Semiconductor Technology from A to Z

Optoelectronics & Photonics

Fabrication of Optoelectronic Devices

Epitaxy: MOCVD and MBE

The fabrication of optoelectronic devices begins with epitaxial crystal growth – unlike many electronic devices, simply doping the substrate isn't sufficient here, since the active region consists of several precisely controlled heterolayers (barriers, quantum wells, DBR stacks) whose thickness and composition directly determine the emission wavelength. Even the smallest crystal defects – dislocations, point defects, impurities – act as non-radiative recombination centers and measurably reduce quantum efficiency. While a MOSFET still tolerates defect densities in the range of 10^4–10^5 cm^{-2}, the same defect densities in an LED active region already cause noticeable efficiency losses; commercial GaN LEDs require dislocation densities well below 10^8–10^9 cm^{-2}, ideally lower still.

MOCVD (Metal-Organic Chemical Vapor Deposition, also called MOVPE) is the industry standard for mass production. Metal-organic precursors – such as trimethylgallium (TMGa), trimethylindium (TMIn), trimethylaluminum (TMAl) – and hydrides (ammonia NH_3 for nitride compounds, arsine AsH_3 or phosphine PH_3 for arsenides/phosphides) are carried in a gas stream (usually hydrogen or nitrogen) over the substrate, heated to 600–1100 °C, where they decompose thermally (pyrolysis), and the released atoms incorporate at the growth front. The process runs at atmospheric or moderate reduced pressure, allows high throughput across multiple wafers simultaneously in planetary reactors, and therefore dominates commercial LED and laser fabrication almost entirely.

MBE (Molecular Beam Epitaxy), by contrast, operates in ultra-high vacuum (below 10^{-10} mbar) with molecular beams from effusion-cell-evaporated elements that condense directly on the substrate without a chemical reaction. The growth rate, typically around one monolayer per second, is markedly lower than MOCVD rates, but layer thickness can be controlled down to the monolayer and monitored in situ via RHEED (Reflection High-Energy Electron Diffraction) – the diffraction pattern shows in real time whether the surface is growing atomically smooth. Due to the lower throughput and higher equipment cost, MBE remains largely reserved for research and specialty applications with particularly demanding heterostructures, such as VCSEL DBR stacks with extremely sharp interfaces.

A central challenge in both processes is lattice matching: growing a material with a differing lattice constant on a substrate creates strain that relaxes through dislocation formation above a critical layer thickness (often just a few nanometers for strong mismatch) – with a directly measurable effect on the

resulting device efficiency. GaN, for instance, is usually grown on sapphire (Al_2O_3) or SiC due to the lack of economically available bulk GaN substrates, despite a substantial lattice mismatch of up to 16% for GaN-on-sapphire. A thin AlN or GaN nucleation layer grown at low temperature (500–600 °C) serves as a buffer layer: it initially grows amorphous to polycrystalline, is then annealed and recrystallizes, relieving part of the strain before the actual single-crystal GaN growth continues at 1000–1100 °C.

Mesa Patterning and Metallization

After epitaxial growth, a continuous layer structure initially covers the entire wafer – for individually addressable devices, it must be divided into isolated mesa structures. Dry etching processes, for III-V materials usually ICP-RIE (Inductively Coupled Plasma Reactive Ion Etching) with chlorine-based chemistry (Cl_2/BCl_3 mixtures, since fluorine chemistry doesn't form sufficiently volatile etch products with Ga and In compounds), remove the layers down below the active region. Etch depth is typically a few micrometers and defines both the lateral extent of each individual device and the electrical isolation of neighboring structures on the same chip – insufficient isolation would cause parasitic leakage currents between neighboring LEDs or lasers on the same wafer.

Contact metallization follows, with fundamentally different requirements on the p- and n-sides. The p-side typically uses Ti/Au or Ni/Au layer stacks, often supplemented with a thin Pt or Pd interlayer as a diffusion barrier; they must achieve as low a specific contact resistance as possible (in the range of 10^{-5}–10^{-6} $\Omega\cdot cm^2$) while also providing good current spreading, which is technologically more demanding than n-contacts due to the lower hole mobility in p-type semiconductors. LEDs therefore often add a transparent, conductive ITO layer (indium tin oxide, sheet resistance typically 10–30 Ω/sq) between the metal contact and p-GaN, to distribute current evenly across the entire chip area without blocking light extraction with large-area, opaque metallization – ITO is over 90% transparent in the visible spectrum. The n-side typically gets a simpler Ti/Al-based contact with lower contact resistance, since n-type semiconductors inherently have lower contact barriers than p-type ones, owing to higher electron mobility and lower effective mass.

Fabrication Flow: From Epitaxy to Mounted Chip

Four key process steps in LED and laser diode fabrication

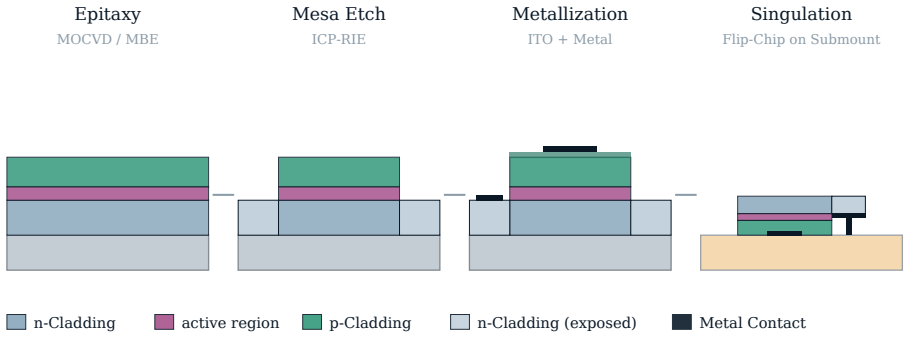

Only the active region (center) must stay defect-poor across the entire process chain – every step can affect quantum efficiency.

Device in Operation: Current Flow and Light Emission

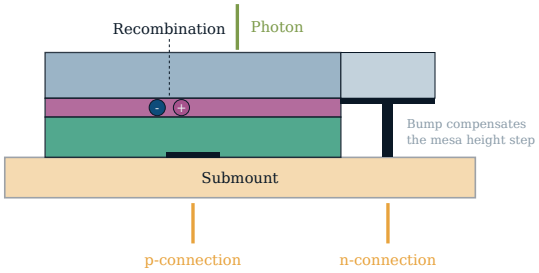

Holes and electrons are injected into the active region via the contacts, recombine radiatively there – the photon exits on the substrate side.

A distinctive feature of VCSEL fabrication is the oxide aperture, which solves one of the central problems of the vertical resonator concept: the need to concentrate the injection current onto the small active area. During epitaxy, a thin, AlAs-rich layer (typically over 98% Al content) is deliberately incorporated into the layer structure near the active region. After mesa etching, this layer is converted to insulating aluminum oxide (Al_xOy) through selective, lateral wet oxidation at elevated temperature in a humid nitrogen atmosphere, starting from the exposed mesa edges – the oxidation front grows from outside in at a material- and temperature-dependent, well-controllable lateral rate, so that only a small central opening a few micrometers in diameter remains unoxidized and thus conductive. This aperture forces the current into a narrowly confined central region of the active area and simultaneously acts as an optical waveguide (due to the oxide's lower refractive index) – both necessary to achieve the current density and transverse mode control needed for lasing despite the VCSEL's short gain length and small device area.

VCSEL Oxide Aperture: Structure and Formation

Lateral wet oxidation confines current and optical field to a small opening

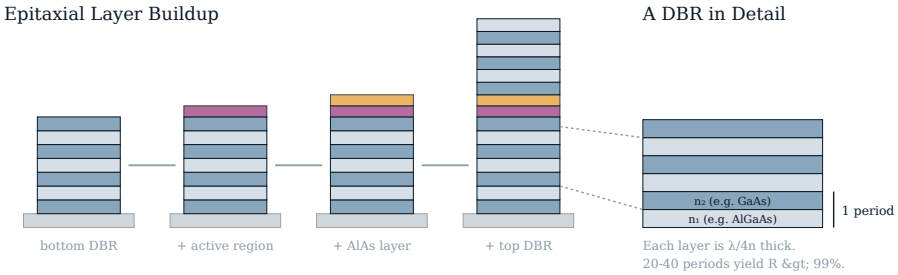

Finished Structure After Mesa Etch, Metallization, and Oxidation

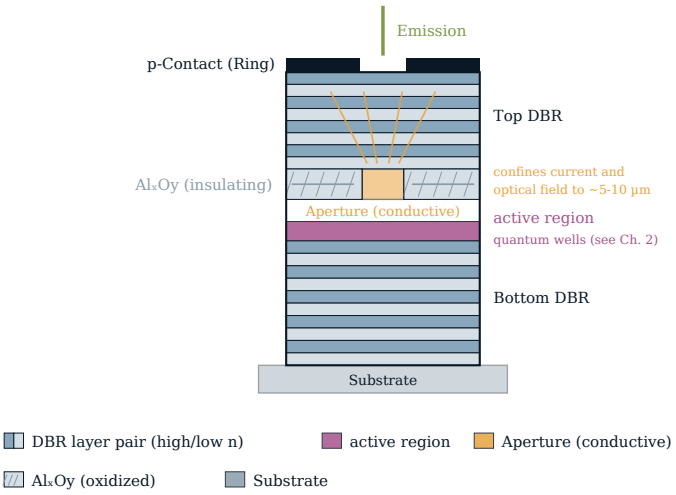

The oxidation front grows inward from the mesa edges, forcing current and optical field through the remaining central opening, just a few µm wide.

Formation of the VCSEL Oxide Aperture (Top View)
Lateral wet oxidation propagates radially inward from the mesa edge

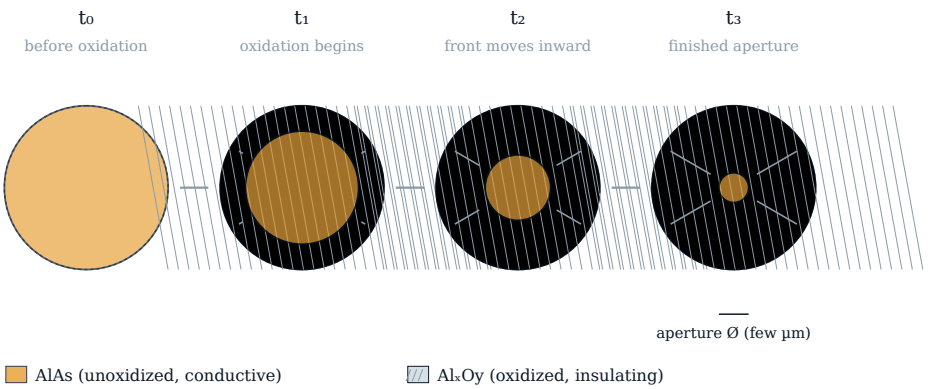

Oxidation time and temperature determine lateral penetration depth and thus the final aperture diameter - central to threshold current and mode behavior.

Facet Formation in Edge Emitters

Edge-emitting laser diodes need two precisely parallel, optically smooth mirror facets at the chip ends for their resonator – unlike planar patterning steps, these can't be produced by etching but are instead created through controlled crystal cleaving along specific low-index crystal planes, such as the (110) plane in GaAs-based structures. The wafer is scored on the backside with a fine, diamond-scribed fracture line; under defined mechanical pressure along this line, the crystal cleaves atomically smooth along the energetically preferred cleavage plane, with no subsequent polishing required – the resulting facet is atomically flat, a prerequisite for low optical scattering losses in the resonator.

The resulting cleaved facets already possess a natural reflectivity of roughly 30% due to the refractive index step between semiconductor ($n \approx 3.5$) and air ($n \approx 1$), per the Fresnel equations – sufficient for simple Fabry-Pérot lasers, but rarely optimal for the desired balance of output power, threshold current, and facet lifetime. The facets are therefore usually additionally coated with dielectric multilayer coatings: a high-reflectivity (HR) coating made of alternating quarter-wavelength layers (e.g., SiO_2/TiO_2 or SiO_2/Ta_2O_5) on the back facet can achieve reflectivities above 95%, maximizing internal resonator gain and thus lowering the threshold current. An anti-reflective (AR) coating on the front facet deliberately reduces reflectivity to a few percent, setting the desired, asymmetric light extraction – the combination of HR back and AR front concentrates the usable output power on one side of the chip. VCSELs skip this entire fabrication step: as described in paragraph 1, the DBR mirrors are

already grown as an integral part of the layer structure during epitaxy, which on one hand reduces the process effort on the finished chip, and on the other hand is precisely what makes the markedly higher mirror reflectivity requirement of over 99% described in Chapter 3 technically achievable in the first place – a dielectric coating applied afterward could not achieve this precision on vertical, epitaxially grown surfaces.

div class="center">

Facet Formation in Edge Emitters
From crystal cleaving to a coated laser facet

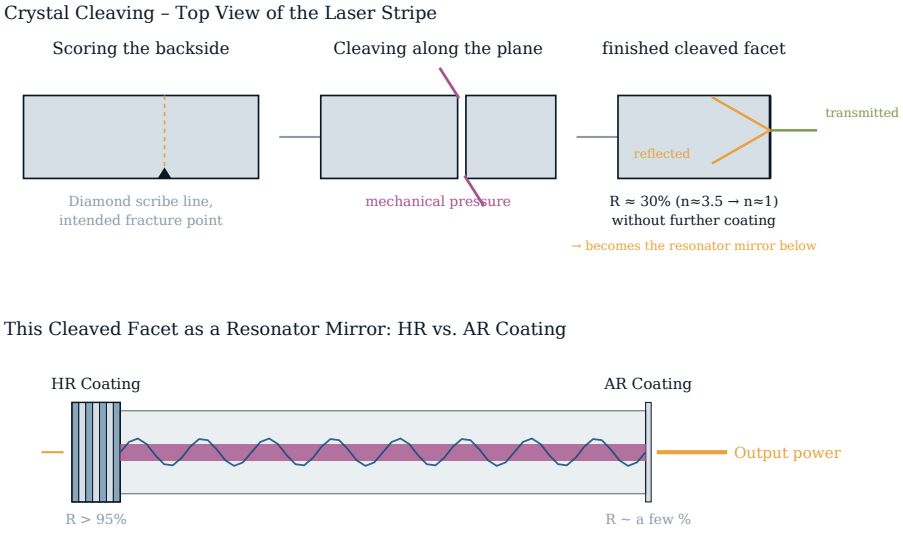

For VCSELs, this entire fabrication step is skipped – the DBR mirrors are already grown epitaxially (see VCSEL chapter).

Crystal cleaving yields atomically smooth facets with a natural reflectivity of ~30%; dielectric coatings then set threshold current and output direction deliberately.

Singulation and Packaging

After patterning and – for edge emitters – facet formation, the wafer still exists as a continuous disc and must be singulated into individual chips. For most III-V devices, especially LEDs and VCSELs, this is done by dicing with a diamond-coated cutting blade, increasingly also by laser scribing followed by breaking along the weakened line, to minimize mechanical damage to the chip edges. Edge emitters, by contrast, are often cleaved a second time along the already-cleaved facet direction, perpendicular to the first cleave, to separate individual laser bars (which initially form as a continuous strip containing many side-by-side lasers) into individual chips.

The singulated chip is then mounted on a submount – usually ceramic (AlN, with high thermal conductivity around 170–200 W/m·K), copper, or diamond for particularly high-power lasers. Two mounting variants dominate, with different thermal and optical consequences: in wire bonding, the chip is glued active-layer-up (usually with conductive silver epoxy or eutectic solder) and the contacts are connected to the package via thin gold wire – simple and low-cost, but with a longer thermal path from the active region through the entire chip substrate to the heat sink. In flip-chip bonding, the chip is inverted and connected directly to the submount via solder bumps (often AuSn solder), placing the active region only a few micrometers from the heat sink – this substantially shortens the thermal path and is standard for high-power LEDs and high-power lasers, as already touched on in the LED chapter's discussion of thermal droop: efficient thermal management is directly decisive for light output and lifetime.

Encapsulation forms the final step, and varies considerably by device type. For white LEDs, the phosphor-silicone mixture (see Chapter 2) is dispensed directly over the chip and thermally cured, often supplemented with a silicone lens for beam shaping and mechanical protection. Laser diodes, by contrast, are often hermetically packaged in a TO can (named after the standardized "transistor outline" form factor) with a protective glass window or integrated collimating lens, to permanently protect the sensitive, uncoated or dielectrically coated facets from moisture, dust, and contamination – organic contamination on the facet can locally burn under high optical power density and permanently damage the facet (COD, Catastrophic Optical Damage), a well-known failure mechanism in high-power lasers.