

Semiconductor Technology from A to Z

Wafer Fabrication

Silicon	3
<i>Properties of Silicon</i>	3
Production of raw silicon	4
<i>Starting Material: Silicon Dioxide</i>	4
<i>Purification of the Raw Silicon</i>	5
<i>Zone Refining</i>	6
Fabrication of the single crystal	7
<i>The Single Crystal</i>	7
<i>Czochralski Crystal Growing Process</i>	7
<i>Crucible-Free Float-Zone Growing</i>	9
The wafer	9
<i>Wafer Slicing and Surface Finishing</i>	9
<i>Historical Development of Wafer Size</i>	12
<i>Why Are Wafers Round?</i>	14
Doping techniques	15
<i>Doping</i>	15
<i>Diffusion</i>	15
<i>Diffusion Methods</i>	16
<i>Ion implantation</i>	18


Silicon

Properties of Silicon

Silicon is the chemical element with atomic number 14 in the periodic table; it belongs to the 4th main group and the 3rd period. Silicon is a classic semiconductor, meaning its conductivity lies between that of conductors and insulators. In nature, silicon (from Latin *silex/silicis*: pebble) occurs exclusively as an oxide: as silicon dioxide (SiO_2) in the form of sand and quartz, or as silicate (compounds of silicon with oxygen, metals, and others). Silicon is therefore, quite literally, as abundant as sand on a beach, making it an extremely inexpensive starting material whose value is only determined through processing. Other semiconductors, such as germanium or the compound semiconductor gallium arsenide, offer in some respects substantially better electrical properties than silicon: charge carrier mobility, and the resulting switching speeds, are significantly higher in germanium and GaAs. However, silicon has decisive advantages over other semiconductors.

This dominant position applies to integrated circuits, not to semiconductor technology as a whole. In power electronics, silicon carbide and gallium nitride have secured a firm place for themselves, since they can withstand higher voltages, higher temperatures, and faster switching. And for anything intended to emit light, silicon is unsuitable: its band transition releases energy as heat rather than as light, which is why light-emitting diodes and lasers are made from compound semiconductors instead.

Bohr atomic model of silicon



Oxide layers can be produced very easily on a silicon crystal; the resulting silicon dioxide is a high-quality insulator that can be deposited on the substrate in a targeted manner. For the semiconductors mentioned above, germanium and


GaAs, producing similarly good insulating layers is, by contrast, very costly. The ability to precisely alter the conductivity of pure silicon through doping is another reason this metalloid is so significant. Other materials are, in some cases, highly toxic, and compounds formed with these elements are not as durable and stable as those based on silicon. A prerequisite for using silicon in semiconductor manufacturing, however, is that it be present in an ultra-pure, single-crystal form. This means that the silicon atoms in the crystal lattice are arranged in a perfectly regular pattern, with no undefined foreign atoms present in the crystal.

In addition to the single-crystal form, there is also polysilicon (poly = many) and amorphous silicon (a-Si). While single-crystal silicon, in the form of wafers, is the foundation of microelectronics, polycrystalline silicon is deposited as a layer on the wafer in various areas to fulfil specific tasks (e.g. as a masking layer, as the gate in a transistor, and others). It can be produced easily and patterned with ease. Polysilicon is made up of many individual single crystals arranged irregularly with respect to one another. Amorphous silicon has no regular lattice structure at all, but rather a disordered one. In chip manufacturing it plays only a minor role; outside of it, however, this is by no means the case: because it can be deposited over large areas and at low temperatures onto glass, the switching elements behind every flat-panel display are built from it, and it is also used in thin-film solar cells.

Production of raw silicon

Starting Material: Silicon Dioxide

The silicon used in semiconductor manufacturing is obtained from quartz. In this process, oxygen — which combines with silicon extremely readily even at room temperature, and which in quartz is likewise present in combination with silicon as silicon dioxide — must be removed. This is done in an arc furnace using carbon, at temperatures that reach up to about 2000 °C in the hottest zone of the furnace. The oxygen separates from the silicon and reacts with the carbon (C) to form carbon monoxide (CO):



This equation summarizes what actually takes place in the furnace over several stages. In the cooler zone, silicon carbide initially forms; this then reacts with additional silicon dioxide in the hotter zone to produce silicon. Silicon carbide is therefore not an unwanted by-product but a necessary intermediate product; the ratio of carbon to quartz is set so that it is fully consumed again by the end of


the process. At these temperatures, the carbon monoxide is gaseous and can easily be separated from the liquid silicon.

The raw silicon obtained in this way (metallurgical-grade silicon) is approximately 98 to 99 % pure. The remainder consists of impurities such as iron, aluminium, calcium, phosphorus, and boron. For use in semiconductor devices, this is far too impure by many orders of magnitude — these substances must be removed in further processes.

Purification of the Raw Silicon


Purification is carried out indirectly via a gaseous compound: the silicon is converted into trichlorosilane, which is purified by distillation, and solid silicon is then deposited from it again. The process is known as the Siemens process, named after its developer, and to this day still supplies the majority of the world's ultra-pure silicon.

Using the trichlorosilane process, many impurities are removed by distillation. The raw silicon reacts with hydrogen chloride (HCl) at around 300 °C to form gaseous trichlorosilane (SiHCl₃) and hydrogen (H₂):




The impurities, which also combine with the chlorine present, only pass into the gaseous state at higher temperatures. This allows the trichlorosilane to be separated out. Only carbon, phosphorus, and boron, which have similar condensation temperatures, cannot be filtered out by this process.

The trichlorosilane process can be reversed, so that the purified silicon is subsequently obtained in polycrystalline form. This takes place with the addition of hydrogen in a quartz vessel containing heated, thin silicon rods (silicon seed rods), at around 1100 °C:



and



The silicon deposits onto the silicon seed rods, which grow over several days into rods roughly 150 to 200 mm thick. These are subsequently broken up and processed further as chunks.

The process is extremely energy-intensive, since the rods must be kept at temperature the entire time while the chamber wall is cooled. As an alternative, fluidized-bed deposition has become established: here, small silicon granules grow within a bed through which gas flows, which requires significantly less energy and operates continuously rather than in batches. However, its purity does not match that of the Siemens process, which is why this material is used

primarily in the solar industry.

This is precisely where an important distinction lies: silicon suitable for solar applications must reach roughly six nines of purity (99.9999 %), whereas silicon suitable for electronics must reach nine to eleven nines. Using the Czochralski process, this polysilicon could already be converted into a single crystal; however, its degree of purity is still not high enough for the fabrication of semiconductor devices.


Zone Refining

To further increase purity, zone refining is used. A coil carrying a high-frequency alternating current is placed around the silicon rods. This melts the silicon inside the rods, and the impurities sink downward due to their high solubility in the melt (the surface tension of the silicon prevents the melt from flowing out). By repeating this process multiple times, the level of impurities in the silicon is reduced to the point where it can be further processed into a single crystal.

Illustration of zone refining

Zone Refining of a Silicon Rod

Surface tension holds the melt between the fixed ends



Impurities remain in the melt and travel along with the zone.

All processes are carried out under vacuum to prevent further contamination.

At the end of these processes, the silicon prepared for semiconductor

manufacturing has a purity of more than 99.9999999 %, corresponding to less than one foreign atom per 1 billion silicon atoms.

Fabrication of the single crystal

The Single Crystal

A single crystal (monocrystal), as required in semiconductor manufacturing, is a regular arrangement of atoms. In addition, there is polycrystalline silicon (composed of many small single-crystal grains) and amorphous silicon (disordered structure). Depending on the orientation of the lattice in space, silicon wafers exhibit different surface structures, which influence various device parameters, such as charge carrier mobility.

Crystal orientations



Chip manufacturing uses almost exclusively (100)-oriented silicon. The reason lies at the interface with the gate oxide: unsaturated bonds remain there after oxidation, and their number is lowest for this orientation. The (111) orientation has the densest arrangement of bonds and therefore the most defects; however, it oxidizes the fastest.

Crystal orientation is also of particular importance in micromechanics. For instance, microchannels with vertical walls can be produced on (110) silicon, whereas (100) orientation results in sidewall angles of 54.74° .

Czochralski Crystal Growing Process

The polycrystalline silicon obtained after zone refining is melted in a quartz


crucible just above the melting point of silicon. Dopants (e.g. boron or phosphorus) can now be added to the melt in order to achieve the desired electrical properties of the single crystal.

A seed crystal (a single crystal) mounted on a rotating rod is brought into contact with the surface of the silicon melt. This seed determines the orientation of the crystal. Upon contact between the seed and the melt, silicon deposits onto the seed and adopts its crystal structure. Because the crucible temperature is only slightly above the melting point of silicon, the deposited silicon solidifies immediately on the seed, and the crystal grows.

The seed is slowly pulled upward while being continuously rotated, maintaining constant contact with the melt. The crucible rotates in the opposite direction to the seed crystal. Constant temperature control of the melt is essential to ensure uniform growth. The diameter of the single crystal is determined by the pulling speed, which ranges from 2 to 25 cm/h. The faster the pulling speed, the thinner the crystal becomes. The entire apparatus is enclosed in a protective gas atmosphere to prevent oxidation of the silicon.

Czochralski Crystal Growth

The seed crystal sets the orientation of the crystal



Diameter is set via pulling speed and melt temperature.

A significant disadvantage of this process is that the melt becomes increasingly enriched with dopants during the process, since the dopants dissolve more readily in the melt than in the solid. As a result, the dopant concentration is not constant along the length of the silicon rod. In addition, impurities or metals can dissolve out of the crucible and become incorporated into the crystal lattice.

The advantages of this process are its low cost and the ability to produce larger wafers than is possible with the float-zone process (see below).

Crucible-Free Float-Zone Growing


In contrast to the Czochralski crystal growing process, crucible-free float-zone growing does not melt the entire polysilicon rod; instead, as with zone refining, only a small region (a few millimetres) is melted.

Here too, a seed crystal, brought into contact with the end of the polycrystalline silicon rod, determines the crystal structure. The polycrystal is melted and adopts the structure of the seed. The heating zone is slowly moved along the rod, and the polycrystalline silicon rod gradually transforms into a single crystal.

Since only a small region of the polycrystalline silicon is melted, very few impurities can accumulate (as with zone refining, the foreign atoms migrate to the end of the silicon rod). Doping is achieved here by adding dopants (e.g. diborane or phosphine) to the protective gas that flows around the apparatus.

Float-Zone Process

Only a few millimeters of the rod are molten




No crucible means no contact with foreign material; doping is introduced via the protective gas.

The wafer

Wafer Slicing and Surface Finishing

The single-crystal rod is first ground down to the desired diameter and then, depending on crystal orientation and doping, given one or two flats. The larger flat is used to precisely align the wafers during production. The second flat identifies the wafer type (crystal orientation, p-/n-doping), but is not always present. For wafers 200 mm in diameter and larger, so-called notches are used instead of flats. These are tiny notches at the edge of the wafer that likewise allow the wafer to be aligned, but take up far less of the wafer's valuable surface area.



Sawing

Historically, the single-crystal rod was cut using an internal diameter (ID) saw, whose cutting edge is coated with diamond fragments. It cuts precisely, but only one wafer at a time, and up to 20 % of the crystal rod is lost to the thickness of the saw blade. Given today's rod diameters and wafer prices, this is no longer acceptable; the standard method is therefore wire sawing, in which several hundred wafers are cut from the rod in a single pass. A long wire, wetted with a slurry of silicon carbide grains and a carrier fluid such as glycol or oil, is guided over rotating rollers. The silicon crystal is lowered into the wire grid and thereby sliced into wafers. The wire moves back and forth at about 10 m/s and is

typically 0.1–0.2 mm thick. Increasingly, instead of a slurry of loose grains, a wire with diamond grains fixed firmly into its surface is used. It cuts faster, produces less waste, and eliminates the need to dispose of spent slurry.


After sawing, the wafers have a roughened surface and, due to the mechanical stress, lattice damage within the crystal. To refine the surface, the wafers pass through several process steps.

Lapping

Using granular abrasives (e.g. aluminium oxide) on a rotating steel plate, 50 μm (0.05 mm) of the wafer surface is removed. The grain size is reduced in stages, but the surface is once again damaged by the mechanical treatment. The flatness after lapping is about 2 μm .

Rounding the wafer edge

In later processes, the wafers must not have any sharp edges, since deposited layers could otherwise flake off. For this reason, the edge of the wafers is rounded off using a diamond grinding tool.



Etching

In a dip-etch step using a mixture of hydrofluoric, acetic, and nitric acid, a further 50 μm is removed. Since this is a chemical process, the surface is not damaged. Crystal defects are finally eliminated.

Polishing

This is the final step toward the finished wafer. At the end of the polishing step, the wafers have a remaining unevenness of less than 3 nm (0.000003 mm). For this, the wafers are treated with a mixture of sodium hydroxide solution (NaOH), water, and silicon dioxide particles. The silicon dioxide removes a further 5 μm

from the wafer surface, while the sodium hydroxide removes oxide and eliminates processing marks left by the silicon dioxide particles.

This is followed by a final cleaning and a measurement of each individual wafer for flatness, thickness, and particles. A considerable proportion of wafers subsequently also receive an **epitaxial layer**: a thin, especially pure, and precisely doped layer of silicon grown onto the polished surface. The actual devices are then formed in this layer, while the wafer beneath serves only as a carrier.

Historical Development of Wafer Size


The manufacture of integrated circuits on silicon wafers began in the mid-1960s using wafers with a diameter of 25 mm. Since then, the diameter has grown in steps; since the late 1990s, the 300 mm wafer has been the largest format used in volume production. Its area is roughly 140 times that of the first 25 mm wafers.

With larger wafers, throughput in chip manufacturing increases considerably, allowing manufacturing costs to be reduced accordingly. For example, at the same feature size, more than twice as many chips can be produced on a 300 mm wafer as on a 200 mm wafer.

Typical wafer data:

Type [mm]	Diameter [mm]	Thickness [μm]	Primary flat [mm]	Bow [μm]
150	150 ± 0.5	~700	55 - 60	25
200	200 ± 0.5	~800	<i>Notch</i>	35
300	300 ± 0.5	~900	<i>Notch</i>	45

Overview of wafer sizes: 25, 38, 51, 75, 100, 125, 150, 200, 300, 450 [mm] (to scale)



The Step to 450 mm That Never Happened

The progression was supposed to continue with 450 mm. Around 2008, several major manufacturers and equipment suppliers joined forces in a shared program; initial equipment and test wafers were produced, and the transition was announced for the early 2010s. It never happened: the program was shut down in the mid-2010s, and 300 mm has remained the largest format to this day.

The reasons for this already lie in the technical difficulties that were being tackled at the time:

- Bow increases with diameter. To prevent stacked wafers from touching each other during transport, spacings, supports, and grippers would have had to be redesigned.
- Film stresses deform a large wafer more easily. It cannot simply be made thicker, since material costs and natural frequency work against this.
- The larger area also has to be processed. Compared with 300 mm, the cost per square centimeter would have had to fall by a factor of 2.25 for the economics to work out at all.
- A 450 mm crystal takes more than twice as long to grow. During this time, the probability of defects being incorporated into the lattice also increases.
- Almost every piece of equipment on the production line would have had to be redeveloped – at an investment level that individual manufacturers could no longer bear on their own.

The decisive point, however, was an economic one. The benefit of a larger wafer lies in spreading the fixed cost per piece of equipment over more chips. This advantage becomes smaller the more expensive the individual piece of

equipment becomes – and lithography equipment became dramatically more expensive during the same period. The effort required for the transition would therefore have increased, while the return from it decreased. Instead, the industry directed its investments toward smaller feature sizes and the third dimension: more devices per unit area rather than more area per wafer.

Why Are Wafers Round?

The question often arises as to why wafers are round, since microchips are, after all, rectangular. This inevitably results in wasted area on the wafer – area where no complete chips fit, and which must ultimately be discarded at the end of semiconductor manufacturing.

Having explained the two manufacturing processes - the [crystal growth process](#) and [zone pulling](#) - this question can easily be answered.

A silicon wafer for microchip manufacturing must be a single crystal. This is only possible with the processes mentioned, and these inherently produce a circular shape.

Even though it would be technically possible to subsequently shape the round silicon crystals into an angular form (e. g. by sawing), the round shape of silicon wafers nevertheless offers several advantages despite the rectangular microchips.

- Straightening the round silicon rods would place additional stress on the material and would inevitably lead to crystal defects that would affect chip quality.
- Round wafers are considerably more robust. Angular wafers could hardly be transported and processed without damage.
- Uniform processing during chip manufacturing using radially symmetric processes (CMP, spin-on, etching) is considerably simpler.
- A narrow edge region would always have to be discarded even with rectangular wafers, since the wafers must be held during processing. Deposited layers would flake off and generate additional particles if they extended all the way to the outermost edge.

As wafer size increases, the amount of wasted area also continues to decrease.

Rectangular wafers, on the other hand, are found in the manufacture of solar cells. Polycrystalline wafers are mostly used here, which can be cast into rectangular shapes. Manufacturing is comparatively simple, so angular wafers can be processed as well. The corners are usually additionally chamfered. For today's common single-crystal solar cells, it is exactly the other way around: they are made from crystals pulled in round form, which are trimmed to a square shape before sawing – with rounded corners, because cutting away the

remaining edge would cost more material than the gain in area would be worth.

Doping techniques

Doping

Doping means introducing a foreign substance into the semiconductor crystal in order to deliberately alter its conductivity through an excess or a deficiency of electrons. In contrast to doping during wafer fabrication itself, where the entire wafer is doped, the doping techniques described on this page allow silicon wafers to be doped partially. The introduction of the foreign substance can be achieved using various methods: [diffusion](#), [implantation](#) (and alloying).

Diffusion


Diffusion is barely used anymore for targeted doping today; it has been superseded by [ion implantation](#). Nevertheless, understanding it remains essential, because it never stops acting: every subsequent high-temperature step causes dopants that have already been introduced to keep migrating. This limits how much heat a wafer may still be subjected to after doping, and is the reason why furnace processes have increasingly been replaced by rapid, second-scale heating.

Diffusion means the spontaneous spreading of a substance within another substance due to a difference in concentration; for example, a drop of ink in a glass of water becomes evenly distributed after a certain amount of time. In a silicon crystal, one finds a solid lattice of atoms through which the dopant must move. This can happen in three ways:

- Vacancy diffusion: the foreign atoms occupy empty sites in the crystal lattice, which can always occur. Similar to hole conduction, an interplay arises between occupied lattice sites and vacancies.
- Interstitial diffusion: the foreign atoms move between the silicon atoms within the crystal lattice.
- Site exchange: foreign atoms located in the crystal lattice swap lattice sites with silicon atoms.

Diffusion in the crystal lattice

Three paths by which the dopant migrates through the lattice



The dopant can continue to move within the semiconductor crystal until either a concentration gradient has been equalized, or the temperature has been lowered far enough that the atoms can no longer move.

The speed of the diffusion process depends on several factors:

- dopant
- concentration difference
- temperature
- substrate
- crystal orientation of the substrate (see [fabrication of single crystals](#))

Diffusion with an Exhaustible Source

Diffusion with an exhaustible source means that only a limited amount of dopant is available. The longer the diffusion process continues, the lower the concentration at the surface becomes; in return, the penetration depth into the substrate increases. The diffusion coefficient of a substance indicates how quickly it moves within the crystal. Arsenic, with a low diffusion coefficient, penetrates the substrate more slowly than, for example, phosphorus or boron.

Diffusion with an Inexhaustible Source

Here, the dopant is available in unlimited quantity. The concentration at the surface therefore remains constant throughout the process, since particles that have penetrated into the substrate are continuously replenished.

Diffusion Methods

In the following processes, the wafers are placed inside a quartz tube that is heated to a specific temperature along its entire length.

Diffusion from the Gas Phase

A carrier gas (nitrogen, argon) is enriched with the desired dopant (likewise in gaseous form, e.g. phosphine (PH_3) or diborane (B_2H_6)) and passed over the silicon wafers, where concentration equalization can take place.

Solid-Source Diffusion

Discs onto which the dopant has been applied are placed between the wafers. As the temperature in the quartz tube rises, the dopant diffuses out of the source discs into the atmosphere. A carrier gas then distributes the dopant evenly throughout the quartz tube, allowing it to reach the surface of the wafers.

Diffusion with a Liquid Source


Boron tribromide (BBr_3) or phosphorus oxychloride (POCl_3) serve as liquid sources. A carrier gas is passed through the liquid and thereby transports the dopant to the wafers, where it then reaches the surface of the silicon discs.

Since the entire wafer is not meant to be doped, certain regions are masked with silicon dioxide. Dopants cannot penetrate this oxide, so no doping takes place at these locations. To avoid stress or even breakage of the wafers, the tube is heated in steps (10 °C per minute) up to approximately 900 °C, at which point the dopant is introduced to the wafers. To initiate the diffusion process, the temperature is then raised to approximately 1200 °C.

Characteristics:

- Since many wafers can be processed simultaneously, this method is quite cost-effective
- If foreign substances from earlier doping steps are already present in the crystal, renewed exposure to heat can cause them to diffuse out again
- Over time, dopants accumulate inside the quartz tube and are additionally transported to the wafers by the carrier gas during subsequent doping steps
- Dopants spread within the crystal not only vertically but also laterally, so the doping windows always end up doped over a larger area than intended

Diffusion through a window in the mask
The dopant also migrates sideways and undercuts the oxide



The lateral spread is roughly eight tenths of the penetration depth.
As a result, the doped region always ends up wider than the window -
this has to be accounted for when designing the masks.

Ion implantation

In the ion implantation charged dopants (ions) are accelerated in an electric field and irradiated onto the wafer. The penetration depth can be set very precisely by reducing or increasing the voltage needed to accelerate the ions. Since the process takes place at room temperature, previously added dopants can not diffuse out. Regions that should not be doped, can be covered with a masking photoresist layer.

An implanter consists of the following components:


- ion source: the dopants in gaseous state (e.g. boron trifluoride BF_3) are ionized
- accelerator: the ions are drawn with approximately 30 kiloelectron volts out of the ion source
- mass separation: the charged particles are deflected by a magnetic field by 90 degrees. Too light/heavy particles are deflected more/less than the desired ions and trapped with screens behind the separator
- acceleration lane: several 100 keV accelerate the particles to their final velocity (200 keV accelerate bor ions up to 2.000.000 m/s)
- Lenses: lenses are distributed inside the entire system to focus the ion beam
- distraction: the ions are deflected with electrical fields to irradiate the desired location
- wafer station: the wafers are placed on large rotating wheels and held into

the ion beam

Illustration of an ion implanter

Ion implanter

The separator lets through only the desired ion species



Penetration depth of ions in the wafer

In contrast to diffusion processes the particles do not penetrate into the crystal due to their own movement, but because of their high velocity. Inside the crystal they are slowed down by collisions with silicon atoms. The impact causes damage to the lattice since silicon atoms are knocked from their sites, the dopants themselves are mostly placed interstitial. There, they are not electrically active, because there are no bonds with other atoms which may give rise to free charge carriers. The displaced silicon atoms must be re-installed into the crystal lattice, and the electrically inactive dopants must be activated.

Recovery the crystal lattice and activation of dopants

Right after the implantation process, only about 5 % of the dopants are bond in the lattice. In a high temperature process at about 1000 °C, the dopants move on lattice sites. The lattice damage caused by the collisions have already been cured at about 500 °C. Since the dopants move inside the crystal during high temperature processes, these steps are carried out only for a very short time.

Channeling

The substrate is present as a single crystal, and thus the silicon atoms are regularly arranged and form "channels". The dopant atoms injected via ion implantation can move parallel to these channels and are slowed only slightly, and therefore penetrate very deeply into the substrate. To prevent this, there are several possibilities:

- wafer alignment: the wafers are deflected by about 7° with respect to the ion beam. Thus the radiation is not in parallel direction to the channels and the ions are decelerated by collisions immediately.
- scattering: on top of the wafer surface a thin oxide is applied, which deflects the ions, and therefore prevents a parallel arrival



Characteristic:

- the reproducibility of ion implantation is very high
- the process at room temperature prevents the outward diffusion of other dopants
- spin coated photoresist as a mask is sufficient, an oxide layer, as it is used in diffusion processes, is not necessary
- ion implanters are very expensive, the costs per wafer are relatively high
- the dopants do not spread laterally under the mask (only minimally due to collisions)
- nearly every element can be implanted in highest purity
- previous used dopants can deposit on walls or screens inside the implanter and later be carried to the wafer
- three-dimensional structures (e.g. trenches) can not be doped by ion implantation
- the implantation process takes place under high vacuum, which must be produced with several vacuum pumps

There are several types of implanters for small to medium doses of ions (10^{11} to 10^{15} ions/cm²) or for even higher doses of 10^{15} to 10^{17} ions/cm².

The ion implantation has replaced the diffusion mostly due to its advantages.

Doping using Alloy

For completeness it should be mentioned that besides ion implantation and diffusion there is an alternative process: doping using alloy. Since this procedure, however, brings disadvantages with it such as cracks in the substrate, it is not used in today's semiconductor technology any more.