SEMICONDUCTOR FABRICATION · PART 2-1

From quartz to a single crystal Purifying silicon and growing the wafer’s foundation

Before a transistor can be made, its silicon foundation needs **exceptional purity** and **a continuous crystal lattice**. Those are different jobs. Follow the chemistry that produces electronic-grade polysilicon, then the physics that turns it into a large single crystal by the Czochralski or float-zone method.

Part 2 · Lesson 2-1 · Intermediate beginner level · Introduction 1-1 is useful background · English edition: 4 November 2026

In this lesson

1. Two transformations before chipmaking can begin
2. Why purity alone is not enough
3. First purification step: take the oxygen out of quartz
4. Second step: make a compound that can be distilled
5. Third step: deposit polysilicon by the Siemens process
6. Czochralski growth: freeze silicon onto a seed
7. Four pieces of physics that determine crystal quality
8. Float-zone growth: remove the crucible from the problem
9. The industry behind the material
10. Connect the chemistry to the crystal
11. Check your understanding
12. A short glossary
13. Sources and reading notes
14. Evidence map and calculation assumptions

How to read this lesson

Begin with the main text. The shaded physics boxes add optional detail and worked estimates. Figures are conceptual and not to scale. Numbered links lead to the supporting sources.

What you will learn

- Explain why purity and a single-crystal structure are separate requirements.
- Follow silicon through reduction, chlorosilane purification and deposition.
- Identify how CZ growth controls dislocations, dopant distribution, oxygen and point defects.
- Compare CZ, MCZ and float-zone growth, including the purpose of neutron transmutation doping.

1. Two transformations before chipmaking can begin

Silicon is abundant. Semiconductor-grade starting material is not. The challenge is to remove unwanted atoms, then arrange the silicon atoms into one continuous crystal.

The route explored here has two main stages. **Purification** converts quartz into metallurgical silicon, turns that silicon into a volatile chemical compound, separates impurities by distillation and deposits high-purity solid silicon. **Crystal growth** melts the purified material and freezes it onto a seed with the required crystal orientation.

One result is a large cylindrical **single-crystal ingot**, from which wafers will later be sliced. Making a pure solid and making a single crystal are distinct achievements: a very pure lump can still contain many differently oriented crystal grains.

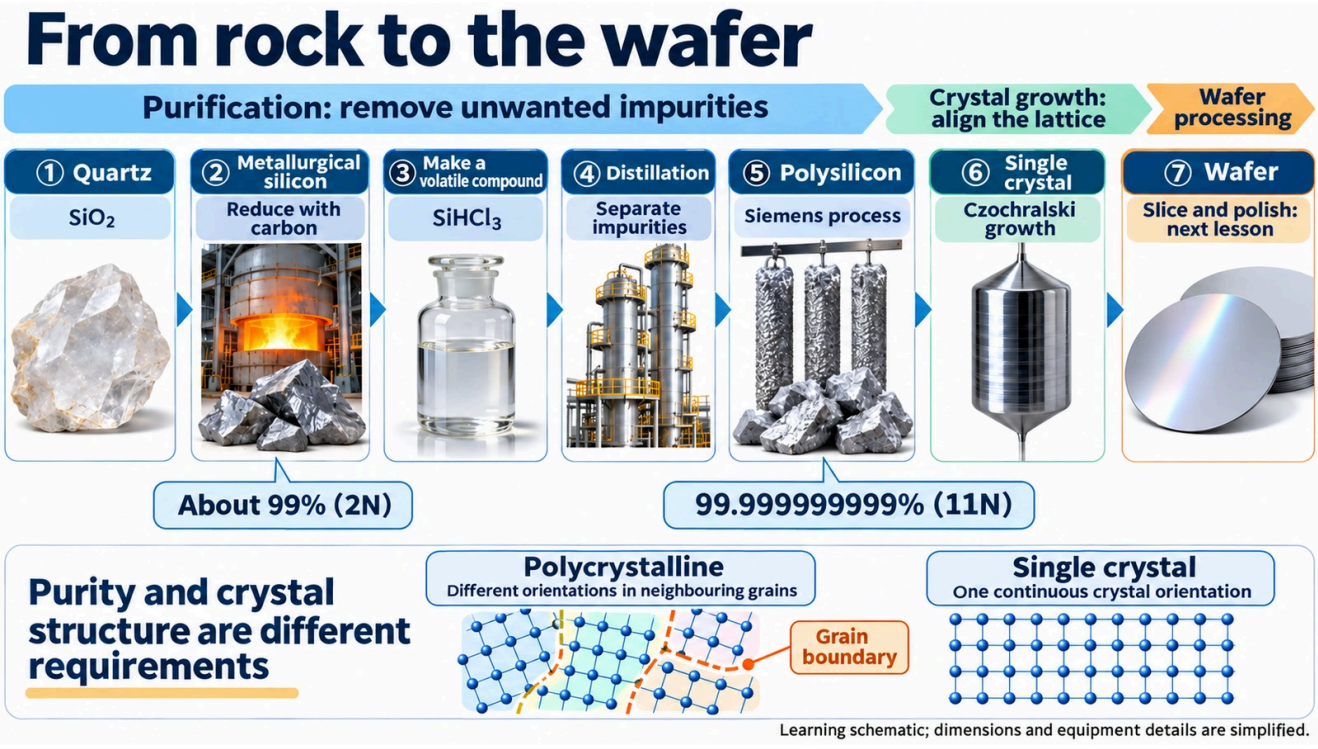

Figure 1 **From quartz to wafers: purity and crystal structure** (schematic)

Read Figure 1 as two different jobs

The early steps remove unwanted impurities. Crystal growth then establishes a continuous lattice orientation. The **11N** label describes a high-purity polysilicon product specification, not a promise that every subsequently grown crystal has the same total impurity level: dopants are deliberately added, and CZ growth also introduces oxygen.

2. Why purity alone is not enough

A semiconductor's electrical behaviour can change dramatically when a tiny fraction of its atoms are replaced by dopants. That is useful when the addition is deliberate. Uncontrolled contamination makes the material much harder to design around.

Consider a simplified atom-counting example. Silicon contains about **5×10^{22} atoms per cm^3**. If a nominal purity is treated as an atomic fraction for this illustration, the difference is enormous:

Illustrative purity	Remaining impurity fraction	Approximate impurity atoms per cm^3
99%, or 2N	10^{-2}	5×10^{20}
99.999999999%, or 11N	10^{-11}	5×10^{11}

Against a deliberate dopant concentration of 10^{16} atoms/cm^3, the first figure is about 50,000 times larger; the second is about 20,000 times smaller. These are teaching calculations. Real purity specifications identify particular contaminants, measurement limits and concentration conventions; the table is not a substitute for a supplier specification.

The crystal structure matters too. **Polycrystalline silicon**, or polysilicon, contains grains with different orientations. At a **grain boundary**, the otherwise regular atomic arrangement is interrupted. Such boundaries can impede charge transport and introduce sites that trap electrons or holes.

A **single crystal** has a continuous lattice orientation throughout the material. That provides a much more consistent foundation for manufacturing large numbers of devices with tightly controlled properties. Purification controls which atoms are present; crystal growth controls how the lattice is organised.

3. First purification step: take the oxygen out of quartz



Quartzite consists largely of quartz, SiO_2, and illustrates the type of silica-rich rock used as a silicon raw material. The photographed pieces were registered as decorative or building material, rather than as a documented semiconductor feedstock.
Image: King Waltz / Wikimedia Commons / CC BY-SA 4.0.

The Royal Society of Chemistry gives silicon as **27.7% of Earth's crust by mass**, second only to oxygen. [5] Its abundance does not mean that ready-to-use elemental silicon can simply be dug out of the ground. It is commonly bound to oxygen in silica and silicate minerals.

The industrial route begins with silica-rich raw material and carbon in a **submerged-arc furnace**. Electrodes extend into the charge, supplying the heat needed for reduction. USGS describes silicon metal production alongside ferrosilicon production in this type of furnace. The US Department of Energy gives an electricity requirement of roughly **10–15 MWh per tonne of metallurgical silicon**. [1][2]

A simplified overall reaction is:

$$SiO_2 + 2C \rightarrow Si + 2CO$$

This balances the atoms when carbon monoxide is the gaseous product. The real furnace chemistry involves several reactions, including intermediates such as SiO and SiC. Educational material gives furnace temperatures of roughly 1,500–2,000 °C, while the IEA report describes about 1,700 °C. These figures describe the process broadly, rather than one uniform temperature throughout a furnace. [3][4]

The result is **metallurgical-grade silicon (MG-Si)**, with a purity of about **99%** in the cited DOE description. Much silicon metal goes into aluminium alloys and chemical products such as silicones; semiconductor and solar applications use only part of the output. [1][2]

Why does reduction need such a high temperature?

Silica is chemically stable. Mixing quartz with carbon at room temperature does not readily release useful silicon.

Thermodynamics expresses the driving force through **$\Delta G = \Delta H - T\Delta S$**. ΔH is the enthalpy change, ΔS the entropy change and T the absolute temperature. A negative ΔG favours the forward reaction under the specified conditions; how quickly it proceeds is a separate question of reaction kinetics.

Removing oxygen from silica has a large energy cost. Producing gaseous CO increases entropy, so raising the temperature makes the $-T\Delta S$ contribution more favourable for this reduction. Removing CO from the reaction region also helps. This explains the need for a hot industrial process without implying that every reaction producing a gas is automatically favourable at high temperature.



Silicon can have a metallic lustre even though its electronic behaviour is that of a semiconductor. The Commons description identifies this as a piece of purified silicon; its appearance alone does not establish a semiconductor-grade purity specification.
Image: Enricoros at English Wikipedia / Wikimedia Commons / public domain.

4. Second step: make a compound that can be distilled

Trying to remove trace contaminants one atom at a time from solid silicon would be extremely difficult. The chemical route takes a more useful detour: convert silicon into a **volatile compound**, purify that compound and then recover solid silicon.

In the route described by PVEducation, powdered metallurgical silicon reacts with hydrogen chloride at about **300 °C** to produce **trichlorosilane**, $SiHCl_3$. [3]

$$Si + 3HCl \rightarrow SiHCl_3 + H_2$$

Trichlorosilane has a normal boiling point of about **31.8 °C**. Other elements in the feed can form different chlorides. Their different volatilities make separation possible. The cited educational description reports electrically active impurities below one part per billion atoms after purification. [3][6]

What distillation actually separates

Molecules can escape from a liquid into the vapour above it. At a given temperature, different substances have different vapour pressures. The normal boiling point is the temperature at which a pure substance's vapour pressure reaches one atmosphere.

Compound	Normal boiling point	Why it matters here
Boron trichloride, BCl_3	12.5 °C	Boron is a p-type dopant
Trichlorosilane, $SiHCl_3$	31.8 °C	The silicon-containing intermediate being purified
Silicon tetrachloride, $SiCl_4$	57.6 °C	Another silicon chloride in the process chemistry
Phosphorus trichloride, PCl_3	About 76 °C	Phosphorus is an n-type dopant

The boiling-point values come from the cited PubChem records. [6–9] Adjacent values differ by roughly 20 °C, but a boiling-point list alone does not specify an industrial separation design.

In a distillation column, repeated contact between rising vapour and descending liquid progressively enriches more volatile components towards the top and less volatile components towards the bottom. Industrial purification uses a sequence of separation stages, rather than one magical tray that removes every contaminant.

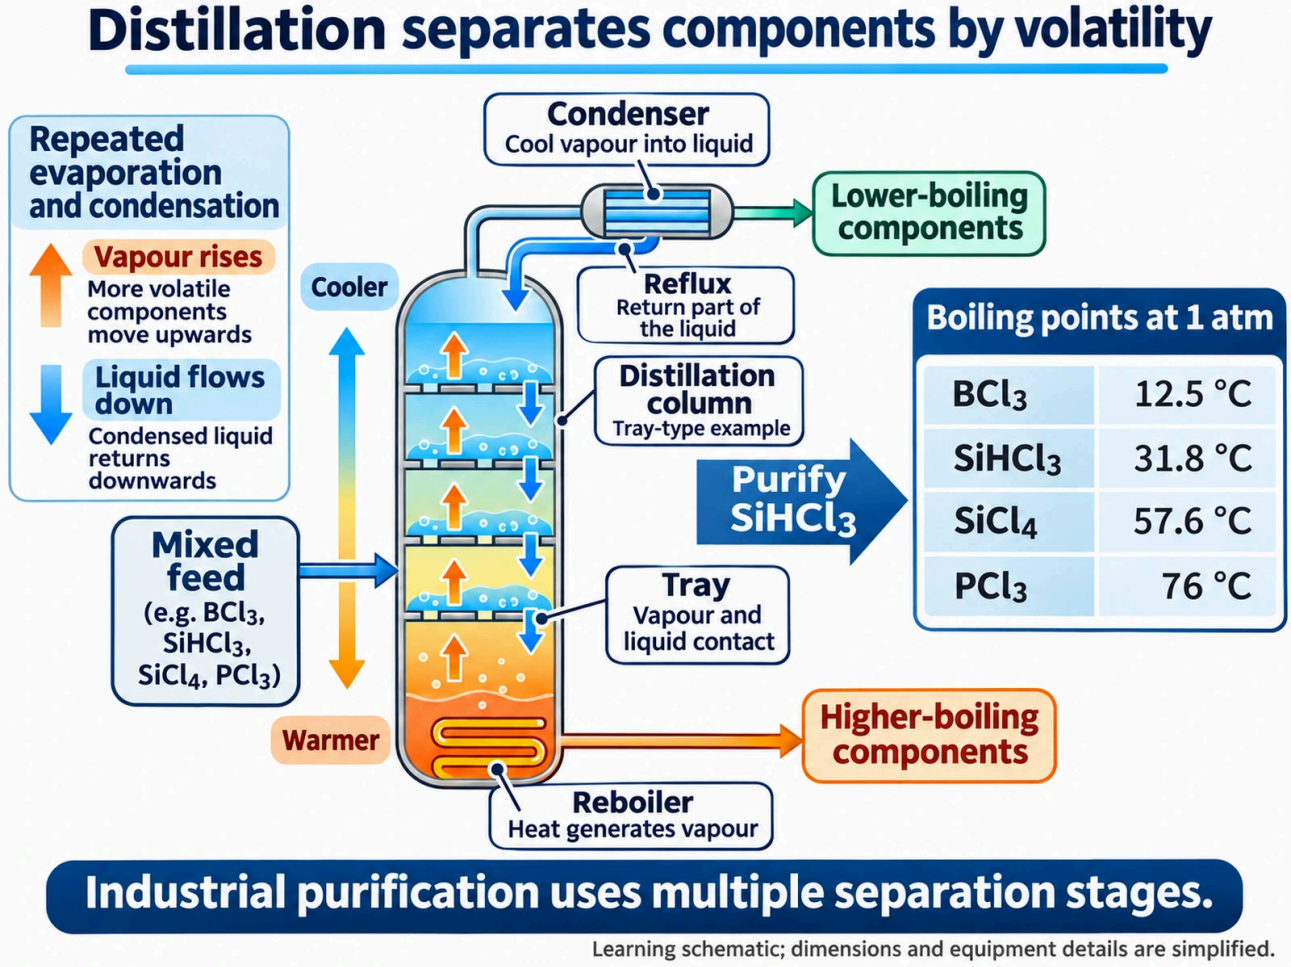

Figure 2 **Distillation uses differences in volatility** (schematic)

The liquid is a silicon compound, not molten silicon

The process does not melt elemental silicon and boil impurities out of it. It first changes silicon chemically into trichlorosilane, which can be handled as a volatile fluid. “Higher-boiling” and “lower-boiling” describe volatility under stated conditions, not simply molecular weight.

5. Third step: deposit polysilicon by the Siemens process

The purified intermediate must now be converted back into solid silicon. In the **Siemens process**, slender silicon starter rods are arranged inside a bell-shaped reactor. Passing current through the rods heats them, and silicon deposits from the process gas onto their surfaces. The rods become thicker as deposition continues. [2][3]

A simplified reaction for the trichlorosilane/hydrogen route is:

$$SiHCl_3 + H_2 \rightarrow Si + 3HCl$$

It is the reverse of the simplified chlorination reaction above. Hydrogen chloride and other process gases can be recovered and reused. This is an example of **chemical vapour deposition (CVD)**: a gas-phase chemical process leaves a solid deposit on a hot surface.

PVEducation gives an illustrative operating description of about **1,100 °C**, **200–300 hours** and rods reaching **150–200 mm diameter**. Actual reactor conditions depend on the technology and supplier. The DOE description explains that the resulting rods are broken into chunks for subsequent processing and packed under an inert atmosphere. [2] [3]



A small section of a high-purity polysilicon rod, approximately 10 cm long and 4 cm in diameter. The rough surface reflects the crystalline grains deposited during growth. Polysilicon becomes feedstock for single-crystal growth.
Image: Warut Roonguthai / Wikimedia Commons / CC BY-SA 3.0.

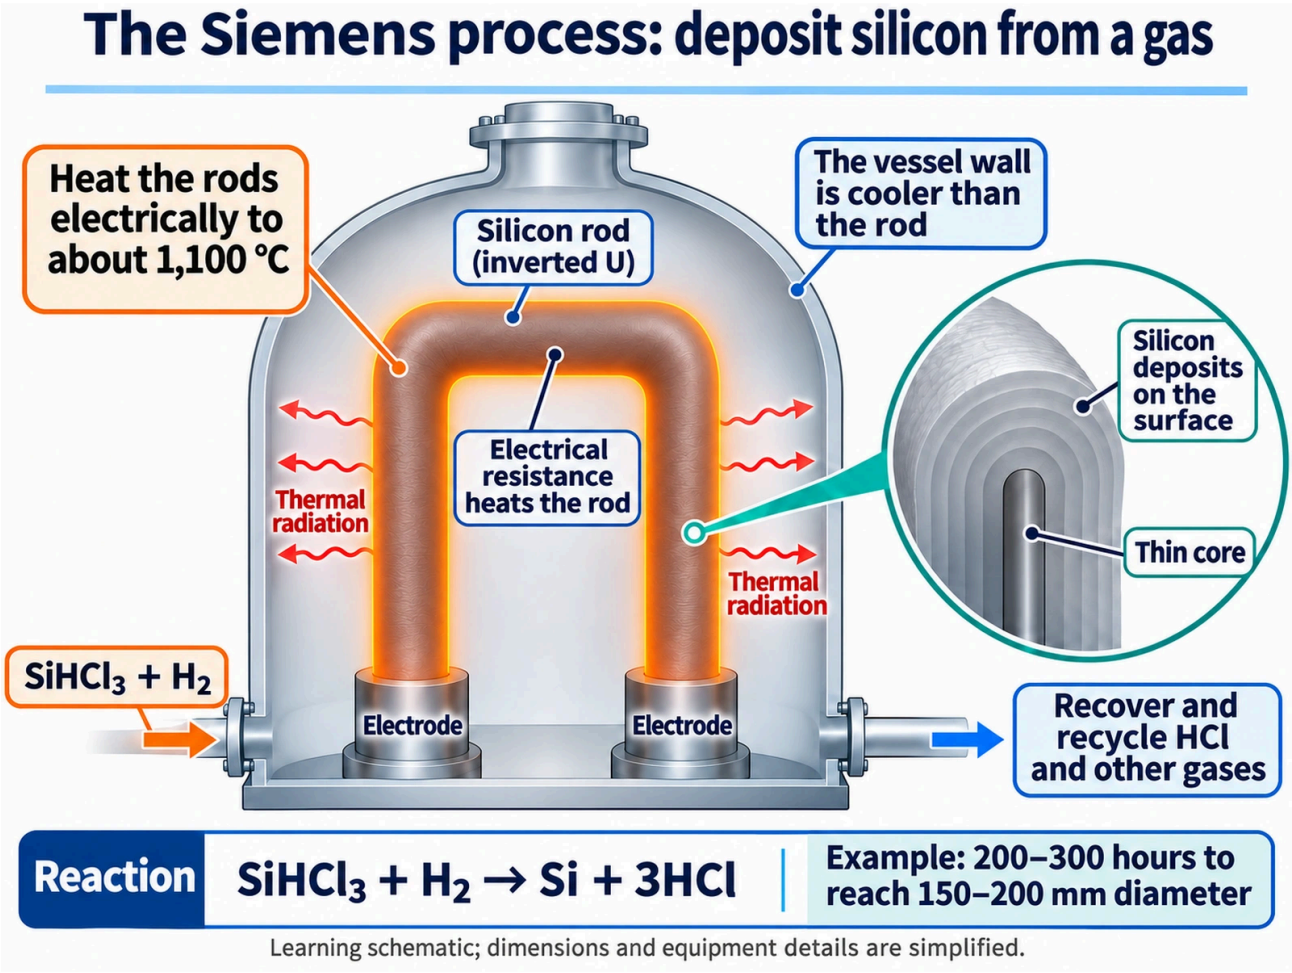

Figure 3 **Inside a Siemens reactor** (schematic)

The rod grows by deposition, not by freezing a melt

Silicon accumulates on the heated rod's surface. The thickened product is still polycrystalline. Producing a single crystal requires another stage: melting this purified feedstock and growing it from a seed.

Why does this process use so much electricity?

The IEA's 2022 report describes a reduction in Siemens-process electricity intensity from over **70 kWh/kg in 2011** to about **50 kWh/kg in 2021**. [4] For comparison, the cited metallurgical-silicon estimate is 10–15 kWh/kg. [2] These are dated process-level figures, not a guaranteed consumption for a particular modern plant.

One reason to examine the thermal design is **radiation**. A hot surface emits thermal radiation with an ideal blackbody flux of σT^4, where $\sigma \approx 5.67 \times 10^{-8}$ W/(m^2·K^4).

At 1,100 °C, or 1,373 K, this gives about **20 W per cm^2** of blackbody surface. Real silicon has an emissivity below one, and net heat loss also depends on its surroundings. The calculation is an illustration, not a measurement of reactor losses. Nevertheless, keeping hot rods inside a cooler enclosure for hundreds of hours makes radiative heat loss a plausible contributor to electricity demand.

Using the example radii and times above gives an approximate radial build-up rate of **0.25–0.5 mm per hour**. This ignores the starter rod radius and treats growth as uniform; it conveys the scale of the process rather than predicting a production cycle.

What "eleven nines" means

Tokuyama describes its polysilicon purity as **99.999999999%**, conventionally called **11N** or "eleven nines". SUMCO describes its wafer feedstock in terms of metallic impurities at the ppb level or below. Those statements refer to particular purity descriptions, so their numbers should not be treated as interchangeable specifications for all contaminants. [10][15]

Wacker's annual report describes its ultra-pure polysilicon in exceptional purity terms and presents its position in the semiconductor market. The important engineering point is that semiconductor applications require extremely tight control of trace contaminants, not merely a shiny-looking solid. [12]

The process is named after Siemens. The cited sources place its development in the 1950s or 1960s, depending on how they describe the history. The IEA's 2022 report gave the Siemens route more than 90% of polysilicon production. That is a historical figure from the report, not a newly measured share for 2026. [2–4]

An alternative: the fluidised-bed reactor

A **fluidised-bed reactor (FBR)** suspends silicon particles in an upward gas flow and deposits more silicon onto them, producing granular polysilicon. The DOE's 2022 report describes a silane/hydrogen route, electricity consumption roughly half that of the Siemens process, and a market share of 3–5% at the time. The IEA report also discusses lower energy use and a purity trade-off. Those shares and comparisons belong to the cited reports' period. [2][4]

The same DOE report gives a broad solar-grade purity range of 8N–11N. Purity alone therefore does not neatly divide all material into two fixed “solar” and “semiconductor” grades; the detailed impurity specification and application requirements matter. [2]

6. Czochralski growth: freeze silicon onto a seed



Jan Czochralski, the Polish chemist whose name is attached to the CZ method, in a photograph from around 1910.
Image: unknown photographer / Wikimedia Commons / public domain.

The **Czochralski method**, usually shortened to **CZ**, turns purified polysilicon into a single crystal. Its essential idea is simple: contact molten material with a seed of the desired crystal orientation, then control freezing as the crystal is withdrawn.

Czochralski developed the method in 1916 while working at AEG in Berlin. It originally investigated metal crystallisation. Gordon Teal and John Little at Bell Laboratories later adapted crystal pulling to semiconductor materials, reporting germanium single-crystal growth in 1950. [31][32]

The basic sequence

SUMCO describes three main steps. First, melt high-purity polysilicon in a **quartz crucible**, with carefully chosen dopant additions such as boron or phosphorus. Next, touch a **seed crystal** to the melt. Finally, rotate and withdraw the seed so that the growing crystal inherits its lattice orientation. SUMCO gives a melt temperature of about **1,420 °C**. [15]

Silicon’s melting point is about **1,414 °C**. The difference between a quoted melting point and an operating temperature is useful: the furnace must maintain and control a melt, not merely sit at a single handbook number. [18]

The crystal and crucible may rotate in opposite directions. GlobalWafers describes one clockwise/anticlockwise arrangement; the diagram uses opposite-direction arrows without turning a viewpoint convention into a universal operating rule. [17]



Two unused crucibles, with pieces of polysilicon feedstock at the bottom. The crucible contains the melt during growth, but also becomes a source of oxygen, as discussed below.
Image: Twisp / Wikimedia Commons / public domain.

An ordinary element can depend on an extraordinary raw material

High-purity quartz for crucibles has an important supply connection to **Spruce Pine, North Carolina**. Reports following Hurricane Helene in 2024 described the roles of Sibelco and The Quartz Corp. One widely reported estimate, attributed to the author of *The World in a Grain*, linked 70–90% of the world's crucibles to quartz from the area. That is an attributed estimate, not an independently established share in this article. [38][39]

The reporting also described crucible lifetimes of roughly 300–400 hours and the disruption after the September 2024 storm. Sibelco resumed production and shipments in October. These are historical reports, not a statement about current operating conditions. Lesson 2-3 returns to wafer supply chains. [38][39]

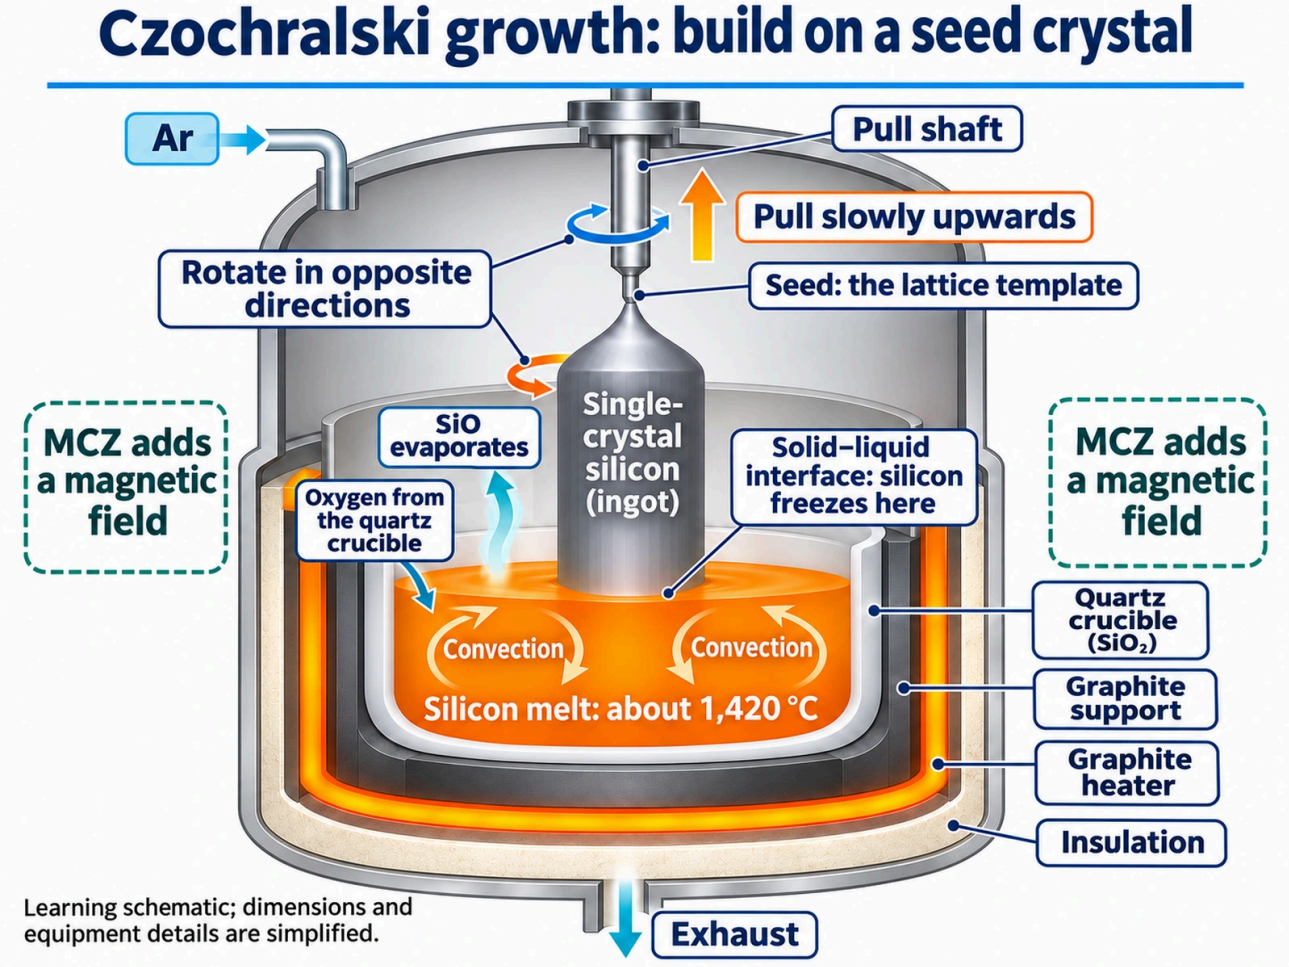

Figure 4 **Inside a CZ crystal puller** (schematic)

The growing end remains connected to the melt

The ingot moves upwards, but new silicon solidifies at its lower end, at the **solid–liquid interface**. The seed provides the lattice template. Heat flow and pull rate must be coordinated to maintain the required diameter and crystal quality.



A Leybold Z141 crystal-pulling installation from 1956, displayed at the Deutsches Museum in Munich. It is a historical machine, not a drawing of current 300 mm production equipment.
Image: Journey234 / Wikimedia Commons / public domain.

- **The ingot has a neck, shoulders, a body and a tail**

At the top is the seed, followed by a deliberately narrow **neck**. The **shoulder** widens towards the target diameter. The constant-diameter **body** supplies the useful wafer material. At the end of growth, a tapered **tail** helps avoid an abrupt separation from the melt and the associated thermal stress.



A silicon single crystal displayed in a Munich mineral museum. The tapered shoulder and narrow neck show that an ingot is more than a simple straight cylinder.
Image: ArticCynda / Wikimedia Commons / CC0.

Siltronic describes 300 mm ingots weighing several hundred kilograms and taking **five to seven days** to grow. A Bonn University museum example is 155 cm tall, 300 mm in diameter and 216 kg in mass. These are examples of scale, not standard dimensions for every product. [19][35]



The exhibit panel describes an ingot weighing about 122 kg and producing wafers of 300 mm diameter. The figure differs from the Bonn example because these are different crystals.
Image: Oleg Alexandrov / Wikimedia Commons / CC BY-SA 3.0.

Why can the crystal not simply be pulled faster?

Freezing releases **latent heat**. The interface can keep solidifying only if heat is removed at the required rate.

A simplified upper-limit estimate, retaining heat conduction through the solid, is:

$$v_{max} \approx k_s G_s / (\rho L)$$

Here k_s is the crystal's thermal conductivity, G_s its temperature gradient near the interface, ρ its density and L the latent heat per unit mass. The model omits other heat-flow contributions and detailed furnace geometry. Its message is that **growth is constrained by heat removal**.

The cited Stanford teaching notes estimate a theoretical maximum of about **35 cm/hour** for a six-inch crystal under their assumptions, with practical rates about half that or less. A 300 mm process has its own thermal constraints; the smaller-crystal estimate should not be used as a production recipe for it. [20]

7. Four pieces of physics that determine crystal quality

- **1. Necking: give dislocations a way out**

When a seed first contacts the melt, thermal shock can generate **dislocations**: line defects in the atomic arrangement. If they continue through the growing material, they compromise the crystal quality required for device manufacturing.

In 1959, W. C. Dash reported a method of growing dislocation-free silicon. **Dash necking** starts growth with a very narrow section. The cited patent example describes a neck 2–4 mm in diameter and pull speeds up to 6 mm/minute during that stage. Those are example process conditions, not a recommendation for all pullers. [21–23][33]

Why does making the neck thin help?

In silicon, dislocations can move on particular crystallographic planes, notably the {111} family. For a common ⟨100⟩ pulling orientation, those planes are inclined to the pulling direction.

Existing dislocations can therefore reach the side surface rather than continuing indefinitely along the crystal. A narrow neck shortens that route and helps limit the strain that drives further dislocation multiplication. Once a dislocation-free section has been established, the crystal can be widened into the shoulder and body. This is the principle illustrated by the red lines terminating at the neck surface. [22]

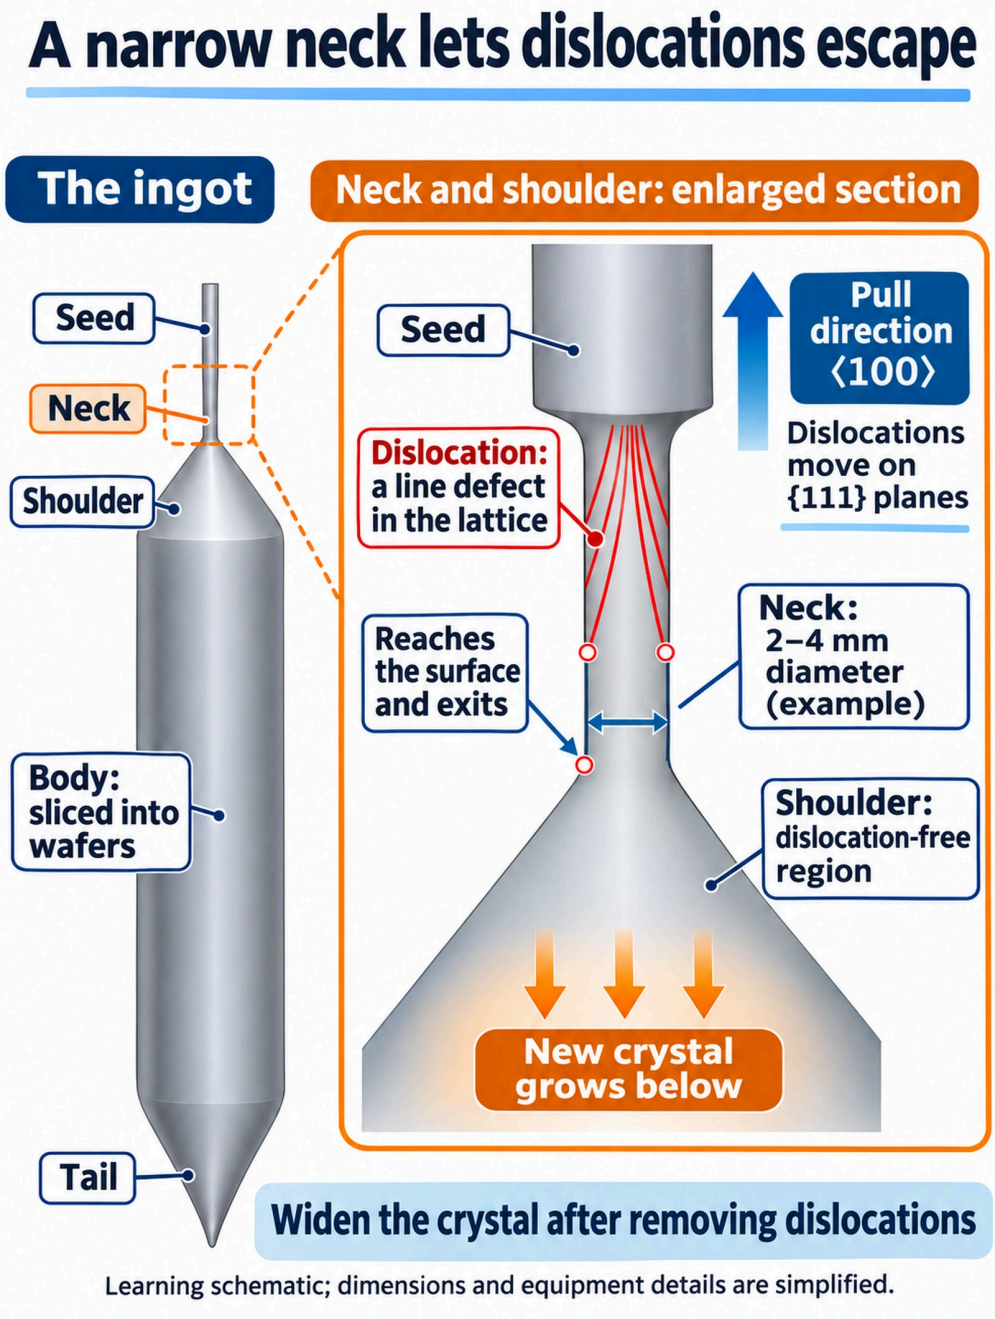

Figure 5 **How a narrow neck removes dislocations** (schematic)

Pulling upwards and growing below are compatible

The finished crystal is withdrawn upwards. Later material freezes onto its lower end. In the enlarged schematic, the narrow neck is formed before the wider shoulder below it. The picture does not show dislocations being mechanically scraped off by a tool.



A single-crystal silicon rod used as the seed for crystal pulling. The seed supplies the crystallographic orientation inherited by the newly solidifying material.
Image: Warut Roonguthai / Wikimedia Commons / CC BY-SA 4.0.

2. Segregation: the melt and crystal do not share impurities equally

When material freezes, an impurity need not enter the solid at the same concentration as in the liquid. The **segregation coefficient**, or partition coefficient, is defined as:

$$k = C_s / C_l$$

C_s and C_l are the concentrations in the solid and liquid at the interface. The Japanese edition uses the following representative values from Trumbore's 1960 compilation. They are reference values for this explanation; actual growth also depends on transport and process conditions. [23]

Element	Role in this discussion	Representative k
Boron, B	p-type dopant	0.80
Phosphorus, P	n-type dopant	0.35
Arsenic, As	n-type dopant	0.3
Antimony, Sb	n-type dopant	0.023
Oxygen, O	Introduced from the crucible in CZ growth	0.5
Copper, Cu	Contaminant	4×10^{-4}
Gold, Au	Contaminant	2.5×10^{-5}
Iron, Fe	Contaminant	8×10^{-6}

For **k < 1**, the impurity favours the liquid. As more silicon freezes, rejected impurity accumulates in the remaining melt. Later-grown crystal can therefore have a higher dopant concentration than the first-grown material. [20]

A worked example with the Scheil model

Assume a fully mixed liquid, negligible diffusion in the solid, a constant segregation coefficient and the model's idealised interface conditions. The **Scheil equation** gives:

$$C_s(f) = kC_0(1 - f)^{k-1}$$

C_0 is the initial melt concentration and f the fraction already solidified. At the start, $C_s(0) = kC_0$. Dividing by that initial crystal concentration gives the following ratios:

Fraction solidified, f	Boron, k = 0.8	Phosphorus, k = 0.35	Antimony, k = 0.023
0	1.00	1.00	1.00
0.5	1.15	1.57	1.97
0.8	1.38	2.85	4.82
0.9	1.58	4.47	9.48

These values are calculations using the stated model, not measurements along a particular production ingot. The simple model should not be extrapolated unquestioningly to the very last drop of melt at f = 1. [20]

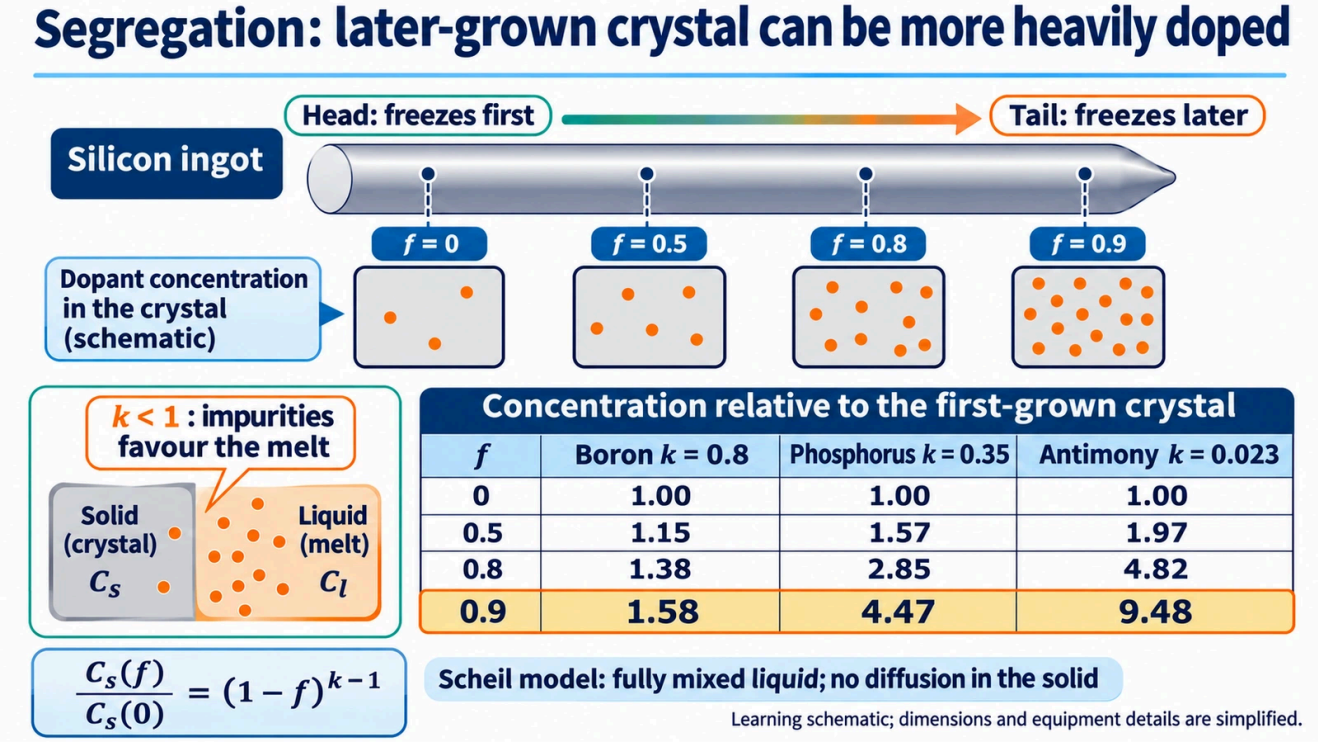

Figure 6 **Segregation changes concentration along an ingot** (schematic)

What does “one times” mean in Figure 6?

The reference is **the first-grown crystal**, $C_s(0)$, not the initial liquid concentration C_0. The dots illustrate a trend; their count is not a measured concentration. A coefficient close to one gives a flatter profile, while a smaller coefficient can produce a much stronger head-to-tail change.

This helps explain why resistivity varies along an ingot and why only the sections meeting a specified range may be selected for a particular product.

Segregation can also help reject contamination. At the representative iron value $k = 8 \times 10^{-6}$, the instantaneous solid concentration is roughly one hundred-thousandth of the adjacent liquid concentration. That is an illustrative interfacial ratio, not a guarantee that the whole ingot has been purified by the same factor.

3. Oxygen: contamination that can also be useful

CZ silicon grows in contact with a quartz crucible. Some oxygen from that crucible enters the melt and is incorporated into the crystal. Some silicon monoxide, SiO, escapes from the melt surface. Oxygen transport is therefore linked to crucible dissolution, melt flow and evaporation, not just to the purity of the starting polysilicon. [17][22][34]

The cited teaching notes give a typical CZ oxygen concentration of roughly **10^{18} atoms/cm^3**. Dividing by 5×10^{22} silicon atoms/cm^3 gives about **20 parts per million atoms (ppma)**, or one in 50,000. GlobalWafers’ cited examples of 11–17 ppma are of the same order. [17][20]

Oxygen is not simply “good” or “bad”. Controlled oxygen can strengthen the wafer against dislocation motion during later processing. Suitable heat treatments can also produce oxygen-related precipitates in the wafer interior that help trap metallic contamination. This is **internal gettering**: moving contaminants away from the near-surface region where devices will be made. GlobalWafers dates the proposal of this use of oxygen precipitates to 1977. [17][20]

The desired oxygen level and precipitation behaviour depend on the product and its later thermal history. A process that benefits one device application may be unsuitable for another.

4. Point defects: balance pull rate against the temperature gradient

Two basic **point defects** occur in a crystal. A **vacancy** is a missing atom at a normal lattice site. A **self-interstitial** is an extra silicon atom occupying a position between the normal sites.

During cooling, these defects can aggregate under suitable conditions. Vacancies can form voids; interstitials can form dislocation loops. A void intersecting a wafer surface can appear as a crystal-originated particle or pit, commonly abbreviated **COP** in this context. [22]

Voronkov’s 1982 work linked the balance between vacancy and interstitial populations to **V/G**: growth or pulling velocity V divided by the temperature gradient G near the solid–liquid interface. [37]

Relative V/G	Tendency in the simplified picture	Possible aggregated defects
Lower	Interstitial-dominated	Dislocation loops
Near a suitable transition range	Aggregation can be suppressed	A low density of the targeted grown-in defects
Higher	Vacancy-dominated	Voids and surface COPs

The cited patent gives an example critical ratio of about **2.1×10^{-5} cm^2/(s·K)**. That is approximately **0.13 mm^2/(min·K)**. With an illustrative temperature gradient of 8 K/mm, it corresponds to a pull rate of roughly 1 mm/minute. The gradient is an assumption for the calculation, not a disclosed production setting. [25]

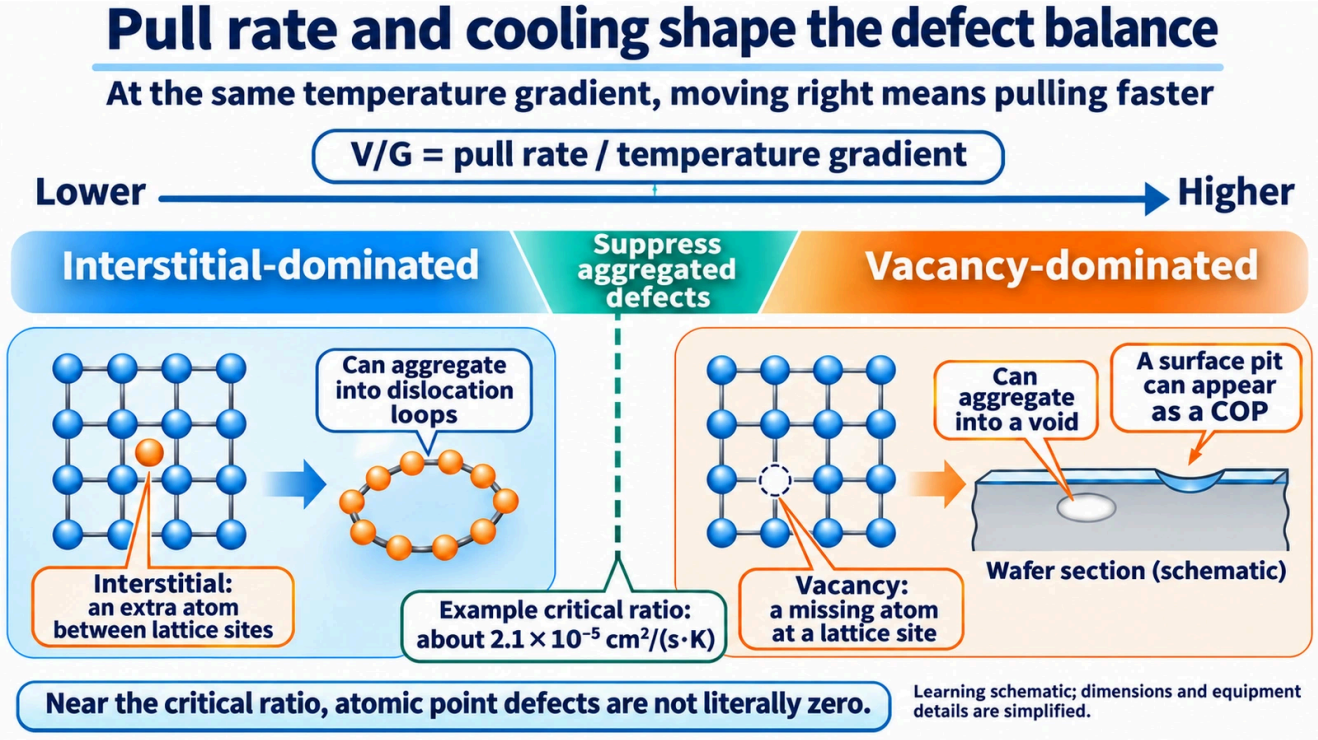

Figure 7 **V/G and the balance between point-defect populations** (schematic)

Low grown-in defect density does not mean no atomic defects

The target is to suppress harmful aggregates such as voids and loops. The transition does not make every atomic point defect vanish. The quoted critical value is an example from a particular source; definitions, thermal history and process conditions matter.

Controlling this balance across an entire large-diameter crystal is difficult because the thermal field can differ between centre and edge. GlobalWafers markets “Perfect Silicon” in terms of the specified absence of COPs and related grown-in defects; the product name should be read in that engineering context. [26]

- **MCZ: use a magnetic field to influence the melt**

The melt moves because of buoyancy-driven convection and the rotation of the growing system. That flow transports heat, dopants and oxygen. **Magnetic Czochralski growth (MCZ)** adds an applied magnetic field to influence the electrically conducting melt. [15][22]

How can a magnet brake a liquid?

Liquid silicon conducts electricity. Motion through a magnetic field can induce electric currents, and the interaction of those currents with the field produces a **Lorentz force**. In an appropriate field arrangement, this damps aspects of the flow rather than simply making the entire melt motionless.

The cited GlobalWafers patent discusses a horizontal field opposing melt motion. The IAEA's 2012 report gives a broad field range of about **50–4,000 gauss** for magnetic-field-assisted CZ processing. These examples explain the principle; field geometry and strength are chosen for the particular growth process. [24][27]

8. Float-zone growth: remove the crucible from the problem



The bright molten region, narrow neck and beginning of the growing crystal show the start of a float-zone process.
Image: Matthias Renner / Wikimedia Commons / CC BY-SA 3.0.

If the crucible is an oxygen source, one solution is to avoid using a crucible. In **float-zone growth (FZ)**, a radio-frequency induction coil heats a small region of a vertical polysilicon feed rod. The molten zone connects the feed to the single crystal being grown from a seed below. As the process advances, feed material melts and then solidifies with the seed's orientation. [16]

The molten zone contacts a controlled atmosphere or vacuum rather than a quartz container. This makes very low oxygen concentrations possible, but it does not mean that every source of oxygen or contamination has literally disappeared. The suspended melt must also remain stable. Surface tension and the electromagnetic/thermal conditions constrain the available geometry; the cited patent describes typical diameters up to about **200 mm**. [27] [28]

FZ material is used where its properties are valuable, including power devices such as **IGBTs**, which handle substantial voltage and current, and some detector applications. The IAEA's 2012 report described CZ material as more than 90% of the silicon market. Treat that percentage as the report's historical description. [16][27]



Figure 8 **Float-zone growth and neutron transmutation doping** (schematic)

Crystal growth and NTD are different operations

The upper panel shows feed material melting and freezing into a single crystal. The lower panel shows a subsequent way of introducing a dopant into a grown crystal. The neutron source is not part of the float-zone growth apparatus.

- **NTD: create phosphorus inside the crystal**

Some FZ silicon is doped by a nuclear transformation rather than by adding dopant to the melt. **Neutron transmutation doping (NTD)** uses the naturally occurring isotope silicon-30. It captures a neutron to form silicon-31, which then undergoes beta decay to stable phosphorus-31. [27]

$$^{30}\mathrm{Si} + \mathrm{n} \rightarrow {}^{31}\mathrm{Si} \rightarrow {}^{31}\mathrm{P}$$

The half-life of **silicon-31** is about **2.62 hours**. After 24 hours, the simple decay calculation leaves about **0.2% of the initial silicon-31 population**. This number concerns that isotope alone; it is not a rule for when an irradiated ingot may be released or used. [27]

Neutrons penetrate the bulk crystal, and silicon-30 is distributed throughout the starting material. With controlled irradiation conditions, the resulting phosphorus can be distributed very uniformly. Subsequent processing includes annealing irradiation-induced lattice damage and confirming the required material properties. This is useful for power-device material whose resistivity must be tightly controlled. The IAEA's 2012 report gave NTD material about 20% of the FZ market at the time. [27]

Feature	CZ	Float zone
Melt container	Quartz crucible	No crucible
Oxygen	Commonly of order 10^{18} atoms/cm^3 in the cited example	Very low oxygen is achievable
Diameter in this overview	300 mm is a major production format	Generally up to about 200 mm in the cited description
Typical applications	Logic, memory and many other mainstream devices	Selected power devices and detectors
Main trade-off highlighted here	Scalable diameter, with crucible-related oxygen to manage	Low oxygen, with molten-zone and diameter constraints

9. The industry behind the material

- **Wafers: large volumes, a concentrated supplier base**

SEMI reported **12,973 million square inches** of silicon-wafer shipments in **2025**, up 5.8%, with revenue of **US$11.4 billion**, down 1.2%. [29] If all that area were converted into equivalent 300 mm wafers, each about 110 square inches, it would represent roughly **120 million wafers**. That is an area-equivalent calculation; the real total includes smaller diameters.

Siltronic's September 2026 presentation estimated that the five leading suppliers—Shin-Etsu, SUMCO, GlobalWafers, Siltronic and SK Siltron—accounted for roughly **75% of the market**, using January–September 2025 revenue data. The number is the company's estimate under its chosen definition, not a new survey conducted for this article. The two largest firms in that presentation are Japanese. Lesson 2-3 explores the supplier landscape in more detail. [30]

A larger **450 mm** wafer format was developed as a possible next step. Industry reporting described the G450C consortium scaling back its activities at the end of 2016. This history does not establish a new timetable for a transition to 450 mm; this lesson makes no such prediction. [36]

- **Polysilicon: solar dominates volume, electronics sets exacting requirements**

The DOE's 2022 report described more than 80% of polysilicon demand as solar-related. The IEA reported China holding 79% of polysilicon production capacity in 2021. Those are historical figures with specific product definitions. [2][4]

USGS's 2026 silicon summary gives China roughly 80% of **2025 silicon-material production** and notes the addition of silicon to the US critical minerals list in 2025. Silicon materials in that commodity context should not be equated automatically with electronic-grade polysilicon. [14]

Only a limited group of suppliers meet the demands of semiconductor polysilicon. Wacker describes itself as the leading supplier in that segment, based on its own analysis. Hemlock describes itself as one of five producers supplying polysilicon for semiconductor wafers. Tokuyama records entering the polysilicon business in 1984. These are attributed company statements rather than a new independent ranking. [11–13]

10. Connect the chemistry to the crystal

Stage	Main task	Physical or chemical idea
Carbothermic reduction	Recover silicon from silica	High-temperature reaction chemistry and CO formation
Chlorination and distillation	Make and purify a volatile silicon compound	Volatility and repeated vapour–liquid contact
Siemens deposition	Recover high-purity solid polysilicon	CVD on electrically heated rods
CZ growth	Freeze a continuous crystal onto a seed	Heat removal, necking, segregation, oxygen and point-defect control
MCZ	Influence the conducting melt	Magnetohydrodynamic damping of flow
Float-zone growth	Grow without a crucible	A stable local molten zone and a seed-controlled lattice
NTD, where used	Form a controlled phosphorus dopant distribution	Neutron capture, beta decay and subsequent material processing

The next lesson, **2-2**, follows the ingot into wafer preparation: slicing, surface finishing, epitaxial growth and silicon-on-insulator structures. Part 2 has four lessons. Lesson **2-3** examines wafer suppliers and the demands created by AI-related applications. Lesson **2-4** turns to non-silicon substrates, including SiC, GaN and InP.

11. Check your understanding

1. What does 11N mean, and how does it compare with the metallurgical-silicon example?

▼ Show answer

11N means 99.999999999% purity in the cited product description. The metallurgical-grade example is about 99%, or 2N. Detailed impurity specifications still matter.

2. Why turn silicon into trichlorosilane instead of trying to distil solid silicon directly?

▼ Show answer

Trichlorosilane is a volatile compound with a normal boiling point of about 31.8 °C. Its different volatility from other components allows staged distillation and purification.

3. Why is a narrow neck deliberately grown below the seed?

▼ Show answer

It helps dislocations reach the side surface and terminate before the crystal is widened into the useful shoulder and body.

4. For phosphorus with k = 0.35, how does the Scheil-model concentration at f = 0.9 compare with the first-grown crystal?

▼ Show answer

It is about 4.47 times as high. The reference is Cs(0), not the initial melt concentration C_0.

5. Name two useful roles of controlled oxygen in CZ silicon.

▼ Show answer

It can improve resistance to dislocation motion during later processing and support internal gettering that traps metallic contamination away from the device region.

6. Why can FZ material have very low oxygen, and what constraint comes with the method?

▼ Show answer

The melt does not contact a quartz crucible. Maintaining a stable molten zone constrains the geometry and typical diameter, given here as up to about 200 mm in the cited description.

12. A short glossary

Annealing
Controlled heat treatment to repair damage or modify material properties.

Carbothermic reduction
High-temperature reduction of an oxide using carbon.

COP
Crystal-originated particle or pit; here, a surface feature from a grown-in void.

Crucible
A container that holds the melt during crystal growth.

Crystal grain
A region of crystal with a particular lattice orientation.

CZ
Czochralski growth: pulling a seeded crystal from a crucible-held melt.

Dash necking
Growing a narrow initial section to help eliminate dislocations.

Dislocation
A line defect in a crystal's atomic arrangement.

Dopant
An intentional impurity used to control electrical properties.

Electronic-grade polysilicon
Polysilicon meeting strict impurity specifications for electronic applications.

Enthalpy
Thermodynamic quantity used to describe heat changes at constant pressure.

Entropy
Thermodynamic measure associated with accessible microscopic states.

FBR
Fluidised-bed reactor: gas suspends particles that grow by deposition.

FZ
Float-zone growth: seeded crystal growth through a local melt, without a crucible.

G450C
Consortium established to develop 450 mm wafer manufacturing.

Gettering
Trapping contaminants away from the device-active region.

Gibbs energy
Thermodynamic potential used to assess the favoured reaction direction.

Grain boundary
The boundary between differently oriented crystal grains.

Ingot
The bulk crystal from which wafers are subsequently prepared.

Interstitial
An atom occupying a position between normal lattice sites.

Latent heat
Heat absorbed or released during a phase change.

Lorentz force
Force on charges or currents in an electromagnetic field.

MCZ
Magnetic-field-assisted Czochralski growth.

Metallurgical-grade silicon
Silicon made by reducing silica; about 99% pure in the cited example.

Neck
The narrow initial section grown below the seed in a CZ ingot.

NTD
Neutron transmutation doping: creating dopants through neutron capture and decay.

Polysilicon
Polycrystalline silicon containing multiple crystal grains.

ppma
Parts per million atoms, an atomic concentration convention.

Quartz
Crystalline silica, SiO_2; a component of silicon raw material.

RF induction coil
Coil producing an RF electromagnetic field for heating.

Scheil equation
Idealised segregation model with fully mixed liquid and negligible solid diffusion.

Seed crystal
A small single crystal that supplies the orientation for newly grown material.

Segregation coefficient
Interfacial solid-to-liquid concentration ratio for a specified impurity.

Siemens process
Silicon deposition from gas onto electrically heated starter rods.

Single crystal
Material with one continuous crystal lattice orientation.

Submerged-arc furnace
Electric furnace with electrodes extending into the raw-material charge.

Trichlorosilane
$SiHCl_3$, a volatile compound used in silicon purification and deposition.

Vacancy
An unoccupied site where an atom would normally sit in the lattice.

Vapour pressure
Equilibrium vapour pressure above a condensed substance.

V/G
Growth velocity divided by the temperature gradient near the interface.

11N
Eleven nines: 99.999999999% purity under the stated specification.

13. Sources and reading notes

Reference numbers match the Japanese edition. Dates, patent-specific examples and attributed company claims remain explicit. Japanese source titles are described in English below; linked originals retain their original language. The final two references are news reports rather than technical primary sources.

1. **US Geological Survey** — Silicon Statistics and Information — Read source
2. **US Department of Energy** — Solar Photovoltaics: Supply Chain Deep Dive Assessment, February 2022 — Read source
3. **C. Honsberg and S. Bowden, PVEducation** — Refining Silicon — Read source
4. **International Energy Agency** — Special Report on Solar PV Global Supply Chains, 2022 — Read source
5. **Royal Society of Chemistry** — Periodic Table: Silicon — Read source
6. **PubChem** — Trichlorosilane: boiling point — Read source
7. **PubChem** — Boron trichloride: boiling point — Read source
8. **PubChem** — Phosphorus trichloride: boiling point — Read source
9. **PubChem** — Silicon tetrachloride: boiling point — Read source
10. **Tokuyama** — Polysilicon product information (Japanese) — Read source
11. **Tokuyama** — Company history (Japanese) — Read source
12. **Wacker Chemie** — Annual Report 2025 — Read source
13. **Hemlock Semiconductor** — Electronics — Read source
14. **US Geological Survey** — Mineral Commodity Summaries 2026: Silicon — Read source
15. **SUMCO** — Silicon wafer manufacturing process (Japanese) — Read source
16. **Siltronic** — Products — Read source
17. **GlobalWafers** — Technology — Read source
18. **PubChem** — Silicon: element 14 — Read source
19. **Siltronic** — Extension to the crystal-pulling hall at Freiberg — Read source
20. **Stanford EE212 teaching notes, hosted by University of Washington EE528** — Chapter 3: Crystal Growth, Wafer Fabrication and Basic Properties of Silicon Wafers — Read source
21. **W. C. Dash** — Growth of Silicon Crystals Free from Dislocations, Journal of Applied Physics 30, 459, 1959 — Read source
22. **H. Föll, Kiel University** — The Science and Art of Si Crystal Growth — Read source
23. **F. A. Trumbore** — Solid Solubilities of Impurity Elements in Germanium and Silicon, Bell System Technical Journal 39, 205, 1960 — Read source
24. **GlobalWafers, US patent US11932962B2** — Horizontal magnetic field and melt-convection control — Read source
25. **MEMC Electronic Materials, US patent US6328795B2** — Critical V/G example — Read source
26. **GlobalWafers** — Products — Read source
27. **International Atomic Energy Agency** — Neutron Transmutation Doping of Silicon at Research Reactors, IAEA-TECDOC-1681, 2012 — Read source
28. **GlobalWafers, US patent US11655559B2** — Float-zone diameter constraints — Read source
29. **SEMI / PR Newswire** — Global silicon wafer shipments in 2025, February 2026 — Read source
30. **Siltronic** — Investor presentation, September 2026 — Read source
31. **IEEE Engineering and Technology History Wiki** — Milestone: Czochralski Process, 1916 — Read source
32. **Computer History Museum** — First grown-junction transistors fabricated — Read source
33. **MEMC Electronic Materials, US patent US5628823A** — Neck diameter and pulling-speed example — Read source
34. **Motorola, US patent US4545849A** — Quartz dissolution and SiO evaporation — Read source
35. **Arithmeum, University of Bonn** — Displayed 300 mm silicon crystal — Read source
36. **EE Times, republished by Design & Reuse** — No Sign of 450mm on the Horizon, January 2017 — Read source
37. **V. V. Voronkov** — The mechanism of swirl defects formation in silicon, Journal of Crystal Growth 59, 625, 1982 — Read source
38. **NBC News** — Spruce Pine and the semiconductor supply chain after Helene, 3 October 2024 — Read source
39. **Associated Press, republished by WJXT** — High-purity quartz production resumes after Helene, 10 October 2024 — Read source

14. Evidence map and calculation assumptions

Topic	Source or assumption	Interpretation
Silicon metal and its uses	USGS [1]; DOE [2]	Metallurgical silicon is distinct from electronic-grade feedstock
Reduction, energy use and chlorosilane/Siemens examples	DOE [2]; PVEducation [3]; IEA [4]	Source-specific conditions and dated energy data
Crustal abundance	RSC [5]	Mass fraction, not the availability of pure silicon
Boiling points	PubChem [6–9]	Normal boiling points; not an industrial column design
High-purity polysilicon and suppliers	Tokuyama [10][11]; Wacker [12]; Hemlock [13]; SUMCO [15]	Attributed product specifications and company statements
Impurity atom counts	5×10^{22} silicon atoms/cm³; purity treated as an atomic fraction	Simplified comparison, not a supplier assay
Thermal radiation and radial build-up	Blackbody σT^4 at 1,373 K; stated radius/time examples	About 20 W/cm² and 0.25–0.5 mm/hour; not measured plant performance
CZ sequence and magnetic/float-zone alternatives	SUMCO [15]; Siltronic [16]; GlobalWafers [17]	General process descriptions
Melting point, ingot scale and historical development	[18][19][31][32][35]	Handbook value, specific exhibits and dated historical reports
Growth-speed estimate	Stanford notes [20]	Simplified heat-removal model and six-inch example
Necking	Dash [21]; Kiel [22]; patent [33]	General principle; patent dimensions are examples
Segregation and Scheil ratios	[20][23]; fully mixed liquid and negligible solid diffusion	Ratios 1.58/4.47/9.48 at f = 0.9 refer to the first-grown crystal
Oxygen, gettering and SiO	[17][20][22][34]	Controlled oxygen can be useful; $10^{18} \div (5 \times 10^{22}) = 20$ ppma
Point defects and V/G	Voronkov [37]; patents [25]; products [26]	Aggregated defects distinguished from atomic point defects; critical value is source-specific
MCZ and float-zone constraints	[15][16][24][27][28]	Field and diameter figures are examples rather than universal limits
NTD	IAEA 2012 [27]	Isotope transformation, annealing and controlled irradiation; historic market figures
NTD decay calculation	$2^{-24/2.62}$	About 0.2% of the starting ^{31}Si remains after a day; not a material-release criterion
Wafer shipments and supplier concentration	SEMI [29]; Siltronic [30]	2025 totals and an attributed supplier estimate
Silicon-material production and 450 mm history	USGS [14]; industry reporting [36]	Keep product categories and dates distinct
Spruce Pine supply-chain example	News reports [38][39]	Attributed estimates and historical disruption, not a current operations update

English edition: 4 November 2026. Fresh English text based on the Japanese edition. Diagrams: OpenAI Image Gen. Sources, dates and photograph credits are provided above.

© 2026 Troy Technical. Original text and illustrations: all rights reserved, subject to applicable legal exceptions. Third-party photographs retain their stated licences. Links to this article are welcome.
Diagrams localised with the built-in OpenAI Image Gen tool for this English edition.