

/technology

Electron lithography

>>simultaneous control of 13,000 electron beams

Evolution in the chip industry closely corresponds with the development of new production equipment.

Optical lithography suffers with both physical and financial constraints. Electron lithography could provide a solution. However, how do you switch every electron beam on and off 2.5 billion times per second?

Lithography is the most important step in the production of chips. This process transcribes the chip pattern onto a silicon wafer. On the wafer is a light-sensitive lacquer that is exposed by a wafer stepper (a type of super projector). The wafer stepper projects a slide (known as a mask), then the wafer is moved up a fraction and another mask is projected onto the wafer, and so on until the entire wafer is full. Subsequently, the lacquer is developed and the wafer is etched.

In the computer-chip-manufacturing industry, there is a constant race to make everything smaller: finer details on a chip mean more functionality and greater speed. Wafer-scanner manufacturers come up with all kinds of ideas to project these finer details using light. In order to write finer details, light of a shorter wavelength is necessary. The next generation of wafer scanners will use extreme ultraviolet (EUV) light. The wavelength of EUV light is close to that of X-rays, at just 13.5nm. This requires a complicated and entirely new machine design, using different materials and even more expensive masks..

Lithography using electron beams

Another way to achieve finer details is to use a different 'light source'. MAPPER Lithography uses electron beams to carry out the 'exposure'. Electron beams can write smaller structures than light beams. This technique works without masks, which results in considerable savings. Using a raster scan, the pattern is written onto the wafer by a large number of electron beams.

Experiments with this technique have been performed since as far back as the 1960s. The main obstacle was the speed. At the time, the electron beam functioned as a plotter that wrote patterns at an exasperatingly slow speed. In order to increase the speed, MAPPER developed a device that uses 13,000 beams, each of which writes part of the pattern.

Blanker and beam

MAPPER creates this large number of beams by aiming one large electron beam at a silicon slice that contains 13,000 holes. Under this slice is the blanker, which is a similar slice containing holes that enables the beams to be switched on and off. The holes in these slices are etched incredibly accurately – to the nanometre – creating 13,000 beams of equal size and therefore equal intensity. Electric lenses bend the beams in one dimension. Every beam can write lines of pixels that are 2µm in length and 3.5nm in width. The wafer moves under this beam, creating a 2µm-wide strip. All 13,000 beams together achieve a maximum chip size of 13,000 x 2µm = 26mm. This is a standard size in the chip industry.

At www.mapperlithography.com/technology, you can find an animated film depicting this process.

Exposure and watercolours

Writing a pattern with electrons can be compared to painting with watercolours. Every electron that hits the lacquer leaves a mark. The beam has a diameter of 25nm and can be positioned to an accuracy of 3.5nm. It is like making a finely detailed sketch with a thick paintbrush. This is possible because you are able to push




down slightly harder on the brush for just an instant, then move it slightly and release the pressure again. Only the sum of a number of exposures provides a dose sufficient to make the lacquer dissolve during development. In this way, it is possible to position the image in the wafer's photolacquer to a degree of accuracy greater than 3.5nm.

>>electron beams can write smaller structures than light beams

Pattern Streamer

The electron beams remain permanently switched on, although they can be interrupted by optically driven switches on the blanker. Every beam can be individually switched on or off. By switching the beam on or off at the right moment, a pattern is created. The switches are operated with pulses of light. These are provided by the pattern streamer – a computer that controls all of the electron beams in real time – via an optical fibre. The mask that is projected onto the wafer is saved in the pattern streamer as a bitmap with a size of 1.2 Gbyte per 2 μ m strip. This original bitmap is corrected for each chip, for example, by shifting or rotating the wafer slightly on the stage. However, the machine itself also has errors, e.g. beams can be slightly crooked. In such cases, the pattern streamer



calculates where the crooked beams touched the wafer and which data from the original bitmap is in the correct place. It moves parts of the original bitmap to create a new bitmap. This process is called resampling – an intensive calculation process with a great deal of data (around 2.5 Gbit per second per strip). The very short time between measurement of the errors and the exposure of the corrected bitmap makes it necessary to carry out these calculations in real time.

Redundancy


All material and process attributes will soon be known. These have already been processed during creation of the mask, even before the data is transmitted to the pattern streamer. The streamer processes only real-time corrections for machine-dependent aspects such as temperature and beam errors. When using 13,000 beams, the natural assumption is that one or two will be defective. This is why the number of beams is 2% higher. The machine can work out for itself which ones are defective. The fact that there will be defective beams is combined with another important design choice: around half of the 13,000 beams are controlled using two steps. For the second step, the wafer is shifted slightly from the position in the first step in order to avoid defective beams from the first step during 'exposure'. An additional advantage is that this halves the size of the pattern streamer, resulting in significant savings.

Calculation power


Resampling is carried out using programmable logic (FPGAs), but they cannot handle a speed of 2.4 Gbit per second. This is why the data pad has been divided into twenty parts per electron beam, each of which operates at 125MHz. One FPGA can control five electron beams. You therefore need an enormous amount of electronics to control 13,000 beams, even when using two steps. The pattern streamer the size of two filing cabinets – it is for this reason that the machine was developed in phases.


One FPGA can control five electron beams. Twelve FPGAs fit onto one board. Twelve of these boards fit on one rack. Four other boards are then added together with hard disks and a computer in order to control the whole thing. The masks of the chip in production are saved on the hard disks and depending on the step, a specific mask is transferred to the RAM memory.

One rack is sufficient for 720 channels. In order to control half of the 13,000 beams per step, a total of eleven racks are needed. Eleven racks actually gives the pattern streamer over 6,500 channels (the minimum number of channels required). The mapping between the channels and the beams is flexible. By correctly assigning the beams to the channels, it possible to create a complete exposure even if the channels of the pattern streamer or the fibres of the connection are defective.



1. Electron source
2. Lens
3. Optical fibre
4. Blaster
5. Wafer





Demo machines

There are already two demo machines in operation. With 110 channels, these machines are capable of writing small test patterns and are large enough to show that a series of separate beams are capable of writing one coherent pattern (beam stitching). One of these machines is at a research lab in Grenoble, and the other is in a production clean room at TSMC. The second machine in particular is generating a great deal of interest – TSMC is using it to develop opportunities to apply electron lithography to its processes.

>>a machine that surpasses the limits of optical lithography

Scale

Speed is an essential aspect in the production of ICs. One of the most important attributes is the transfer speed of machines (the number of wafers that can be processed per hour). At the moment, MAPPER is working on a machine that can expose ten wafers per hour, named Matrix 10.1. In order to limit initial costs, a 1-wafer-per-hour machine – Matrix 1.1 – is currently being built. This machine's data pad (the pattern streamer and the blanker) has been limited to 1,300 beams. By exposing every wafer 20 times using this machine, a 1-wafer-per-hour machine is created. Eventually, Matrix 10.10 will consist of ten Matrix 10.1 machines side-by-side, each of which will expose a combined total of 100 wafers per hour. Even with all of the peripheral equipment, this is still smaller than an EUV machine.

System design

The pattern streamer's eleven racks for the Matrix 10.1 are roughly the same size as the rest of the e-beam machine. Therefore, with ten machines, you have ten pattern streamers. They will be stored in the service area above or below the clean room, where a less strict chemical classification applies. A thick bundle of 7,800 optical fibres will run through the floor to the machine.

The other electronic equipment required to operate the machine must be kept as close as possible: if the distance between them is too great, the signals from these measuring and regulating devices are delayed. Distance does not affect the pattern streamer, as all of the signals from the pattern streamer are subject to the same delay, which the machine's main control unit corrects. This tells the wafer stage and the pattern streamer what to do. A precise clock provides the basis for the actions: you begin at this time, and you begin ten milliseconds later.

Multi-stage development

Multi-stage development spreads the risk. Furthermore, an entire machine with all of these parts would be too expensive. Costly FPGAs and memories are currently essential to achieving these speeds. However, the developers know that over time, prices fall and performance increases: a universal law of nature in the chip industry. By the time that MAPPER starts building Matrix 10.10, other FPGAs will have been developed that can do more and cost less. And for the fully developed machine, you could even use ASICs (application-specific integrated circuits). By using a step-by-step approach, this innovative technology will help develop a machine that transcends the limitations of optical lithography.