Biomedical Ultrasound Fundamentals of Imaging and Micromachined Transducers

TA : Anujkumar D Prajapati, Course Instructor : Dr. Hardik J. Pandya

Department of Electronic Systems Engineering

Indian Institute of Science, Bangalore

Lecture - 52

Hello everyone, and welcome to today's TA session. Our focus today is on wet etching and Miller indices, with particular emphasis on how the understanding of Miller indices aids in achieving the desired shapes during the wet etching of silicon. Let's begin by first understanding what etching means in the context of microfabrication.

Etching refers to the process of removing unwanted materials from a surface to create patterns or structures, which is a crucial step in microfabrication. During this process, various materials are deposited onto a substrate, and etching helps remove excess material to form the desired pattern. For instance, if you're looking to create an interconnect between two points say, point A and point B, you would start by depositing a blanket layer of conductive material such as polysilicon or a metal like gold. After the blanket deposition, a patterning technique like lithography would be used, and then the excess material is etched away, leaving a conductive path between the two points.

Etching isn't only applied to conductive layers but also to insulating layers. For example, a blanket layer of silicon nitride or silicon oxide could serve as insulation between two layers, but etching can be used to remove this insulating material where needed to establish a conductive path. Etching is also used to remove sacrificial layers in microstructures, such as creating hollow spaces. In this case, a material like SiO_2 or polysilicon would be deposited and later removed using etching, leaving behind the hollow space once the structural layers are in place.

In addition to these, etching can involve etch masks, etch stop layers, or structural layers, depending on the application. The material selected for etching depends on various factors like the specific function, the mechanical and electrical properties needed for the device, and how the etched layer interacts with preceding or subsequent layers.

The chemicals used in the etching process are known as etchants, and choosing the right one is critical. The selection of an etchant depends on factors like the type of material being etched, the specific application, and the mechanical and electrical properties of the layers involved. Additionally, you must ensure that the etching process doesn't negatively affect subsequent fabrication steps.

One key factor to consider when selecting an etchant is the etch rate, which refers to how quickly the etching process occurs. If an etchant works too quickly, it becomes difficult to control, and even small timing errors could lead to under-etching (removal of too little material) or over-etching (removal of too much material). For instance, if the etch process is completed in only 5 seconds, even a millisecond delay could produce inaccurate results.

In conclusion, wet etching is a critical process in microfabrication, allowing for precise removal of materials to create the intended structure or pattern. Understanding the selection of the right etchant and managing the etch rate effectively is crucial to achieving high precision. Up next, we'll explore the wet etching of silicon in more detail and discuss how Miller indices can guide us in shaping the material accurately.

The first important point to consider when selecting an etchant is the etch rate. The etching process should not be too fast, as rapid etching can make it difficult to control, leading to potential issues with under-etching or over-etching. Conversely, the etch rate should not be too slow either. If you have a device and need to etch a layer, the process should not take hours to complete. There should be an optimal etch rate, and this balance is critical when selecting your etchant.

Next, the etchant should not corrode or react with the protective mask. Imagine you have a layer that needs to be etched, and you apply the etchant directly. If there's no protective mask, the entire area will be etched, which is not what you want. To prevent this, an etch mask is used to protect specific areas from the chemicals. If your etchant corrodes the protective mask, the process will fail, which is why it's crucial that the etchant does not affect the protective material.

Additionally, the etchant should not react with the existing layers. Let's say you want to etch layer B, but before that, you had deposited or patterned layer A. If the etchant also etches away layer A, all your previous steps will be wasted. Therefore, another critical criterion is ensuring that the etchant does not react with or damage the existing or previous layers.

The concept of selectivity is also crucial. While ideally, the etchant should not etch the protective mask, in reality, that may not always be the case. Selectivity refers to the ratio of the etch rate of the material you want to etch to the material you do not want to etch, such as the protective mask or previous layers. High selectivity means the etchant primarily targets the material you intend to etch while minimizing the impact on other materials.

Etching can be classified into two broad categories: wet etching and dry etching. In wet etching, the entire sample is submerged in a liquid chemical bath, and the etching reaction occurs in all directions. This process is isotropic, meaning the etchant attacks the material equally from all sides. In contrast, dry etching uses a mixture of free radicals and plasma instead of liquid chemicals. This process is highly directional, with the etching focused in the direction of the plasma, resulting in more precise, straight walls. Techniques like RIE (Reactive Ion Etching) and DRIE (Deep Reactive Ion Etching) are commonly used for achieving highly directional etching in dry etching processes.

Now, focusing more on wet etching, there are different types of wet etching processes. One such process is isotropic etching, where the etching occurs equally in all directions. As I mentioned earlier, in wet etching, the entire sample is submerged in a chemical bath, allowing the chemicals to attack the material from every angle. If you observe an example of isotropic etching, you will see that the material being etched—such as a sacrificial SiO_2 layer—is attacked uniformly, creating a rounded profile.

When silicon dioxide (SiO_2) is placed in an etchant or the desired chemical, the chemical penetrates from all directions, attacking the material wherever possible and completely removing it. This is referred to as isotropic etching, where etching occurs uniformly in all directions, which often leads to the formation of suspended or hanging structures, as shown in certain figures. In addition to isotropic etching, wet etching can also exhibit anisotropic etching, where the etching is not uniform in all directions. While dry etching is typically highly anisotropic, in wet etching, crystal planes can influence the process, resulting in anisotropic etching as well.

In the ideal case of isotropic wet etching, the etching would produce a nearly spherical hole. However, when etching silicon, the presence of crystal planes can lead to the formation of a pyramidal shape. The specific angle of this pyramidal structure depends on the crystal orientation, and we will explore how this angle is formed in more detail.

Now, let's focus on etch anisotropy. In isotropic etching, the etch rate is uniform in all directions, meaning the lateral etch rate (parallel to the surface) is roughly the same as the vertical etch rate (perpendicular to the surface). For instance, if a protective mask is applied to a sample, the chemicals will attack from an open window in both the vertical and lateral directions at an equal rate. This results in uniform etching, and the process does not depend on the orientation of the mask.

In anisotropic etching, however, the etch rate varies depending on the crystal plane orientation. This means the lateral etch rate can be either greater or smaller than the vertical etch rate. When the lateral and vertical etch rates are equal, the etching is isotropic, but if they differ, the etching is anisotropic. We will delve deeper into how the orientation of the mask edges relative to the crystal planes influences the final etched shape, and how we can achieve the desired edge profile by carefully designing the mask in accordance with the crystal orientation.

Although anisotropic etching may initially seem like a disadvantage due to its complexity, it can actually be advantageous if harnessed correctly. This technique allows the creation of intricate and precise shapes. However, improper design can lead to unexpected results. To prevent this, standard shapes, recipes for mask design, and standard wafer orientations are typically used to produce consistent outcomes.

When discussing the anisotropic etching of silicon, we need to examine the various planes in the silicon crystal. Different crystal planes have distinct activation energies for etching reactions, meaning that depending on the orientation of the plane facing the etchant, the etch rate will vary. This results in the anisotropic effect. For example, in the etching of silicon using potassium hydroxide (KOH), the reaction rate varies based on the specific crystal planes exposed to the etchant, which leads to a characteristic etching profile.

The etching process in KOH is not diffusion-limited, as the chemical diffuses in all directions. Instead, it is reaction rate-limited, primarily due to the activation energies of the different crystal planes. In the case of KOH etching, the most prominent plane we encounter is the (111) plane, which has a significantly lower etch rate. As a result, the etching process tends to stop at the (111) plane, making it useful as an etch-stop plane in anisotropic etching of silicon. In contrast, the (100) and (110) planes exhibit much higher etch rates. This characteristic of the (111) plane is why it is often used for defining precise geometries in silicon etching processes.

When we examine the intersection of the various (111) planes, which can number up to eight, the result is the formation of complex shapes like V-grooves, pyramidal pits, and cavities. These geometries are intrinsic to the anisotropic etching of silicon and are determined by the crystallographic planes exposed to the etchant. For example, when etching silicon and observing it from a top-down view, the protective mask (represented in blue) shields specific areas from being etched. If we take a cross-sectional view, instead of observing an isotropic spherical shape, we find a pyramidal structure formed by the (111) planes, showing how anisotropic etching results in sharp, well-defined features.

Now, let's transition to Miller indices, a concept introduced by William Miller in 1839. These indices are integer values that define the directions and families of planes within a crystal structure. Miller indices are essential in crystallography, especially when describing the orientation of crystal planes. For 2D crystals, we use two integers in the Miller indices set, while 3D structures require three integers. Miller indices play a crucial role in identifying and describing the planes that dictate the etching behavior in various materials, such as silicon.

In a crystal, the smallest repeating unit is called the unit cell, analogous to an atom in a material. Just as atoms repeat to form a material, the unit cells repeat along the x, y, and z axes to form the crystal. The Miller indices describe these unit cells and their orientations. To illustrate, consider the simplest crystal structure: the simple cubic unit cell. If you imagine a cube aligned along the x, y, and z axes, Miller indices help specify the orientation of the planes within that cube. For instance, the origin lies at the intersection of these axes, and the Miller indices define how the planes within this cube are oriented with respect to the crystallographic axes.

We begin by identifying a point in the crystal structure using Cartesian coordinates (0,0,0), known as the origin. If we move along the x-axis, we label this distance as a; along the y-axis, it is labeled as b; and along the z-axis, it is called c. For instance, if we move along the x-axis by one unit, the coordinates become (a=1, b=0, c=0). Similarly, if we move to another point, such as (1,1,1), it means we've moved one unit along each axis. These distances along the crystallographic axes form the basis for Miller indices, which are used to describe the orientations of planes and directions within the crystal.

All points in this crystal, marked by red dots in the illustration, are known as lattice points. Now, let's focus on determining Miller indices for directions. To calculate the Miller indices for a direction, follow these steps: first, subtract the tail coordinates from the head coordinates; next, reduce the result to the smallest integer values; and finally, enclose them in square brackets to represent the direction. For example, in the direction shown on the left, the head coordinates are (1, 0, 1) and the tail coordinates are (0, 0, 0). Subtracting the tail from the head, we get (1, 0, 1), and when enclosed in square brackets, we obtain [1 0 1] as the Miller indices for that direction.

It's important to note that when denoting a point, it should be separated by commas and enclosed in round brackets. Let's take another example: if the tail coordinates are (1, 0, 1) and the head coordinates are (0, 1, 0), subtracting the tail from the head yields (1, -1, 1). For negative values, we replace the minus sign with an overbar, resulting in [1 $\bar{1}$ 1], which represents the Miller indices for that direction.

Now, let’s move on to calculating Miller indices for planes. This process is similar to finding directions, but instead of subtracting coordinates, we look at where the plane intersects the x, y, and z axes. The distances at which the plane intersects these axes are labeled a, b, and c. Once we know these values, we calculate the reciprocals of a, b, and c to find h, k, and l, which are the Miller indices for that plane. If the reciprocals yield fractions, we convert them to the lowest integer values. For negative values, we again use the overbar notation.

The final Miller indices for planes are enclosed in round brackets, like (h k l). However, if we're referring to a family of planes with similar orientations, we enclose them in curly braces, like {h k l}. This distinction helps clarify whether we're discussing a specific plane or a set of equivalent planes within the crystal structure.

Let's take a closer look at the figure on the right. Here, we have a cube, and from this perspective, we can see three of its faces. Let’s focus on the blue face. If you examine it, you’ll notice that this blue plane intersects the z-axis at a point labeled C. In this case, since it’s a cube, let’s assume that the length of each side is 1 unit, meaning C = 1. Now, one thing to observe about this cube is that all the faces are parallel, and the opposite faces are parallel to each other. This means that the blue plane is parallel to the XY-plane.

When a plane is parallel to something, it never intersects it. Therefore, since the blue plane is parallel to the xy-plane, it will never intersect the x-axis or the y-axis. In mathematics, something that never intersects is said to intersect at infinity. So, we can say that for this blue plane, $a = \infty$ and $b = \infty$ because it doesn’t intersect the x or y axes, and it intersects the z-axis at $c = 1$.

To calculate the Miller indices (h k l), we take the reciprocals of a, b, and c. Therefore, the reciprocals are:

$1/\infty$, $1/\infty$, and $1/1$, which simplifies to 0, 0, and 1. These three integers give us the Miller indices for the blue plane, which is denoted as (0 0 1).

Now, let’s consider the red plane. The red plane intersects only the y-axis, and it does not intersect the x-axis or the z-axis. So, in this case, $a = \infty$, $b = 1$, and $c = \infty$. Taking the reciprocals of these values, we get 1/∞, 1/1, and 1/∞, which simplifies to 0, 1, and 0. Therefore, the Miller indices for the red plane are (0 1 0).

For the green plane, it intersects only the x-axis, while not intersecting the y-axis or z-axis. Using the same process, the Miller indices for this plane are (1 0 0).

One important thing to note about all three of these planes is that they are the faces of the cube, and if you carefully observe the Miller indices, you'll notice that each plane has one value of 1 and two values of 0. These planes can be represented as a family of planes, indicated by using curly braces: {1 0 0}. This family of planes is particularly significant in silicon crystals, and most silicon wafers you encounter will have a {1 0 0} plane orientation.

Let's now consider a more complex example. In this case, rather than having simple unit vectors like a = 1, b = 1, and c = 1, let's say the unit vectors in the crystal structure are A, B, and C, where A is in the x-direction, B in the y-direction, and C in the z-direction. Let's assume that the plane intersects the x-axis at 3A, the y-axis at 2B, and the z-axis at 2C.

To find the Miller indices for this plane, we first take the reciprocals of these values:

1/3, 1/2, and 1/2. Since Miller indices must be integers, we multiply each reciprocal by a common factor to eliminate the fractions. In this case, multiplying all by 6 (the least common denominator), we get 2, 3, and 3. Therefore, the Miller indices for this plane are (2 3 3).

Let's go over the process step by step. When dealing with Miller indices, one efficient way to maintain the ratio is to multiply by the least common multiple (LCM). For example, if you have intercepts at 3, 2, and 2, the LCM of these numbers is 6. By multiplying through by 6, you get 2, 3, and 3, which represent the Miller indices for this plane: (2 3 3).

Now, let's look at the next example of Miller indices for a different plane. If we examine the plane in question, we can mark the points of intersection along the axes. In the x and y directions, the plane intercepts at a, and in the z direction, it does not intercept, meaning it intersects at infinity. So, for the z-axis, the intercept is infinity. If we write down the intercepts, we have a, a, and infinity. Taking the reciprocals, we get 1/a, 1/a, and 1/infinity. In this case, the LCM would be a, and simplifying gives us (1 1 0) as the Miller indices for this plane.

Similarly, in another example, if the intercepts are all at equal distances, say A, A, and A, it's a straightforward process. The reciprocals of these intercepts will be 1/A, 1/A, and 1/A, and after taking the LCM, we get the Miller indices (1 1 1). This plane forms a triangular shape, which is often referred to as the (1 1 1) plane, and it is another crucial plane in silicon crystal structures.

Let's consider another random plane for fun. In this case, if you observe the intercepts, there is no z-intercept, meaning the plane is parallel to the z-axis. Therefore, the intercept for z is infinity. In the x direction, the intercept is a/2, and in the y direction, it's a. So, writing down the intercepts, we have a/2, a, and infinity. Taking reciprocals gives 2/a, 1/a, and 1/infinity, and after finding the LCM, the Miller indices for this plane become (2 1 0).

Now, let's discuss the notation for Miller indices. If the values are enclosed in round brackets and separated by commas, it represents a point. If the values are enclosed in round brackets without commas, it represents the Miller indices for a plane. When the values are enclosed in curly braces, they represent a family of planes. For example, {1 0 0} represents the family of planes with Miller indices (1 0 0), (0 1 0), and (0 0 1). When enclosed in square brackets, the values represent a direction, and if they are in angle brackets, they denote a family of directions. For instance, <1 0 0> represents a family of directions along the (1 0 0) axis.

Understanding these notations is vital, particularly in wet bulk micromachining, because the etching characteristics heavily depend on the crystal's orientation, as described by the Miller indices. In order to fabricate microstructures with precision, it's essential to align the mask edges correctly with respect to the crystal directions. Otherwise, the resulting shape might differ from the intended design. The anisotropic etching of silicon is highly dependent on crystallographic planes. In wet etching with KOH, which is one of the most common etching methods, the etch rate is fastest on the (1 0 0) planes, followed by the (1 1 0) planes, while the (1 1 1) planes exhibit the slowest etch rates.

We'll stop here for now. In the next session, we'll dive deeper into the concept of Miller indices, how to visualize these indices within silicon crystals, and how this knowledge can help us create complex microstructures through wet etching.