SEMICONDUCTOR FABRICATION · PART 2-2

Making a wafer fit for a chip From slicing and polishing to epitaxy and SOI

A single-crystal ingot is only the starting point. It must become a thin disc with **tightly controlled shape and surface quality**. Some products then need a new crystal layer; others need an insulating layer buried beneath the devices. Follow the processing steps and the reasons behind them.

Part 2 · Lesson 2-2 · Intermediate beginner level · Read Lesson 2-1 first · English edition: 4 November 2026

In this lesson

1. The wafer is an engineered starting material
2. Why does flatness matter so much?
3. Shape the ingot and slice it into discs
4. Round the edges, control the faces and remove damage
5. Polishing creates the final surface
6. Clean the wafer and measure the result
7. Epitaxy: grow a new layer with the required properties
8. SOI: put an insulating layer under the devices
9. Bonding and thinning also enable backside power
10. Put the processing steps together
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

- Follow an ingot through slicing, damage removal, polishing, cleaning and inspection.
- Explain why lithography needs careful control of local wafer flatness.
- Distinguish a polished wafer, an epitaxial wafer and an SOI wafer.
- Read the two-stage epitaxial growth model and the Smart Cut layer-transfer sequence.

1. The wafer is an engineered starting material

Growing a single crystal solves one problem: it gives the material a continuous crystal orientation. It does not yet provide the shape, thickness or surface needed for chip production.

The next job is to cut thin discs, remove the damage left by cutting and grinding, and create smooth, flat surfaces. Cleaning removes processing residues; inspection checks whether the result meets the required specification.

After that, different applications take different routes. A **polished wafer** can be used as the starting substrate. An **epitaxial wafer** adds a newly grown single-crystal layer whose electrical properties can be selected separately. An **SOI wafer** places a thin silicon layer above an insulator. These are different ways of engineering the material on which the devices will be made.

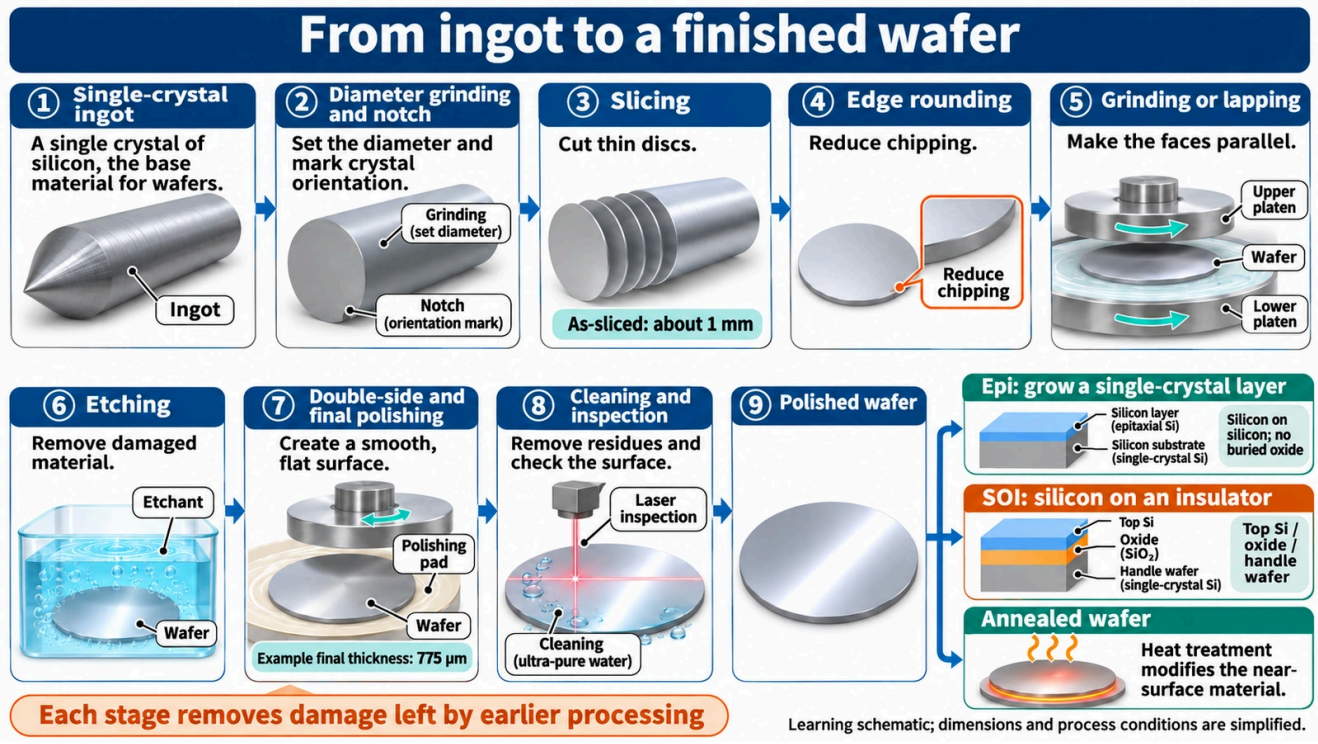

Figure 1 **From ingot to finished wafer products** (schematic)

Read the flow as a sequence of problems being solved

Slicing creates discs but also leaves damage. Grinding or lapping controls thickness and parallelism. Etching removes damaged material chemically. Polishing develops the final geometry and surface quality. Actual flows vary with the product; the diagram summarises the roles rather than prescribing one mandatory recipe.

2. Why does flatness matter so much?

A wafer can look like a mirror and still be unsuitable for a demanding lithography step. A mirror-like finish tells us little about whether the surface stays at the required height across a device area.

Lithography projects a pattern into a light-sensitive film called **photoresist**. If parts of that film sit too far above or below the intended focal plane, the printed pattern can change. The wafer therefore has to provide a controlled foundation for later layers and exposures.

Fine resolution comes with a shallow focus range

A camera lens with a wide aperture can blur the background. Lithography faces a related trade-off: collecting light over a larger angular range helps resolve fine features but reduces the useful depth of focus.

A simplified description uses two relations: [7]

$$\text{Resolution } W = k_1 \lambda / NA$$

$$\text{Depth of focus DOF} = k_2 \lambda / NA^2$$

λ is wavelength and **NA**, numerical aperture, describes the optical system's angular collection capability. The coefficients depend on the pattern, illumination, resist and process requirements. For the following estimates, take k_2 = 0.5–1.0; this is an illustrative assumption rather than a universal process window.

Example optical system	λ / NA^2	Estimated DOF for k_2 = 0.5–1.0
ArF immersion: 193 nm, NA 1.35	About 106 nm	About 53–106 nm
EUV example: 13.5 nm, NA 0.33	About 124 nm	About 62–124 nm

The optical parameters come from ASML; the DOF values are calculations using the stated assumptions. The EUV row refers specifically to the 0.33-NA example. Actual focus tolerances depend on the imaging task. [8][9]

The important feature is **NA squared in the denominator**. Increasing NA to improve resolution makes control of focus increasingly demanding.

This focus budget has to accommodate more than the starting wafer's shape. Tool levelling errors and the topography of films already on the wafer also use part of it. The numbers above therefore should not be read as an allowance for that much starting-wafer unevenness.

Measure the relevant area, not just the whole disc

An exposure system works with local fields and scans rather than exposing an entire wafer in one operation. ASML's cited immersion system supports a **26 mm × 33 mm field**. [9] Wafer geometry is likewise evaluated over defined local **sites**.

One measure is **SFQR**, the range of surface-height deviations from a least-squares reference plane fitted within a site. It captures local shape after fitting that reference plane; it is different from simply measuring the wafer's absolute height above a table. [38]

The **2007 ITRS roadmap** gave the following targets for 300 mm wafers evaluated with **26 mm × 8 mm sites**. The site size in this historical table is different from the exposure field quoted above. [10]

Roadmap year	2007	2010	2015	2019	2022
SFQR target	≤65 nm	≤45 nm	≤25 nm	≤16 nm	≤11 nm
Reference DRAM half-pitch	65 nm	45 nm	25 nm	16 nm	11 nm

These were forward-looking targets published in 2007, not evidence that every product made in those later years met one common specification. They show how wafer requirements were expected to tighten alongside device scaling.

Siltronic uses vivid scale comparisons for 20 nm of wafer topography: a leaf floating on Lake Chiemsee, or a bacterium on a tennis court, after scaling the wafer to that size. The point of the analogy is the extraordinary relative flatness, not a literal surface-inspection method. [6]

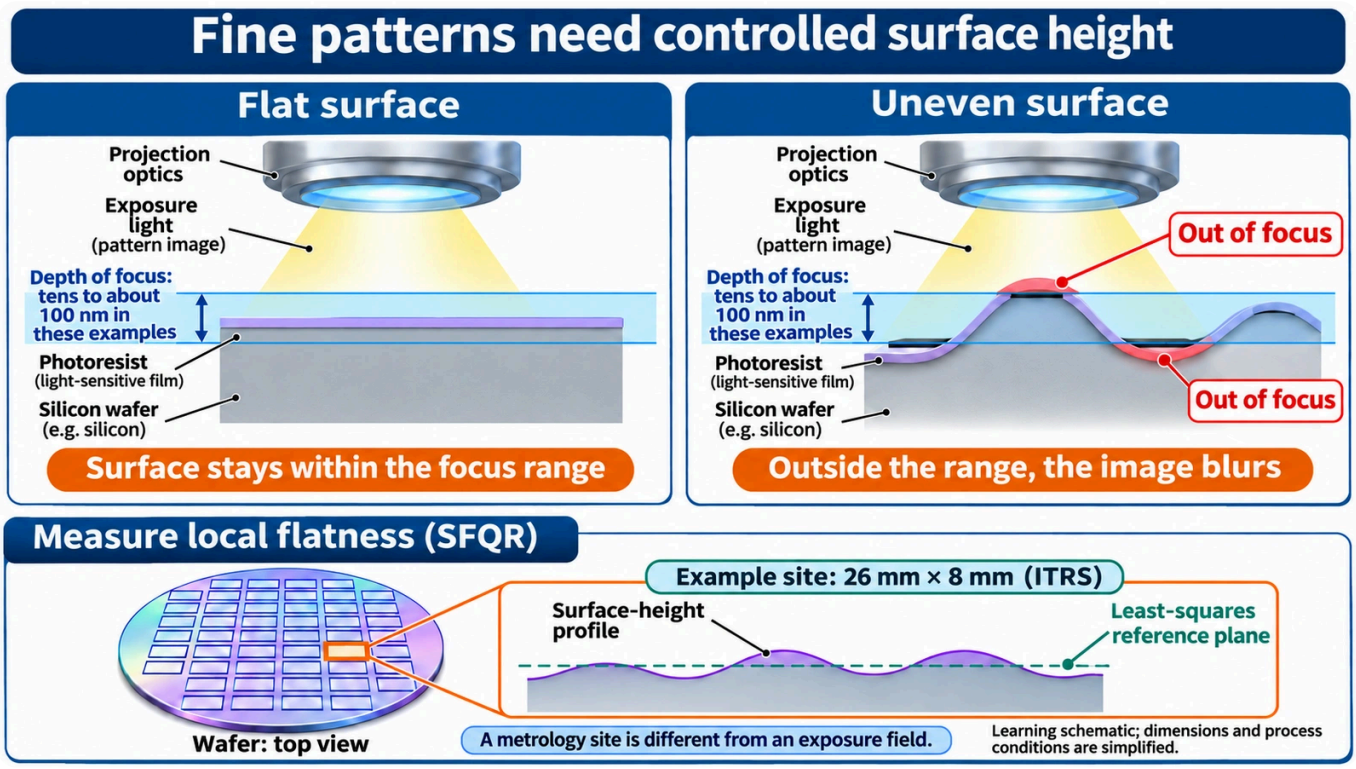

Figure 2 **Surface height, focus and local flatness** (schematic)

Two ideas in one figure

The upper panels show whether the resist surface lies within a usable focus range. The lower panel shows how local flatness can be described relative to a fitted plane. Focus control and flatness measurement are connected, but they are not the same measurement.

3. Shape the ingot and slice it into discs

Set the diameter and mark the orientation

The as-grown ingot is first ground around its circumference to control its diameter. [1] An orientation mark helps later equipment align the wafer with its crystal axes. Smaller formats may use a straight **orientation flat**; large wafers commonly use a small **notch**.

This is more than a convenient handling mark. Silicon's behaviour depends on crystal direction: cleavage and some etching processes, for example, respond differently to different orientations.

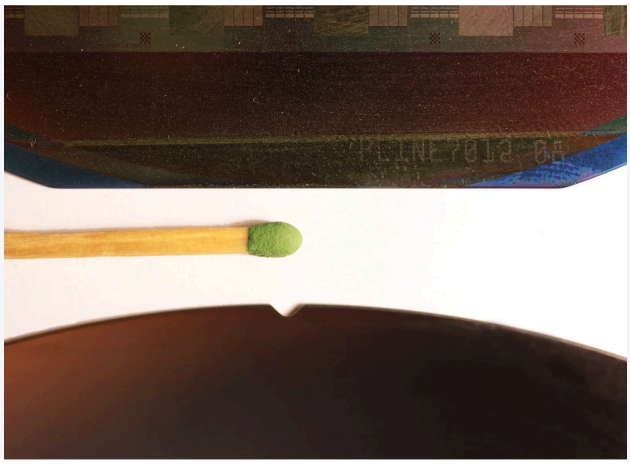

The upper 150 mm wafer has a straight flat; the lower 200 mm wafer has a notch. A matchstick provides scale. The upper wafer already carries processed patterns, so the photograph illustrates the orientation marks rather than two newly polished blank wafers.
Image: Kuebi = Armin Kübelbeck / Wikimedia Commons / CC BY-SA 3.0.

A wire web makes many cuts in parallel

SUMCO describes slicing with either an **inner-diameter saw** or a wire saw, producing discs of roughly **1 mm thickness**. [1] A multi-wire saw creates many parallel cutting segments from one long wire looped around grooved rollers. The ingot feeds slowly into this moving web.

The key geometry is simple: **the cutting wires cross the ingot's long axis**. Each cutting plane separates a disc from its neighbour. Wires running along the axis would not create the wafer slices shown in the illustration.

The cited patent describes 300–400 wire loops as an example. The wire moves back and forth while the ingot is fed into it, allowing many wafers to be cut in one operation. [3][13]

In a **loose-abrasive slurry** process, liquid carries abrasive particles, such as silicon carbide, into the cut. The moving wire brings the particles into contact with silicon, which they remove mechanically. The wire and abrasive together determine the cut; it is not a thermal melting process.

One Shin-Etsu patent discusses moving from a 0.16 mm wire with 11.5 µm abrasive to a **0.14 mm wire with 8 µm abrasive**. Another gives a **20–25 hour** cutting example for a 300 mm ingot and also discusses fixed-abrasive wire with a core around 0.1 mm in diameter. In a fixed-abrasive process, diamond particles are attached to the wire rather than carried freely in slurry. These are source-specific examples, not a claim about every present-day production line. [11][12]

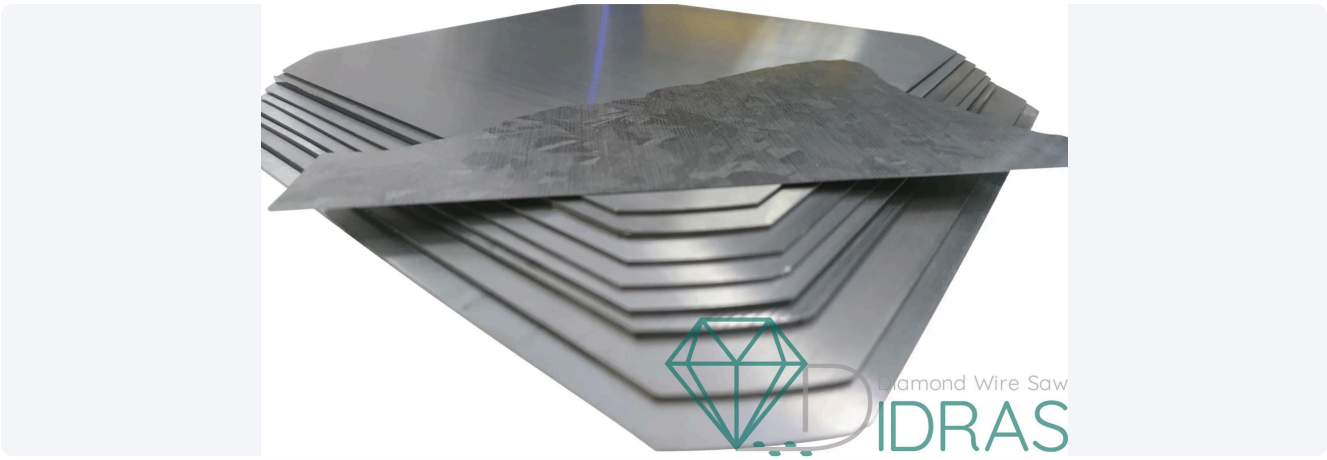

Diamond-wire-sawn silicon wafers, as described on Commons. Their cropped square outline differs from the circular wafers used in this lesson's IC examples, but the photograph shows the unpolished appearance after slicing. The visible equipment logo belongs to the original photograph.

Image: Prof M. Vogel / Wikimedia Commons / CC BY-SA 4.0.

Kerf loss: purified silicon becomes cutting waste

A saw does not merely separate two surfaces. It removes material between them. The width of that removed material is the **kerf**, and the lost silicon contributes to **kerf loss**.

For the loose-abrasive example, the cited patent estimates kerf as wire diameter plus roughly two to three abrasive diameters. [11] Using a 0.14 mm wire and 8 µm abrasive gives:

$$\text{Kerf} \approx 0.14\ \text{mm} + (2\text{–}3) \times 0.008\ \text{mm}$$

The result is **0.156–0.164 mm**, or about 0.16 mm. If each useful as-sliced disc is 1 mm thick, the fraction of ingot length consumed by the kerf is approximately **0.16 / 1.16 ≈ 14%**.

Ignoring the ends and taking a constant 1 mm slice thickness, a one-metre ingot would yield roughly **860 slices**, rather than 1,000. These are geometric teaching calculations; actual counts depend on the process and usable ingot length.

Lesson 2-1 followed this material through purification and days of crystal growth. Reducing cutting waste preserves some of that hard-won value.

Slicing also leaves a mechanically disturbed **damage layer**. A 2002 ECS conference abstract reported about **11 µm** of damage for its wire-sawn samples and about **22 µm** for its inner-diameter-sawn samples. The comparison belongs to those samples and cutting methods; it is not a universal damage-depth specification. [14]

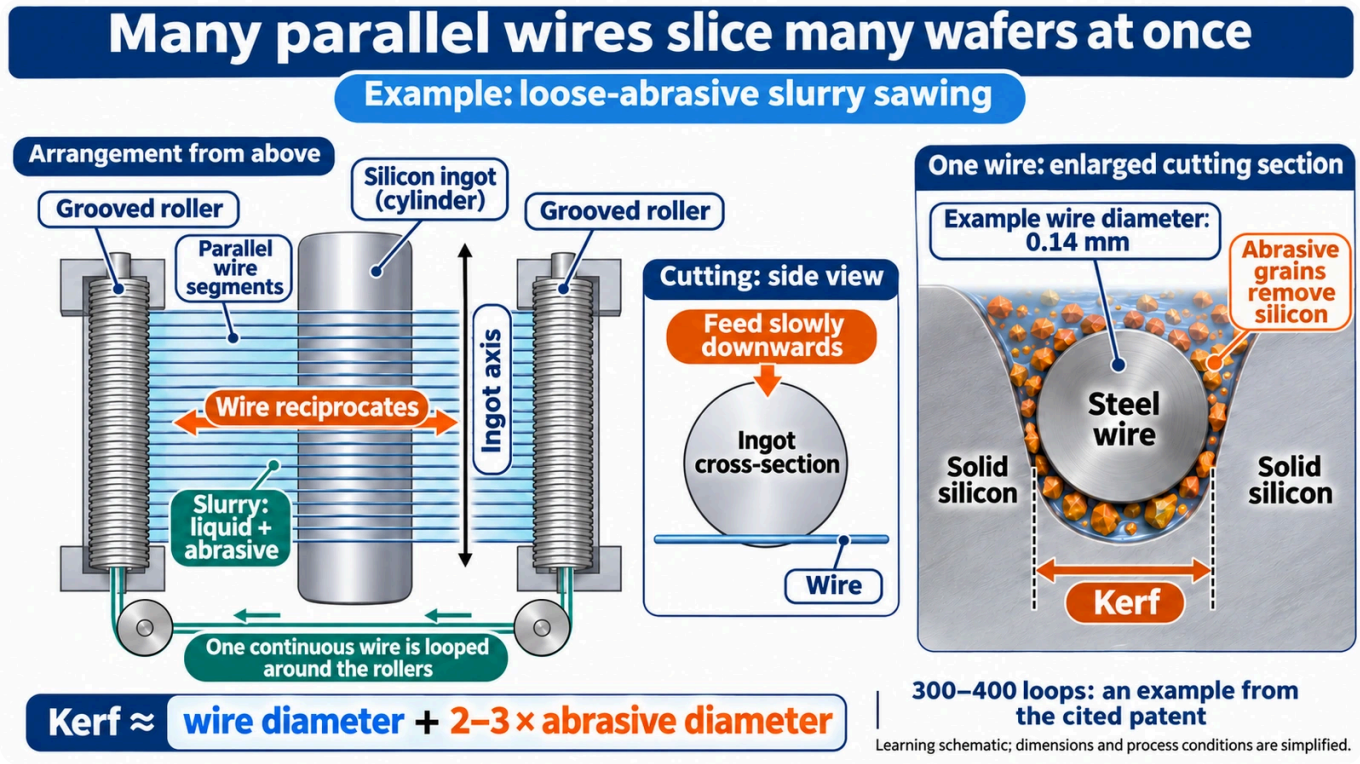

Figure 3 **The corrected geometry of multi-wire sawing** (schematic)

Read the top view and side view separately

The arrangement on the left shows the ingot and wire directions from above. The small side view shows the downward feed into a cutting wire. The enlarged section on the right explains how loose abrasive creates a kerf. Keeping these views separate avoids the misleading impression that the wires run along the ingot.

4. Round the edges, control the faces and remove damage

- **Edge rounding prevents small defects from becoming bigger problems**

Single-crystal silicon is brittle. The edge of a sliced wafer is rounded or **chamfered** to reduce chipping. [3] A chipped edge can shed particles, and an edge flaw can become a starting point for fracture during later thermal or mechanical loading.

- **Grinding and lapping establish thickness and parallelism**

The two faces must be brought towards the required thickness and made suitably parallel. **Lapping** removes material with abrasive between plates. SUMCO describes an alumina-abrasive route, while Siltronic gives an example removing about **10 µm per side**. [1][3]

Some flows use grinding instead of lapping. A Siltronic patent gives a 300 mm example combining wire sawing, edge rounding, sequential grinding of the two faces, chemical etching, edge polishing and double-side polishing. The flow demonstrates how the operations can be combined; it is not the only valid sequence. [15]

Chemical etching removes the disturbed layer

More mechanical removal can create fresh damage as it removes old material. **Etching** uses chemical reactions to remove the damaged layer without another cutting or grinding pass. [1]

The cited Shin-Etsu patent compares acid and alkaline etches. [16]

Route	Useful feature	Trade-off described in the source
Hydrofluoric/nitric acid mixture	Can produce a smoother surface	Larger removal can degrade flatness; gaseous reaction products require management
Alkali, such as NaOH or KOH	Can preserve the geometry established by lapping	Orientation-dependent etching can leave pits and a rougher surface

Acid and alkali remove silicon in different ways

In the acid-mixture explanation, nitric acid oxidises silicon and hydrofluoric acid dissolves the oxide. Coupled oxidation and dissolution remove material without the strong crystallographic selectivity associated with alkaline silicon etching.

Alkaline etching depends much more strongly on crystal orientation. Different atomic arrangements react at different rates, allowing facets and pits to develop. That local roughness can coexist with useful preservation of the wafer's larger-scale geometry.

The cited patent gives an example combining approximately **20 µm of alkaline removal** with **10 µm of acid removal**. Those are example removal amounts, not a processing instruction or a universal recipe. [16]

5. Polishing creates the final surface

The aim is to combine good geometry with a smooth, low-defect surface. SUMCO describes **mechanochemical polishing** using **colloidal silica**: very small silica particles dispersed in liquid. [1]

In double-side polishing, wafers sit in openings in a carrier between upper and lower polishing pads. Motion and slurry act on both faces. A cited Siltronic patent example processes **15 wafers of 300 mm diameter**, using a slurry containing **3 wt% silica at pH 11.5**, with about **30 µm removal** and a final thickness of **775 µm**. Its local-flatness target was SFQRmax no greater than 0.13 µm for 26 mm × 8 mm sites. These values describe that patent's example and target. [15]

A later **final polish** concentrates on surface quality while limiting further removal. Another Siltronic patent describes 10–20 µm removal in a geometry-forming polish, followed by no more than 1.5 µm in a haze-free finishing step, with a softer pad and slurry in the pH 9.0–10.5 range. [17]

Why chemistry and mechanics work better together

Purely mechanical abrasion can leave scratches and subsurface damage. Chemical removal alone does not necessarily reduce high and low areas at the different rates needed to flatten the surface.

In the simplified polishing picture, the slurry changes the near-surface material chemically. Pad contact and particles then remove material, with raised regions generally experiencing stronger contact. Controlling particle size, agglomeration, pad behaviour and chemistry helps remove material while limiting damage.

A useful empirical relation is the **Preston relation**:

$$\text{Removal rate} \approx K \times P \times V$$

P is pressure, V relative velocity and K a process-dependent coefficient. It helps explain why contact conditions affect removal, but it does not replace a detailed polishing model or equipment control.

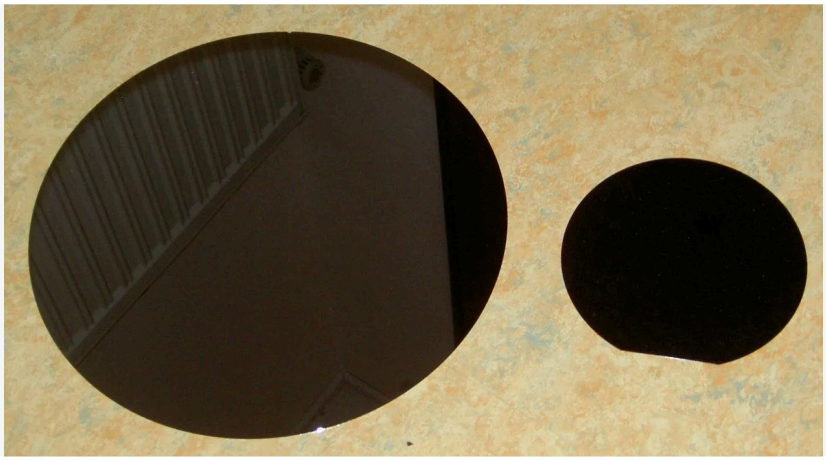

Two unpatterned silicon wafers. Reflections in the larger surface illustrate its mirror-like finish; the smaller wafer has an orientation flat. A photograph alone cannot establish its flatness or particle specification.
Image: Hebbe at German Wikipedia / Wikimedia Commons / public domain.

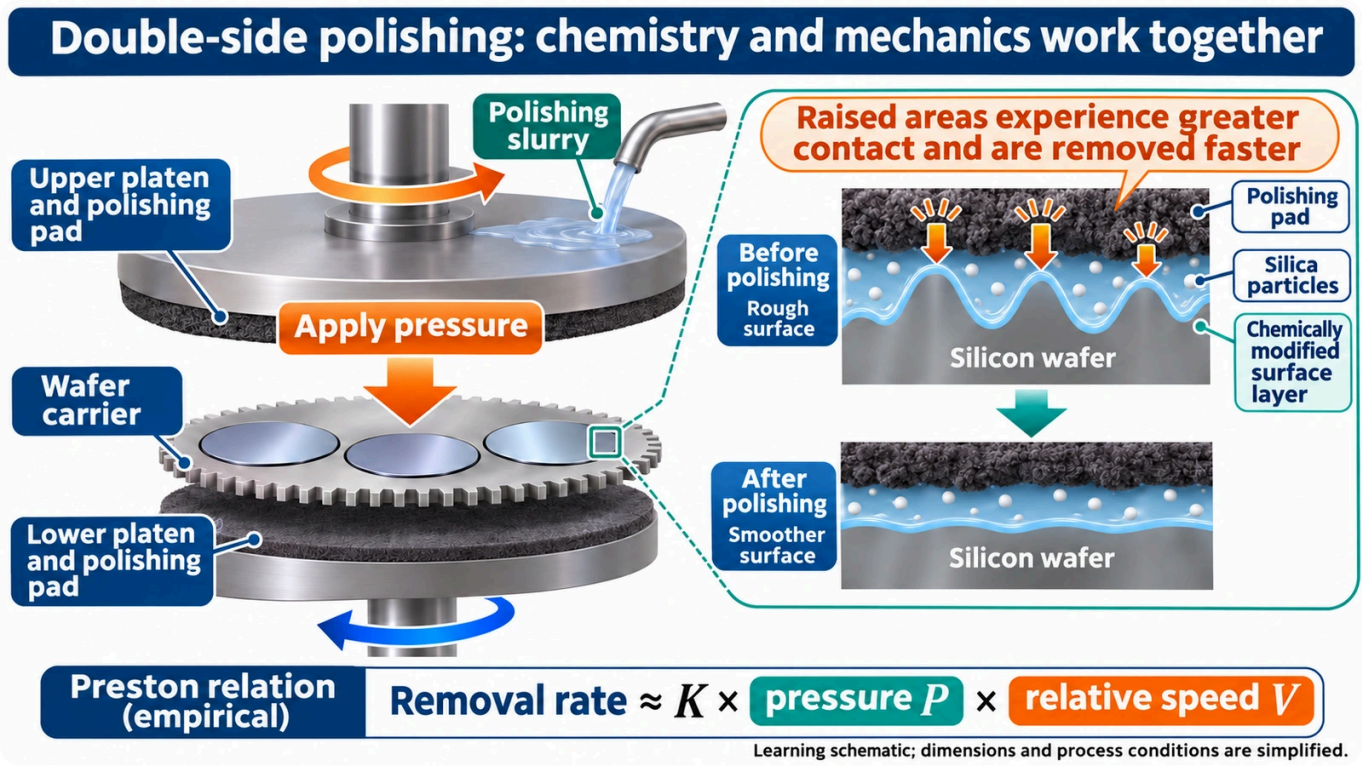

Figure 4 **Double-side polishing and surface removal** (schematic)

A smooth surface and parallel faces are different achievements

The equipment controls wafer geometry while the local interactions develop surface quality. The chemically modified layer is drawn much thicker than reality so that the mechanism is visible. Pressure and speed matter, but so do slurry, particles, pad condition and process history.

6. Clean the wafer and measure the result

- **RCA cleaning**

Polishing leaves residues that must be removed: slurry particles, organic contamination and metallic impurities. A foundational approach is **RCA cleaning**, developed by Werner Kern and colleagues. Kern's retrospective describes research beginning in 1961, internal use of the SC-1 and SC-2 sequence, and publication in 1970. [18]

SC-1 uses an alkaline hydrogen-peroxide solution containing ammonia. The cited historical account gives 75–85 °C and discusses oxidation of organic contamination and complexing of metals. **SC-2** uses a hydrochloric-acid/hydrogen-peroxide solution to address remaining inorganic contamination. These descriptions explain the roles of the cleaning families; actual cleaning flows have evolved with materials and device requirements. [18]



Equipment for batch processing six-inch wafers, including cleaning and etching operations, at LAAS-CNRS in France. The photograph documents a facility; the detailed chemistry of wafer cleaning is covered in Lesson 3-1.
Image: Guillaume Paumier (user:guillom) / Wikimedia Commons / CC BY-SA 3.0.

What inspection needs to establish

A wafer is checked before shipment. [1] Different measurements answer different questions:

Measurement	What it tells us
Thickness and TTV	Total thickness variation: the maximum minus minimum thickness across the defined wafer area
Local flatness, including SFQR	Surface shape within defined sites
Nanotopography	Small-amplitude surface variation over millimetre-scale distances
Edge roll-off	Geometry changes near the wafer perimeter
Surface particles and defects	Features detected, for example, through scattered laser light

The 2007 ITRS roadmap also projected increasingly small particles to be counted: from 65 nm and above in its 2007 target to 16 nm and above in its 2021–2022 targets. Again, those are historical roadmap entries, not a statement of universal inspection performance in those years. [10]

Annealed wafers: modify the near-surface material with heat

A polished wafer can receive a high-temperature treatment in hydrogen or argon to improve the near-surface material. SUMCO describes **annealed wafers** in terms of removing oxygen near the surface and improving crystal perfection. [2]

The underlying idea is to prepare a good device region while retaining useful gettering capability deeper in the wafer. As Lesson 2-1 explained, controlled oxygen-related structures can help trap metals away from the active region. The 2011 ITRS discussion noted that shorter, lower-temperature device thermal cycles made it harder to develop the oxygen precipitates needed for some internal-gettering schemes. Read that as a dated process-integration concern. [19]

7. Epitaxy: grow a new layer with the required properties

The lattice continues across the interface

An **epitaxial wafer** has a newly grown single-crystal layer on a prepared crystalline substrate. The layer inherits the underlying lattice orientation. In the silicon-on-silicon example here, the crystal continues across the interface: **there is no buried oxide between the epi layer and substrate**. [4]

SUMCO describes silicon epitaxy using gases such as silicon tetrachloride or trichlorosilane at high temperature, giving about 1,200 °C in its general process explanation. [1] The simplified trichlorosilane reaction is familiar from Lesson 2-1:

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

The reactor supplies gas to a crystalline wafer surface so that silicon can join the existing lattice. The same broad chemistry can produce polysilicon on a different growth surface; the substrate and growth conditions make the crystallographic difference.

A discussion of single-wafer reactors describes infrared-lamp heating and a cooler quartz enclosure. [25] A SUMCO patent example gives a 300 mm wafer, **1,130 °C**, a growth rate of **2.2 µm/minute** and a **2 µm layer**. A Siltronic patent discusses preferred thicknesses above 0.5 µm and up to 20 µm for its stated control and productivity considerations. These are cited examples, not one thickness or temperature required for all epitaxial products. [21] [22]



A molecular-beam epitaxy system at LAAS-CNRS. MBE grows crystals in vacuum using directed molecular or atomic beams. The production silicon epitaxy discussed in this lesson uses gas-phase chemistry and different equipment; the photograph illustrates the wider family of epitaxial techniques.
Image: Guillaume Paumier (user:guillom) / Wikimedia Commons / CC BY-SA 3.0.

Why add crystal to an already crystalline wafer?

The purpose is to make the **surface material suited to the device**, while choosing the bulk substrate for other requirements.

The cited material describes epitaxial surfaces without substrate-originated COPs and layers whose electrical properties can be selected separately. Those benefits depend on the substrate and growth process; epitaxy does not automatically eliminate every possible defect. [5][22]

SUMCO discusses multilayer structures with different resistivities. Siltronic describes applications including integrated circuits, power devices and image sensors, with different wafer diameters serving different markets. [1][4]

An important example is a **lightly doped p− epitaxial layer on a heavily doped p+ substrate**. The device region can have one resistivity while the substrate underneath provides a lower-resistance region. In CMOS, unintended parasitic bipolar structures can interact and sustain a large current: **latch-up**. The cited SEH America patent explains how the p/p+ arrangement can reduce the gain of a parasitic path and help suppress that behaviour. Suitable structures deeper in the substrate can also provide gettering sites. [23]

SEMI's report on 2025 wafer shipments linked growth to demand for advanced logic epitaxial wafers and polished wafers for HBM. That is the market context in the dated report, rather than a universal choice of substrate for every logic or memory design. [26]

Two stages set the growth rate: the Grove model

A precursor has to pass through two stages: reach the surface, then react and become part of the growing material. The simplified model expresses the fluxes as: [24]

$$\text{Transport: } J_1 = h_g(C_g - C_s)$$

$$\text{Surface reaction: } J_2 = k_s C_s$$

C_g is the precursor concentration in the bulk gas; C_s is its concentration near the surface. h_g describes transport and k_s the surface reaction. At steady state, the two fluxes are equal. For a fixed bulk precursor concentration, the effective coefficient is:

$$h_g k_s / (h_g + k_s)$$

This resembles two resistances in series. If reaction is slow, it dominates: **reaction-limited growth** is strongly temperature-sensitive. If reaction is sufficiently fast, delivery becomes the bottleneck: **transport-limited growth** is less temperature-sensitive in this simplified picture.

The cited SUMCO patent discusses trichlorosilane growth at **1,100–1,130 °C** in a transport-limited regime. Such a regime can reduce sensitivity to small temperature variations across a large wafer, but uniform delivery and flow still need to be controlled. [20]

Autodoping: the substrate can influence the new layer

At growth temperature, dopants can move. **Out-diffusion** carries dopants from the substrate into the growing material. **Autodoping** describes a gas-mediated route: dopant released from exposed substrate surfaces enters the gas and is incorporated into the layer.

This is particularly relevant when a lightly doped layer is grown on a heavily doped substrate. The cited reactor discussion describes approaches including single-wafer processing, alternative lower-temperature chemistry and backside coatings. [25]

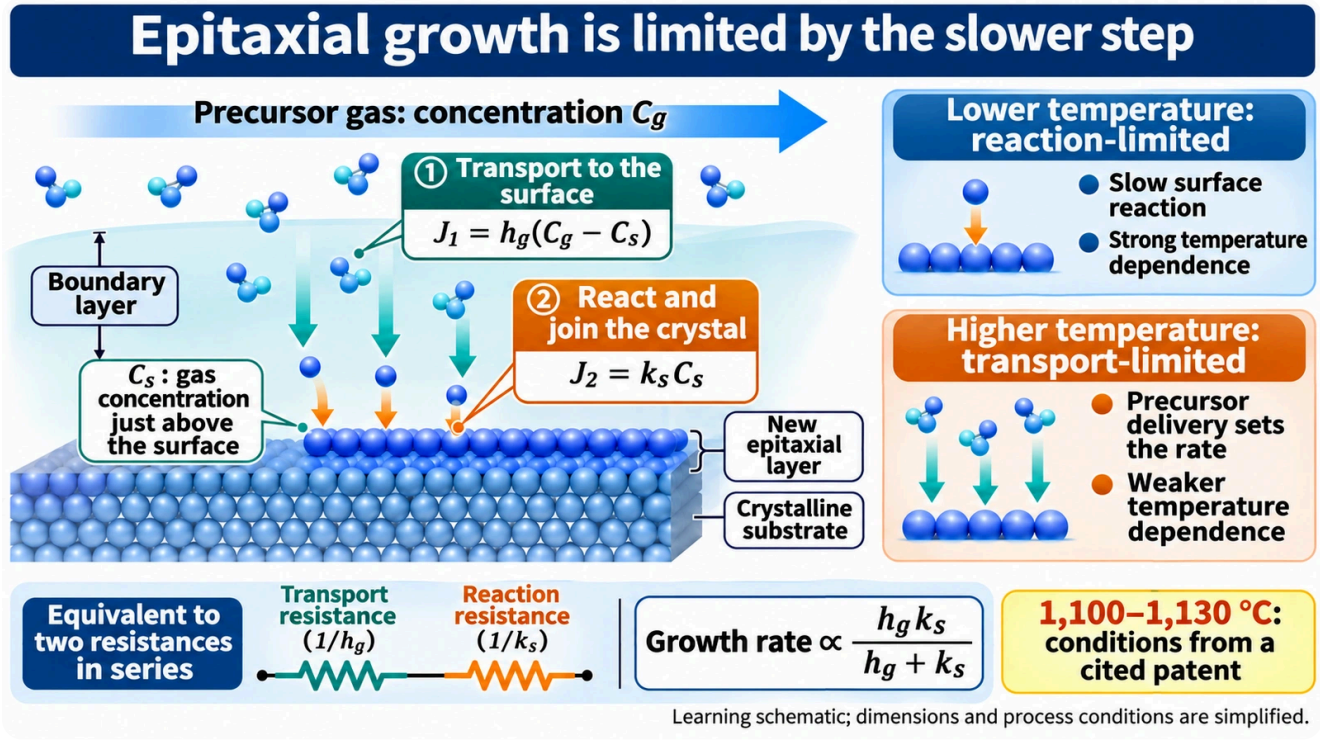

Figure 5 **Transport and reaction: two different bottlenecks** (schematic)

Transport-limited does not mean flow is unimportant

It means transport is setting the rate, so precursor concentration, flow and the boundary layer require careful control. The diagram's proportionality assumes other quantities, including bulk concentration, are held fixed. Its C_s is a different quantity from the solid dopant concentration denoted by the same symbol in Lesson 2-1's segregation model.

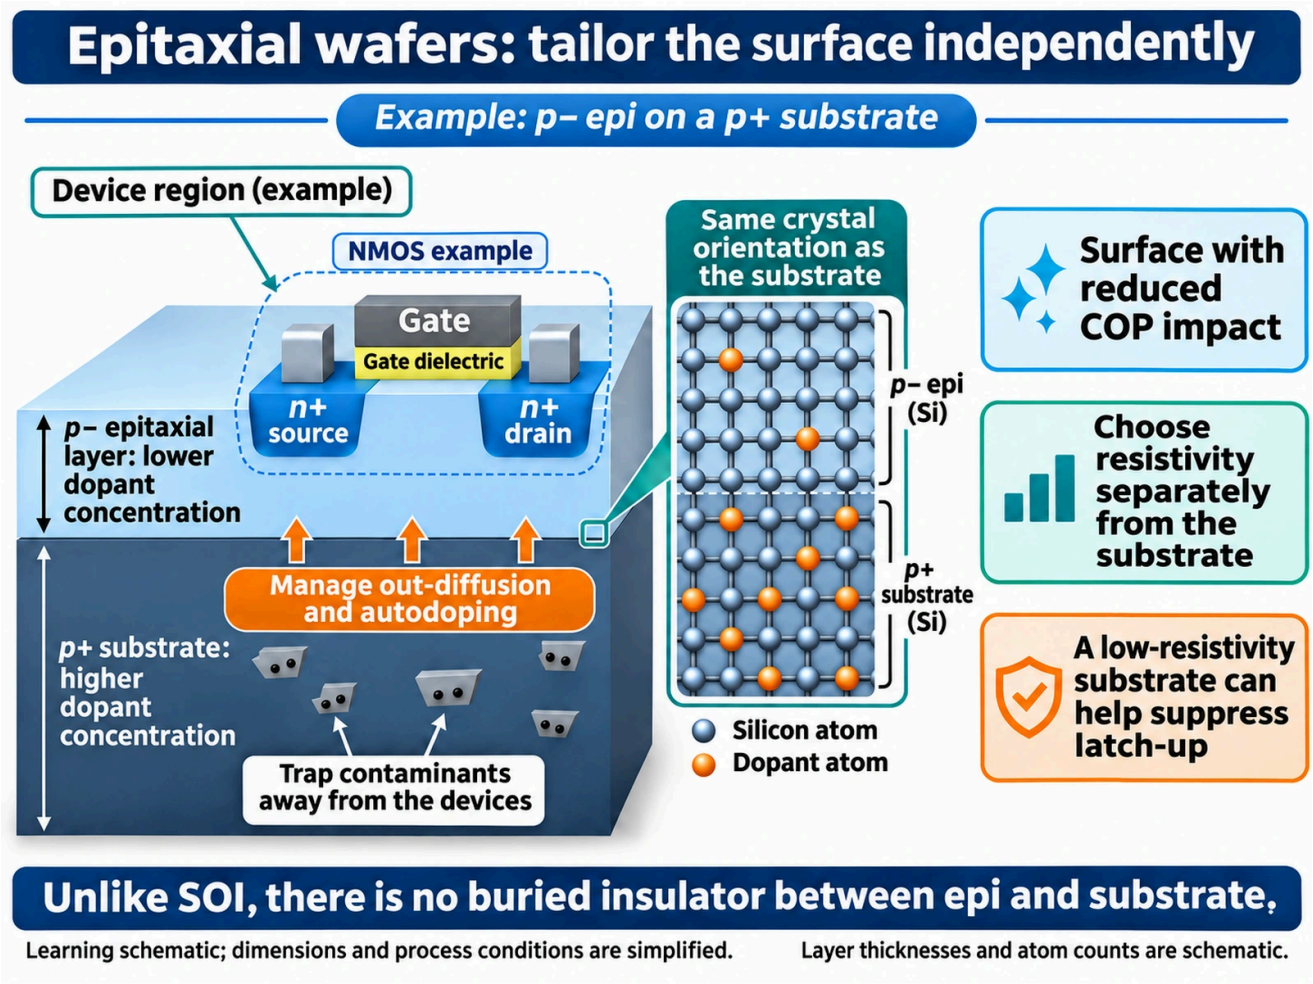

Figure 6 **A p− epitaxial layer on a p+ substrate** (schematic)

Same crystal orientation, different electrical properties

Both regions are silicon. Different doping allows different resistivities without inserting an insulating layer between them. The atom counts and thicknesses are exaggerated for explanation. The NMOS device is an example, not a complete CMOS latch-up circuit, and the gettering symbols represent a function rather than literal manufactured boxes.

8. SOI: put an insulating layer under the devices

An **SOI wafer**, silicon on insulator, has a thin single-crystal silicon layer over a **buried oxide (BOX)**, supported by a thicker handle substrate. Devices are made in the upper silicon.

The buried insulator changes the electrical relationship between devices and substrate. Soitec discusses reduced source/drain capacitance and leakage, and the removal of the conventional bulk-CMOS latch-up path. SUMCO describes the resulting opportunities for power, speed, density and reliability. The actual benefit depends on device design and application. [1][28]

SOI family	Structure highlighted in the source	Example application
FD-SOI, fully depleted SOI	Very thin device silicon and BOX; one cited 28 nm-generation example uses about 12 nm Si and 25 nm BOX	Low-power logic
RF-SOI	Device silicon, BOX, a trap-rich layer and a high-resistivity substrate in the cited arrangement	Smartphone and other RF circuits
Photonics-SOI	The top silicon can form optical waveguides	High-speed optical communication

Soitec's cited FD-SOI paper describes a 300 mm example with a **0.3 nm range of device-layer thickness**. Taking an illustrative silicon-plane spacing of about 0.14 nm, 0.3 nm is only around two atomic-layer spacings. The comparison depends on orientation and how layers are counted. [27]

FD-SOI also supports **back biasing**: a voltage applied from beneath the BOX changes transistor behaviour. Soitec contrasts its usable bias range with roughly ±300 mV for the bulk example in the cited paper. That comparison belongs to the structures and constraints discussed there. [28]

Soitec attributes widespread smartphone use to its RF-SOI products, including a claim of presence in all 5G smartphones. This is the company's claim, not a survey performed for this article. Its photonics-SOI material addresses optical-communication applications. [29][31]

Three ways to create the buried-insulator structure

SIMOX creates buried oxide by implanting oxygen into silicon and applying the necessary subsequent processing. Izumi and colleagues reported CMOS devices on oxygen-implanted buried oxide in 1978. [32]

Bonding and thinning starts with two wafers. Oxide is formed on one or both bonding surfaces; the wafers are joined and heat-treated, then the upper active wafer is ground and polished down to the required thickness. SUMCO describes this route. Removing most of a donor by grinding, however, consumes material and makes very thin uniform layers challenging. [1]

Smart Cut transfers a thin layer by combining implantation and bonding. Originating at CEA-Leti and industrialised by Soitec, it allows the active layer to be prepared separately from the mechanical support. The donor remainder can be refreshed and reused. [30]

Smart Cut: decide where the donor will split

The basic patent describes implanting light ions, bringing the donor surface into contact with a support and applying thermal treatment to separate a thin layer. [33] A practical teaching sequence is:

1. Form oxide on a silicon donor wafer.
2. Implant hydrogen so that a hydrogen-rich region develops at a controlled depth within the donor silicon.
3. Flip the donor and bond its oxidised face to a handle wafer.
4. Split at the implanted region, leaving thin silicon above oxide on the handle.

The patent gives a hydrogen dose around **5×10^{16} cm^{-2}** and a silicon splitting-temperature example of about **500 °C**. A cited university paper describes bonding near room temperature, separation at 400–600 °C, higher-temperature bond strengthening and final surface finishing. These figures describe examples; implantation and thermal conditions are chosen for the process. [33][34]

Why does fracture follow the implanted region?

Implanted hydrogen accumulates over a depth distribution governed partly by implantation energy. It creates defects and thin microcavities in a buried region. Heating allows hydrogen-containing cavities to evolve, and gas pressure and stress help drive fracture along that region.

Bruel's account identifies gas pressure in the microcavities and stress in the layer as driving factors. Without suitable bonding, surface blistering can occur. Bonding changes the mechanical constraints and helps a continuous split propagate along the implanted region. This is a fracture process, not a conventional blade cutting through the wafer. [35]

The implantation conditions help determine the transferred-layer thickness. Subsequent processing establishes the final interface and surface quality. Reclaiming the remaining donor makes repeated layer transfer possible. [30]

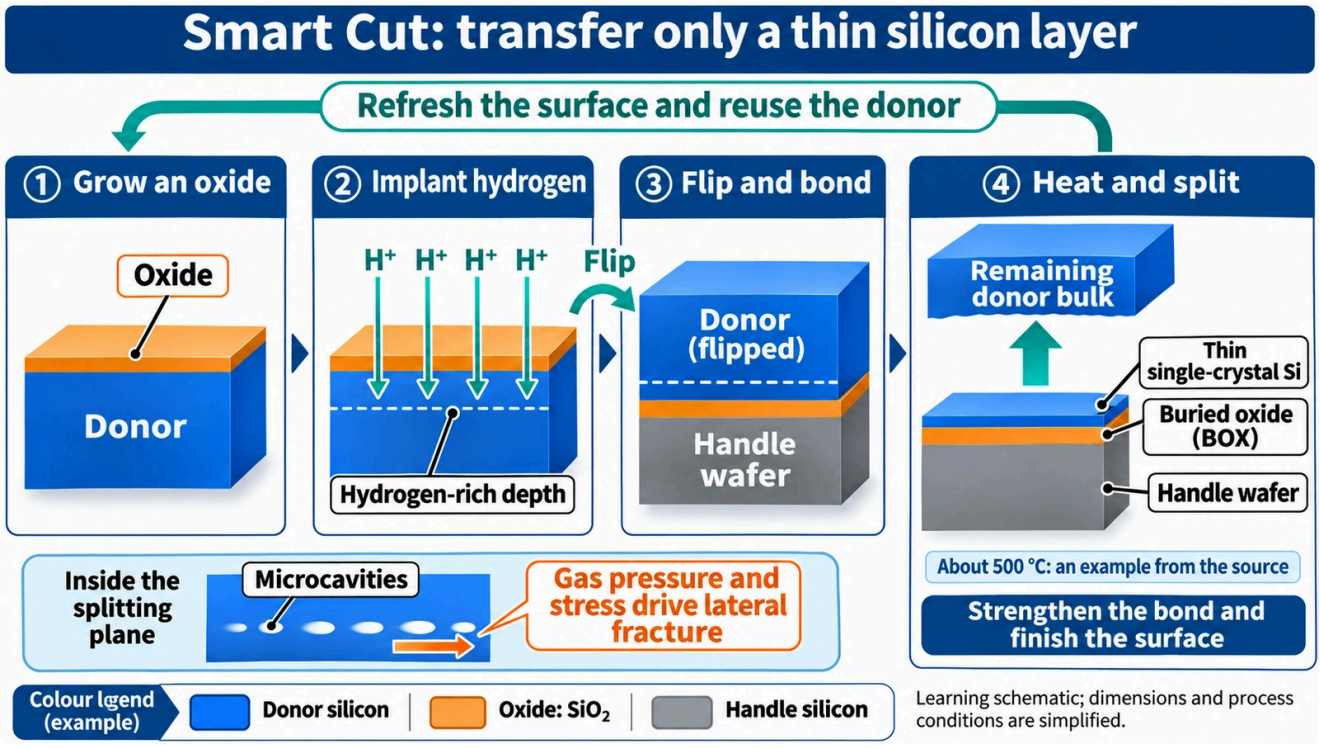

Figure 7 **Smart Cut: preserve the layer order through each step** (schematic)

Follow the colours through the flip

Before bonding, the oxide is on the donor's upper surface. After flipping, it faces the handle. Splitting leaves **thin donor silicon on top, oxide beneath it, and the handle at the bottom**. Most of the donor lifts away for potential reuse. Its bulk is not the transferred device layer.

Shin-Etsu's relationship with the technology includes a licence from Soitec in 1997, as described in the cited company history. Lesson 2-3 examines the supplier landscape rather than treating a process name as exclusive to one wafer supplier. [37]

9. Bonding and thinning also enable backside power

Wafer preparation techniques now appear inside advanced device-fabrication flows as well. **Backside power delivery**, introduced in the earlier AI lesson, is one example.

Imec describes bonding a processed wafer's front side to a carrier, then removing silicon from the back until the remaining silicon is only a few hundred nanometres thick. Connections can then be made from the backside. The carrier supports material that would otherwise be extremely difficult to handle. [36]

Starting from 775 µm and retaining a few hundred nanometres leaves roughly one part in a few thousand of the original thickness, depending on the final value, depending on the chosen final value. The point is the scale of the thinning, not one universal final thickness.

Bond quality, surface geometry and thickness control therefore matter beyond the wafer supplier's own production line. Their consequences reach into later device integration and yield. Lesson 3-12 returns to the backside-power structure.

10. Put the processing steps together

Step	Main purpose	Key idea
Diameter grinding and orientation marking	Set shape and provide alignment information	Crystal direction matters
Slicing	Produce thin discs	Parallel cutting planes and kerf loss
Edge rounding, grinding and lapping	Reduce chipping and control thickness/parallelism	Manage mechanical geometry and damage
Etching	Remove the disturbed layer chemically	Reaction chemistry and crystallographic selectivity
Double-side and final polishing	Develop geometry and surface quality	Chemical modification plus mechanical removal
Cleaning and inspection	Remove residues and verify the result	Chemistry, geometry and defect measurement
Epitaxial growth	Tailor the device-layer material	Transport, reaction and dopant control
SOI fabrication	Put device silicon above an insulator	Bonding, thinning or layer transfer

The next lesson, **2-3**, examines the wafer industry: supplier positions, the strengths of Japanese manufacturers, quality requirements and demand linked to AI. The raw material has travelled a long way from quartz. Its value now depends on controlled structure, geometry, electrical properties and a reliable manufacturing process.

11. Check your understanding

1. What happens to depth of focus when NA increases in DOF = $k_2\lambda/NA^2$?

▼ **Show answer**

It becomes smaller. Better resolution through higher NA makes surface-height and focus control more demanding.

2. In the loose-abrasive example, what sets the kerf width?

▼ **Show answer**

Approximately the wire diameter plus two to three abrasive diameters. The stated relation belongs to that process example; fixed-abrasive cutting has its own details.

3. Why etch after slicing or grinding instead of relying only on more mechanical removal?

▼ **Show answer**

Chemical etching can remove the disturbed material without another mechanical pass creating a fresh damage layer.

4. Why combine chemical and mechanical action in polishing?

▼ **Show answer**

Chemistry changes the near-surface material, while pad and particle interactions preferentially remove raised regions. Process control helps develop flatness while limiting damage.

5. Give two reasons for using an epitaxial wafer.

▼ **Show answer**

It can provide a suitable low-defect device surface and allow the device layer’s resistivity to be chosen separately from the substrate. A p−/p+ example can also help suppress latch-up and support gettering away from devices.

6. Why can a transport-limited epitaxial regime help thickness uniformity?

▼ **Show answer**

Its growth rate is less sensitive to small temperature differences in the simplified model. Precursor delivery and flow still need careful control.

7. What changes when the device silicon sits above a buried insulator in SOI?

▼ **Show answer**

The electrical connection to the bulk substrate changes, reducing some parasitic effects and breaking the conventional bulk-CMOS latch-up path. FD-SOI can also support a useful back-bias range.

8. After Smart Cut splitting, which layers remain on the handle wafer, from top to bottom?

▼ **Show answer**

Thin single-crystal silicon transferred from the donor, then buried oxide (BOX), then the handle wafer. Most donor material separates above the transferred layer.

12. A short glossary

Annealed wafer
Wafer heat-treated to modify its near-surface crystal properties.

Autodoping
Gas-mediated incorporation of dopants released from the substrate during growth.

Back bias
Substrate-side voltage used to adjust transistor behaviour.

Boundary layer
Near-surface region through which precursor transport must occur.

BOX
Buried oxide, the insulating oxide layer within an SOI structure.

Carrier wafer
A support bonded to a wafer during processes such as extreme thinning.

Chamfering
Rounding or shaping the edge to reduce chipping.

Colloidal silica
Fine silica particles dispersed in liquid, used in polishing.

Damage layer
Material disturbed by mechanical processing near the surface.

Double-side polishing
Polishing both wafer faces in one process.

DOF
Depth of focus: the axial range giving acceptable image quality.

Edge roll-off
A change in surface geometry near the wafer perimeter.

Epitaxial growth
Crystal growth that inherits the substrate's lattice orientation.

Epitaxial wafer
Wafer carrying a deliberately grown single-crystal layer.

FD-SOI
Fully depleted SOI, using sufficiently thin device silicon for full depletion.

Final polish
Limited finishing removal aimed at surface quality.

Fixed-abrasive wire
Cutting wire with abrasive particles attached to it.

Gettering
Trapping contaminants away from the device-active region.

Grove model
Simplified growth model combining transport and surface reaction.

Inner-diameter saw
Saw using the inner edge of an annular blade.

Kerf
Width of material removed by a cut.

Kerf loss
Material lost as cutting waste.

Lapping
Abrasive removal used to control wafer thickness and parallelism.

Latch-up
Sustained unwanted current through interacting parasitic device structures.

Mechanochemical polishing
Polishing combining chemical action with mechanical removal.

Multi-wire saw
Saw creating many cutting segments from a long wire.

NA
Numerical aperture, a measure of an optical system's angular capability.

Nanotopography
Small-amplitude surface shape over millimetre-scale distances.

Notch
Small edge cut-out providing an orientation mark.

Orientation flat
Straight edge section used as an orientation mark.

Out-diffusion
Dopant migration from the substrate into adjacent material.

Polished wafer
Wafer finished to specified geometry and surface quality.

Preston relation
Empirical link between removal rate, pressure and relative speed.

RCA clean
Cleaning families commonly described as SC-1 and SC-2.

Reaction-limited growth
Growth controlled mainly by slow surface reaction.

RF-SOI
SOI engineered for radio-frequency applications.

SFQR
Range of surface deviations from a site's fitted least-squares plane.

SIMOX
SOI formation using implanted oxygen and subsequent processing.

Single-wafer reactor
Equipment processing one wafer at a time.

Slurry
Liquid carrying particles used in cutting or polishing.

Smart Cut
Layer transfer combining implantation, bonding and controlled splitting.

SOI
Silicon on insulator: silicon above an insulating layer.

Transport-limited growth
Growth controlled mainly by precursor delivery.

TTV
Total thickness variation over the specified wafer area.

13. Sources and reading notes

Reference numbers follow the completed Japanese edition. Patent values describe examples or discussions in those documents, not independently confirmed present-day production settings. Roadmap entries retain their publication dates. Japanese titles are described in English below.

1. **SUMCO** — Silicon wafer manufacturing process (Japanese) — Read source
2. **SUMCO** — Product line-up: annealed wafers (Japanese) — Read source
3. **Siltronic** — Products and wafer processing — Read source
4. **Siltronic** — Epitaxial wafers — Read source
5. **GlobalWafers** — Technology — Read source
6. **Siltronic** — Investor presentation, September 2026, page 19 — Read source
7. **H. Kirchauer, TU Wien doctoral thesis** — Some Fundamental Considerations: resolution and depth of focus — Read source
8. **ASML** — EUV lithography systems — Read source
9. **ASML** — TWINSCAN NXT:2050i — Read source
10. **ITRS** — 2007 Front End Processes, Tables FEP2a/2b — Read source
11. **Shin-Etsu, US patent US9260643** — Wire diameter, abrasive size and kerf — Read source
12. **Shin-Etsu, US patent US10011046** — Sawing time and fixed-abrasive wire examples — Read source
13. **Shin-Etsu, US patent US9776340** — Wire-loop count — Read source
14. **GRINM and co-authors** — ECS 201st meeting abstract 565, 2002: sawing damage comparison — Read source
15. **Siltronic, US patent US6861360** — Wafer processing and double-side polishing example — Read source
16. **Shin-Etsu, US patent US6239039** — Acid and alkaline etching comparison — Read source
17. **Siltronic, US patent US6530826** — Polishing and haze-free finishing — Read source
18. **W. Kern** — Citation Classic, 1983: retrospective on Kern and Puotinen, RCA Review, 1970 — Read source
19. **ITRS** — 2011 Front End Processes — Read source
20. **SUMCO, US patent US8753962** — Temperature regime for 300 mm epitaxial growth — Read source
21. **SUMCO, US patent US9631297** — Epitaxial growth example — Read source
22. **Siltronic, US patent US7875115B2** — COPs and epitaxial layer thickness — Read source
23. **SEH America, US patent US6346460B1** — p/p+ structure, latch-up and gettering — Read source
24. **MIT OpenCourseWare 6.152J, 2005** — Chemical Vapor Deposition lecture — Read source
25. **Semitracks** — March 2017 newsletter: epitaxial reactors and autodoping — Read source
26. **SEMI** — 2025 worldwide silicon wafer shipments and revenue, 10 February 2026 — Read source
27. **Soitec** — FD substrate white paper — Read source
28. **Soitec** — Planar fully-depleted silicon technology for 28 nm and beyond — Read source
29. **Soitec** — RF-SOI — Read source
30. **Soitec** — Smart Cut technology — Read source
31. **Soitec** — Photonics-SOI — Read source
32. **K. Izumi, M. Doken and H. Ariyoshi** — CMOS on buried SiO_2 formed by oxygen implantation, Electronics Letters 14, 593, 1978 — Read source
33. **CEA, US patent US5374564A, inventor M. Bruel** — Basic Smart Cut patent — Read source
34. **N. M. Forhan and co-authors, University of São Paulo** — Hydrogen-implantation SOI layer transfer — Read source
35. **M. Bruel** — History, Physics, and Applications of the Smart-Cut Process, MRS Bulletin 23(12), 35, 1998 — Read source
36. **imec** — How to power chips from the backside — Read source
37. **Soitec** — Company activity and licensing history — Read source
38. **Patent EP1335420A1** — Wafer shape evaluation: SFQR definition — Read source

14. Evidence map and calculation assumptions

Topic	Source or assumption	How to interpret it
Wafer-processing route and product families	SUMCO [1][2]; Siltronic [3][4]; GlobalWafers [5]	Representative operations and products; flows vary
Focus relations and optical parameters	[7–9]	k_2 = 0.5–1.0 assumed for the DOF estimates; EUV row is the 0.33-NA example
SFQR and historic geometry targets	[10][38]	Fitted-site reference plane; 2007 roadmap targets, not later measured results
Scale analogy	Siltronic [6]	Illustrative comparison for 20 nm topography
Wire, abrasive and cutting examples	Patents [11–13]	Specific descriptions rather than universal current line settings
Kerf calculation	0.14 mm wire, 8 µm abrasive, 1 mm slice	Kerf 0.156–0.164 mm; approximately 14% loss and 860 slices/metre, ignoring ends
Slicing damage	ECS abstract [14]	About 11 µm versus 22 µm for the reported samples
Geometry-forming and final polish	Patents [15][17]	Example removals, slurry conditions and 775 µm finished thickness
Acid and alkaline etching	Patent [16]	Source-specific comparison and combined-removal example
RCA cleaning	Kern retrospective [18]	Historical development and cleaning roles
Annealed-wafer integration	[2][19]	Near-surface modification and a dated gettering concern
Epi conditions, layer properties and p/p+ benefits	[1][4][5][20–23]	Controlled-growth examples; no claim that every defect disappears
Grove model and autodoping	[24][25]	Steady-state simplified model and gas-mediated dopant incorporation
Market context	SEMI [26]	Reported drivers of 2025 shipments
FD-SOI, RF-SOI and photonics-SOI	[27–31]	Application-specific structures and attributed company claims
Atomic-layer comparison	0.3 nm / 0.14 nm	Roughly two spacings under the stated illustrative assumption
SIMOX and Smart Cut	[32–35]	Historical report, implantation/bonding/splitting mechanism and example conditions
Licensing	Soitec [37]	Attributed company history
Backside power and extreme thinning	imec [36]	Example integration route, not a universal residual thickness

English edition: 4 November 2026. Fresh English text based on the completed Japanese edition. Diagrams: OpenAI Image Gen. Historical targets and patent examples retain their context; photograph credits accompany the images.

© 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.