

# Electron-beam and emerging lithography for the magnetic recording industry

A. A. G. Driskill-Smith\*

Hitachi Global Storage Technologies, 5600 Cottle Road, San Jose, CA 95193

## ABSTRACT

Today in 2004, the areal density in magnetic recording systems stands at approximately 100 Gbit/in<sup>2</sup>. It is projected to increase by 30–60% per year for the foreseeable future and reach 1 Tbit/in<sup>2</sup> around the end of the decade. The corresponding rapid reduction in bit dimension poses significant challenges for lithography, both now in the thin-film head, and in the future on the disk, with the possible transition to patterned media. In thin-film head production, the critical dimension is now less than 100 nm and by 2006 will be less than 50 nm. To meet these requirements, the magnetic recording industry in recent years has been turning to direct-write electron-beam lithography. However, advances in both electron-beam systems and chemically-amplified resists will be required as areal density approaches the Terabit per square inch regime. No lithography is presently used on the disk or media, but this could change within a few years if discrete track media or patterned media is introduced. The key challenge will be establishing a low cost manufacturing capability at the very high resolution bit dimensions required. One new technology that shows considerable potential is nanoimprint lithography, but significant tool development will be required to improve throughput and provide sufficiently low cost per disk. High-resolution electron-beam systems and advanced resists will also be required to fabricate the 1× master templates. Further in the future, a combination of lithography and self-assembly of magnetic nanoparticles may provide a path to areal densities as high as 40 Tbit/in<sup>2</sup>.

**Keywords:** Magnetic recording, giant magnetoresistive thin-film head, patterned media, electron-beam lithography, nanoimprint lithography

## 1. INTRODUCTION

Throughout the 1990s, propelled by the introduction of the magnetoresistive (MR) recording head in 1991 and the giant magnetoresistive (GMR) recording head in 1997, the areal density of magnetic recording systems grew at a faster rate than ever before, nearly doubling each year<sup>1</sup>. In 1991, areal densities were less than 100 Mbit/in<sup>2</sup>; today, they have reached approximately 100 Gbit/in<sup>2</sup>, an increase of three orders of magnitude in the space of little more than a decade. The corresponding rapid reduction in bit dimension poses significant challenges for lithography, both now in the thin-film head, and in the future on the disk, with the possible transition to patterned media.

In thin-film head fabrication, the critical dimension (CD) has traditionally been significantly larger than that of transistors in integrated circuits, and the magnetic recording industry has therefore been able to rely on optical lithography systems developed for the semiconductor industry. However, this is no longer the case — the CD of thin-film read heads is now less than 100 nm and continues to shrink by 20–30% every year — and in recent years the magnetic recording industry has been turning to electron-beam lithography to fulfil its requirements. Electron-beam lithography is a viable option because the relatively low wafer volume in the industry and the very small area (<0.01%) of critical features on thin-film head wafers combine to make its throughput limitations surmountable in this application. However, further advances in both electron-beam systems and resists will be required as thin-film head production approaches the Terabit per square inch regime.

Presently the bits on the disk are not defined by lithography, but this could change within a few years if patterned media is introduced into products. Multi-grain bits on continuous-film media would be replaced by arrays of lithographically-patterned, single-domain magnetic elements, in which each element stores one bit. This extends the limit placed on bit size by thermal instability effects (the superparamagnetic effect) without increasing the media noise level. The key

---

\* alexander.driskill-smith@hgst.com; tel 408 717 6549; fax 408 717 9130; www.hgst.com

challenge here is low cost fabrication at the very high resolution bit dimensions required, and one new technology that shows considerable potential is nanoimprint lithography.

This paper addresses some of the challenges that must be met if electron-beam lithography and other emerging lithography techniques, such as nanoimprint, are to be successfully applied in the magnetic recording industry in the future.

## 2. AREAL DENSITY AND LITHOGRAPHY ROADMAPS

Areal density growth is the magnetic recording industry's equivalent to Moore's Law in the semiconductor industry. Packing greater numbers of bits into the same area lowers the cost per gigabyte of a hard disk drive. It permits greater drive capacities, yet at the same time allows smaller disks and fewer heads and disks per drive. From the introduction of the first hard disk drive in the 1950s through to 1990, the average compound annual growth rate in areal density remained at approximately 25%. However, it increased dramatically in the 1990s, as a result of new sensor technology. Annual areal density growth increased first to 60% with the introduction of the MR recording head in 1991 and then to 100% with the introduction of the GMR recording head in 1997, as shown in Figure 1(a). A direct result of this growth has been the rapid decrease in cost of storage from over \$1 per megabyte in 1991 to less than \$1 per gigabyte in 2004. If the 100% annual growth in areal density of the late 1990s had continued, areal densities would have reached 1 Tbit/in<sup>2</sup> by 2006. However, the growth rate has slowed considerably in the past 3 years to around 50% and will probably remain in the range 30–60% for the foreseeable future. With areal densities in 2004 standing at approximately 100 Gbit/in<sup>2</sup>, the introduction of Terabit per square inch recording systems should occur between 2009 and 2013.

In order to translate the areal density roadmap into a lithography roadmap, some description of a magnetic recording system is required. In a hard disk drive, the read–write head (Figure 2) is mounted on a slider arm and placed in close contact with the disk or recording medium (Figure 3), which stores data in circular tracks. The data bits are magnetized along the direction of the track in the plane of the disk, a configuration that is known as longitudinal recording. As the disk rotates, the separation between head and disk reaches a stable value known as the fly height. Early inductive heads used a simple two-pole geometry both to read and to write the data. However, since the introduction of the MR sensor in the early 1990s, the read and write functions have been performed by separate elements in the head. The write head, as before, consists of two poles separated by a write gap. The write field in the disk, which must exceed the disk's coercivity in order to store data, depends on the saturation magnetization of the write pole material, the write gap, and the fly height. The read head consists of a sensor, surrounded by two shields, that in today's products utilizes the GMR effect to detect magnetic fields perpendicular to the surface of the disk that arise at the transitions between bits along the track. As areal densities move towards the Terabit per square inch regime, other technologies may be introduced, such as tunneling magnetoresistive (TMR) read sensors, which provide higher sensitivity during read operations, and perpendicular recording, in which the bits are magnetized perpendicular to the disk.

The areal density of data on the disk is equal to the reciprocal of the product of track pitch and linear bit pitch along the track. If patterned media is introduced in the future, the bits will have to be defined lithographically on the disk. However, this is not the case at present: the bits are defined by the read–write head during recording and depend on the geometry of the head and the fly height. The track pitch is related to the read sensor and write pole widths, while the linear bit pitch is related to the write gap (the distance between the write poles) and the read gap (the distance between the shields) in the head. The track pitch is greater than the linear bit pitch by a factor known as the Bit Aspect Ratio (BAR); that is, the width of the track is greater than the bit spacing along the track. One reason for this is because the linear bit spacing is controlled by thin-film depositions, whereas the track pitch is limited by lithography during the thin-film head wafer process.

In today's products, the areal density is approximately 100 Gbit/in<sup>2</sup> and the BAR is between 6 and 8. The track pitch is approximately 200 nm and the linear bit pitch is approximately 33 nm (Table 1). However, in order to achieve any given track pitch on the disk, the write pole width must be approximately 10% smaller than the width of the track, to minimize adjacent-track writing, and the read sensor width must be approximately 50% smaller, to minimize side-reading from adjacent tracks. The most aggressive CD in the head is therefore the read sensor width, which today stands at approximately 100 nm. Two other points are worth noting. First, the BAR has steadily been falling from over 15 a decade ago to between 6 and 8 today and is likely to continue falling over the coming years. This means the largest

contribution to areal density growth since the early 1990s has come from reducing track pitch, which is controlled by lithography in the head, rather than reducing linear bit pitch. And second, the fly height must scale with all other dimensions in the recording system in order to achieve higher areal density, although not necessarily at the same rate. While the read sensor width decreased by a factor of 100 between 1990 and 2004, the fly height fell by a factor of only 10 in the same period, from approximately 100 nm in 1990 to less than 10 nm in 2004.

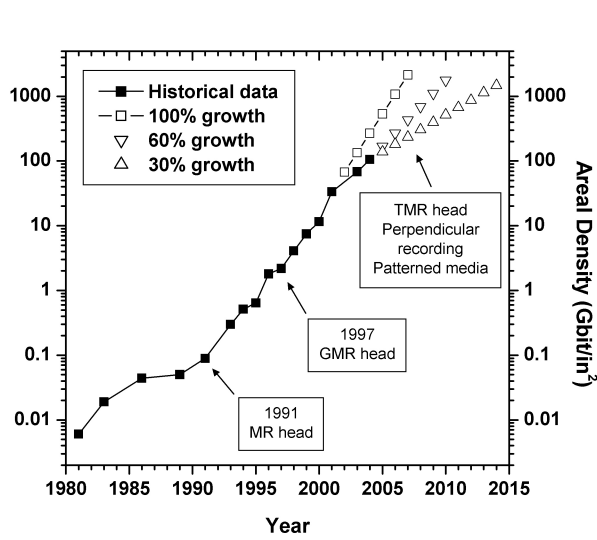


Fig. 1(a). Evolution of areal density in the magnetic recording industry between 1981 and 2004. Also shown are projections for 2005 and beyond at 30% and 60% compound annual growth rates, and a hypothetical projection for 2002 and beyond if the 100% annual growth of the 1990s had continued.

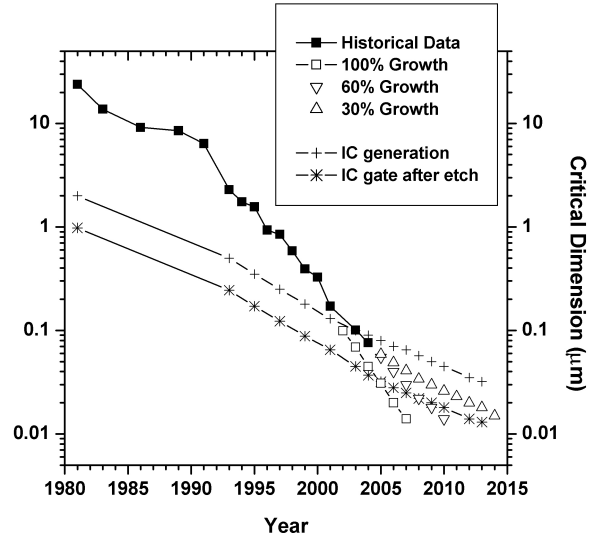


Fig. 1(b). Evolution of read sensor trackwidth in the magnetic recording industry between 1981 and 2004. Also shown for comparison are historical and projected data for the IC industry (IC generation and IC gate width after etch).

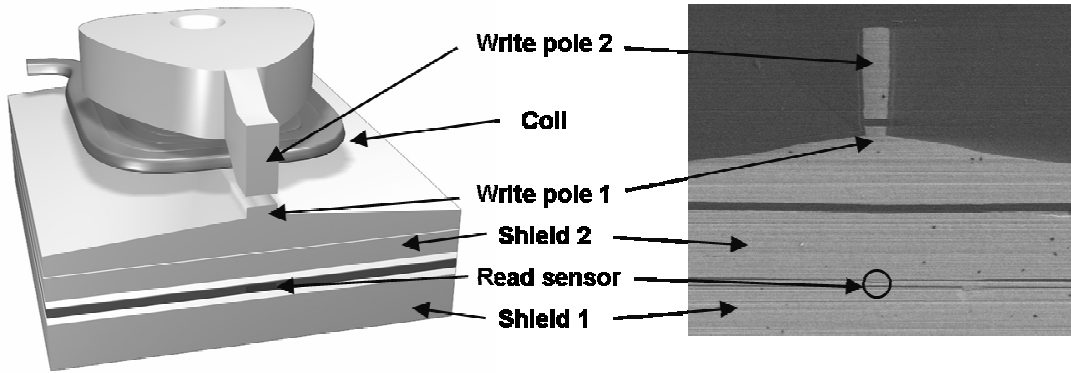


Fig. 2. Cross-sectional views at the air-bearing surface (ABS) of a thin-film magnetic recording head. The critical read sensor and write pole widths are defined lithographically during the wafer process. After completion, the wafers are diced into individual heads. Each head is then rotated 90 degrees and bonded to a slider arm. In its final configuration, the ABS of the head is parallel to and in close contact with the disk surface.

Although the read sensor width in 2004 is approximately the same as the IC generation level, this has historically not been the case (Figure 1(b)). Prior to 1990, the magnetic recording industry's CD lagged that in the semiconductor industry by an order of magnitude, and it was only after the introduction of MR recording heads in the 1990 that the read sensor width began to catch up. If the 100% annual areal density growth of the late 1990s had continued, with the corresponding 30% or greater annual decrease in read sensor trackwidth, the CD requirement in thin-film heads would

have exceeded that of the IC gate width after etch by 2005. Faced with the possibility that optical lithography development for the semiconductor industry roadmap may no longer satisfy the requirements of its own lithography roadmap, the magnetic recording industry began evaluating direct-write electron-beam lithography in the late 1990s as a potential successor to optical lithography for CD definition in thin-film heads. Although areal density growth rates have slowed since 2001, postponing any crossover with the IC roadmap, there are additional cost reasons for making a transition to electron-beam lithography in the low wafer volume environment of a thin-film head wafer fab.




Fig. 3. Schematic view of a longitudinal recording system, showing a thin-film recording head in close contact with one track of a disk.

Areal Density (Gb/in <sup>2</sup> )	"Square Bit" pitch BAR = 1	Bit Aspect Ratio (BAR) = 6	
		Track Pitch (nm)	Linear Bit Pitch (nm)
1	800	2000	330
10	250	600	100
100	80	200	33
1000	25	60	10
10000	8	20	3

Table 1. Bit size in nanometers versus areal density in gigabits per square inch. Bit Aspect Ratio (BAR) refers to the ratio between track pitch and linear bit pitch.

### 3. THIN-FILM HEADS

#### 3.1 Advanced lithography options

Direct-write electron-beam lithography is a viable alternative to optical lithography in thin-film head production because the relatively low wafer volume in the industry and the very small coverage of critical features on thin-film head wafers combine to make its throughput limitations surmountable. Currently, 250 million disk drives are produced per year, containing 700 million heads. With approximately 25,000 heads fabricated on each wafer, this corresponds to just 100 wafer starts per day for the entire magnetic recording industry. More realistically, there are probably between 200 and 400 wafers starts per day, depending on yield, wafer layout, and other factors, but these figures are still a very small fraction of the total wafer starts per day in the semiconductor industry. The wafer layout of Figure 4 shows that the single critical feature at each head corresponds to a pattern density of less than 0.01%. This critical feature can be exposed by electron-beam lithography, leaving other non-critical features to be exposed by mix-and-match optical lithography in the same resist.

In order to estimate the throughput of thin-film heads in an electron-beam system, the first step is to calculate the "beam-on" time, that is, the time when the beam is actually on and exposing resist. Using realistic assumptions for the operating conditions of Variable Shaped Beam (VSB) systems and Gaussian (spot) Beam (GB) systems, the calculated beam-on time in both cases is approximately 5 minutes, which equates to a maximum possible throughput of 12 wafers per hour. For the VSB system, the assumptions are 2 A/cm<sup>2</sup> current density and 10 exposure shots per device, and for the GB system, 1 nA beam current, 10 nm beam-step size, and 25 MHz exposure clock speed. In both cases, the exposure area per device is 30 μm<sup>2</sup> and the exposure dose is 40 μC/cm<sup>2</sup> for chemically-amplified resist. In reality, however, the actual throughput is substantially lower than 12 wafers per hour, since it is dominated by other overheads that include wafer exchange time, beam calibration, alignment, and, most important, stage travel time.

Calculations for the stage travel overhead depend on whether the stage is of the continuously-moving or step-and-repeat variety (Figure 5). For a continuously-moving stage, the stripe width is critical: wider stripes result in shorter total stage travel and thus higher throughput. For a step-and-repeat stage, the deflection field size is critical: larger deflection fields result in fewer stage moves and thus higher throughput. In VSB systems with continuously-moving stages, the exposure takes place in parallel with the stage movement and throughput of 3–4 wafers per hour is achievable (Table 2). At high stage speed, the stage overhead becomes less than the beam-on time and further increases in throughput can be achieved only by increasing the current density and therefore reducing the beam-on time (Figure 6). In contrast, GB systems use

step-and-repeat stages and the exposure and stage moves take place serially. The achievable throughput of these systems is typically lower than that of VSB systems by a factor of two and there is a strong dependence on both stage step-and-settle time and deflection field size (Table 3 and Figure 7). One way to close the throughput gap between GB and VSB systems would be to replace the step-and-repeat stages in GB systems with continuously-moving stages.




Fig. 4. Typical layout for a thin-film head wafer process. The exposure area per device for each critical level is only a few tens of square microns, corresponding to a pattern density of less than 0.01%.

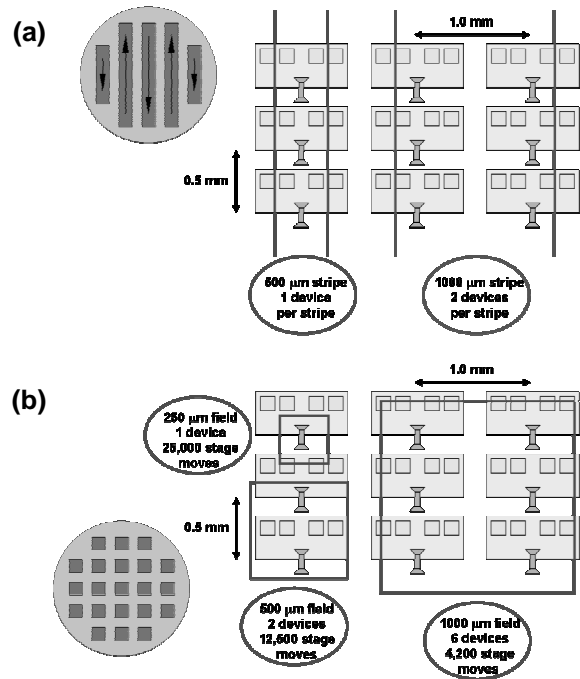


Fig. 5. Wafer exposure strategy for (a) a Variable Shaped Beam system with continuously-moving stage, and (b) a Gaussian Beam system with step-and-repeat stage.

The advanced lithography options available to the magnetic recording industry are listed in Table 4. For isolated features, the minimum achievable resolution with 193 nm optical lithography and phase-shift masks is approximately 50 nm. VSB systems are limited by beam-edge slope and have comparable resolution capability for isolated features. In contrast, GB systems with <10 nm beam diameter provide superior resolution and therefore multiple-generation capability. Optical lithography is the clear leader in terms of throughput and provides at least an order of magnitude lower cost per wafer pass, provided the tools are operated at full capacity. However, in the low wafer volume environment of a thin-film head wafer fab (25–100 wafers per day), cost per wafer pass is driven by utilization of the tool's available throughput and optical lithography loses its advantage. For example, in a small MR wafer fab that starts 25 wafers per day and exposes two critical levels (the read and write trackwidths), only 50 wafers would pass through the 193 nm optical lithography tool per day. Operating at a throughput of 100 wafers per hour, the exposures would be completed in half-an-hour, leaving the tool idle for the remainder of the day. In this situation, the cost per wafer pass is twice that of either electron-beam lithography option.

As thin-film head fabrication becomes more complex, the number of critical levels per wafer will increase. In addition to the critical read and write trackwidths, other levels will likely require exposure on an advanced lithography tool, due to resolution, overlay or aspect ratio requirements. Figure 8 shows the effect of an increase in the number of critical levels per wafer on the cost per wafer pass and the number of tools required. For the 193 nm optical lithography option, only one tool is required and the cost per wafer pass falls as the number of critical levels (and tool utilization) is increased. For the two electron-beam lithography options, the cost per wafer pass reaches a plateau when the tools attain full utilization. Additional tools are required when the number of critical levels increases, but the cost per wafer pass remains the same, provided all tools are fully utilized. In general, the cost per wafer pass through a Variable Shaped Beam system is less than that through a Gaussian Beam system due to its higher throughput.

Stage speed (mm/s)	Devices per stripe	Stage overhead (mins)	Current density (A/cm <sup>2</sup> )	Beam-On Time (mins)	Throughput (wafers per hour)
10	1	21	2	5	1.9
	2	10			2.9
30	1	7			3.5
	2	4			4.0
100	1	2			4.0
	2	1			4.0

Table 2. Throughput estimates for a Variable Shaped Beam system with continuously-moving stage. Beam-on time is fixed at 5 minutes and other overheads not shown in the table (wafer load, unload, calibration, alignment) are assumed to total 10 minutes.

Stage step & settle time (ms)	Devices per e-beam deflection field	Total stage overhead per wafer (minutes)	Maximum throughput (wafers per hour)
500	2	104	0.5
	6	35	1.2
200	2	42	1.1
	6	14	2.1
100	2	21	1.7
	6	7	2.7

Table 3. Throughput estimates for a Gaussian Beam system with step-and-repeat stage. Beam-on time is fixed at 5 minutes and other overheads not shown in the table (wafer load, unload, calibration, alignment) are assumed to total 10 minutes.

		193 nm optical	Variable Shaped EB	Gaussian (Spot) EB
Minimum resolution for isolated features (nm)		50	50	10
Throughput (wafers per hour)		100	3–4	1–2
Cost per wafer pass (\$)	At full tool capacity	10	100	200
	Large MR wafer fab (100 wafers per day, 2 critical levels)	100 (1 tool)	100 (2 tools)	200 (4 tools)
	Small MR wafer fab (25 wafers per day, 2 critical levels)	400 (1 tool)	200 (1 tool)	200 (1 tool)

Table 4. Lithography tool options for thin-film head manufacturing. The figures for cost per wafer pass assume that the cost of an electron-beam lithography tool (Variable Shaped or Gaussian) is approximately one half the cost of a 193 nm optical lithography tool.

Comparing the two electron-beam lithography options, the key differentiating factor for Variable Shaped Beam systems is their higher throughput. They are the most attractive option for smaller thin-film head wafer fabs, due to their lower cost per wafer pass, especially when compared with 193 nm optical lithography. Installing multiple tools per fab will be a viable option as the number of critical levels having dimensions in the range 50–120 nm increases, but resolution improvements will be required if they can continue to be used for the most critical read sensor trackwidth when it falls below 50 nm. The key differentiating factor for Gaussian Beam systems is their minimum resolution and hence multiple-generation capability. The adoption of Gaussian Beam systems by all thin-film head wafer fabs is probably inevitable when the read sensor width falls below 50 nm, but their use may be limited to one or two critical levels per wafer due to

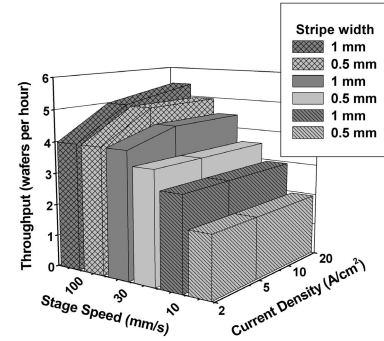


Fig. 6. Throughput estimates for a Variable Shaped Beam system with continuously-moving stage. Other overheads are assumed to total 10 minutes.

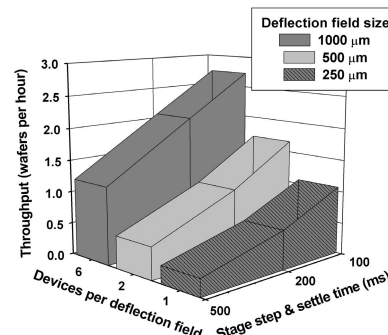


Fig. 7. Throughput estimates for a Gaussian Beam system with step-and-repeat stage. Beam-on time is fixed at 5 minutes and other overheads are assumed to total 10 minutes.

the higher cost per wafer pass. Throughput improvements will be required in order to cut the cost per wafer pass, increase the number of tools per fab, and therefore expand their use beyond a small number of critical levels per wafer. Finally, it is important that continual improvements in CD uniformity, beam placement and overlay are made to both types of electron-beam system. Their performance in these areas must match or exceed the semiconductor industry's roadmap for optical lithography tools, otherwise the main argument in their favor, their lower cost per wafer, will have little bearing.

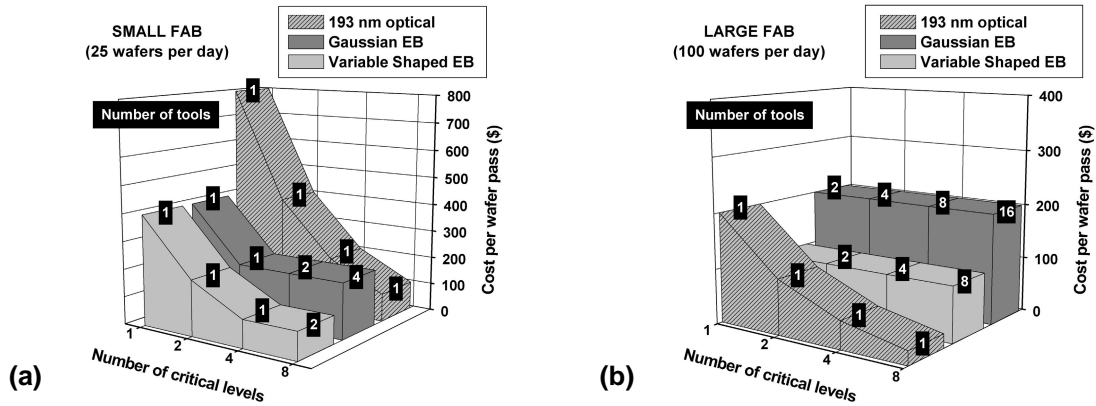


Fig. 8. Cost per wafer pass for the three lithography tool options as the number of critical levels per wafer is increased from 1 to 8 in (a) a small and (b) a large thin-film head wafer fab. Also shown is the number of tools required for each option.

### 3.2 Thin-film head processing

Thin-film heads have a three-dimensional structure, as shown in Figure 2, and the wafer process involves many layers. Tight alignment is required between layers in the read head and the write head, yet their vertical separation can be as large as 10  $\mu\text{m}$ . In terms of resolution, the two most critical lithography levels are the write pole and the read sensor. Although the write pole fabrication is complicated by its high aspect ratio, it is the read sensor that has the most aggressive CD requirement. Today, it is less than 100 nm, and by 2006 it will be substantially less than 50 nm, at which point the manufacturing process may become limited by the availability of suitable resist materials.




Fig. 9. GMR read sensor processing. (a) Fabrication sequence (A expose negative-tone electron-beam resist followed by RIE pattern transfer, B ion mill GMR sensor materials followed by deposition of seed, hard bias, and lead layers, C resist liftoff). Transmission Electron Micrographs of (b) a 32 nm thin-film read sensor fabricated with a standard, chemically-amplified, negative electron-beam resist and (c) a 13 nm thin-film read sensor fabricated with HSQ, an inorganic, high-resolution, negative electron-beam resist.

The read sensor process utilizes a negative-tone electron-beam resist, RIE pattern transfer, then ion mill, deposition and liftoff steps, as shown in Figure 9(a). Using standard, chemically-amplified, negative-tone electron-beam resists and

Gaussian electron-beam lithography, it is possible to fabricate read sensors having trackwidths down to approximately 30 nm (Figure 9(b)). However, the process below 50 nm is challenging, because it must operate at or close to the resolution limit of the resist, which is typically between 30 and 50 nm for current-generation, chemically-amplified materials. In contrast, by using Hydrogen Silsesquioxane (HSQ), an inorganic, high-resolution, negative-tone electron-beam resist, it is possible to fabricate read sensor dimensions down to approximately 10 nm (Figure 9(c))<sup>2</sup>. In this case, the process is limited not by the resist material, but by the resolution limit of the Gaussian Beam exposure tool, which has a beam diameter of less than 10 nm. HSQ also has the advantage of a high silicon content, which makes it very resistant to etch during the RIE pattern transfer step; however, it is not suitable for manufacturing due to shelf-life issues and its insensitivity to optical exposure.

Availability of suitable resist materials will therefore be a key factor in successful read sensor fabrication when the trackwidth falls below 50 nm. New, negative-tone, chemically-amplified resists will be required that are sensitive to both electron-beam exposure for critical features and optical exposure for non-critical features. Ideally, these resists would have high silicon content for pattern transfer and would have a minimum resolution of less than 30 nm, as shown in Table 5. In the longer term, a minimum resolution of 10 nm will be required, but this may not be achievable with chemically-amplified systems.

Resist property	Standard resist	HSQ resist	Ideal resist
Negative tone	Yes	Yes	Yes
Si containing	No	Yes	Yes
Chemically amplified	Yes	No	Yes
Mix and match	Yes	No	Yes
Minimum resolution	30–50	< 10 nm	< 30 nm
Shelf life	Good	Poor	Good

Table 5. Comparison of resist properties for thin-film read sensor processing.

## 4. MEDIA

At present, no lithography is used on the disks in a magnetic recording system. The data bits are defined on continuous-film media by the read-write head during formatting and recording, and each bit consists of between 100 and 1,000 individual magnetic grains. Today in 2004, at an areal density of approximately 100 Gbit/in<sup>2</sup>, the bit size is approximately 200 × 30 nm<sup>2</sup>. However, as areal densities continue to increase, the associated reduction in bit size will be affected by two issues. The first is statistical noise. If the bit size is reduced without a commensurate reduction in grain size, the number of grains per bit decreases. This in turn adversely affects the signal-to-noise ratio, which is proportional to the square root of the number of grains in the bit. The second issue is thermal instability. If the grain size is reduced below approximately 8 nm, the stored magnetic energy in the grain, which is proportional to its volume, becomes comparable to its thermal energy and makes the grain thermally unstable. This is known as the superparamagnetic effect.

It is becoming increasingly difficult to reconcile these two constraints using conventional media<sup>3</sup>. Strategies for delaying the onset of superparamagnetic effects that have been implemented in recent years include increasing the media coercivity, reducing the bit aspect ratio, and introducing antiferromagnetically-coupled media. Other options for the future include perpendicular recording, in which the grains are magnetized perpendicular to the disk, and discrete track recording, where the tracks are patterned lithographically. However, the ultimate solution will be the introduction of fully patterned media, in which the individual bits are lithographically-patterned and each consists of a single-domain magnetic element. The key challenge that must be addressed here is establishing a low cost manufacturing capability at the very high resolution bit dimensions required.

### 4.1 Patterned servo: “Magnetic lithography”

Before turning to patterned media, it is worth mentioning the servo write issue in a hard disk drive. Currently, the magnetic servo pattern has to be recorded on all disks prior to drive shipment in a serial process that is both slow and expensive. If the throughput of this process could be improved, there is potential for significant cost savings. In addition



to servo write, two other applications would benefit from a fast process for copying magnetic data: drive formatting and distribution of digital content, such as software and music. What is needed is a magnetic equivalent to the stamping process that is used in high volume distribution of optical media like CDs and DVDs. Track-by-track recording in a hard disk drive is very slow in comparison and scales unfavorably with increases in drive capacity.

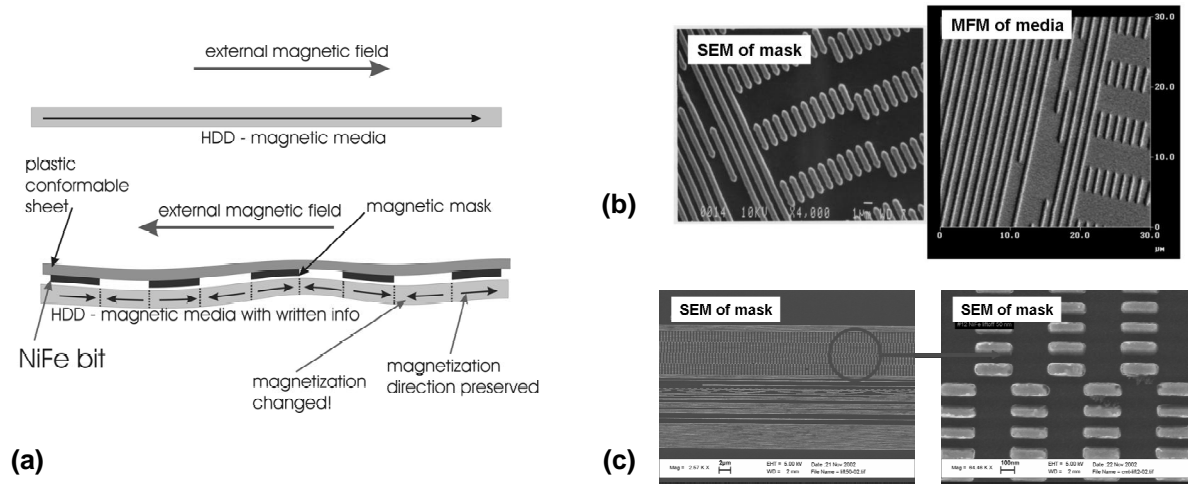


Fig. 10. (a) Principle of “magnetic lithography”. (b) Scanning Electron Micrograph of a CoFe mask having a bit length of 640 nm and Magnetic Force Micrograph of the same pattern after transfer to disk. (c) Scanning Electron Micrographs of a NiFe mask having a bit length of 78 nm.

One potential solution can be described as “magnetic lithography” and the principles behind this process are shown in Figure 10(a)<sup>4</sup>. First, the disk is magnetized uniformly by an external magnetic field. Then, a “magnetic mask”, patterned with a soft magnetic material such as NiFe, is placed in close contact with the disk and the direction of the external magnetic field is reversed. The magnetic mask selectively shields the external field, thus transferring the mask pattern into a magnetic pattern on the disk. The process has been successfully demonstrated at relatively large bit lengths (640 nm), as shown in Figure 10(b). However, when the bit length is reduced below 100 nm (Figure 10(c)), it becomes difficult to reverse the magnetization of the NiFe mask elements: the small particle size prevents domain wall formation and increases the field required to set the magnetization. Since the linear bit density in current products is already close to 30 nm, it is therefore not clear that magnetic lithography can be successfully implemented in practice, at least for leading-edge products.

#### 4.2 Patterned tracks and patterned media

Returning to conventional lithography, the two main applications for disk patterning are shown in Figure 11: discrete track media and patterned media. Both also offer the option of patterning servo features at the same time. In discrete track media, only the tracks on the disk are patterned lithographically. As with conventional media, each bit consists of many grains and the bit transitions along the track continue to be defined by the read–write head during recording. The main advantage of discrete track media is that it provides higher signal-to-noise ratio than conventional media for the same track pitch. Noise from the track edges is reduced, there is no exchange between adjacent tracks, and the read sensor width can be increased, providing higher signal output<sup>5</sup>. In fully patterned media, the situation is different. Both bit dimensions, the track pitch and the linear bit pitch, are patterned lithographically. Each patterned element is a single magnetic domain and stores one bit. The main advantage of patterned media is the improved thermal stability of the bits. The superparamagnetic effect applies to the whole bit, since each one consists of a single magnetic domain, not up to one thousand individual grains as in conventional media. For the same areal density, the volume, and hence stored magnetic energy, of a patterned bit is much larger than that of the individual grains in continuous media. This means that a patterned media system can be scaled to much higher areal densities than a conventional media system. Other advantages of patterned media include reduced track edge noise, reduced transition noise, and improved signal-to-noise ratio, because there is no statistical noise due to averaging of grains. Furthermore, the media coercivity can be reduced, enabling lower write fields. Patterned media also provides a path to very high areal densities of greater than 1 Tbit/in<sup>2</sup>,

since the stored energy in a magnetic element is proportional not just to its volume, but also to its anisotropy. By choosing appropriate materials and controlling the properties of the magnetic elements, the bit size can potentially be reduced below the ~8 nm size limit that exists for grains in today's conventional media.



Fig. 11. Schematics of (a) Discrete Track media and (b) Patterned Media. In Discrete Track media, only the tracks on the disk are patterned lithographically, while the bit transitions along the track continue to be defined by the thin-film head during recording. In fully patterned media, both bit dimensions on the disk are defined lithographically.

Year (60% growth)	Areal Density (Gb/in <sup>2</sup> )	Discrete Track Recording Bit Aspect Ratio (BAR) = 4		Fully Patterned Media BAR = 1	
		Track Pitch (nm)	Track half-pitch (nm)	“Square Bit” pitch (nm)	Half-pitch (nm)
2004	100	160	80	80	40
2006	300	92	46	46	23
2009	1,000	50	25	25	12.5
2011	3,000	30	15	15	7.5
2014	10,000	16	8	8	4

Table 6. Lithography requirements for Discrete Track Recording and Patterned Media. At any given areal density, the lithography half-pitch for fully patterned media is smaller than that for discrete track media by a factor of two, assuming a Bit Aspect Ratio of 4 for discrete track media.

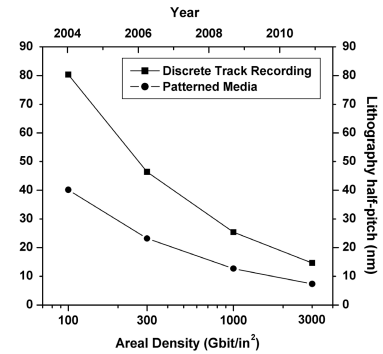


Fig. 12. Lithography requirements for Discrete Track Recording and Patterned Media.

The challenges involved in patterning tracks or bits on a manufacturing scale are significant. Assuming that neither option can be introduced for at least two years, the areal density requirement will be at least 300 Gbit/in<sup>2</sup> (Table 6 and Figure 12). For Discrete Track Recording (DTR), assuming a Bit Aspect Ratio (BAR) of 4, this corresponds to a track pitch of 92 nm or a lithography half-pitch of 46 nm. For fully patterned media, in which both bit dimensions are defined lithographically and the BAR is 1, the lithography half-pitch is 23 nm, twice as aggressive as the DTR case for the same areal density. And by the time areal densities reach 1 Tbit/in<sup>2</sup>, probably at the end of this decade, the lithography half-pitch requirement will be 25 nm for DTR and 12.5 nm for patterned media. In addition to the high resolution, the patterned features must have circular symmetry, and the manufacturing method must have high throughput and low cost. Several techniques, such as focused-ion-beam<sup>6</sup> and electron-beam patterning<sup>7</sup>, have been used to fabricate patterned media test structures over small areas, but these methods are not suitable for manufacturing due to their very low throughput. Other methods, such as X-ray<sup>8</sup>, interferometric lithography<sup>9</sup>, and ion-beam projection<sup>10</sup>, solve the throughput problem, but are generally deficient in terms of cost, resolution, or the ability to fabricate patterns with circular symmetry. However, one very promising technique that has emerged within the last decade is nanoimprint lithography. It has the potential to pattern very high resolution features at high throughput and low cost.

### 4.3 Nanoimprint lithography for patterned media

Remarkable advances have been made in nanoimprint lithography (NIL) in the past decade. In the original, thermal imprint process, a nickel or silicon master template is pressed at relatively high temperature and pressure into a resist layer covering the substrate surface<sup>11</sup>. The resist conforms to the features on the template and retains this shape after the template is released. The pattern is then transferred to the substrate, typically by reactive-ion etch (RIE). The template can be re-used many times and printed features of less 10 nm have been demonstrated with this method: it appears that the minimum resolution is limited only by the resolution of the template. More recently, a photo-polymerization nanoimprint method has been developed<sup>12</sup>. It requires a transparent template material, such as quartz, and uses a low-

viscosity, photo-polymerizable resist layer. In contrast to the thermal imprint process, which requires high temperature and pressure, this method uses UV light irradiation through the template to cure the resist during the imprint step. Its throughput is therefore greater and, since the template is transparent, it is also easier to align the template to the substrate.




Fig. 13. Nanoimprint lithography (NIL) for patterned media. (a) Electron-beam lithography and (b) RIE pattern transfer to form the 1× master template (c) NIL, either by thermal imprint or by photo-polymerization, to form the replica pattern in resist on the disk, (d) RIE to transfer the resist pattern into the disk substrate, and (e) magnetic multilayer deposition to create the decoupled, magnetic bits.

A NIL process for fabricating patterned media is shown in Figure 13. The template is fabricated by electron-beam lithography and RIE pattern transfer and then stamped onto the disk by either of the nanoimprint methods. The magnetic domains can be formed by sputtering Co/Pd multilayers with perpendicular anisotropy over the entire disk, resulting in isolated, single-domain pillars<sup>13</sup>. The total thickness of the magnetic multilayers must be less than the height of the pillars so that the bits remain decoupled from each other. This is an important consideration: as the lateral bit dimensions decrease, the height of the pillars must decrease for reasons of mechanical strength, which in turn imposes a maximum thickness on the magnetic multilayers and reduces the signal. Other considerations include minimizing the variation across the disk in the field required to reverse the magnetization direction of the bits, which is known as the switching-field distribution, and optimizing the coercivity of the multilayers. As the size of the bits decreases, the coercivity of the multilayer stack increases, but this can be controlled by adjusting the deposition conditions and the composition of the multilayers; however, the switching field distribution also increases for reasons that are not yet fully understood.

NIL throughput (disks per hour)	Number of NIL tools required	Cost per disk (\$) (NIL contribution only)
60	2300	0.83
720	200	0.07

Table 7. Cost per disk (NIL tool contribution only) as a function of NIL tool throughput. The calculation assumes 700 million patterned disk surfaces per year, \$1 million tool cost, 4 year tool depreciation, 80% tool uptime, 70% yield.

Resolution, uniformity and throughput will all be important elements of a nanoimprint process for patterned media manufacturing, but overlay is not a critical requirement in this application. At its earliest possible introduction, discrete track recording will require a lithography half-pitch of 46 nm and fully patterned media will require 25 nm. More likely, the process would be introduced later, by which time the lithography requirement will be even more aggressive. The uniformity requirement for the patterned bits will likely be 10% of CD and it will be important to minimize edge roughness, for this minimizes the switching field distribution. Local placement errors must also be minimized in order to reduce position jitter noise as the head travels along the track. Furthermore, the cost per disk must be substantially less than \$1, which means the throughput will need to be in the neighborhood of 5 s per stamp, or 720 stamps per hour (Table 7). This level of performance has not been demonstrated by any available nanoimprint tool, but given the rapid pace of advances in NIL tooling over recent years, there are grounds for optimism that the technical hurdles can be overcome.

#### 4.4 NIL template fabrication by electron-beam lithography

The main disadvantage of NIL is the need for a  $1\times$  master template. Indeed, fabricating the templates could be the most challenging part of implementing a NIL process for patterned media. A combination of high-performance electron-beam systems and high resolution, high sensitivity resists will be necessary to define the template pattern at the required resolution and with acceptable throughput. The ideal electron-beam system would probably combine a Gaussian or spot beam, in order to define the sub-50 nm feature sizes, with a rotating stage on a linear slider, in order to achieve the circular symmetry of the pattern. In this configuration, the beam remains on-axis throughout the exposure, thus minimizing beam diameter variation and stitching errors that arise from beam deflection. One track is defined per complete stage rotation and successive tracks are defined by moving the stage along a linear slider. Rotating stages are already used in combination with electron-beam systems to make master templates for optical disks, but the feature sizes in that application are at least one order of magnitude larger and Gaussian Beam systems have not been used. For high-resolution, patterned media applications, several issues associated with a rotating stage will have to be addressed, such as how to make accurate stage position measurements, achieve uniform track spacing, and make measurements of local height variations, all while the stage rotates.

The electron-beam exposure time for a patterned media NIL template depends heavily on both the sensitivity of the resist used and the beam current (Table 8 and Figure 14). Assuming a pattern density of 50% and a disk diameter of 2.5 inches, the exposure time can be reduced below 24 hours only by using a high sensitivity chemically-amplified resist (CAR) or a high beam current, or both. The issue is whether future developments in CAR performance will offer sufficient resolution for this application. As described in Sections 4.2 and 4.3, discrete track recording will require a lithography half-pitch of substantially less than 50 nm at its earliest possible introduction and fully patterned media will require less than 25 nm. At 1 Tbit/in<sup>2</sup>, the requirements for DTR and patterned media will be 25 nm and 12.5 nm, respectively. If CAR proves not to be an option, a high-resolution, non-chemically-amplified resist, such as ZEP, would have to be used, in which case the exposure dose and exposure time could increase by a factor of up to 10. Note that the actual exposure dose for any particular resist depends on the accelerating voltage of the electron-beam system: higher beam voltages require higher exposure doses. There will also be a tradeoff involved between beam current and beam diameter. Between 10 and 100 nA, the beam diameter of Gaussian beam systems typically increases rapidly, which may preclude using high beam currents when the bit size becomes very small.

A further requirement will be placed on the exposure clock frequency, or beam-blanking frequency. For discrete track media, this may be a minor issue, because the tracks are continuous and the beam would remain on, or un-blanked, for much of the exposure, except when exposing servo features. However, this is not the case for fully patterned media: the beam must be un-blanked and blanked to expose each bit around the tracks. As shown in Table 9 and Figure 15, the exposure frequency required for a 10 hour template exposure at 300 Gbit/in<sup>2</sup> is 60 MHz. The stage rotation speed during this exposure would be 800 rpm. As the areal density increases, both the linear bit density along the track and the track density increase, which means that a higher beam-blanking frequency is required to keep the exposure time constant. At 1 Tbit/in<sup>2</sup>, the exposure frequency for a 10 hour exposure increases to 200 MHz, and the associated stage rotation speed is 1500 rpm.

The lifetime of the NIL templates will be critical, since the contribution to the cost per disk from fabricating templates is strongly related to the template usage, that is, the number of stamps that can be made from each master template before it degrades. For patterned media manufacturing to be viable, the cost per disk must be substantially less than \$1, which means the usage of each template must be on the order of 100,000, as shown in Table 10. If such template usage proves to be unrealistic, another option may be to fabricate "daughter" templates from the original master template by NIL and pattern transfer. If this could be done with minimal degradation of the template uniformity, it would reduce the cost per disk and relax the throughput requirement of the electron-beam master template exposure.

No electron-beam system is currently available that combines a Gaussian beam with a rotating stage and no chemically-amplified, electron-beam resist is currently available that would meet the requirements of fully patterned media at 300 Gbit/in<sup>2</sup> areal density, let alone 1 Tbit/in<sup>2</sup>. Significant cost and effort would be required to develop such systems and materials for patterned media applications. It is therefore likely that the fabrication of suitable NIL templates will be the key challenge in patterned media manufacturing.

Resist material	Exposure dose ( $\mu\text{C}/\text{cm}^2$ )	Exposure time (hrs) for 50% coverage		
		1 nA	10 nA	100 nA
PMMA	1000	4200	420	42
ZEP	200	830	83	8
CAR	20	83	8	0.8

Table 8. Electron-beam exposure time for a NIL template as a function of resist sensitivity and beam current. The calculation is for an electron-beam system having a rotating stage and linear slider, and assumes 50% coverage of a 2.5 inch diameter disk.

Areal Density ( $\text{Gb}/\text{in}^2$ )	Exposure frequency (MHz)		
	24 hr exposure	10 hr exposure	1 hr exposure
100	8	20	200
300	24	60	600
1000	80	200	2000

Table 9. Electron-beam exposure (beam-blanking) frequency for a NIL template as a function of areal density and exposure time. The calculation is for an electron-beam system having a rotating stage and linear slider.

NIL template usage	Number of EB tools required	Cost per template (\$)	Cost per disk (\$) (EB contribution only)
1,000	2000	2500	3.54
10,000	200	2500	0.35
100,000	20	2500	0.04

Table 10. Cost per disk (electron-beam tool contribution only) as a function of NIL template usage. The calculation assumes 700 million patterned disk surfaces per year, \$5 million tool cost, 10 hour exposure time per template, 4 year tool depreciation, 80% tool uptime, 70% yield.

#### 4.5 Beyond lithography: Self-assembly

Further into the future, self-assembly of arrays of magnetic particles may provide the path to areal densities as high as 40 Tbit/in<sup>2</sup>. One promising candidate material is FePt, since it is more chemically stable than other hard magnetic materials and forms nanoparticles whose diameter can be controlled in the range 4–10 nm (Figure 16)<sup>14</sup>. The size distribution of the nanoparticles is very narrow (<5%), especially when compared with the grain size distribution in continuous-film media (>25%), and it forms highly regular arrays of particles when self-assembled on a substrate. Provided the FePt nanoparticles are in the ordered  $L1_0$  phase, they have high anisotropy and large coercivity at room temperature. However, several issues remain to be solved before they could be used as a data storage medium. First, the as-deposited nanoparticles are in a disordered fcc phase and are superparamagnetic at room temperature. The ordered phase is formed only after high temperature annealing, which results in sintering. Second, there is currently no method for aligning the easy axis of the particles in the array. And third, the self-assembled nanoparticles form ordered arrays only on a length scale of microns: there is no long-range order over the surface of an entire disk.

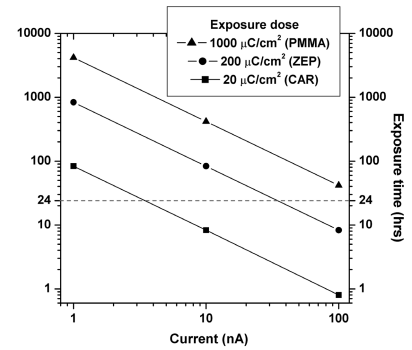


Fig. 14. Electron-beam exposure time for a NIL template as a function of resist sensitivity and beam current.

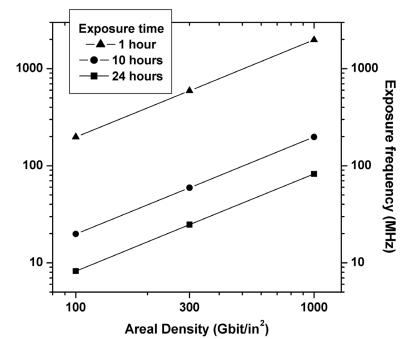


Fig. 15. Electron-beam exposure (beam-blanking) frequency for a NIL template as a function of areal density and exposure time.

One solution to the long-range order issue could be to combine lithography with self-assembly, for example by printing circular tracks on the disk by NIL, then forming several rows of nanoparticles within each track by self-assembly<sup>15</sup>. Either the nanoparticles could be magnetic, as in the case of FePt above, or they could act as a mask for subsequent etch of an underlying magnetic layer. This latter approach has been demonstrated using a PMMA-polystyrene diblock copolymer solution, which separates during anneal into PMMA nanoparticles surrounded by a matrix of polystyrene<sup>16</sup>. The PMMA nanoparticles are selectively removed by oxygen plasma treatment and the voids are then filled with spin-on-glass, which serves as the etch mask. The size of the nanoparticles is related to the molecular weight of the polymer and diameters down to 10 nm have been shown. By adjusting the width of the lithographically-defined tracks, the number of nanoparticle rows per track can be controlled. However, it is clear that excellent control of the track-edge roughness will be required at the NIL step, otherwise defects in the nanoparticle ordering will occur during the subsequent self-assembly process.




Fig. 16. Scanning Electron Micrograph of an array of 4 nm diameter FePt magnetic nanoparticles created by self-assembly, corresponding to an areal density of 40 Tbit/in<sup>2</sup>.

## 5. CONCLUSIONS

In thin-film heads, the most aggressive CD is the read sensor width. It is now less than 100 nm and continues to shrink by 20–30% every year. By 2006, it will be less than 50 nm. In order to meet these requirements, the magnetic recording industry in recent years has been turning to direct-write electron-beam lithography. The rising cost of advanced optical lithography systems, the relatively low wafer volume in the industry, and the very small coverage of critical features on thin-film head wafers mean that electron-beam lithography can provide a cost-effective solution, despite its lower throughput. Variable Shaped electron-beam systems provide higher throughput, and therefore lower cost per wafer pass, than Gaussian electron-beam systems. However, Gaussian beam systems provide better minimum resolution and therefore multiple-generation capability. Looking to the future, Variable Shaped beam systems will require improved resolution if they can continue to be used when the read sensor width falls below 50 nm and Gaussian beam systems will require improved throughput if their use can be expanded beyond a small number of critical levels per wafer. In addition, both systems types will require continual improvements in CD uniformity, beam placement and overlay. Availability of suitable resist materials will also be a key factor in successful read sensor fabrication when the trackwidth falls below 50 nm. New, negative-tone, chemically-amplified resists will be required that are sensitive to both electron-beam exposure for critical features and optical exposure for non-critical features. Ideally, these resists would have high silicon content for pattern transfer and would have a minimum resolution of less than 30 nm.

Presently the bits on the disk are not defined by lithography, but this could change within a few years if discrete track media or patterned media is introduced into products. Discrete track media, in which only the tracks on the disk are patterned lithographically, provides higher signal-to-noise ratio than conventional media for the same track pitch. However, fully patterned media, in which both bit dimensions, the track pitch and the linear bit pitch, are patterned lithographically, provides the ultimate solution to the thermal stability problem, potentially extending areal densities well into the Terabit per square inch regime. The key challenge that must be addressed is establishing a low cost manufacturing capability at the very high resolution bit dimensions required. At its earliest possible introduction, discrete track media would require a lithography half-pitch of 46 nm, and fully patterned media would require 23 nm, twice as aggressive as the discrete track case for the same areal density. One new technology that shows considerable potential for patterned media manufacturing is nanoimprint lithography. However, significant tool development will be

required to improve throughput and provide sufficiently low cost per disk. Resolution and uniformity will also be important elements of the nanoimprint process, but overlay is not a critical requirement in this application. The main disadvantage of nanoimprint lithography is the need for a  $1\times$  master template, and fabricating the templates could be the most challenging part of implementing an imprint process for patterned media. A combination of high-performance electron-beam systems and high resolution, high sensitivity resists will be necessary to define the template pattern at the required resolution and with acceptable throughput. The ideal electron-beam system would probably combine a Gaussian or spot beam, in order to define the sub-50 nm feature sizes, with a rotating stage on a linear slider, in order to achieve the circular symmetry of the pattern. No chemically-amplified, electron-beam resist is currently available that would meet the requirements of fully patterned media at 300 Gbit/in<sup>2</sup> areal density, let alone 1 Tbit/in<sup>2</sup>.

The final question is whether areal densities in the Terabit per square inch regime can be reached by lithography alone, or whether an entirely new technique such as self-assembly will be required. At 1 Tbit/in<sup>2</sup>, the lithography half-pitch for fully patterned media will be just 12.5 nm, close to the resolution limit of electron-beam lithography. Self-assembly of magnetic nanoparticles has been demonstrated, but one important issue in a magnetic recording system will be achieving long-range order over an entire disk. Perhaps the solution will be to combine lithography with self-assembly, for example by defining circular tracks on the disk by lithography, then forming rows of nanoparticles within each track by self-assembly. Using 4 nm diameter FePt nanoparticles could extend areal densities to as high as 40 Tbit/in<sup>2</sup>.

### ACKNOWLEDGEMENTS

The author gratefully acknowledges Robert Fontana and Bruce Terris for helpful discussions during the preparation of this paper. A large team has contributed to the work presented, including I. McFadyen, R. Tiberio, J. Katine, D. Druist, K. Lee, S. MacDonald, M.-C. Cyrille, E. Dobisz, F. Dill, J. Li, M. Morijiri, M. Arasawa, T. Okada, H. Kimura, A. Chiu, N. Robertson, I. McFadyen, P. Kasiraj, T. Thomson, M. Albrecht, L. Folks, C. Rettner, Z. Bandic, M. Rooks, S. Raoux, M. Best, G. Hu, G. McClelland, T. Albrecht, T. Nishida, H. Ikekami, T. Suzuki, T. Ando, M. Ogino, and A. Miyauchi.

### REFERENCES

1. R. E. Fontana, J. Katine, M. Rooks, R. Viswanathan, J. Lille, S. MacDonald, E. Kratschmer, C. Tsang, S. Nguyen, N. Robertson, and P. Kasiraj, *IEEE Trans. Magn.* **38**, 95 (2002)
2. A. A. G. Driskill-Smith, J. A. Katine, D. P. Druist, K. Y. Lee, R. C. Tiberio, and A. Chiu, accepted for publication in *Microelectron. Eng.* (2004)
3. D. Weller and A. Moser, *IEEE Trans. Magn.* **35**, 4423 (1999)
4. Z. Z. Bandic, H. Xu, Y. M. Hsu, and T. R. Albrecht, *IEEE Trans. Magn.* **39**, 2231 (2003)
5. Y. Soeno, M. Moriya, K. Ito, K. Hattori, A. Kaizu, T. Aoyama, M. Matsuzaki, and H. Sakai, *IEEE Trans. Magn.* **39**, 1967 (2003)
6. C. T. Rettner, M. E. Best, and B. D. Terris, *IEEE Trans. Magn.* **37**, 1649 (2001)
7. C. Haginoya, S. Heike, M. Ishibashi, K. Nakamura, K. Koike, T. Yoshimura, J. Yamamoto, and Y. Hirayama, *J. Appl. Phys. Lett.* **85**, 8327 (1999)
8. F. Rousseaux, D. Decanini, F. Carcenac, E. Cambril, M. F. Ravet, C. Chappert, N. Bardou, B. Bartenlian, and P. Veillet, *J. Vac. Sci. Technol. B* **13**, 2787 (1995)
9. C. A. Ross, H. I. Smith, T. Savas, M. Schattenburg, M. Farhoud, M. Hwang, M. Walsh, M. C. Abraham, and R. J. Ram, *J. Vac. Sci. Technol. B* **17**, 3168 (1999)
10. A. Dietzel, R. Berger, H. Loeschner, E. Platzgummer, G. Stengl, W. H. Bruenger, and F. Letzkus, *Adv. Mater.* **15**, 1152 (2003)
11. S. Y. Chou, P. R. Krauss, and P. J. Renstrom, *Science* **272**, 85 (1996)
12. M. Colburn, S. Johnson, M. Stewart, S. Damle, T. Bailey, B. Choi, M. Wedlake, T. Michaelson, S. V. Sreenivasan, J. Ekerdt, and C. G. Willson, *Proc. SPIE, Emerging Lithographic Technologies III*, 379 (1999)
13. G. M. McClelland, M. W. Hart, C. T. Rettner, M. E. Best, K. R. Carter, and B. D. Terris, *Appl. Phys. Lett.* **81**, 1483 (2002)
14. S. H. Sun, C. B. Murray, D. Weller, L. Folks, and A. Moser, *Science* **287**, 1989 (2000)
15. J. Y. Cheng, C. A. Ross, E. L. Thomas, H. I. Smith, R. G. H. Lammertink, and G. J. Vancso, *IEEE Trans. Magn.* **38**, 2541 (2002)
16. K. Naito, H. Hieda, M. Sakurai, Y. Kamata, and K. Asakawa, *IEEE Trans. Magn.* **38**, 1949 (2002)