

UNIVERSITY OF WYOMING

Department of Mechanical Engineering

ME/ERS 4203/5203

Nuclear Materials

Reading Notes

Units 1–7

Instructor: Prof. Sharma

Fall 2026

Contents

Unit 1: Introduction to Nuclear Materials 3

Unit 2: Crystal Structure and Defects 10

Unit 3: Radiation Damage 26

Unit 4: Defect Evolution and Transport 42

Unit 5: Mechanical Properties after Irradiation 58

Unit 6: Nuclear Fuel Materials and Performance 74

Unit 7: Corrosion and Environmental Degradation 91

Unit 1: Introduction to Nuclear Materials

1.1 Nuclear reactor systems 4

1.2 Components and neutron absorption 5

1.3 Mechanical, thermal, and chemical requirements 6

1.4 Radiation damage and degradation 7

Questions 9

Learning Objectives

1. Identify the functions of fuel, cladding, coolant, moderator, and control materials.
2. Use simple heat balances and neutron-absorption calculations.
3. Explain how heat transfer, stress, corrosion, and irradiation affect material selection.

Key quantities

Quantity	Symbol	Common units
Thermal and electric power	P_{th}, P_e	W
Microscopic absorption cross section	σ_a	barn
Macroscopic absorption cross section	Σ_a	cm^{-1}
Fluence	Φ	cm^{-2}

1.1 Nuclear reactor systems

A nuclear reactor produces heat in its fuel. The coolant removes this heat, and a power plant converts part of it into electricity. Materials must carry loads, transfer heat, and retain radioactive material while their properties change during service. A room-temperature strength value alone cannot establish whether a material is suitable.

In a **pressurized water reactor (PWR)**, water serves as both coolant and neutron moderator. The primary coolant transfers heat to a separate water circuit in a steam generator. High primary pressure keeps the bulk coolant liquid; local nucleate boiling can still occur at the fuel-cladding surface. In a **boiling water reactor (BWR)**, water boils in the reactor vessel, and the steam goes to the turbine. These arrangements create different pressure, flow, and water-chemistry conditions (NRC, *Reactor Concepts Manual*).

A moderator slows neutrons mainly by scattering. Thermal reactors use the resulting low-energy neutrons to sustain the chain reaction. Fast reactors avoid strong moderation. Sodium- and lead-cooled fast reactors therefore have different coolant-compatibility and irradiation requirements from water-cooled reactors. Helium-cooled reactors can use a thermal or fast spectrum, depending on their design. A coolant name alone does not define the neutron spectrum. Molten-salt concepts also include different fuel arrangements and spectra. The six Generation IV families are sodium-cooled fast, lead-cooled fast, gas-cooled fast, very-high-temperature, supercritical-water-cooled, and molten-salt reactors. These names describe system families; their material requirements depend on the selected temperature, coolant chemistry, spectrum, and component function.

Fusion joins light nuclei. Deuterium–tritium fusion produces helium and a neutron with about 14.1 MeV of energy. Plasma-facing components experience heat and particle loads; surrounding structures also experience neutron damage. The plasma temperature is not the temperature of the wall. After shutdown, radioactive decay continues to release heat. Spent fuel and waste therefore remain materials problems: containers must remove decay heat and resist corrosion and mechanical damage under the specified storage or disposal conditions (Murty and Charit, 2013, Ch. 1).

For simple energy balances, thermal efficiency is $\eta = P_e/P_{\text{th}}$, where P_e is electric power and P_{th} is thermal power. Rejected heat is $P_{\text{th}} - P_e$. Energy delivered over time t at constant power is $E = P_e t$. For a single-phase coolant with constant heat capacity,

$$P_{\text{th}} = \dot{m} c_p (T_{\text{out}} - T_{\text{in}}).$$

Here $\dot{m}$ is mass flow rate and c_p is specific heat capacity. MW is power; MWh is energy. This coolant equation excludes boiling and changes in kinetic energy.

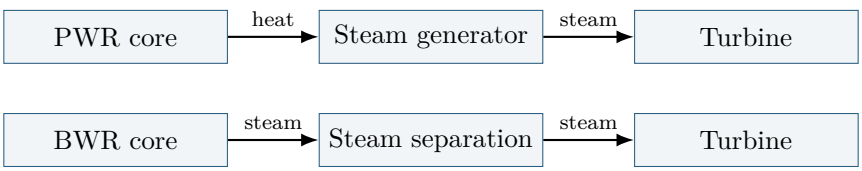


Figure 1.1: Heat and steam paths. PWR turbine steam comes from a separate water circuit; BWR steam comes from the reactor vessel. Return flows are omitted.

Example: Heat removal

Assume $\dot{m} = 1000$ kg/s, $c_p = 5$ kJ/(kg K), and a 20 K coolant rise. The removed heat is $1000(5)(20) = 100{,}000$ kJ/s $= 100$ MW. At $\eta = 0.30$, electric power is 30 MW and rejected heat is 70 MW.

1.2 Components and neutron absorption

Each component has a different function. Fuel produces heat and retains part of the fission-product inventory. Cladding separates fuel from coolant and provides another barrier to release. A pressure vessel carries pressure; its surface cladding improves corrosion resistance. Internals support and position components. Control materials absorb neutrons deliberately. A material useful in a control rod may be undesirable around the fuel because of the neutrons it absorbs.

Component	Typical material	Main requirements
LWR fuel	UO_2 ceramic	Heat generation, stability, fission-product retention
Fuel cladding	Zirconium alloys	Low neutron absorption, corrosion resistance, strength
Pressure vessel	Low-alloy steel	Strength and fracture resistance
Core internals	Stainless steels	Structural support and coolant compatibility
Steam-generator tubes	Nickel alloys	Heat transfer and resistance to corrosion/cracking
Control absorbers	B_4C; Ag–In–Cd in selected designs	Neutron absorption and dimensional stability

These are examples, not interchangeable specifications. Composition, processing, welds, irradiation history, and operating conditions affect performance (Murty and Charit, 2013, Ch. 1).

The microscopic absorption cross section σ_a describes the probability of neutron absorption by a nucleus at a specified neutron energy. It is not the geometric area of the nucleus. One barn equals 10^{-24} cm^2. Natural-element values depend on isotopic composition. At a neutron speed of 2200 m/s, NIST lists approximately 0.185 barn for natural Zr and 2.56 barn for natural Fe. This difference helps explain the use of zirconium alloys as cladding.

For an elemental solid of density ρ and molar mass M, the atomic number density is $N = \rho N_A/M$, with $N_A = 6.022 \times 10^{23}$ mol^{-1}. Using g/cm^3 and g/mol gives atoms/cm^3. The macroscopic absorption cross section is

$$\Sigma_a = N\sigma_a, \qquad \Sigma_a^{\text{alloy}} = \sum_j N_j \sigma_{a,j}.$$

Σ_a has units of cm^{-1}. An average microscopic cross section uses *atomic* fractions: $\bar{\sigma}_a = \sum_j x_j \sigma_{a,j}$. Mass fractions must first be converted using $x_j = (w_j/M_j)/\sum_k (w_k/M_k)$.

Example: Absorption in a layer

Assume $N = 4.0 \times 10^{22}$ atoms/cm^3 and $\sigma_a = 0.20$ barn. Then $\Sigma_a = 0.0080$ cm^{-1}. For straight neutron paths through a uniform, absorption-only layer of thickness h, the absorbed fraction is $1 - e^{-\Sigma_a h}$. At $h = 0.05$ cm it is 0.00040, or 0.040%. Actual reactor transport also includes scattering and an energy spectrum.

1.3 Mechanical, thermal, and chemical requirements

Material selection starts with component function and service conditions. The selected material must meet several requirements together. Low neutron absorption does not compensate for inadequate strength, and high strength does not prevent corrosion. Manufacturing, joining, inspection, cost, and available test data also affect selection.

Strength describes resistance to plastic deformation or failure under a specified test. **Ductility** measures plastic deformation before fracture. **Fracture toughness** describes resistance to crack extension under specified conditions. A strong material can still fracture with little plastic deformation. **Creep** is time-dependent deformation under sustained load; **fatigue** involves damage under repeated loading. Temperature and irradiation can alter all these responses (Murty and Charit, 2013; Was, 2007).

For a straight specimen under axial force F, nominal stress is $\sigma = F/A$, using initial area A. Since 1 N/mm^2 = 1 MPa, this form is convenient for simple calculations. For a thin cylindrical tube with internal-to-external pressure difference Δp, mean radius r, and wall thickness $h \ll r$, the approximate hoop stress is $\sigma_\theta = \Delta p\, r/h$. It neglects local stress concentrations and is not a pressure-vessel design rule.

Heat must pass from fuel through the gap and cladding to the coolant. A low thermal conductivity raises the temperature difference required to carry a given heat flux. For steady one-dimensional conduction across a flat layer with constant conductivity k and no internal heat generation,

$$\Delta T = \frac{q'' h}{k}, \qquad R'' = \frac{h}{k}.$$

Here q'' is heat flux in W/m^2, h is thickness in m, and R'' is thermal resistance per unit area in m^{2}K/W. Temperature drops add across layers carrying the same heat flux. This flat-layer result approximates a thin tube wall; fuel with internal heat generation needs a different equation.

Corrosion removes material or changes the surface. Localized corrosion can penetrate deeply even when average mass loss is small. Coolant chemistry and temperature affect oxide stability and transport. Reducing wall thickness improves conduction but can increase stress and reduce the amount of corrosion that a component can tolerate. Surface condition and material properties must therefore be evaluated together.

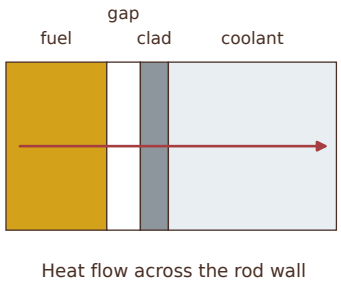


Figure 1.2: Fuel, gap, cladding, and coolant form the radial heat-transfer path.

Example: Thickness trade-off

Assume $q'' = 0.4$ MW/m^2, $h = 0.50$ mm, and $k = 20$ W/(m K). Then $\Delta T = 10$ K. With $r = 5$ mm and $\Delta p = 4$ MPa, the tube estimate gives $\sigma_\theta = 40$ MPa. Doubling h doubles the conduction temperature drop and halves the estimated hoop stress.

1.4 Radiation damage and degradation

An energetic particle can transfer energy to a lattice atom and displace it. The displaced atom may cause further collisions. This sequence is a **displacement cascade**. A vacant lattice site is a vacancy; an extra atom between normal sites is an interstitial. Some vacancies and interstitials recombine. Others move to surfaces, grain boundaries, and dislocations or join defect clusters. Thus the defects left in the material differ from the total number initially produced.

Neutron **flux** ϕ measures the neutron population multiplied by speed, summed over directions and energies. Fluence is its time integral; at constant flux, $\Phi = \phi t$. Common units are cm^{-2}s^{-1} for flux and cm^{-2} for fluence. **Displacements per atom (dpa)** is a model-based measure of displacement exposure. At constant dpa rate, $D = \dot{D}t$. It counts displacement events relative to the reference number of atoms, not the fraction of atoms currently vacant. Two exposures with the same dpa can produce different defects because temperature, dose rate, spectrum, and gas production differ (Was, 2007).

Defect clusters can hinder dislocation motion and raise yield strength. The same microstructural changes can reduce ductility or shift the ductile-to-brittle transition temperature of ferritic steels upward. A transition shift is measured using a stated test criterion; it is not a universal fracture temperature. For a simple comparison, define $M_T = T_{\text{service}} - T_{\text{tr}}$. A smaller M_T indicates less temperature separation from that criterion, but does not establish component safety.

Voids can cause **swelling**, $S = (V - V_0)/V_0$. For small isotropic swelling, the linear strain is approximately $S/3$; exactly, $L/L_0 = (1+S)^{1/3}$. Irradiation creep involves stress-dependent strain during irradiation. Irradiation growth can change shape without an applied stress, particularly in anisotropic materials. These are distinct contributions to dimensional change.

Corrosion and cracking can occur alongside irradiation. Stress corrosion cracking requires a susceptible material, a suitable environment, and tensile loading, including residual stress. At

Davis-Besse in 2002, leakage of borated water associated with vessel-head nozzle cracking caused severe corrosion of the head steel. The case shows why inspection of leakage, deposits, wall loss, and cracks must accompany material selection (NRC Bulletin 2002-01).

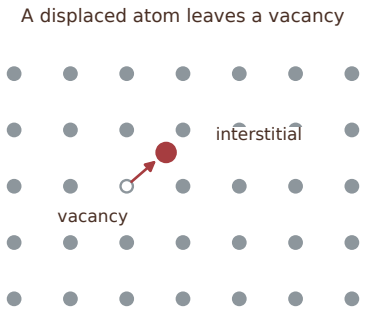


Figure 1.3: A displaced lattice atom can leave a vacancy and become an interstitial.

Example: Exposure and dimensions

At an assumed $\dot{D} = 2 \times 10^{-7}$ dpa/s for 10^7 s, $D = 2$ dpa. If a separate measurement gives $S = 0.006$, the small-strain estimate is $S/3 = 0.002$. A 100 mm specimen length increases by about 0.20 mm. The swelling was measured, not deduced from dpa alone.

References. K. L. Murty and I. Charit, *An Introduction to Nuclear Materials*, Wiley-VCH (2013), Ch. 1; G. S. Was, *Fundamentals of Radiation Materials Science*, Springer (2007). NRC, *Reactor Concepts Manual*, and Bulletin 2002-01. NIST NCNR, neutron cross-section table (2200 m/s).

Unit 1 questions

Conceptual questions

1.C1 Where is steam produced in a PWR and in a BWR?

1.C2 How does a moderator differ from a coolant?

1.C3 Why can a good neutron absorber be a poor cladding material?

1.C4 Why can a thinner wall improve heat transfer but worsen mechanical performance?

1.C5 Why is dpa not the surviving vacancy fraction?

Numerical problems

1.N1 A plant produces 900 MW electric from 3000 MW thermal. Find its efficiency.

1.N2 For $N = 5.0 \times 10^{22}$ cm^{-3} and $\sigma_a = 2.0$ barns, calculate Σ_a.

1.N3 Find fluence for $\phi = 10^{14}$ cm^{-2}s^{-1} over 2×10^6 s.

Combined problems

1.P1 Assume $\dot{m} = 2000$ kg/s, $c_p = 5$ kJ/(kg K), and $\Delta T = 25$ K. Find thermal power, then electric power at 32% efficiency.

1.P2 A tube has $r = 5$ mm, $h = 0.40$ mm, $k = 20$ W/(m K), and $\Delta p = 4$ MPa. At $q'' = 5 \times 10^5$ W/m^2, find its hoop stress and wall temperature drop.

Unit 2: Crystal Structure and Defects

2.1 Lattices, unit cells and metallic structures 11

2.2 Crystallographic directions and planes 13

2.3 X-ray diffraction and crystal geometry 15

2.4 Interstitial sites and ceramic structures 17

2.5 Point defects and equilibrium vacancies 19

2.6 Dislocations, grain boundaries and stacking faults 21

Questions 25

Learning Objectives

1. Count atoms and sites in unit cells, and relate crystal geometry to density.
2. Index directions and planes, and use diffraction peaks to determine lattice spacing.
3. Distinguish point defects, dislocations, grain boundaries, and stacking faults.
4. Estimate equilibrium vacancy populations and simple dislocation dimensions.

Key quantities

Quantity	Symbol	Common units
Lattice parameter and plane spacing	a, d_{hkl}	nm or Å
Vacancy fraction	C_v	dimensionless
Formation and migration energies	E_f, E_m	eV
Burgers-vector magnitude	b	nm

2.1 Lattices, unit cells and metallic structures

A crystal has long-range periodic order: the atomic arrangement repeats through translations. An amorphous solid can have well-defined local neighbours without this long-range repetition. A lattice is a periodic set of equivalent points. A basis is the atom or group of atoms attached to every lattice point. Their combination specifies the crystal structure. The basis matters because identical geometry with different species can give different properties.

Definition: Lattice, basis and crystal

A **lattice** gives equivalent positions related by translation. A **basis** specifies the atoms attached to each position. A **crystal structure** combines the two. Translating one complete basis to every lattice point reproduces both geometry and chemical identity.

A unit cell reproduces the crystal when translated along its three edge vectors. A primitive cell contains one lattice point; it need not contain one atom. A conventional cell may contain several lattice points because its larger shape displays the symmetry clearly. Three lengths a, b, c and three angles describe the cell: α lies between b and c, β between a and c, and γ between a and b. A cubic cell has equal lengths and three right angles.

There are seven crystal systems and fourteen Bravais lattices. Cubic, tetragonal and orthorhombic conventional cells have three right angles; their lengths are respectively all equal, two equal, and independent. Hexagonal axes have $a = b$, $\gamma = 120°$ and $\alpha = \beta = 90°$. Rhombohedral axes for the trigonal system have equal lengths and equal angles. Monoclinic cells have two right angles; triclinic lengths and angles are independent. Symmetry defines the system, so accidental equality of measured lengths is not sufficient to establish it. The centring symbols P, I, F, C and R mean primitive, body-centred, face-centred, base-centred and rhombohedral.

2.1.1 Atom counting and contact geometry

An atom on a corner of a cubic cell contributes $1/8$, a face atom contributes $1/2$, and an interior atom contributes one. Body-centred cubic (BCC) therefore contains $8/8 + 1 = 2$ atoms per conventional cell. Face-centred cubic (FCC) contains $8/8 + 6/2 = 4$. Each BCC atom has eight nearest neighbours; each FCC atom has twelve. In an equal hard-sphere model, BCC atoms touch along a body diagonal and FCC atoms along a face diagonal:

$$\text{BCC: } \sqrt{3}a = 4r, \qquad \text{FCC: } \sqrt{2}a = 4r.$$

Here r is the model atomic radius. Real atoms do not have rigid spherical boundaries, but these relations are useful descriptions of ideal geometry.

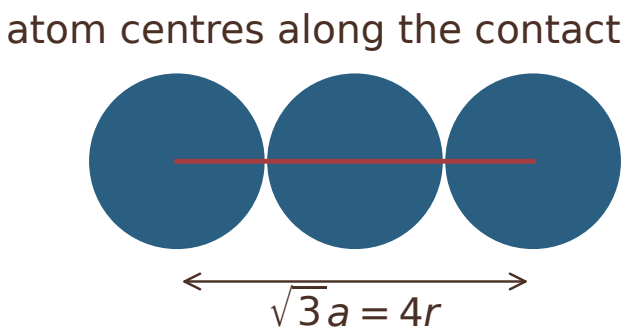


Figure 2.1: BCC contact occurs along the body diagonal, not along a cube edge.

Follow the diagonal from one corner through the centre to the opposite corner. Along this line, the centre-to-centre distances total four radii. Along an edge the corner atoms do not touch in this model. The contact path therefore determines the radius–lattice-parameter relation before the sphere volumes are counted.

The atomic packing factor is the fraction of cell volume occupied by the model spheres:

$$\mathrm{APF} = \frac{n(4\pi r^3/3)}{V_{\mathrm{cell}}}.$$

Using the matching atom count and cell volume gives $\pi\sqrt{3}/8 = 0.680$ for BCC and $\pi/(3\sqrt{2}) = 0.740$ for FCC. BCC has dense $\{110\}$ planes but no close-packed plane. FCC has close-packed $\{111\}$ planes.

Key point: count and volume must match

A larger choice of unit cell contains proportionally more atoms and volume. It leaves density and packing fraction unchanged. Changing the atom count without changing the corresponding volume does not describe a different physical crystal; it is a bookkeeping error.

2.1.2 Hexagonal packing and density

Hexagonal close-packed (HCP) layers follow ABAB, whereas FCC layers follow ABCABC. Ideal HCP has twelve nearest neighbours and the same 0.740 packing fraction as FCC. Its ideal geometry gives $a = 2r$ and $c/a = \sqrt{8/3} = 1.633$. Real metals can depart from this ratio. HCP is a hexagonal lattice with a two-atom basis. Its primitive parallelepiped contains two atoms; the large hexagonal prism contains six and has three times the volume. Mixing these cell conventions gives an incorrect density.

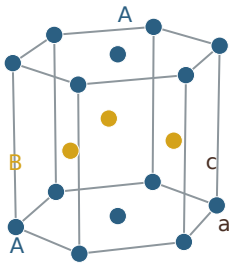


Figure 2.2: The large HCP prism contains six atoms after sharing is included.

In the stacking picture, B-layer atoms sit in one set of hollows of an A layer. A third layer can return to the A positions, producing HCP, or occupy the other set of hollows, producing the ABC sequence of FCC. Both sequences preserve the same local close packing while changing the longer-range repetition.

Mass density follows from mass per cell divided by cell volume:

$$\rho = \frac{nM}{N_A V_{\text{cell}}}, \qquad N_A = 6.022 \times 10^{23}\ \text{mol}^{-1}.$$

For an elemental crystal, M is molar atomic mass. Use M in g/mol and V in cm^3 for ρ in g/cm^3; 1 nm $= 10^{-7}$ cm. For a compound, use the number of formula units and the molar formula mass. Packing fraction alone does not rank the mass densities of different materials.

Example: Density from an FCC radius

Given. An ideal elemental FCC crystal has $r = 0.140$ nm and molar mass $M = 60.0$ g/mol. Find its theoretical density.

Model. Use four atoms per conventional cell and contact along the face diagonal: $a = 2\sqrt{2}r$. The numerator in $\rho = nM/(N_A a^3)$ is the mass of those four atoms.

$$a = 2\sqrt{2}(0.140) = 0.396\ \text{nm} = 3.960 \times 10^{-8}\ \text{cm},$$

$$\rho = \frac{4(60.0)}{(6.022 \times 10^{23})(3.960 \times 10^{-8})^3} = 6.42\ \text{g cm}^{-3}.$$

Meaning. This is the density of the ideal structure under the stated assumptions. The length conversion is cubed with the cell edge. Using two atoms with this FCC volume would halve the answer, while using an edge-contact relation would give the wrong volume.

2.2 Crystallographic directions and planes

A direction describes an arrow, while a plane describes an orientation and position within a family of parallel planes. Square brackets $[uvw]$ denote one direction and angle brackets $\langle uvw \rangle$ its symmetry-related family. Parentheses (hkl) denote a plane or indexed set of parallel planes;

braces $\{hkl\}$ denote symmetry-related plane orientations. A bar indicates a negative index: $[\bar{1}10]$ has components $(-1, 1, 0)$. These distinctions prevent a line, a plane and a family from being confused.

To index a direction, subtract its start coordinates from its end coordinates, expressing both in cell-edge units. Clear fractions and reduce the resulting integers. The location of the arrow is irrelevant; its displacement determines the direction. For example, an arrow from $(1, 0, 0)$ to $(0, 1, 0)$ is $[\bar{1}10]$. Reversing it reverses all signs. Direction indices do not specify a length. In a cubic crystal the actual displacement $a(u, v, w)$ has length $a\sqrt{u^2 + v^2 + w^2}$; a prefactor must be retained if the physical displacement is shorter.

Key point: notation identifies the object

$[uvw]$ specifies one direction; $\langle uvw \rangle$ collects equivalent directions. (hkl) specifies a plane orientation with an indexed spacing; $\{hkl\}$ collects equivalent orientations. For cubic crystals, [100] and [010] belong to $\langle 100 \rangle$, but they are different arrows.

2.2.1 Miller indices and spacing

To obtain Miller indices, write the plane intercepts in units of a, b, c, take their reciprocals and clear fractions. An infinite intercept has reciprocal zero, so a zero index means that the plane is parallel to that axis. A negative intercept gives a negative index. If the plane contains the origin, choose a parallel plane or move the origin before taking reciprocals. A finite triangular patch in a drawing represents part of an infinite plane.

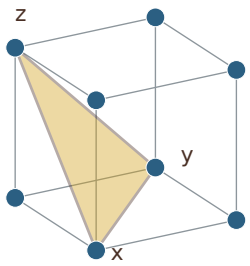


Figure 2.3: A cubic (111) plane intercepts all three positive axes at one cell edge.

The triangular patch joins the three intercepts; it is not the boundary of the physical plane. Extending that patch gives the same orientation throughout the crystal. A family of parallel planes is obtained by translating it along its normal, which is [111] for this cubic example.

The planes (100) and (200) have the same orientation, but the indexed spacings are different. In diffraction, do not reduce (200) to (100): the former has twice the reciprocal spacing. For a cubic crystal,

$$d_{hkl} = \frac{a}{\sqrt{h^2 + k^2 + l^2}}.$$

Here d_{hkl} is a perpendicular spacing, not the distance measured obliquely along an arbitrary axis. Increasing the magnitude of the indices decreases this spacing for fixed a. Thus $d_{100} = a$, $d_{110} = a/\sqrt{2}$, $d_{111} = a/\sqrt{3}$ and $d_{200} = a/2$.

In cubic crystals $[hkl]$ is normal to (hkl). A direction $[uvw]$ lies within a plane (hkl) if

$$hu + kv + lw = 0.$$

This is especially useful when checking slip geometry. For example, $[1\bar{1}0]$ lies in (111) because $1 - 1 + 0 = 0$. It does not follow that every close-packed direction belongs to every close-packed plane.

The angle between cubic plane normals follows from the vector dot product:

$$\cos\psi = \frac{h_1h_2 + k_1k_2 + l_1l_2}{\sqrt{h_1^2 + k_1^2 + l_1^2}\sqrt{h_2^2 + k_2^2 + l_2^2}}.$$

For unoriented planes, use the acute angle between their normals when requested. The same-index normal rule and this simple Euclidean index formula must not be transferred unchanged to a noncubic lattice, whose axes can have unequal lengths or non-right angles.

2.2.2 Hexagonal notation

Hexagonal crystals commonly use three basal axes a_1, a_2, a_3 separated by 120°, together with a perpendicular c axis. Four-index planes are $(hkil)$ with $i = -(h + k)$; directions are $[uvtw]$ with $t = -(u + v)$. The redundant basal index displays hexagonal symmetry. The basal plane is (0001) and the c direction is $[0001]$. A plane $(10\bar{1}0)$ is prismatic. Four-index components use the hexagonal basis, so they cannot be treated as four mutually perpendicular Cartesian components.

Example: From intercepts to spacing

Given. A plane in a cubic crystal with $a = 0.500$ nm has intercepts $(1, 3, \infty)$ in cell-edge units. Find its indices, spacing and a direction in the plane.

Model. Miller indices are proportional to reciprocal intercepts. The infinite intercept makes the third index zero.

$$(1, 3, \infty) \longrightarrow (1, \tfrac{1}{3}, 0) \longrightarrow (310), \qquad d_{310} = \frac{0.500}{\sqrt{3^2 + 1^2 + 0^2}} = 0.158 \text{ nm}.$$

For $[1\bar{3}0]$, the plane test is $hu + kv + lw = 3(1) + 1(-3) + 0(0) = 0$, so this direction lies in the plane. In a cubic crystal $[310]$ is instead normal to it.

Meaning. The plane is parallel to the z axis, but its perpendicular spacing is not simply its x or y intercept. The indices carry orientation information; the lattice parameter supplies the length scale.

2.3 X-ray diffraction and crystal geometry

X-rays scattered by different atoms retain phase relationships. At some angles their contributions reinforce one another; at others they cancel. Diffraction therefore measures periodic structure rather than simply the total number of atoms illuminated. A peak position provides a spacing.

Its intensity also depends on the positions and types of atoms, orientation distribution and measurement conditions. A tall peak alone does not establish that its planes are the most densely occupied.

2.3.1 Bragg's law

Consider rays scattered from adjacent parallel planes separated by perpendicular distance d. The extra path travelled by the lower ray is $2d\sin\theta$, where θ is measured from the plane. Constructive interference requires this path difference to be an integer number of wavelengths:

$$n\lambda = 2d\sin\theta.$$

The order n is a positive integer; λ and d must use matching length units. Most indexed powder-pattern calculations here use $n = 1$, with higher indexed reflections carrying the smaller spacings explicitly. The incident and diffracted beams enclose the scattering angle 2θ. An XRD graph normally plots intensity against 2θ, so its horizontal coordinate must be halved before substitution.

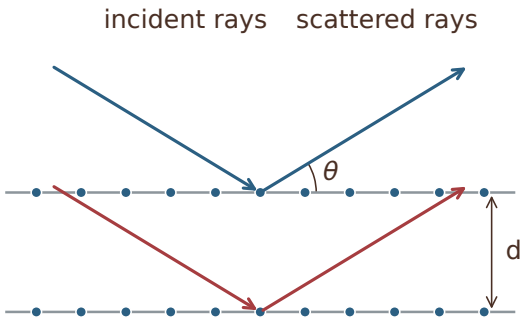


Figure 2.4: The path difference from neighbouring planes is $2d\sin\theta$; the scattering angle is 2θ.

A measured peak can be converted to a spacing using $d = \lambda/(2\sin\theta)$. If a cubic reflection is known to be (hkl), then $a = d\sqrt{h^2 + k^2 + l^2}$. A single peak with unknown indices generally does not determine a unique lattice parameter: several assignments are possible. Several consistent reflections and appropriate selection rules make the interpretation stronger.

At fixed wavelength and fixed reflection indices, lattice expansion increases d and shifts the peak to smaller 2θ. Contraction does the opposite. This conclusion follows directly from Bragg's law. It does not require an assumption about the peak height. A proposed reflection also must satisfy $n\lambda \leq 2d$, since a real sine cannot exceed one.

Key point: peak position and intensity

Position enters Bragg's law and determines a spacing once λ is known. **Intensity** depends on interference within the unit cell as well as the measurement. An absent reflection can therefore reveal the basis even when its spacing is geometrically possible.

2.3.2 Structure factors and systematic absences

Meeting Bragg's geometric condition is necessary but not sufficient for a reflection. Contributions from different atoms within a unit cell may cancel. The structure factor is their complex sum:

$$F_{hkl} = \sum_j f_j \exp\left[2\pi i(hx_j + ky_j + lz_j)\right].$$

Here f_j is the atomic scattering factor and (x_j, y_j, z_j) are fractional coordinates. The phase term records where each atom sits relative to the indexed planes. Cancellation gives $F_{hkl} = 0$, a systematic absence in the ideal structure. This is distinct from a weak allowed peak falling below an instrument's detection limit.

For identical atoms at the BCC corner and body-centre positions, the phase difference gives $F = f[1 + (-1)^{h+k+l}]$. If the sum is even, $F = 2f$; if it is odd, $F = 0$. Thus (110), (200) and (211) are allowed, whereas (100) and (111) are absent. For conventional elemental FCC, indices must be all even or all odd. Consequently (111), (200) and (220) are allowed, while (100) and (110) are absent. Zero counts as even. These rules describe the specified lattice and basis. A multi-atom basis can add further absences and can change relative intensities; elemental rules alone do not fully describe every compound.

Example: Interpreting an indexed diffraction peak

Given. Cu Kα_1 radiation has $\lambda = 1.5406$ Å. A peak at $2\theta = 60.0°$ is assigned to elemental FCC (200), using $n = 1$.

Model. First obtain d from Bragg's law, then use the cubic spacing relation to obtain a.

$$\theta = 30.0°, \qquad d = \frac{1.5406}{2\sin 30.0°} = 1.5406 \text{ Å},$$

$$a = d\sqrt{2^2 + 0^2 + 0^2} = 3.0812 \text{ Å}.$$

Interpretation. The all-even (200) indices satisfy the FCC selection rule. An FCC (110) assignment would not, even though inserting its indices into the spacing formula would produce a numerical lattice parameter. Geometry and the basis must both agree with an assignment. Further allowed peaks should give the same a within measurement uncertainty.

2.4 Interstitial sites and ceramic structures

An interstitial site is a geometrical position between normal atom sites. It is not automatically a defect. A perfect elemental crystal can have empty interstitial sites, while a perfect compound can have one species regularly occupying sites defined relative to another sublattice. An interstitial defect instead means an extra atom relative to the specified perfect structure. This distinction is essential when moving from metal structures to ceramics.

A tetrahedral site is surrounded by four atoms; an octahedral site has six surrounding atoms. A conventional FCC cell has eight tetrahedral and four octahedral sites. A conventional BCC cell has twelve tetrahedral and six octahedral sites. BCC sites are distorted: surrounding distances

are not all equivalent as in an ideal regular polyhedron. A geometric hole size can guide a hard-sphere picture, but bonding, electronic structure and lattice relaxation determine the defect energy. The largest geometrical hole need not host the most stable self-interstitial. A relaxed split configuration may be preferred instead.

2.4.1 Species, occupancy and coordination

In fluorite UO_2, uranium forms an FCC sublattice and oxygen occupies all eight tetrahedral sites of the conventional cell. The net contents are four U atoms and eight O atoms, or four UO_2 formula units. Each U has eight O neighbours; each O has four U neighbours. The two coordination numbers differ because the populations differ. Counting the same U–O links from either end gives $4(8) = 8(4) = 32$ links per cell.

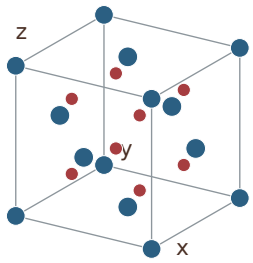


Figure 2.5: Fluorite UO_2: four uranium atoms (blue) and eight oxygen atoms (red) per conventional cell.

The corner and face-centred uranium positions must be shared with adjacent cells before counting. The eight oxygen positions lie within the cell. These two counting rules explain why a drawing with many visible uranium spheres still represents only four uranium atoms per conventional cell.

Coordination describes neighbours around one atom; site count describes the number of available positions in a cell. Neither is interchangeable with the other. For instance, saying that oxygen occupies tetrahedral sites means that each oxygen has four uranium neighbours. It does not mean that uranium has four oxygen neighbours. The atom counts and neighbour-link check provide independent ways to avoid reversing this relationship.

Rock salt has one species on an FCC sublattice and the other on all four octahedral sites. Its conventional cell contains four of each species, giving MX and four formula units. Each species has six unlike nearest neighbours. UC and UN are nuclear-material examples of this structure type. Zinc blende instead fills half of the eight tetrahedral positions, again giving four M and four X atoms but fourfold coordination. Equal chemical ratios do not require equal structures.

Antifluorite reverses the species roles of fluorite, giving M_2X, as in Li_2O. Wurtzite is a hexagonal tetrahedral structure; zinc blende is cubic. SiC occurs in multiple polytypes, so the stacking sequence must be specified. Other relevant formula types are perovskite ABX_3 and spinel AB_2X_4. A formula alone does not supply all atomic coordinates or establish which structure is stable.

Key point: composition is a ratio

Cell contents M_4X_8 reduce to the formula MX_2, but the cell still contains four formula units. Coordination counts neighbours around an individual site; it does not reduce or multiply the formula. Keep composition, cell contents and coordination as separate quantities.

2.4.2 Composition and density

For a compound with Z formula units per cell and molar formula mass M_f,

$$\rho = \frac{ZM_f}{N_A V_{\text{cell}}}.$$

Do not insert total atom count in place of Z while still using M_f; that would count the mass repeatedly. For mixed species, the equivalent numerator is the sum of each species' atom count times its molar atomic mass.

Nonstoichiometry means departure from the ideal chemical ratio. UO_{2-x} is oxygen deficient and UO_{2+x} oxygen rich. Oxygen deficiency can involve oxygen vacancies; excess can involve oxygen interstitials and clusters. Uranium charge states and defect populations adjust to preserve overall charge balance. Counting missing oxygen alone does not fully specify the charge compensation or the equilibrium defect arrangement.

Example: Formula-unit counting in fluorite

Given. An ideal MX_2 fluorite cell has $a = 0.600$ nm and molar formula mass $M_f = 120$ g/mol. Find its density, then examine 75% occupancy of the X sites.

Model. Four M atoms and eight X atoms make four MX_2 formula units, so $Z = 4$.

$$\rho = \frac{4(120)}{(6.022 \times 10^{23})(6.00 \times 10^{-8})^3} = 3.69 \text{ g cm}^{-3}.$$

At fractional occupancy 0.75, the mean X count is $0.75(8) = 6$. With all four M sites occupied, the ratio is M_4X_6, or $MX_{1.5}$.

Meaning. The calculated density applies to the fully occupied MX_2 cell. A density for the deficient cell would need its actual mass and volume; the original formula mass cannot simply be reused. Fractional occupancy describes an average across cells. The count alone does not determine whether the proposed deficient structure and its charge compensation are stable.

2.5 Point defects and equilibrium vacancies

A point defect is localized near one or a few atomic sites. A vacancy is an unoccupied normal site; neighbouring atoms relax around it. An atom can jump into the vacancy, leaving its old site vacant. Thus atom and vacancy displacements are opposite in an individual exchange. A vacancy is a missing-atom configuration, not an extra particle moving through the crystal.

A substitutional solute replaces the usual species on a normal site. An interstitial solute lies between normal sites. A self-interstitial is an additional atom of the host species, and its relaxed arrangement need not coincide with an ideal tetrahedral or octahedral hole. In a dumbbell configuration two atoms share the region around one normal lattice site. BCC iron commonly provides a $\langle 110 \rangle$ dumbbell example, but preferred orientations are material-dependent. One BCC metal's stable defect configuration is not a universal rule for all BCC metals.

A Frenkel pair consists of a vacancy and a self-interstitial created when an atom leaves its normal site but remains in the crystal. Total atom count is conserved. Vacancy and interstitial counts initially increase equally, although subsequent recombination or capture at different sinks can change the surviving populations. Energetic irradiation can produce these defects far above thermal equilibrium concentrations.

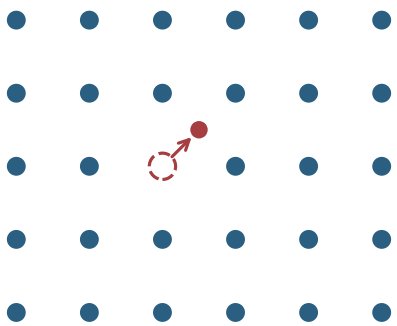


Figure 2.6: A displaced atom leaves a vacancy and becomes a self-interstitial; the total atom count is unchanged.

Definition: Formation and migration

Formation compares a defect-containing state with a specified reference and atom reservoir. **Migration** describes the barrier between two positions of an existing defect. A population calculation needs the former; a jump-rate calculation needs the latter.

2.5.1 Formation and migration

Formation energy E_f measures the energy cost of making a defect with its atom bookkeeping specified. Migration energy E_m is the barrier for an existing defect to move between configurations. A low migration barrier does not mean a low cost of creating the defect. Likewise, a large equilibrium population does not by itself establish rapid transport. Population and mobility are different physical quantities.

For dilute, noninteracting vacancies in thermal equilibrium, the vacancy fraction is approximately

$$C_v^{\mathrm{eq}} = \frac{N_v}{N_s} \simeq \exp\left(\frac{S_f}{k_B}\right) \exp\left(-\frac{H_f}{k_B T}\right).$$

Here N_v is the number of vacant normal sites, N_s the total number of normal sites, H_f formation enthalpy, S_f formation entropy, T absolute temperature and k_B Boltzmann's constant. The exponent is dimensionless. Use $k_B = 8.617 \times 10^{-5}$ eV/K when formation energies are in eV. Kelvin, rather than degrees Celsius, belongs in the exponential.

For illustrative estimates we set the entropy prefactor to one and replace H_f by E_f, giving $C_v \approx \exp[-E_f/(k_B T)]$. Increasing temperature increases this fraction, while increasing formation energy decreases it. The number of vacancies follows from $N_v = C_v N_s$; the reciprocal $1/C_v$ is the average number of sites per vacancy. Neither number is a nearest-neighbour distance. The formula assumes equilibrium and cannot be used as the total defect population during irradiation without considering production and loss.

The temperature dependence is exponential rather than linear. With the same prefactor at temperatures T_1 and T_2,

$$\frac{C_v(T_2)}{C_v(T_1)} = \exp\left[\frac{E_f}{k_B}\left(\frac{1}{T_1} - \frac{1}{T_2}\right)\right].$$

If $T_2 > T_1$, the exponent is positive. This ratio is a useful way to compare temperatures without introducing the sample volume: it changes the fraction first, and the normal-site count then converts that fraction into a vacancy number.

2.5.2 Stoichiometric vacancy sets

A Schottky set removes ions in the compound's stoichiometric ratio to a reservoir such as the surface, leaving vacancies in the interior. In monovalent MX, one cation vacancy and one anion vacancy form a neutral set. In ideal UO_2, one U vacancy and two O vacancies form a neutral set. Relative to the occupied sites their effective charges are -4 and $+2$, so $-4 + 2(+2) = 0$. A Schottky set therefore does not always contain two vacancies. Unlike an interior Frenkel rearrangement, it changes the interior atom count while preserving the formula ratio of the removed ions.

Example: Equilibrium vacancy population

Given. Take $E_f = 0.80$ eV, $T = 900$ K, $N_s = 2.0 \times 10^{12}$ normal sites and unit entropy prefactor.

Model. For the dilute equilibrium approximation, $C_v = \exp[-E_f/(k_B T)]$ and $N_v = C_v N_s$.

$$k_B T = (8.617 \times 10^{-5})(900) = 0.07755 \text{ eV},$$

$$C_v = \exp\left(-\frac{0.80}{0.07755}\right) = 3.31 \times 10^{-5}, \qquad N_v = (3.31 \times 10^{-5})(2.0 \times 10^{12}) = 6.62 \times 10^{7}.$$

Meaning. Only about 33 sites per million are vacant, although this region contains tens of millions of vacancies. The reciprocal fraction is about 3.02×10^4 sites per vacancy, not a spatial separation. The exponent is negative and dimensionless, and the small fraction is consistent with the dilute approximation. Actual counts fluctuate around this equilibrium mean.

2.6 Dislocations, grain boundaries and stacking faults

A dislocation is a line defect separating slipped and unslipped regions. Its motion lets atoms rearrange locally, allowing plastic deformation at much lower stresses than a rigid displacement

of an entire perfect plane at once. An edge dislocation can be pictured as the termination of an extra half-plane. Its Burgers vector $\mathbf{b}$ is perpendicular to the line direction $\boldsymbol{\xi}$. A screw dislocation has $\mathbf{b}$ parallel to $\boldsymbol{\xi}$ and planes that form a helical ramp around the line.

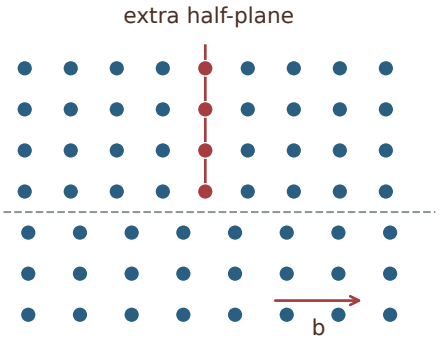


Figure 2.7: An extra half-plane ends at an edge dislocation, viewed along its line.

The disturbance is concentrated around the end of the half-plane. As the dislocation glides, neighbouring bonds rearrange locally and the slipped region grows behind it. The half-plane picture shows the edge character; it does not mean that a complete plane must move rigidly across the sample at once.

The Burgers vector gives the direction and magnitude of the lattice displacement carried by the dislocation. A Burgers circuit compares lattice steps around a dislocation with the same steps in a perfect reference; the closure failure defines $\mathbf{b}$ using a stated sign convention. Along a curved dislocation with no junction, $\mathbf{b}$ stays constant while the local line direction changes. The character can therefore vary continuously from screw through mixed to edge.

Common perfect-dislocation Burgers vectors are $(a/2)\langle 111\rangle$ in BCC, of magnitude $a\sqrt{3}/2$, and $(a/2)\langle 110\rangle$ in FCC, of magnitude $a/\sqrt{2}$. HCP a-type dislocations have $(1/3)\langle 11\bar{2}0\rangle$, of magnitude a in the hexagonal basis. Direction indices alone do not supply the displacement magnitude; the prefactor and lattice metric matter.

2.6.1 Slip and defect-assisted motion

A slip system specifies both a plane and a direction within that plane. The Burgers vector is along the slip direction. The cubic test $hu + kv + lw = 0$ checks compatibility. FCC commonly slips on $\{111\}\langle 110\rangle$, giving twelve geometric systems. BCC commonly uses $\{110\}$, $\{112\}$ and $\{123\}$ planes with $\langle 111\rangle$ directions, although it has no close-packed plane. HCP can activate basal, prismatic and pyramidal systems. Which systems operate depends on mobility, resolved stress, temperature and material; system count alone does not determine ductility.

Glide moves a dislocation within its slip plane. Edge-dislocation climb changes the terminating half-plane by absorbing or emitting point defects and moves the line out of its glide plane. Screw segments can cross-slip between allowed slip planes containing the segment; they cannot choose an arbitrary plane. These mechanisms connect point-defect transport to plastic deformation without making the two processes identical.

Key point: line direction and Burgers vector

The line direction tells where a dislocation runs; **b** tells the displacement it carries. Perpendicular vectors give edge character and parallel vectors give screw character. On a curved line, the local character can change while the same Burgers vector is retained.

2.6.2 Planar defects

A grain boundary joins crystals of different orientation. It can impede dislocation motion and absorb mobile point defects. Its effect on irradiation response depends on boundary structure, stability and defect capture. A simple low-angle tilt boundary can be represented by an array of edge dislocations. For small misorientation,

$$\theta \simeq \frac{b}{L},$$

where b is the relevant Burgers-vector magnitude and L the array spacing. The angle must be in radians. Closer dislocations produce a larger misorientation, but the small-angle relation is not a general high-angle boundary model.

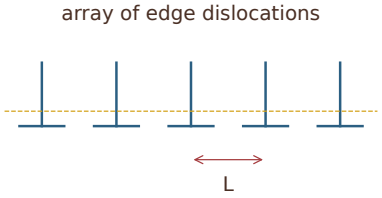


Figure 2.8: An edge-dislocation array accommodates a small tilt; decreasing spacing increases misorientation.

A stacking fault interrupts the layer sequence: perfect FCC has ABCABC stacking, whereas ABCABABC contains a local ABAB segment. A coherent FCC twin reverses stacking around its plane, as in ABCACBA, relating the two orientations by a twin operation. Partial dislocations, with translations smaller than a full lattice translation, can bound a stacking fault.

Stacking-fault energy, $\gamma_{\mathrm{SF}} = E_{\mathrm{excess}}/A$, is excess energy per area in $\mathrm{J/m^2}$ or $\mathrm{mJ/m^2}$. It favours a narrow ribbon between partials; elastic repulsion favours separation. Lower fault energy generally gives wider separation under comparable conditions. The energy per area describes the fault itself; it is not the energy of an entire dislocation or the stress required to move it. For a ribbon of width w and length ℓ, its planar contribution is $E_{\mathrm{SF}} = \gamma_{\mathrm{SF}} w \ell$. Increasing the width raises this contribution, while the elastic interaction between the bounding partials favours their separation. The observed spacing reflects that balance and can also depend on stress and local composition.

Grain boundaries and dislocations can act as sinks for mobile vacancies and interstitials. A sink removes a defect from the freely migrating population when it captures and incorporates it. Capture does not necessarily leave the sink unchanged: an edge dislocation can climb, and a boundary can change its local structure. Consequently, knowing the initial number of defects is not enough to predict their later distribution. Radiation swelling also depends on defect production, transport and sinks.

Example: A tilt boundary and a fault-energy estimate

Given. An FCC crystal has $a = 0.380$ nm. A simple tilt boundary has dislocation spacing $L = 12.0$ nm. Separately, a stacking fault has $A = 2.0 \times 10^{-16}$ m^2 and $\gamma_{\mathrm{SF}} = 0.030$ J/m^2.

Model and substitution. Use a perfect FCC Burgers vector and the small-angle tilt relation:

$$b = \frac{a}{\sqrt{2}} = 0.269 \text{ nm}, \qquad \theta \simeq \frac{0.269}{12.0} = 0.0224 \text{ rad} = 0.0224\frac{180}{\pi} = 1.28^\circ.$$

The separate planar-fault energy is

$$E_{\mathrm{excess}} = \gamma_{\mathrm{SF}} A = (0.030)(2.0 \times 10^{-16}) = 6.0 \times 10^{-18} \text{ J}.$$

Meaning. The small angle is consistent with the array approximation. Halving L would double this predicted angle while the small-angle model remains suitable. Fault energy scales with area: doubling the fault area at fixed γ_{SF} doubles its excess energy.

References. G. S. Was, *Fundamentals of Radiation Materials Science* (Springer, 2007), crystal structure and defects. Cambridge DoITPoMS, *Lattice Planes and Miller Indices* and *Introduction to Dislocations*. AFLOW Prototype Encyclopedia, fluorite AB2_cF12_225_a_c-001.

Unit 2 questions

Conceptual questions

2.C1 Distinguish a lattice from a basis.

2.C2 What distinguishes [110], ⟨110⟩, (110) and {110}?

2.C3 Why can a geometrically possible reflection still be absent?

2.C4 Explain why fluorite has coordination numbers eight and four.

2.C5 How do edge and screw dislocations differ in the relation between **b** and line direction?

Numerical problems

2.N1 Find a for a BCC material with $r = 0.130$ nm.

2.N2 Index a plane with intercepts $(2, 1, \infty)$ in cell units.

2.N3 For a simple tilt boundary with $b = 0.250$ nm and $L = 10.0$ nm, calculate θ in radians and degrees.

Combined problems

2.P1 An elemental BCC (110) peak occurs at $2\theta = 60.0^\circ$ with $\lambda = 1.5406$ Å and $n = 1$. Calculate d, then a, and confirm the selection rule.

2.P2 A BCC region has $a = 0.300$ nm and volume 1.00×10^{-18} m^3. Find its normal-site count and equilibrium vacancy count for $E_f = 0.80$ eV at 1000 K, using unit entropy prefactor. Use $k_B = 8.617 \times 10^{-5}$ eV/K.

Unit 3: Radiation Damage

3.1 Energy transfer and displacement thresholds 27

3.2 Stopping power and damage energy 29

3.3 Collision cascades and primary defects 32

3.4 Displacement models and prompt survival 34

3.5 Recoil spectra and displacement cross sections 36

3.6 Neutron dose and irradiation comparisons 38

Questions 41

Learning Objectives

1. Follow neutron energy through recoil production, stopping, and a displacement cascade.
2. Distinguish displacement estimates from defects that survive after a cascade.
3. Calculate damage-weighted cross sections, fluence, and dpa.
4. Identify the information needed to compare different irradiation conditions.

Key quantities

Quantity	Symbol	Common units
Incident neutron and recoil energies	E_n, T	eV or keV
Damage energy and displacement threshold	T_d, E_d	eV
Displacement cross section	σ_D	barn
Displacement dose	$\mathcal{D}$	dpa

3.1 Energy transfer and displacement thresholds

Radiation damage starts with an energy-transfer event. A neutron can scatter from a nucleus or cause a nuclear reaction. A charged particle loses energy through both atomic collisions and electronic excitation. Electrons and photons interact strongly with electrons, but sufficiently energetic electrons or secondary particles can also displace atoms. The present displacement calculation follows the first recoiling lattice atom, called the *primary knock-on atom* (PKA), and the recoils that it subsequently produces.

Definition: Primary knock-on atom

The **PKA** is the lattice atom that first receives kinetic energy from the incident radiation event being followed. Its initial energy T is an input to the subsequent collision cascade. Secondary recoils receive energy from that moving atom or its descendants.

In an elastic neutron collision, kinetic energy and momentum are conserved. For a stationary target and nonrelativistic motion,

$$E_n = E'_n + T, \qquad \mathbf{p}_n = \mathbf{p}'_n + \mathbf{p}_R.$$

Here E_n and E'_n are the incident and outgoing neutron energies; T is the initial PKA kinetic energy. Inelastic scattering instead leaves some energy in nuclear excitation. Capture and particle-emission reactions may change the isotope or element and introduce additional energetic products. The elastic formulas below describe one specified scattering channel.

Let m be the neutron mass, M the target mass and $A = M/m$. Conservation gives

$$\Lambda = \frac{4mM}{(m+M)^2} = \frac{4A}{(1+A)^2}, \qquad T_{\max} = \Lambda E_n,$$

$$T = \frac{\Lambda E_n}{2}(1 - \cos\theta_{\mathrm{CM}}).$$

The angle is the *neutron scattering angle in the centre-of-mass frame*; it is not the laboratory recoil angle. Forward scattering gives zero transfer and backscattering gives the maximum. Equal masses allow complete transfer. For heavy targets, only a small fraction of the neutron energy can be transferred in a single elastic event. This fraction alone does not determine defect production: even a small fraction of a MeV can greatly exceed a displacement threshold.

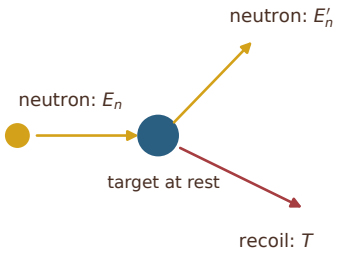


Figure 3.1: Elastic scattering transfers part of the neutron energy to a PKA. The target is initially at rest.

The arrows show momentum directions, while the energy equations determine how much kinetic energy each outgoing particle carries. The maximum-transfer event is a limiting collision geometry. A measured or calculated spectrum includes many other scattering angles and therefore many smaller recoil energies.

An upper bound is not an average. Only when scattering is isotropic in the centre-of-mass frame is $\cos\theta_{\mathrm{CM}}$ uniformly distributed. The resulting recoil energies are uniform between zero and $T_{\max}$, and $\langle T\rangle = T_{\max}/2$. A different angular distribution changes the mean while leaving the kinematic endpoint unchanged. Nearly equal projectile masses can therefore give similar endpoints but different recoil yields.

Key point: an upper bound is not a typical recoil

$T_{\max}$ is fixed by elastic kinematics. The probability of obtaining energies near it depends on the angular scattering distribution. Use $T_{\max}/2$ as the mean only when the stated centre-of-mass scattering is isotropic.

3.1.1 Threshold displacement energy

A Frenkel pair consists of one vacant lattice site and one self-interstitial. The displacement threshold E_d is the minimum recoil energy, for a stated direction and condition, that leaves a pair after prompt relaxation. An atom may move substantially yet return to its site, leaving no pair. A primary pair that does survive can still disappear during later annealing.

The threshold is not the pair formation energy. Formation energy compares relaxed configurations; E_d concerns a dynamic path in which energy spreads into neighbouring atoms and vibrations. Crystal direction affects encounters and return paths, so E_d generally depends on direction. A displacement model often adopts one effective value. The value must be stated with the model; 40 eV below is an assumed effective input, not a universal threshold for every direction in iron.

For a compound, a simple high-energy average is possible only under an explicit allocation rule. If every atom receives the same damage energy and species i has atomic fraction s_i and threshold $E_{d,i}$, the count proportional to $1/E_d$ gives

$$\frac{1}{E_d^{\mathrm{eff}}} = \sum_i \frac{s_i}{E_{d,i}}, \qquad \sum_i s_i = 1.$$

Different recoil species or energy allocations require a more detailed treatment.

Example: The energy available to a recoil

Given. Assume a mass ratio $A = 3$, neutron energy $E_n = 800$ eV and a centre-of-mass scattering angle of 90°. Use an effective displacement threshold of 30 eV.

Model. For a stationary target in an elastic collision,

$$\Lambda = \frac{4A}{(1+A)^2} = \frac{12}{16} = 0.75, \qquad T_{\max} = 0.75(800) = 600 \text{ eV}.$$

At the stated angle,

$$T = \frac{600}{2}(1 - \cos 90^\circ) = 300 \text{ eV}, \qquad E_n' = 800 - 300 = 500 \text{ eV}.$$

Meaning. The recoil exceeds the assumed threshold by a factor of ten, while remaining below its 600 eV upper bound. Energy conservation is satisfied by the outgoing neutron and recoil. Dividing T by E_d does not yet give a cascade count: electronic losses, energy sharing and the damage model remain to be specified.

The effective-threshold expression above is a harmonic weighting, because the high-energy count varies as $1/E_d$. It gives more weight to the species with the lower threshold than an arithmetic mean would. For a real compound, the recoil population and the energy deposited on each sublattice must justify the assumed allocation before this shortcut is used.

3.2 Stopping power and damage energy

A moving recoil repeatedly transfers energy to its surroundings. Its instantaneous kinetic energy E decreases along its path coordinate x. The stopping power is the positive energy loss per unit path length,

$$S(E) = -\frac{dE}{dx} = S_n(E) + S_e(E).$$

The subscripts denote nuclear and electronic stopping. Units such as eV/nm or keV/nm express energy per length. If an energy-versus-distance graph slopes down, its negative slope is positive stopping power. Over an interval Δx, the average is $(E_{\text{in}} - E_{\text{out}})/\Delta x$. This average need not equal the local derivative at either endpoint.

Nuclear stopping transfers energy in interactions with atoms or nuclei. A sufficiently energetic transfer may launch another recoil and extend the cascade. Smaller transfers can remain below the displacement threshold and become atomic vibrations. Consequently, nuclear stopping is not a direct measure of energy stored permanently in defects. Much of the transferred energy ultimately becomes heat even when a cascade leaves defects.

Electronic stopping transfers energy into excitation and ionization of electrons. In a recoil displacement budget this removes energy from the atomic collision sequence. Electron–lattice coupling can subsequently return energy to the atoms as heat. Strong electronic excitation can itself cause structural changes in some materials, so it is inappropriate to state that electronic

energy deposition can never damage a solid. Here it is separated from the energy assigned to the nuclear collision cascade.

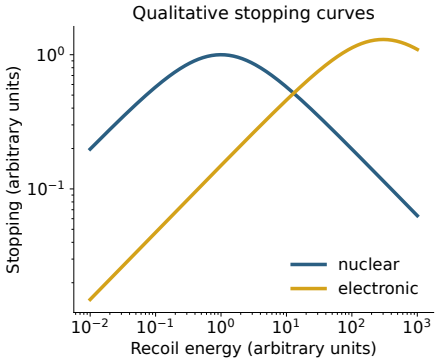


Figure 3.2: Nuclear and electronic stopping change with energy. The arbitrary axes show trends, not material-specific stopping data.

A point on each curve refers to a recoil at one instantaneous energy. As that recoil slows, it samples other parts of the curves. Secondary recoils also sample their own energy ranges. The full partition therefore accumulates many transfers; reading the relative curve heights at the starting energy does not perform that accumulation.

The relative stopping contributions depend on energy, moving species and target. A single universal crossover energy does not exist. Moreover, the moving recoil slows down and generates secondary recoils with different energies. An instantaneous ratio at the initial recoil energy cannot by itself describe the energy allocation of the entire cascade.

Definition: Stopping power and damage energy

Stopping power has units of energy per length and describes loss along a moving particle's path. **Damage energy** has units of energy and describes the cascade budget after electronic loss. One is a rate of loss with distance; the other is an allocated total.

3.2.1 Lindhard energy partition

The damage energy T_d is the part of the initial recoil energy T assigned to atomic motion after electronic losses. Write

$$T_d = f_n(T)T, \qquad E_{\text{electronic}} = T - T_d, \qquad 0 \leq f_n(T) \leq 1.$$

The dimensionless fraction f_n represents a cascade energy partition. It is not generally the local ratio $S_n/(S_n + S_e)$ evaluated at the initial energy. A specified partition model accounts for how energy is shared as recoils slow down. When f_n is supplied, use it directly with the initial recoil energy and carry the resulting T_d into the chosen damage-count model.

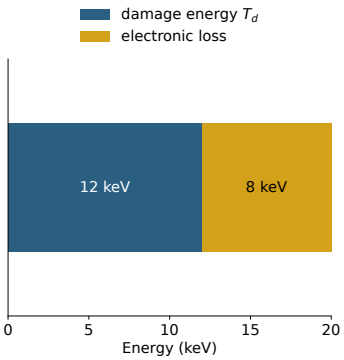


Figure 3.3: An assumed 20 keV recoil supplies 12 keV of damage energy and loses 8 keV electronically. Damage energy is not retained-defect energy.

The two parts of the diagram add to the initial recoil energy. The 12 keV atomic-motion share can create recoils and vibrations, and much of it eventually becomes heat. No third "surviving-defect" energy slice can be inferred by multiplying this budget by a defect-survival fraction without an additional physical model.

Damage energy and surviving defects are different levels of description. The energy fraction f_n acts on T. A survival factor, introduced below, acts on a model defect count after the energy partition has already been applied. Applying a survival factor in place of f_n, or applying the electronic loss twice, changes the physical meaning of the result.

Energy conservation provides a useful check: $0 \le T_d \le T \le T_{\max}$ for the elastic event considered. Keep energy units consistent before subtraction or multiplication. A partition fraction may vary among recoil bins, so an average fraction must have the appropriate energy weighting if it is used to represent several events. Two equal-energy recoils with different f_n values need not supply equal damage energy.

Example: Stopping over a path and partition over a cascade

Given. A recoil slows from 18 to 12 keV over 3 nm. A separate full-cascade partition calculation gives $f_n = 0.70$ at the initial 18 keV.

Model and substitution. The interval-average stopping power uses only the energy lost on that path interval:

$$\overline{S} = \frac{18 - 12}{3} = 2 \text{ keV nm}^{-1}.$$

The cascade energy budget uses the initial energy and its supplied partition:

$$T_d = 0.70(18) = 12.6 \text{ keV}, \qquad E_{\text{electronic}} = 18 - 12.6 = 5.4 \text{ keV}.$$

Meaning. The 6 keV loss over the first path interval contains whatever nuclear and electronic losses occur there. It is not automatically the full electronic loss. The full partition conserves the original energy, $12.6 + 5.4 = 18$ keV. Damage energy can exceed the 12 keV remaining at the end of the selected interval because these quantities describe different stages of the accounting.

3.3 Collision cascades and primary defects

A PKA can transfer enough energy to several neighbouring atoms that they become secondary recoils. Those recoils can generate further motion. A branching collision picture follows the energy transfers, but a real cascade also has a spatial structure: collision paths overlap, nearby atoms act on one another, and the excited region exchanges energy with the surrounding crystal. Thus a cascade is not simply a list of independent permanent displacements.

The initial *ballistic stage* contains rapid, directed collisions and commonly occupies a sub-picosecond interval. The number of atoms far from their original sites rises quickly. As directed motion becomes more collective and disordered, a dense excited region may develop. An energetic cascade can have a temporarily liquid-like core, often called a thermal spike. The earliest motion is strongly out of equilibrium, so assigning it a single local temperature may be misleading.

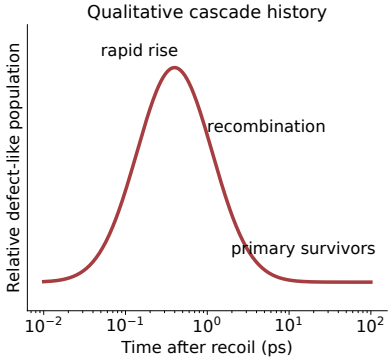


Figure 3.4: The transient population away from lattice sites can greatly exceed the primary population after cooling. The curve is schematic.

Read the descending part as recovery of lattice occupancy while energy leaves the excited region. It does not mean that the incident energy is recovered by the radiation source. The final plotted level is a chosen observation stage, rather than a guarantee that those defects remain unchanged indefinitely.

Energy flows out of the excited region. As it cools, many atoms return toward lattice positions and nearby vacancies and interstitials annihilate. Prompt relaxation commonly spans picoseconds to tens of picoseconds, with the duration depending on the material and event. Defects remaining at its end are *primary survivors*. Longer-time diffusion, clustering and removal at sinks belong to subsequent defect evolution and can change the population again.

The spatial distribution helps explain recombination. Atoms expelled from a central region can leave a vacancy-rich core and interstitial content farther away. Separation can help a pair survive prompt relaxation. Individual cascades, however, are irregular; they need not form perfect concentric shells. At higher recoil energies, an atom may travel between distinct dense collision regions, creating separated subcascades. Their overlap and cooling differ from those of one compact region. No universal recoil energy marks their onset in all materials.

Key point: state the counting time

A cascade has a rapidly changing population during its early motion. **Primary survivors** are counted after prompt cooling under a specified criterion. A later retained population also reflects diffusion, reactions and sinks. Comparing counts requires both the same defect definition and a comparable counting stage.

3.3.1 Crystal directions and defect counting

Crystal rows can focus momentum transfer through successive collisions. In a replacement sequence, one atom occupies the site of its neighbour and passes energy along the row. Many atoms can change sites while the final defect content consists of only a vacancy at one end and an interstitial at the other. This is why atom replacements, atomic mixing and surviving Frenkel pairs are distinct observables.

Channeling concerns motion between rows or planes, where a particle can avoid close nuclear encounters. It can change penetration and the locations at which energy is deposited. Focusing transmits collisions along rows; channeling guides motion through relatively open directions. Both show why orientation matters, especially when interpreting an ion irradiation experiment.

Molecular dynamics follows atomic motion using a specified interatomic potential. A calculation must state the PKA energy and direction, starting temperature, electronic-loss treatment, boundary conditions and time at which defects are counted. Empty reference lattice sites identify vacancies; excess occupancy identifies interstitial content. A cluster count also needs a stated rule for deciding which defects belong together. A drawing or a peak population alone does not establish the number of primary survivors.

For an atom-conserving crystal with no sinks, a Frenkel-pair count N_F corresponds to N_F vacancies and N_F interstitials, or $2N_F$ individual point defects. This equality follows from site and atom bookkeeping. Surfaces, sinks or atom injection can alter the simple balance, so the assumptions matter. Mass, crystal structure, stopping, defect energetics and mobility all influence the result; mass alone cannot rank primary damage in different materials.

Example: Pairs, individual defects and recombination

Given. A closed, atom-conserving simulation contains 45 vacancies and 45 self-interstitials after prompt cooling. Later, 15 vacancy–interstitial pairs recombine.

Model. One Frenkel pair contributes one defect of each kind. Recombination removes one of each, so

$$N_F = 45, \qquad N_{\text{point}} = 45 + 45 = 90,$$

$$N_F^{\text{later}} = 45 - 15 = 30, \qquad N_{\text{point}}^{\text{later}} = 30 + 30 = 60.$$

Meaning. The later pair count is two-thirds of the primary count. This ratio describes the specified later recombination; it is not the prompt NRT survival factor. The equality of vacancy and interstitial counts follows from the closed-system assumption. Neither count determines how many atoms exchanged positions during the event, since replacements can leave no vacant site.

3.4 Displacement models and prompt survival

Displacement models convert an energy budget into an expected count. The count can be noninteger because it represents an average over events. It is not an exact inventory of atoms in one cascade. The Kinchin–Pease (KP) model approximates a cascade by binary elastic collisions between equal-mass atoms with a sharp displacement threshold. It neglects the detailed crystal geometry, many-body relaxation and prompt recombination.

For the simplified model with an electronic cutoff $E_c > 2E_d$,

$$\nu_{\mathrm{KP}}(T) = \begin{cases} 0, & T < E_d, \\ 1, & E_d \le T < 2E_d, \\ T/(2E_d), & 2E_d \le T < E_c, \\ E_c/(2E_d), & T \ge E_c. \end{cases}$$

Below threshold there is no displacement. The next interval produces one. At higher energies, energy sharing gives a linear mean count until the idealized cutoff. Above E_c, the simplified model assigns the excess energy to electrons. The factor $2E_d$ is not a thermodynamic formation energy per pair. It results from the collision model and the termination of branches that fall below threshold. Matching the linear branch to one at $T = 2E_d$ fixes its slope.

The piecewise form makes a physical distinction between three outcomes: insufficient energy, one displacement, and enough energy for multiplication through further recoils. The cutoff is an additional simplifying assumption of this KP version. A calculation should identify the interval before manipulating the linear formula; otherwise an apparently smooth estimate can contradict the model's threshold behaviour.

3.4.1 Norgett–Robinson–Torrens model

The NRT model uses damage energy after electronic losses and a coefficient of 0.8 in its linear branch:

$$\nu_{\mathrm{NRT}}(T_d) = \begin{cases} 0, & T_d < E_d, \\ 1, & E_d \le T_d < 2E_d/0.8, \\ 0.8T_d/(2E_d), & T_d \ge 2E_d/0.8. \end{cases}$$

The upper threshold is $2.5E_d$, ensuring that the linear branch starts at one. With assumed $E_d = 40$ eV, it begins at 100 eV. Applying the linear formula below this boundary incorrectly produces fractional counts in the one-displacement interval. The 0.8 coefficient belongs to the damage-model calibration; it is not a universal statement that 80% of primary defects survive.

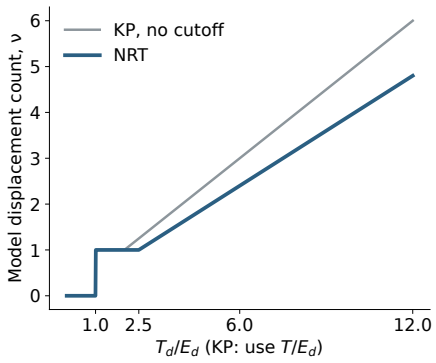


Figure 3.5: The complete NRT function has zero, one-displacement and linear branches. Branch selection precedes substitution.

For the assumed $E_d = 40$ eV, a 30 eV damage event lies below threshold and contributes zero. A 60 eV event lies on the one-displacement plateau. A 200 eV event gives $0.8(200)/80 = 2$ on the linear branch. These values illustrate why a single straight line cannot represent all recoil energies.

Energy partition must precede counting. In the linear regime, $\nu_{\mathrm{NRT}} = 0.4 f_n T/E_d$. Using the full recoil energy when $f_n < 1$ overestimates this count. In contrast to the simplified KP cutoff, the specified damage-energy partition describes electronic losses continuously. Both models still estimate conventional production counts rather than a later retained population.

Key point: three separate model inputs

E_d sets the threshold scale. The partition f_n sets the fraction of recoil energy assigned to atomic motion. The survival factor η describes primary survival relative to the chosen production count. Their numerical values and physical roles are independent inputs unless a stated model relates them.

3.4.2 Athermal recombination correction

Define a prompt-survival factor relative to the NRT count by

$$\eta(T_d) = \frac{\nu_{\mathrm{surv}}(T_d)}{\nu_{\mathrm{NRT}}(T_d)}, \qquad \nu_{\mathrm{arc}}(T_d) = \eta(T_d)\nu_{\mathrm{NRT}}(T_d).$$

The ratio is used where the denominator is nonzero. ARC means athermal recombination corrected: it accounts for primary cascade survival, not all later thermal evolution. Its parameters depend on the material and damage model. One high-energy form is

$$\eta(T_d) = (1 - c)\left(\frac{T_d}{2E_d/0.8}\right)^b + c, \qquad T_d \geq 2E_d/0.8,$$

with $b < 0$ and a limiting factor c. It gives $\eta = 1$ at the branch boundary and approaches c at high energy. The zero and one-displacement branches are retained below it. A single assumed survival factor is suitable only for the recoil conditions for which it was specified. These distinctions follow the NRT and ARC formulations of Norgett, Robinson and Torrens (1975) and Nordlund et al. (2018).

A primary surviving pair count is still not a prediction of hardening, swelling or a long-time defect concentration. Those properties require subsequent defect evolution. Likewise, atomic mixing can remain after sites refill, so neither NRT nor an ARC pair count alone measures all atom replacements.

Example: From recoil energy to primary survivors

Given. Take $T = 10$ keV, $f_n = 0.8$, $E_d = 40$ eV and an assumed prompt-survival factor $\eta = 0.25$ relative to NRT.

Model. Apply the energy partition first, select the NRT branch, and then apply the stated survival factor.

$$T_d = 0.8(10) = 8 \text{ keV} = 8000 \text{ eV}, \qquad 2.5E_d = 100 \text{ eV}.$$

Since $8000 > 100$ eV, the linear branch applies:

$$\nu_{\text{NRT}} = \frac{0.8(8000)}{2(40)} = 80, \qquad \nu_{\text{surv}} = 0.25(80) = 20 \text{ Frenkel pairs}.$$

Meaning. The two fractions act on different quantities: f_n reduces an energy budget, whereas η reduces a conventional count. In an atom-conserving region, 20 surviving pairs mean 20 vacancies and 20 interstitials. Subsequent diffusion and reactions can reduce or rearrange that primary population.

3.5 Recoil spectra and displacement cross sections

A recoil spectrum describes a population of events distributed over energy and, in general, recoiling species. It must state whether the horizontal coordinate is initial recoil energy T or damage energy T_d. When bins are specified in T, the partition $T_d(T)$ is applied before counting damage. A distribution already expressed in damage energy must not have the same electronic loss subtracted again.

Let n_j be the number of recoils in damage-energy bin j, represented by $T_{d,j}$. The event fraction is $n_j / \sum_j n_j$, whereas the fraction of total damage energy carried by that bin is

$$w_j = \frac{n_j T_{d,j}}{\sum_j n_j T_{d,j}}.$$

Rare high-energy events may contribute strongly to the energy total even when most events lie in low-energy bins. An event-weighted average can hide this tail. The meaning of a mean or normalized curve therefore depends on what was counted in its numerator.

For discrete damage energies, the cumulative energy fraction at upper limit T_* is

$$W(T_*) = \frac{\sum_{T_{d,j} \le T_*} n_j T_{d,j}}{\sum_j n_j T_{d,j}}.$$

The upper limit is included. Thus $1 - W(T_*)$ is the fraction from energies *strictly above* that

limit. For a nonzero total, the curve ends at one. A count cumulative instead sums n_j in both numerator and denominator. It also ends at one, but answers a different question. Between discrete bin energies the cumulative is constant because the assumed population contains no events there.

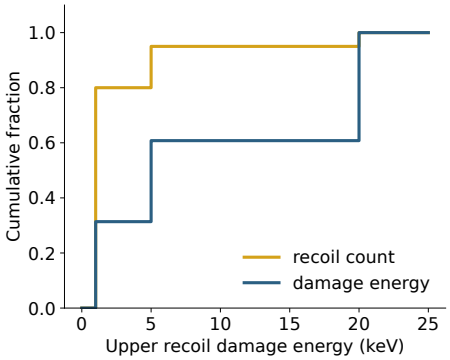


Figure 3.6: Count and energy cumulative curves differ for the same spectrum. Both include the stated upper limit and use their own normalizing total.

Example: Event weighting and energy weighting

Given. Twelve recoils have damage energy 2 keV and three have damage energy 12 keV. Find the contribution of each bin and the cumulative fractions at 2 keV.

Model. Count weighting uses n_j; energy weighting uses $n_j T_{d,j}$.

$$N = 12 + 3 = 15, \qquad E_{\text{damage}} = 12(2) + 3(12) = 60 \text{ keV}.$$

The higher-energy bin therefore contributes $3/15 = 0.20$ of the events and $36/60 = 0.60$ of the damage energy. Including the low-energy bin at the upper limit gives

$$F_{\text{count}}(2 \text{ keV}) = \frac{12}{15} = 0.80, \qquad W(2 \text{ keV}) = \frac{24}{60} = 0.40.$$

Meaning. Most events are low-energy events, but most damage energy comes from the three larger recoils. Both cumulative curves reach one at 12 keV because each uses its own total. Normalization removes the absolute event number or energy, so matching endpoints do not imply matching irradiation rates.

Key point: choose the quantity being accumulated

For event fractions, sum events. For energy fractions, sum energy. For displacement response, weight each event by its model count. A rare energetic recoil can have little influence on the first total and substantial influence on the other two.

3.5.1 Weighting collisions by their damage

At a fixed incident neutron energy, let σ_j be the collision cross section integrated over recoil bin j, and ν_j the corresponding mean model count. The displacement cross section is the count-weighted response

$$\sigma_D \simeq \sum_j \sigma_j \nu_j, \qquad \sigma_{\text{coll}} = \sum_j \sigma_j.$$

Both have area units, commonly barns, where 1 b $= 10^{-24}$ cm^2. They have different meanings: the first counts model displacements, the second collision events. Their ratio is the collision-weighted mean count. Since one event can produce many displacements, σ_D can exceed σ_{coll} without violating probability conservation.

If the supplied input is a differential cross section, bin width is essential:

$$\sigma_j \simeq \left(\frac{d\sigma}{dT}\right)_j \Delta T_j, \qquad \sigma_D(E_n) = \int_0^{T_{\max}} \frac{d\sigma(E_n, T)}{dT} \nu[T_d(T)]\, dT.$$

The differential cross section has area per energy units. This integral keeps incident energy fixed while summing recoil energies for one channel; additional reaction channels and species require additional contributions. Counts must use the same model throughout. An ARC-weighted response generally cannot be obtained by multiplying the NRT response by an arbitrary unweighted mean of η.

Damage-energy weighting and displacement weighting are proportional only when every contributing event uses the same linear count law and threshold. Near threshold, the zero and constant branches break that proportionality. Energy-dependent survival corrections can break it again. This is why a cumulative energy distribution is useful but cannot always replace the full count-weighted response.

Example: A collision-weighted response

Given. Two recoil bins have integrated collision cross sections 0.30 and 0.10 b and NRT counts 10 and 50, respectively.

Model and substitution. Weight each cross section by its own count:

$$\sigma_D = 0.30(10) + 0.10(50) = 8.0 \text{ b}, \qquad \sigma_{\text{coll}} = 0.40 \text{ b},$$

$$\langle \nu \rangle_{\text{coll}} = \frac{8.0}{0.40} = 20.$$

Meaning. The high-count bin provides 5/8 of the displacement response but only 1/4 of the collisions. An unweighted average of 10 and 50 would give 30 and would misrepresent this event population. The response exceeds the collision cross section because each collision produces multiple model displacements.

3.6 Neutron dose and irradiation comparisons

The scalar neutron flux measures the particle population available for interactions. For monoenergetic neutrons of number density n and speed v, $\phi = nv$, with units cm^{-2}s^{-1}. Scalar flux is not the net current through a surface: contributions from opposite directions add to flux even when their currents cancel. For a single-direction beam, it equals the number crossing a perpendicular unit area per second. A spectrum requires summing the energy contributions.

Fluence is time-integrated flux, $\Phi = \int \phi(t)\, dt$, in cm^{-2}. At constant flux, $\Phi = \phi t$. For a uniform

exposure and matching model displacement cross section,

$$N_{\mathrm{disp}} = N_{\mathrm{target}}\Phi\sigma_D, \qquad d = \frac{N_{\mathrm{disp}}}{N_{\mathrm{target}}} = \Phi\sigma_D, \qquad \dot{d} = \phi\sigma_D.$$

Here d is displacements per atom (dpa), a dimensionless accumulated model count per target atom; $\dot{d}$ has units dpa/s. Reporting NRT-dpa or ARC-dpa identifies the underlying count. One dpa does not imply that each distinct atom moved once. Some atoms can be involved repeatedly while others are untouched. Nor is dpa the fraction of atoms remaining as defects after irradiation.

Definition: Flux, fluence and dose

Flux sets an interaction rate per area; **fluence** accumulates that exposure over time. Multiplying fluence by a matching **displacement cross section** gives model dpa. A residual defect fraction requires the additional production, survival and evolution history.

3.6.1 Incident-energy and time dependence

The incident-energy groups used for dose are different from the recoil bins used to construct each response. For group-integrated fluxes ϕ_i and consistently averaged responses $\sigma_{D,i}$,

$$\dot{d} = \sum_i \phi_i\sigma_{D,i}, \qquad \overline{\sigma}_D = \frac{\sum_i \phi_i\sigma_{D,i}}{\sum_i \phi_i}.$$

A simple arithmetic average is justified only for equal flux weights. With a differential flux $\phi(E,t)$ in $\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{eV}^{-1}$, the corresponding expression is

$$d = \int_0^{t_f} dt \int_0^{\infty} \phi(E,t)\sigma_D(E)\, dE.$$

The response must correspond to the material and damage model. With constant spectrum and flux, this reduces to $d = \phi_{\mathrm{total}}\overline{\sigma}_D t$. Equal fluence can give different doses when the spectra produce different averaged responses.

Example: A constant-flux displacement dose

Given. Assume an NRT displacement response of 300 b, constant flux $\phi = 2 \times 10^{14}$ $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and duration $t = 5 \times 10^5$ s.

Model. Use 1 b $= 10^{-24}$ cm^2, $\dot{d} = \phi\sigma_D$ and $d = \dot{d}t$:

$$\dot{d} = (2 \times 10^{14})(300 \times 10^{-24}) = 6.0 \times 10^{-8}\ \text{NRT–}\ \mathrm{dpa\,s}^{-1},$$

$$d = (6.0 \times 10^{-8})(5 \times 10^5) = 0.030\ \text{NRT–}\ \mathrm{dpa}.$$

As a check, the fluence is $\Phi = \phi t = 1.0 \times 10^{20}$ cm^{-2}, giving the same $d = \Phi\sigma_D$.
Meaning. The area units cancel, leaving a dimensionless model count per target atom. Doubling the exposure time at this fixed flux and response doubles the dose. It does not necessarily double the retained defect concentration because defects also evolve during the exposure.

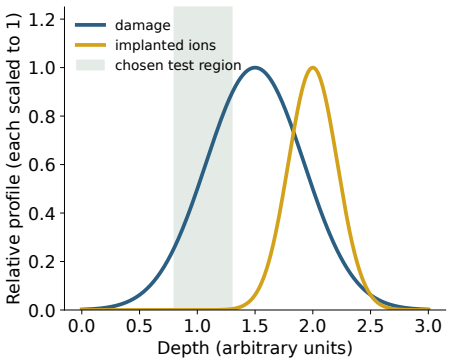


Figure 3.7: Ion damage and implanted species vary with depth. Each illustrative profile is normalized to its own maximum.

The horizontal coordinate selects the depth being sampled. Because each curve has its own normalization, equal plotted heights do not imply equal physical concentrations or equal deposited energy. An experimental measurement that averages over a layer should be compared with the dose and implanted-species distribution over that same layer.

3.6.2 Dose, gases and material state

Equal dpa is useful for comparisons but does not guarantee equal microstructure. Temperature controls defect motion; dose rate changes the time between production events; recoil spectra affect cascade sizes and clustering. Starting composition, sinks and microstructure also matter. Ion irradiation adds depth gradients, injected species and nearby surface sinks. A comparison with neutron irradiation should state the sampled depth and these conditions as well as model dose.

Nuclear transmutation changes chemistry. For example, $^{10}\mathrm{B} + n \rightarrow {}^{7}\mathrm{Li} + {}^{4}\mathrm{He}$ produces energetic charged products; some branches include nuclear excitation and gamma emission. Nickel can produce helium through $^{58}\mathrm{Ni}(n,\gamma)^{59}\mathrm{Ni}$ followed by $^{59}\mathrm{Ni}(n,\alpha)^{56}\mathrm{Fe}$. The second route depends on the isotope inventory created by the first, so spectrum and irradiation history matter.

Helium can bind to vacancies and collect in bubbles, affecting cavity evolution. Gas content is therefore an additional variable, not a consequence determined by dpa alone. One atomic part per million (appm) is an atomic fraction of 10^{-6}. If a constant production ratio g in appm/dpa is explicitly assumed, the produced concentration is $C = gd$. Escape can make retained concentration smaller. A quoted gas ratio must specify its material and irradiation conditions rather than being transferred to all environments.

References. G. S. Was, *Fundamentals of Radiation Materials Science*, Springer (2007). G. H. Kinchin and R. S. Pease, "The displacement of atoms in solids," *Reports on Progress in Physics* 18 (1955), 1–51. M. J. Norgett, M. T. Robinson and I. M. Torrens, "A proposed method of calculating displacement dose rates," *Nuclear Engineering and Design* 33 (1975), 50–54, doi:10.1016/0029-5493(75)90035-7. K. Nordlund et al., "Improving atomic displacement and replacement calculations with physically realistic damage models," *Nature Communications* 9 (2018), 1084, doi:10.1038/s41467-018-03415-5.

Unit 3 questions

Conceptual questions

3.C1 Explain why displacement threshold and Frenkel-pair formation energy differ.

3.C2 Why is it incorrect to call T_d the energy stored in surviving defects?

3.C3 Describe the ballistic stage, thermal spike and prompt cooling in physical terms.

3.C4 Distinguish f_n, η and a later retained-defect fraction.

3.C5 Distinguish fluence, dpa and a residual defect fraction.

Numerical problems

3.N1 For $T = 8$ keV and assumed $f_n = 0.75$, calculate damage energy and electronic loss.

3.N2 A closed crystal contains 28 vacancies and 28 interstitials after cooling. Find the pair count and individual point-defect count.

3.N3 An assumed NRT response is 200 b and fluence is 5×10^{20} cm^{-2}. Find NRT-dpa.

Combined problems

3.P1 Assume $T = 6$ keV, $f_n = 0.50$, $E_d = 30$ eV and $\eta = 0.40$. Find T_d, NRT count and primary surviving pairs.

3.P2 Two recoil bins have collision cross sections 0.4 and 0.1 b and damage energies 1 and 4 keV. Use NRT with $E_d = 40$ eV. Find both counts, σ_D and the collision-weighted mean count.

Unit 4: Defect Evolution and Transport

4.1 Diffusion and conservation 43

4.2 Atomic jumps and defect diffusivity 45

4.3 Sinks and point-defect balances 47

4.4 Clusters and dislocation loops 49

4.5 Cavities, helium and swelling 52

4.6 Radiation-enhanced diffusion and segregation 54

Questions 57

Learning Objectives

1. Connect atomic jumps to diffusion and concentration changes.
2. Use production, recombination, and sink removal to explain defect populations.
3. Relate loops and cavities to stored defects and swelling.
4. Distinguish radiation-enhanced diffusion from radiation-induced segregation.

Key quantities

Quantity	Symbol	Common units
Diffusivity	D	m^2/s
Number concentration	c	$\mathrm{atoms}/\mathrm{m}^3$
Defect fraction	C	dimensionless
Sink strength	k^2	m^{-2}
Swelling	S	fraction or percent

4.1 Diffusion and conservation

A displacement cascade leaves a population of vacancies, self-interstitials and clusters. The later structure depends on what these defects do after the rapid collision sequence ends. Mobile defects may recombine, reach a surface or interface, or join a cluster. Transport connects the initial defect population to loops, cavities and local changes in alloy composition. A defect that exists but cannot move contributes differently from one that can explore many lattice sites during the exposure.

Diffusion is the cumulative result of many atomic jumps. Individual jumps occur in both directions, including in a crystal with uniform composition. A concentration gradient produces a net flux because the more populated side offers more opportunities for jumps across a chosen plane. In this simple description, concentration is the only driving force; stress, electric fields and other driving forces are absent. Zero net flux therefore does not mean that individual atoms have stopped moving.

4.1.1 Flux and concentration gradients

For one-dimensional transport, Fick's first law is

$$J = -D\frac{\partial c}{\partial x}. \tag{4.1}$$

Here J is the number of atoms crossing unit area per second, in atoms $\mathrm{m}^{-2}\ \mathrm{s}^{-1}$; D is diffusivity in $\mathrm{m}^2\ \mathrm{s}^{-1}$; c is number density in atoms m^{-3}; and x is distance in metres. A positive J points in the positive x direction. The minus sign makes the flux point down the concentration gradient. If concentration decreases as x increases, its derivative is negative and J is positive. The first law specifies flux at an instant; it does not require the concentration profile to be steady.

For a linear profile between c_L and c_R, separated by length L, use $\partial c/\partial x = (c_R - c_L)/L$. Multiplying J by area A gives a signed transport rate in atoms s^{-1}. Multiplying again by time gives the net number transported when the flux remains constant. Area is not needed to calculate the flux itself.

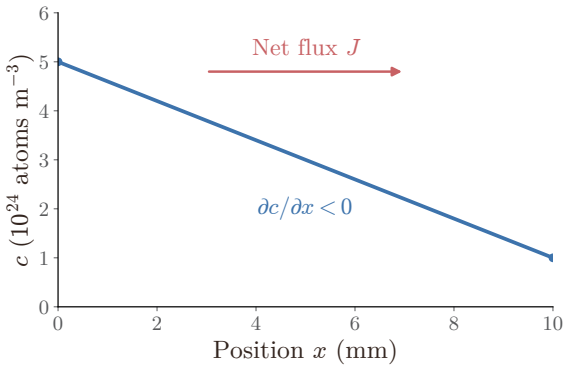


Figure 4.1: A falling concentration profile produces flux in the positive direction. The plotted separation is 10 mm.

Key point

Flux is transport; accumulation is a difference in transport. Equal incoming and outgoing fluxes maintain a steady concentration even while atoms cross the specimen. A curved profile changes because neighbouring slices carry different fluxes.

4.1.2 Accumulation and spreading

Consider a thin slice with no internal source or sink. Its atom count increases if more atoms enter than leave. Dividing this balance by the slice volume and taking the thin-slice limit gives

$$\frac{\partial c}{\partial t} = -\frac{\partial J}{\partial x} = \frac{\partial}{\partial x}\left(D\frac{\partial c}{\partial x}\right). \tag{4.2}$$

If D is constant in space, this becomes Fick's second law,

$$\frac{\partial c}{\partial t} = D\frac{\partial^2 c}{\partial x^2}. \tag{4.3}$$

The slope determines flux; the curvature determines local accumulation. At the centre of a smooth concentration peak, curvature is negative and concentration falls. At a valley, curvature is positive and concentration rises. A linear profile has zero curvature. With fixed boundary concentrations, equal flux enters and leaves every slice, so a nonzero flux can coexist with a time-independent profile.

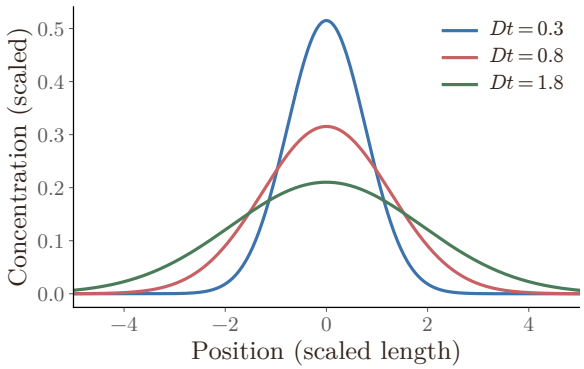


Figure 4.2: A concentration peak spreads and falls while the integrated amount is conserved. Axes use scaled units.

A spreading profile must also satisfy its initial and boundary conditions. Fick's second law alone does not specify the amount present or whether atoms can leave a specimen. The illustrated Gaussian profiles represent spreading without loss of total amount. A specimen with an absorbing boundary can instead lose atoms or mobile defects through that boundary. Internal irradiation production, recombination and trapping require source or loss terms in the balance; they are not hidden inside the constant-D second law.

Example: Flux through a small area

Given. A linear profile falls from 8×10^{23} to 2×10^{23} atoms m^{-3} over 2 mm. Take $D = 3 \times 10^{-13}$ m^2 s^{-1} and area $A = 2 \times 10^{-6}$ m^2.

Model and substitution. Convert the separation to 2×10^{-3} m and use the signed slope:

$$\frac{\partial c}{\partial x} = \frac{(2-8) \times 10^{23}}{2 \times 10^{-3}} = -3 \times 10^{26} \text{ atoms m}^{-4}.$$

Thus $J = -D\partial c/\partial x = 9 \times 10^{13}$ atoms m^{-2} s^{-1}, and the net rate is $JA = 1.8 \times 10^{8}$ atoms s^{-1}.

Meaning. The positive sign means transport toward the lower concentration on the right. If the boundaries maintain this linear profile, the same rate enters and leaves each slice. A nonzero transport rate therefore coexists with zero local accumulation.

4.2 Atomic jumps and defect diffusivity

A defect migrates when atoms rearrange between locally stable configurations. The initial and final configurations are energy minima separated by a higher-energy saddle. The migration energy E_m is the barrier measured from the initial state. Formation energy creates a defect; migration energy moves one that already exists. Irradiation can supply defects without requiring their thermal formation, but their later thermally activated motion still depends on migration barriers.

For a simple activated jump process,

$$\Gamma = \Gamma_0 \exp\left(-\frac{E_m}{k_B T}\right), \qquad D = D_0 \exp\left(-\frac{Q}{k_B T}\right). \tag{4.4}$$

T is absolute temperature in K and $k_B = 8.617 \times 10^{-5}$ eV K^{-1} when energies are in eV. The total successful jump frequency Γ and prefactor Γ_0 have units s^{-1}. Allowed jump directions are already included in this convention for Γ_0. The diffusivity prefactor D_0 has units m^2 s^{-1}. For an existing mobile defect, the simple activation energy Q is its migration energy. Increased temperature raises the fraction of attempts that cross the barrier; it does not remove the barrier.

4.2.1 Random walks and transport distance

For independent jumps of fixed length ℓ, the mean-square displacement after n jumps is $\langle r^2 \rangle = n\ell^2$. With $n = \Gamma t$, comparison with isotropic three-dimensional diffusion gives

$$\langle r^2 \rangle = 6Dt, \qquad D = \frac{\ell^2 \Gamma}{6}, \qquad r_{\text{rms}} = \sqrt{6Dt}. \tag{4.5}$$

The factor is 6 in three dimensions and 4 in an isotropic two-dimensional walk. Correlated jumps require additional factors. The root-mean-square displacement measures statistical spreading from the starting position, not the total path length $n\ell$. Many steps partly cancel. Consequently,

making the exposure four times longer doubles r_{rms} at constant D; it does not quadruple it. Capture at a sink interrupts the unrestricted walk and limits the distances actually explored.

Definition: Formation and migration energies

Formation energy is the energy needed to create the specified defect from a consistent reference state. **Migration energy** is the barrier for a jump of a defect already present. Irradiation changes the defect supply, so the activation energy of equilibrium atomic diffusion need not describe irradiated transport.

4.2.2 Defect mobility and atomic self-diffusion

A vacancy moves when a neighbouring atom fills the empty site. The vacancy and that atom move in opposite directions during the same exchange. One vacancy can exchange with many different atoms, whereas one tagged atom must wait for suitable vacancy encounters. In a dilute, simple vacancy mechanism,

$$D_a^v \simeq f C_v D_v. \tag{4.6}$$

D_v is vacancy diffusivity, D_a^v is the vacancy contribution to atomic self-diffusivity, C_v is vacant-site fraction and f is a dimensionless correlation factor. Thus a large vacancy mobility need not imply a large atomic diffusivity when vacancies are scarce. A self-interstitial is an extra host atom; its migration can involve several atoms moving through an interstitialcy sequence. A small interstitial solute can have a different migration mechanism.

At thermal equilibrium, the dilute vacancy fraction is

$$C_v^{\mathrm{eq}} = \exp(S_f/k_B)\exp[-E_f/(k_B T)], \tag{4.7}$$

where E_f is formation energy and S_f is formation entropy. If the prefactors and f are temperature independent, equilibrium vacancy-mediated atomic diffusion contains both energies: $D_a^v \propto \exp[-(E_f + E_m)/(k_B T)]$. The diffusivity of an existing vacancy contains only the migration term. These are different quantities even in the same crystal.

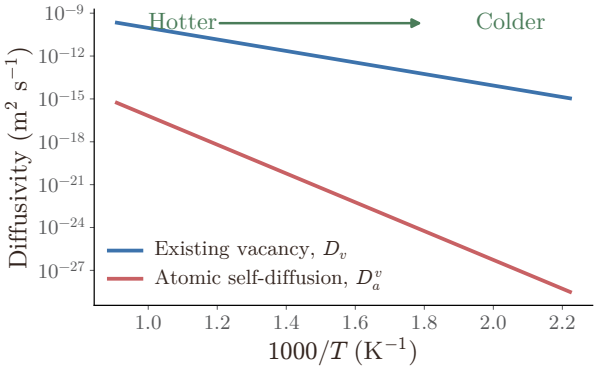


Figure 4.3: Atomic self-diffusion falls more steeply than vacancy diffusion when cooling also reduces the equilibrium vacancy population.

An Arrhenius plot makes this distinction visible. Plotting $\ln[D/(1\,\mathrm{m^2\,s^{-1}})]$ against $1/T$ gives slope $-Q/k_B$. Against $1000/T$, the numerical slope is $-Q/(1000k_B)$. Moving right on either inverse-temperature axis means cooling. The reference unit makes the logarithm dimensionless. A

steeper negative slope means a larger activation energy, provided the same mechanism describes the fitted range.

Example: Distance travelled and net displacement

Given. An existing defect makes independent three-dimensional jumps of length $\ell = 0.30$ nm at total frequency $\Gamma = 2 \times 10^6$ s^{-1} for $t = 0.5$ s. Ignore capture.

Model and substitution. First convert $\ell = 3.0 \times 10^{-10}$ m:

$$D = \frac{\ell^2 \Gamma}{6} = \frac{(3.0 \times 10^{-10})^2 (2 \times 10^6)}{6} = 3.0 \times 10^{-14} \text{ m}^2\,\text{s}^{-1}.$$

The rms displacement is $\sqrt{6Dt} = 3.0 \times 10^{-7}$ m, or 300 nm. The number of jumps is $\Gamma t = 10^6$, giving a total path length $n\ell = 0.30$ mm.

Meaning. Random changes of direction make net displacement much smaller than the accumulated path. A sink 300 nm away is not guaranteed to capture this particular defect: rms displacement describes a population spread, not a deterministic trajectory.

4.3 Sinks and point-defect balances

Rate theory tracks average mobile-defect populations. Here $C_v = N_v/N_{\text{sites}}$ and $C_i = N_i/N_{\text{sites}}$ are dimensionless site fractions. Number density and fraction are related by $c = C/\Omega$, where Ω is atomic volume in m^3 per atom. A source G measured in defects per site per second has units s^{-1}. One useful description is $G = \xi \dot{d}$, where $\dot{d}$ is a specified displacement dose rate and ξ is the fraction supplying the mobile population represented by the model. Its value must match that dose definition and population; it is not a universal survival factor.

Recombination occurs when an interstitial fills a vacancy, removing one of each defect while retaining the atom. In a dilute continuum capture model,

$$R_{iv} = K_{iv} C_i C_v, \qquad K_{iv} \simeq \frac{4\pi r_c (D_i + D_v)}{\Omega}. \tag{4.8}$$

r_c is capture radius in m and $D_i + D_v$ is relative diffusivity. With site fractions, both K_{iv} and R_{iv} have units s^{-1}. A number-density convention would require a different coefficient. The encounter model is an approximation to atomic-scale capture, not a detailed trajectory calculation.

4.3.1 Capture at microstructural sinks

Surfaces, grain boundaries, dislocations and particle interfaces can remove defects from the mobile population. Capture changes the sink or stores defects there; it does not destroy atoms. A surface can change its local position, a boundary can rearrange, and a dislocation can climb by absorbing point defects. Interface structure, temperature and chemistry affect capture efficiency. Diagram arrows indicate net transport, not straight paths or a universal force pulling every defect to a sink.

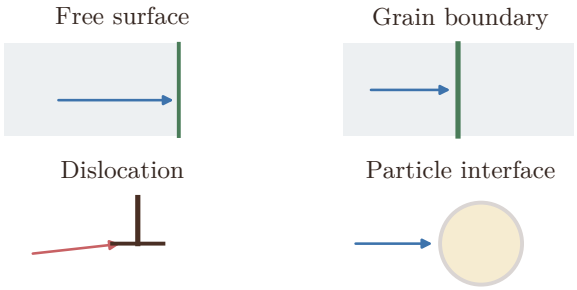


Figure 4.4: Several microstructural features can absorb mobile point defects through different atomic rearrangements.

The forward absorption rate is $R_s = k^2 DC$. Sink strength k^2 has units m^{-2} and describes effective capture geometry and efficiency. The product $k^2 D$ is a removal frequency in s^{-1}. Approximate descriptions include

$$k_d^2 \simeq Z\rho_d, \qquad k_{\mathrm{GB}}^2 \propto d_g^{-2}, \qquad k_{\mathrm{cavity}}^2 \simeq 4\pi r N. \tag{4.9}$$

ρ_d is dislocation line length per volume in m^{-2}; Z is a dimensionless capture factor; d_g is grain size; and dilute ideal spherical absorbers have radius r and number density N. Similar boundary geometry and efficiency are required for the grain-size scaling. Approximately independent sink strengths can be added. Vacancy and interstitial strengths should be treated separately if their capture efficiencies differ. In the usual dislocation-bias model, $Z_i > Z_v$, but actual absorption rates also depend on diffusivity and concentration.

Key point

The three quantities answer different questions: k^2 measures effective sink strength, $k^2 D$ sets the capture frequency, and $k^2 DC$ gives the loss rate for the population present. Comparing k^2 alone cannot determine which defect species is absorbed faster.

4.3.2 Time dependence and limiting balances

A basic homogeneous coupled model is

$$\frac{dC_v}{dt} = G - K_{iv} C_i C_v - k_v^2 D_v (C_v - C_v^{\mathrm{eq}}), \tag{4.10}$$

$$\frac{dC_i}{dt} = G - K_{iv} C_i C_v - k_i^2 D_i C_i. \tag{4.11}$$

Equal source terms assume equal mobile production. The vacancy term includes equilibrium emission; the thermal interstitial population is neglected. Clustering, additional sources and spatial variation require additional terms. Equal production alone does not imply equal mobile concentrations because removal frequencies can differ.

With constant D and k^2, negligible emission and no recombination, sink-only decay after production stops is $C = C_0 e^{-t/\tau_s}$, where $\tau_s = 1/(k^2 D)$. About 37% remain after one lifetime.

With constant production and $C(0) = 0$,

$$C(t) = C_{\mathrm{ss}}(1 - e^{-t/\tau_s}), \qquad C_{\mathrm{ss}} = \frac{G}{k^2 D} = G\tau_s. \tag{4.12}$$

Steady concentration means balanced arrival and removal, not stopped motion. In a different, recombination-only limit with the additional assumption $C_i = C_v = C$, $C = \sqrt{G/K_{iv}}$. Thus sink-dominated concentration scales linearly with G, whereas this recombination-only concentration scales as its square root. Temperature alone does not establish which loss dominates.

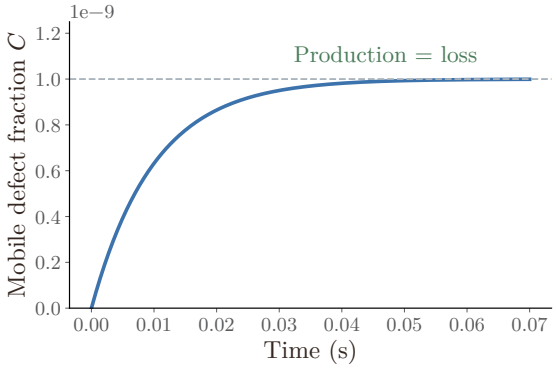


Figure 4.5: A constant source builds a mobile population toward a steady value while capture continues.

Example: A sink-controlled mobile population

Given. Let dislocations dominate with $\rho_d = 2 \times 10^{13}$ m^{-2}, $Z = 2$, $D = 5 \times 10^{-13}$ m^2 s^{-1} and $G = 6 \times 10^{-8}$ s^{-1}. Initially no excess mobile defects are present. Neglect recombination and emission.

Model and substitution. The effective strength is $k^2 = Z\rho_d = 4 \times 10^{13}$ m^{-2}. Hence

$$k^2 D = 20 \text{ s}^{-1}, \qquad \tau_s = \frac{1}{20} = 0.05 \text{ s}, \qquad C_{\mathrm{ss}} = \frac{6 \times 10^{-8}}{20} = 3 \times 10^{-9}.$$

At $t = 0.15$ s, three lifetimes have elapsed, so $C = C_{\mathrm{ss}}(1 - e^{-3}) = 2.85 \times 10^{-9}$, about 95% of the steady value.

Meaning. The small steady fraction is maintained by continuous replacement. At steady state, capture removes $k^2 D C_{\mathrm{ss}} = G$ defects per site per second. If production stops, the population decays on the same 0.05 s timescale within this fixed-sink model.

4.4 Clusters and dislocation loops

Mobile defects can meet by diffusion and bind into clusters. Cascades can also create clusters directly. A cluster is energetically bound when its formation energy is below that of the same defects far apart. For two defects A and B,

$$E_b = E_f(A) + E_f(B) - E_f(AB). \tag{4.13}$$

Positive E_b is favourable binding: separating the pair into these isolated states costs energy. Formation and binding energies must use compatible reference states. Binding does not make a cluster permanent. Thermal emission can remove defects, and entropy contributes to stability at finite temperature. Cluster size, shape, defect type and material all affect the energetics.

A vacancy cluster need not remain an open hole. Depending on crystal structure, energetics, gas content and temperature, vacancy aggregates can form cavities, collapsed dislocation loops or stacking-fault tetrahedra. A stacking-fault tetrahedron is bounded by four faulted faces and is especially relevant to FCC crystals. Interstitial clusters can form an extra patch of an atomic plane. The physical structure matters because different clusters store defects, migrate and interact with moving dislocations differently.

4.4.1 Loop geometry and character

The edge of an extra plane patch is a dislocation line. Because the patch has a closed perimeter, the dislocation is a loop. It can be circular, faceted or irregular; the word loop describes closure, not a required circular shape. For a vacancy loop, a missing patch of plane collapses as neighbouring planes close together. The remaining perimeter is a dislocation. A section cutting across the loop encounters its perimeter twice. A void, in contrast, retains an open three-dimensional cavity.

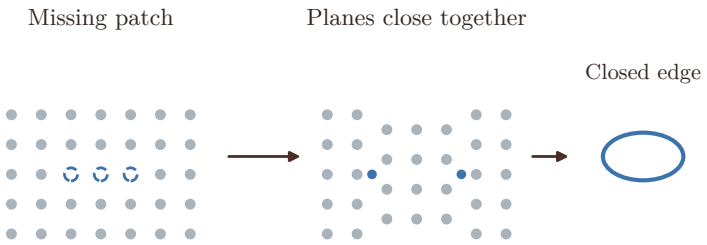


Figure 4.6: Collapse of a missing plane patch leaves a dislocation loop rather than an open cavity.

The Burgers vector **b** describes the lattice displacement associated with a dislocation. In FCC material, a Frank loop has $\mathbf{b} = a\langle 111\rangle/3$ on a $\{111\}$ plane, where a is lattice parameter. It encloses a stacking fault and cannot glide conservatively like a perfect dislocation. An FCC perfect loop has $\mathbf{b} = a\langle 110\rangle/2$ and no enclosed stacking fault. A perfect loop is not automatically freely mobile: geometry, solutes, pinning and other obstacles still matter. Common BCC loop families include $a\langle 111\rangle/2$ and $a\langle 100\rangle$. HCP basal and prismatic directions are not equivalent, and several loop families are possible.

4.4.2 Stored defects, growth and observations

For a planar loop enclosing area A, the number n of stored excess or missing atoms is related to the normal Burgers-vector component by

$$n\Omega = A|b_\perp| = \pi r^2 |b_\perp| \quad \text{for a circular loop.} \tag{4.14}$$

Use the component normal to the loop plane, not automatically the full magnitude of **b**. The atomic volume Ω and product $A|b_\perp|$ both convert an atom count into volume. At fixed loop character, $n \propto r^2$: doubling radius requires four times the stored defects. A radius inferred from this formula is a geometric radius, not necessarily the width of an electron-microscopy contrast feature.

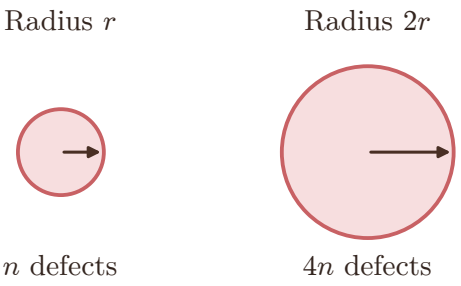


Figure 4.7: Loop area counts stored point defects through the displacement normal to its plane.

In the area sketch, the normal displacement gives the thickness of material represented by the loop. Multiplying that thickness by the enclosed area gives the equivalent excess or missing volume. The dislocation line is the perimeter of this volume, so counting line length alone would not count stored point defects.

An interstitial loop grows by net interstitial absorption and shrinks when vacancy absorption removes the excess. A vacancy loop has the opposite response. Emission must also be included when important. Increasing loop size does not require increasing loop number density: existing loops can collect additional defects without new loops nucleating. When N_L identical loops per m^3 each store n defects, the stored-defect density is $N_L n$, and the corresponding stored-defect fraction is $N_L n \Omega$. This population is distinct from the mobile point-defect fraction in a rate balance.

Transmission electron microscopy reveals loops through diffraction contrast. Changing diffraction conditions can alter the appearance of the same region; a sequence of such images is not necessarily a growth sequence. Marks used to locate small defects are not their physical edges. Both number density and size must be measured to characterize a population. Loops can impede dislocation motion, so their evolution can produce irradiation hardening even without appreciable swelling.

Example: Counting defects stored in a loop

Given. A circular loop has $r = 4$ nm, $|b_\perp| = 0.25$ nm and atomic volume $\Omega = 0.020$ nm^3. Its character remains fixed as it absorbs additional defects of the same type.

Model and substitution. Consistent nanometre units are convenient:

$$n = \frac{\pi r^2 |b_\perp|}{\Omega} = \frac{\pi(4)^2(0.25)}{0.020} = 628 \text{ defects, approximately.}$$

Absorbing another 628 defects doubles n. Since $n \propto r^2$,

$$\frac{r'}{r} = \sqrt{\frac{n'}{n}} = \sqrt{2}, \qquad r' = 5.66 \text{ nm.}$$

Meaning. Doubling stored defects increases radius by only 41%. The loop stores defects over an area; it is not a spherical cavity whose volume scales with r^3. Emission or a change of loop character would require revising this calculation.

4.5 Cavities, helium and swelling

A cavity is an open three-dimensional volume within the solid. Its formation and growth can increase the specimen volume. Void swelling is the volume change relative to the original material volume, $S = \Delta V/V_0$. It is a fractional quantity; multiply by 100 to express it as a percentage. Volume swelling is not the same as a change in one length. The relevant microstructural quantity is total cavity volume, which depends on both cavity size and number density.

For dilute, equal-sized spherical cavities,

$$S \simeq \frac{4\pi}{3} N r^3. \tag{4.15}$$

Here N is cavity number density in m^{-3} and r is radius in m. The dilute approximation allows the reference-volume distinction in number density to be neglected. For a size distribution, add $\sum_j (4\pi/3) N_j r_j^3$ over populations. Because volume scales as r^3, a few large cavities can contribute strongly to swelling. Using a diameter as though it were radius overestimates the volume by a factor of eight. A two-dimensional micrograph alone does not directly give the full three-dimensional volume fraction; each image also needs its own scale and suitable size and number-density information.

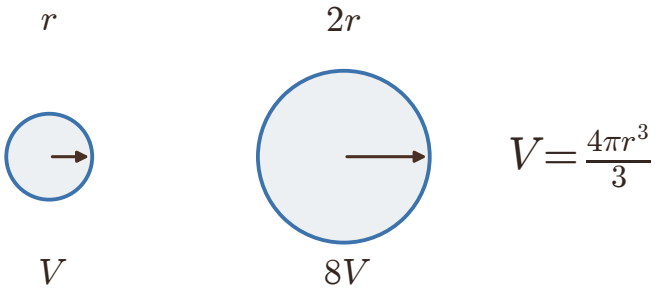


Figure 4.8: Doubling cavity radius at fixed number density multiplies the cavity volume and dilute swelling by eight.

4.5.1 Nucleation and gas stabilization

A new cavity creates internal surface, which costs energy. Collecting vacancies from a supersaturated population can offset that cost. Define vacancy supersaturation $S_v = C_v/C_v^{\mathrm{eq}}$, distinct from swelling S. In the classical spherical capillarity model,

$$\Delta G = 4\pi r^2\gamma - \frac{4\pi r^3}{3\Omega}k_BT\ln S_v, \qquad r^* = \frac{2\gamma\Omega}{k_BT\ln S_v}, \quad S_v > 1. \tag{4.16}$$

γ is surface energy in J m^{-2}, Ω is atomic volume and r^* is critical radius. Use $k_B = 1.381\times10^{-23}$ J K^{-1} with these SI surface-energy units. The positive surface term scales as area, while the supersaturation term scales as volume. The free-energy maximum separates subcritical cavities, which tend to shrink in this model, from larger cavities favoured to grow. Increasing supersaturation lowers r^* when the other quantities remain fixed.

This description omits gas pressure, elastic effects and the discrete structure of very small clusters. It should not be interpreted as an exact atomic-scale prediction. Helium produced by nuclear reactions has low solubility in many metals and can bind to vacancy clusters. Gas can stabilize small cavities and alter their nucleation and growth. Helium is not required for every void, and helium production per dpa depends on composition and irradiation spectrum.

4.5.2 Growth and the defect balance

For one cavity, a simple defect-count balance is

$$\frac{dV_c}{dt} = \Omega(\dot{N}_v - \dot{N}_i - \dot{N}_{\mathrm{em}}), \qquad \frac{dV_c}{dt} = 4\pi r^2\frac{dr}{dt} \quad \text{for a sphere.} \tag{4.17}$$

Each dotted N is a number of defects per second: vacancies captured, interstitials captured and vacancies emitted, respectively. A retained vacancy adds approximately one atomic volume. Interstitial capture or vacancy emission reduces the cavity volume. These are count rates for one cavity, not site-fraction rates. A positive vacancy arrival rate alone does not establish growth; it must exceed both opposing terms.

Dislocation bias can favour swelling because dislocations may capture interstitials more efficiently, leaving a net vacancy supply to cavities. Equal pair production does not require equal capture at

each type of sink. At low temperature, slow vacancy transport can limit growth. At intermediate temperatures, vacancies can reach cavities. At higher temperatures, emission and reduced supersaturation can limit growth. Any swelling peak depends on material, dose rate, exposure and sink structure; it is not a universal operating-temperature limit.

An incubation period with little measurable swelling can precede cavity growth. A steady mobile-defect concentration can sustain a continuous cavity flux, so swelling can continue after that concentration stabilizes. Grain refinement or stable interfaces may shorten capture paths, but their stability, segregation and gas accumulation also matter. The location and type of retained damage must be considered along with its amount.

Example: Cavity radius and swelling

Given. Dilute spherical cavities have number density $N = 2 \times 10^{21}$ m^{-3} and radius $r = 3$ nm. Compare them with the same number of cavities at 6 nm radius.

Model and substitution. The cavity volume per unit specimen volume is

$$S \simeq \frac{4\pi}{3} N r^3 = \frac{4\pi}{3}(2 \times 10^{21})(3 \times 10^{-9})^3 = 2.26 \times 10^{-4}.$$

Multiplying the fraction by 100 gives 0.0226%. Doubling radius multiplies swelling by $2^3 = 8$, giving $S' = 1.81 \times 10^{-3}$, or 0.181%.

Meaning. Size and number density play different roles: doubling N only doubles swelling. If total cavity volume were instead held fixed during coarsening, doubling radius would require reducing the number density to $N/8$.

4.6 Radiation-enhanced diffusion and segregation

Irradiation supplies defects in excess of their thermal populations. If they survive and remain mobile, these defects create additional opportunities for atoms to move. A simple self-diffusion estimate is

$$D_a \simeq f_v D_v C_v + f_i D_i C_i. \tag{4.18}$$

The D values describe defect mobility, the C values are mobile site fractions, and dimensionless f_v, f_i include correlation and mechanism factors. Both mobility and population matter. A large stored population of immobile clusters is not interchangeable with the mobile fraction in this expression. Species-specific jump rates and coupling are needed for alloying-element diffusion; this estimate is not an exact formula for every solute.

4.6.1 Finite exposure and temperature

The sink balance provides a useful vacancy-only example. Let $C_{\mathrm{ex}} = C_v - C_v^{\mathrm{eq}}$, initially zero. Hold each temperature fixed, assume constant production and sink strength, and omit recombination, trapping and clustering. Then

$$\frac{dC_{\mathrm{ex}}}{dt} = G - k^2 D_v C_{\mathrm{ex}}, \qquad C_{\mathrm{ex}}(t) = \frac{G}{k^2 D_v}\left(1 - e^{-k^2 D_v t}\right). \tag{4.19}$$

The atomic diffusivity is estimated as $D_a = f_v D_v (C_v^{\mathrm{eq}} + C_{\mathrm{ex}})$. At long time the excess approaches $G/(k^2 D_v)$. For a short exposure or very low mobility, $1 - e^{-x} \simeq x$, so $C_{\mathrm{ex}} \simeq Gt$. It does not diverge as mobility approaches zero at fixed finite time. In that limit the excess atomic diffusion is still small because D_v is small.

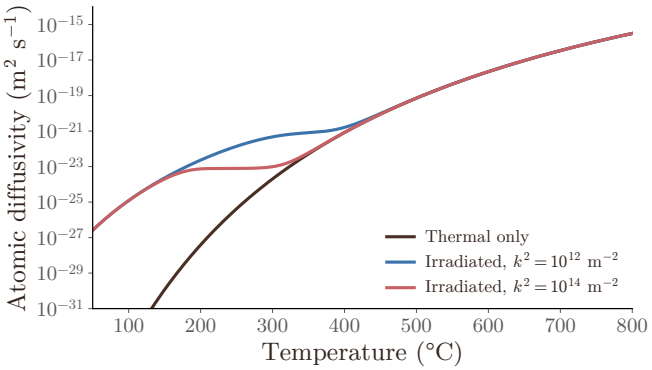


Figure 4.9: A finite-time vacancy model illustrates increased atomic diffusion under irradiation and reduced enhancement with stronger sinks.

Stronger sinks shorten the excess-defect lifetime and reduce the radiation contribution in this model. At intermediate temperatures, irradiation can sustain a population well above equilibrium while defects are mobile. At high temperature, thermal populations and thermal diffusion can become dominant, reducing the relative enhancement. These regimes depend on exposure time, dose rate, material and sinks. Direct atomic relocation in a cascade is displacement mixing; it differs from the subsequent thermally activated motion described by these diffusion coefficients.

Key point

Enhanced diffusion describes faster atomic transport. **Segregation** describes a change in local composition. Faster transport can accelerate redistribution, but a composition change also requires unequal transport of the alloy constituents.

4.6.2 Composition changes near sinks

Radiation-induced segregation (RIS) occurs when defect transport couples to unequal transport of alloying elements. Local composition then changes near a sink. It is a redistribution: enrichment in one region must be supplied by another region in a closed specimen. For an illustrative binary alloy, $C_A(x) + C_B(x) = 1$, where these C symbols are atomic fractions of alloy constituents, not the point-defect fractions above. Enrichment and depletion are measured relative to the bulk composition. Atomic fractions and weight percentages are different composition conventions and must not be mixed in a calculation.

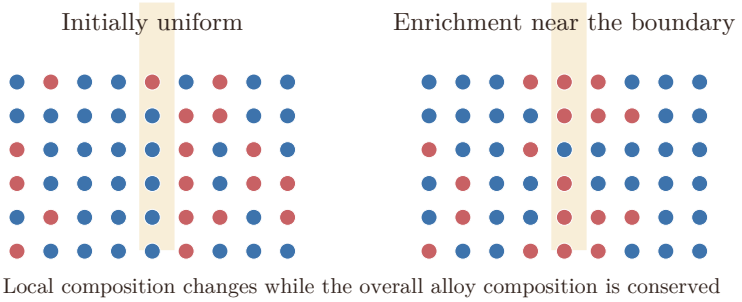


Figure 4.10: Defect-coupled transport redistributes solute near a boundary while preserving the overall number of solute atoms.

In a simplified vacancy-mediated inverse Kirkendall mechanism, vacancies move toward a sink while exchanging atoms move oppositely. If constituent A is transported away more strongly than B, A becomes depleted near the sink. Interstitial-associated transport can instead carry some solutes toward sinks and enrich them there. The contributions can compete, so atomic size alone does not establish the direction of segregation for an element. The alloy, irradiation conditions and sink must be specified.

At low mobility, RIS can develop slowly. Mobile defects can sustain unequal solute fluxes, while thermal diffusion acts to smooth the resulting concentration gradients. Intermediate-temperature peaks can therefore occur, but neither their position nor segregation direction is universal. Grain-boundary chemistry can affect material performance; for example, chromium depletion in irradiated stainless steel can reduce corrosion resistance. Cracking also depends on stress, environment and microstructure. Improving defect capture does not by itself guarantee improved performance if harmful chemistry or gas accumulation develops at those sinks.

Example: Excess defects and enhanced atomic diffusion

Given. Use $D_v = 2 \times 10^{-12}$ m^2 s^{-1}, $k^2 = 5 \times 10^{12}$ m^{-2}, $G = 3 \times 10^{-8}$ s^{-1} and $f_v = 0.75$. Apply the vacancy-only steady model above.

Model and substitution. The removal frequency is $k^2 D_v = 10$ s^{-1}, so the excess fraction and its diffusion contribution are

$$C_{\text{ex}} = \frac{G}{k^2 D_v} = 3 \times 10^{-9}, \qquad \Delta D_a = f_v D_v C_{\text{ex}} = 4.5 \times 10^{-21} \text{ m}^2\,\text{s}^{-1}.$$

Substitution also gives $\Delta D_a = f_v G / k^2$: doubling k^2 halves the steady excess contribution to 2.25×10^{-21} m^2 s^{-1}.

Meaning. Cancellation of D_v occurs because a more mobile vacancy has a shorter sink lifetime. It applies to this steady limiting balance. At finite time, the factor $1 - e^{-k^2 D_v t}$ retains the dependence on mobility; slowly moving defects may never reach the assumed steady state during exposure.

References. G. S. Was, *Fundamentals of Radiation Materials Science: Metals and Alloys*, Springer, 2007, Chs. 4–8 and 12, DOI: 10.1007/978-3-540-49472-0. K. L. Murty and I. Charit, *An Introduction to Nuclear Materials: Fundamentals and Applications*, Wiley-VCH, 2013, Ch. 6.

Unit 4 questions

Conceptual questions

4.C1 Distinguish the questions answered by Fick's first and second laws.

4.C2 Explain why vacancy diffusivity and vacancy-mediated atomic diffusivity differ.

4.C3 Distinguish sink strength k^2 from removal frequency $k^2 D$.

4.C4 Distinguish a collapsed vacancy loop from a void.

4.C5 Explain how vacancy motion toward a sink can produce depletion of an alloying element near that sink.

Numerical problems

4.N1 For $D = 6 \times 10^{-16}$ m^2 s^{-1}, find r_{rms} after 100 s, neglecting capture.

4.N2 Find n for a circular loop with $r = 5$ nm, $|b_\perp| = 0.20$ nm and $\Omega = 0.020$ nm^3.

4.N3 Find fractional and percent swelling for dilute spherical cavities with $N = 10^{21}$ m^{-3} and $r = 4$ nm.

Combined problems

4.P1 Assume dislocations dominate with $Z = 1.5$, $\rho_d = 4 \times 10^{13}$ m^{-2}, $D = 10^{-12}$ m^2 s^{-1} and $G = 1.2 \times 10^{-7}$ s^{-1}. Neglect recombination and emission. Find k^2, τ_s and C_{ss}.

4.P2 Use the vacancy-only steady model with $G = 8{\times}10^{-8}$ s^{-1}, $k^2 = 2{\times}10^{13}$ m^{-2}, $D_v = 2{\times}10^{-12}$ m^2 s^{-1} and $f_v = 0.5$. Find C_{ex} and the excess atomic diffusivity. With $C_v^{\text{eq}} = 10^{-9}$, find the ratio of total to thermal atomic diffusivity.

Unit 5: Mechanical Properties after Irradiation

5.1 Tensile properties and localized deformation 59

5.2 Radiation hardening 61

5.3 Fracture and the ductile–brittle transition 64

5.4 Reactor pressure vessel steels 66

5.5 Irradiation creep, swelling and growth 68

5.6 Helium effects and materials selection 70

Questions 73

Learning Objectives

1. Interpret changes in tensile properties and dislocation motion after irradiation.
2. Relate defect size and number density to radiation hardening.
3. Distinguish transition-temperature measurements from fracture-mechanics quantities.
4. Compare irradiation creep, swelling, growth, and helium effects.

Key quantities

Quantity	Symbol	Common units
Yield strength	σ_y	MPa
Obstacle number density	N	m^{-3}
Stress intensity	K_I	$\mathrm{MPa}\sqrt{\mathrm{m}}$
Creep strain	ε_c	dimensionless
Helium content	He	appm

5.1 Tensile properties and localized deformation

A radiation-damaged microstructure changes the way a material carries load. Defect clusters, loops, cavities and solute-rich features can obstruct dislocations. A higher applied stress may then be needed to start plastic flow. This increase in strength is useful only if the component also retains enough ductility, crack resistance and dimensional stability for its job. These properties describe different responses and require different measurements.

In a tensile test, a specimen of initial cross-sectional area A_0 and gauge length L_0 is pulled by a force F. The engineering stress and engineering strain are

$$\sigma_e = \frac{F}{A_0}, \qquad e = \frac{L - L_0}{L_0} = \frac{\Delta L}{L_0}.$$

Stress is force per area: 1 N/mm^2 = 1 MPa. Strain is dimensionless; $e = 0.02$ means a 2% extension. Engineering quantities use the original dimensions throughout the test. In the initial elastic range, $\sigma_e \simeq Ee$, where E is Young's modulus. Most of this extension disappears on unloading. Plastic deformation remains.

Definition: Elastic and plastic strain

Elastic strain is recovered on unloading. **Plastic strain** remains after the load is removed. In a small-strain tensile description, the total strain is the sum of these contributions. Yield marks the onset of appreciable plastic deformation; it does not mark fracture.

Yield strength marks the start of appreciable plastic flow. When the transition from elastic to plastic response is gradual, a stated proof-stress convention, such as a 0.2% offset, defines the measured strength. Ultimate tensile strength is the maximum engineering stress, not necessarily the stress at fracture. Uniform elongation is the extension accumulated before necking starts. Total elongation includes subsequent localized extension up to fracture and depends on the chosen gauge length.

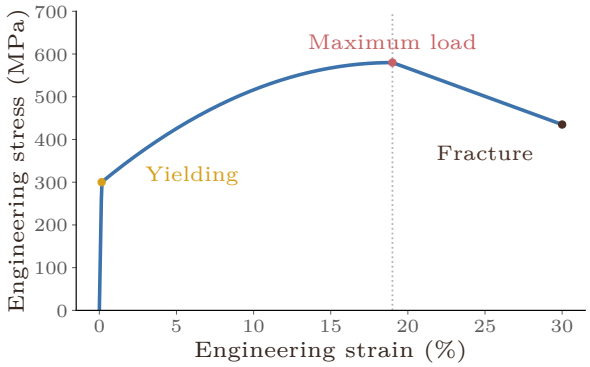


Figure 5.1: A tensile curve separates yielding, uniform extension, necking and fracture.

Read the tensile figure from left to right as increasing extension, not increasing time at a prescribed rate. The maximum engineering stress identifies the largest force divided by the original area. Beyond this point, the falling engineering curve can coexist with a rising local true stress because the neck area is shrinking. The graph therefore needs both stress definitions and

the location of deformation to be interpreted correctly.

Irradiation often raises yield strength while reducing the extra strain hardening available after yield. Strain hardening means that continued plastic deformation raises the stress needed for further flow. It must compensate for the shrinking cross-section of a stretched specimen. If it cannot do so, a slightly thinner region stretches faster and becomes a neck. High initial yield strength therefore does not guarantee large uniform elongation.

True stress σ_t uses the instantaneous area, while true strain ε_t measures accumulated logarithmic extension. For rate-independent uniaxial tension, neglecting elastic strain, the onset of necking satisfies

$$\frac{d\sigma_t}{d\varepsilon_t} = \sigma_t.$$

Before this crossing, sufficient hardening can sustain uniform extension. After it, area loss wins. This is an ideal necking criterion; it does not specify the microscopic fracture mechanism or replace a measured ductility value.

Key point

Strength measures resistance to plastic flow. **Strain hardening** measures how that resistance changes as plastic strain accumulates. Raising the starting resistance without retaining enough subsequent hardening can shorten the uniform part of a tensile test.

Dislocation channeling provides a second reason for nonuniform strain. The first dislocations can remove or rearrange irradiation defects along their path. Later dislocations find that path easier, so deformation concentrates in a narrow band. A channel has fewer effective obstacles; it need not be completely free of defects. The surrounding material can remain relatively hard while the channel carries much of the plastic strain. A channel is a deformation feature, not a track left by an incident neutron.

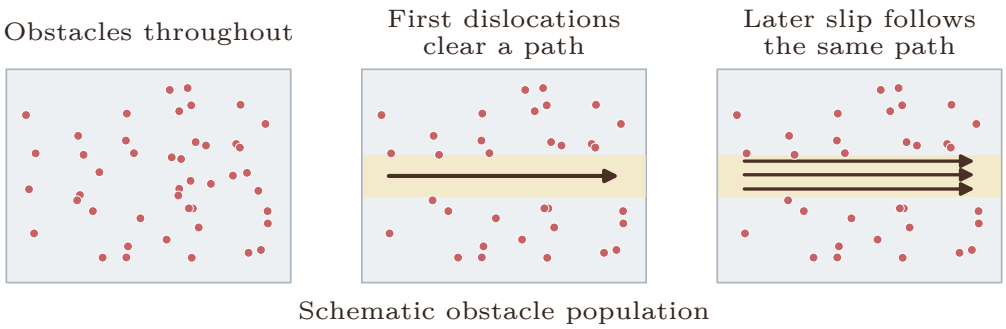


Figure 5.2: Obstacle removal can make a narrow path easier for later dislocations to traverse.

Example: Reading a tensile measurement

Given. A specimen has initial area $A_0 = 12$ mm^2 and gauge length $L_0 = 40$ mm. At one instant the force is 3.6 kN and the measured extension is 2 mm.

Model and substitution. Engineering quantities use the original dimensions. Convert 3.6 kN to 3600 N:

$$\sigma_e = \frac{3600}{12} = 300 \text{ N/mm}^2 = 300 \text{ MPa}, \qquad e = \frac{2}{40} = 0.05 = 5\%.$$

Meaning. These values specify a point on the engineering stress–strain curve. They do not establish whether the specimen has yielded or begun necking without the surrounding curve. If extension is uniform, the gauge length changes from 40 to 42 mm throughout its length; after necking, much of the additional extension can occur in a small region.
For comparison of irradiated and reference specimens, keep test temperature, strain rate and geometry consistent. Irradiation temperature controls the retained microstructure; test temperature controls its response during loading.

5.2 Radiation hardening

Plastic flow requires dislocations to move or new dislocation loops to be emitted by sources. Irradiation can hinder both processes. Friction hardening describes resistance encountered during motion. Source hardening raises the stress required to activate a dislocation source, for example when defects shorten its available length. These contributions depend on the initial dislocation structure as well as the damage introduced by irradiation.

A dislocation pinned at two strong obstacles bows under shear stress. Applied force per unit length is τb, where τ is shear stress and b is the Burgers-vector magnitude. Line tension $\mathcal{T}$ resists curvature, giving $\tau b \simeq \mathcal{T}/R$. With $\mathcal{T} \sim Gb^2/2$ and critical bow radius $R \sim L/2$, the characteristic stress is $\tau \sim Gb/L$. Here G is shear modulus and L is pin spacing. Closer obstacles require a sharper bow and a higher stress. Numerical factors depend on geometry and interaction strength.

A dislocation may cut, absorb or rearrange an obstacle, or bypass it. A strong particle can force bypass and leave a loop around the particle; that is not the outcome of every encounter. Obstacle structure, size, temperature and dislocation character therefore matter in addition to obstacle count.

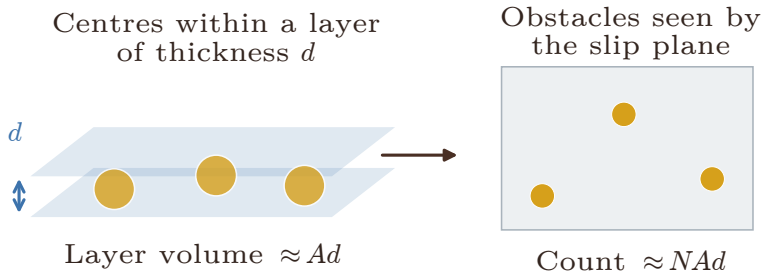


Figure 5.3: A slip plane intersects obstacles with centres in a layer about one obstacle diameter thick.

The obstacle-intersection drawing connects a three-dimensional population to a two-dimensional slip plane. Centres just above or below the plane still contribute if the obstacle extends across it. Larger obstacles are more likely to intersect the plane, which is why the product Nd, rather than N alone, enters the spacing estimate. The drawing describes a statistical spacing, not a regular array of equally spaced defects.

For random obstacles of number density N in m^{-3} and characteristic diameter d in metres, an area A intersects approximately NAd obstacles. The areal density is consequently $N_A \approx Nd$, in m^{-2}. One obstacle per area of order L^2 then gives

$$L \approx (Nd)^{-1/2}, \qquad \Delta\sigma_y = M\alpha Gb\sqrt{Nd} = \frac{M\alpha Gb}{L}.$$

This dispersed barrier hardening equation predicts an *increment* in tensile yield strength. Add it to a specified reference strength to estimate the final strength. The dimensionless factor α represents effective barrier strength. The Taylor factor M converts a shear increment to a polycrystal tensile increment: $\Delta\sigma_y = M\Delta\tau$. An ideal random FCC aggregate commonly uses $M \approx 3.06$, but texture and active slip systems can change it. Neither factor is a universal material constant.

Example: From obstacle population to hardening

Given. Take $N = 4 \times 10^{22}$ m^{-3}, diameter $d = 10$ nm, $M = 3$, $\alpha = 0.3$, $G = 80$ GPa and $b = 0.25$ nm. The reference yield strength is 250 MPa.

Model and substitution. In the dispersed barrier model,

$$Nd = (4 \times 10^{22})(10 \times 10^{-9}) = 4 \times 10^{14}\ \text{m}^{-2}, \qquad L = (Nd)^{-1/2} = 5.0 \times 10^{-8}\ \text{m} = 50\ \text{nm}.$$

Using $G = 80 \times 10^{9}$ Pa and $b = 0.25 \times 10^{-9}$ m,

$$\Delta\sigma_y = \frac{M\alpha G b}{L} = \frac{3(0.3)(80 \times 10^{9})(0.25 \times 10^{-9})}{50 \times 10^{-9}} = 360\ \text{MPa}.$$

The estimated irradiated yield strength is $250 + 360 = 610$ MPa.
Meaning. The 360 MPa is an increment, so it must be added to the reference strength. Although L and d are both lengths, the obstacle spacing here is five times its diameter. These assumed model inputs describe an example, not a measured alloy dataset.

Increasing N fourfold at fixed other inputs halves L and doubles hardening. Coarsening can have the opposite effect. For spherical obstacles, volume fraction scales as Nd^3. Doubling d while reducing N to $N/8$ preserves this volume fraction but reduces Nd to one quarter, so hardening halves if α is unchanged. Constant obstacle volume does not mean constant obstacle resistance.

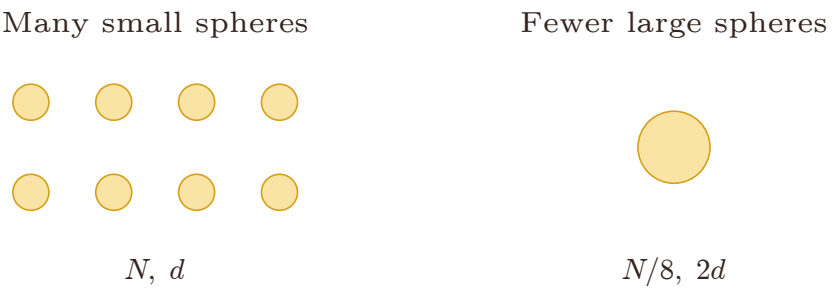


Figure 5.4: At fixed obstacle volume, coarsening can increase slip-plane spacing.

Key point

A hardening estimate requires **size, number density and interaction strength**. Two populations with equal stored defect content can have different spacings and different values of α. Microstructural images supply part of the evidence; they do not determine the full strength increment without an interaction model.

Different obstacle populations require a superposition rule supported for their interactions. Linear addition gives $\Delta\sigma = \sum_i \Delta\sigma_i$; root-sum-square (RSS) addition gives $\Delta\sigma = (\sum_i \Delta\sigma_i^2)^{1/2}$. RSS is not automatically the more accurate choice. A dislocation network also resists slip, approximately as $\tau \sim \alpha_d G b \sqrt{\rho_d}$, where ρ_d is line length per volume in m^{-2}. Its interaction factor α_d need not equal α. Recovery, coarsening, recombination and loss to sinks can produce a hardening plateau with dose, but neither a plateau nor its onset dose is universal. Heating can remove obstacles or promote additional solute-related hardening.

5.3 Fracture and the ductile–brittle transition

Ductility is the amount of plastic deformation before fracture. Fracture toughness measures resistance to extension of a sharp crack. These quantities cannot be inferred from yield strength alone. A material can withstand a large stress before general plastic flow and still have poor tolerance of an existing flaw. The fracture mode also matters. Ductile tearing often involves cavity nucleation, growth and joining with substantial local plasticity. Cleavage can propagate with much less macroscopic deformation. If boundaries lose resistance, a crack can instead follow an intergranular path.

A Charpy impact test breaks a notched specimen with a pendulum and measures absorbed energy in joules. Repeating the same test at different temperatures gives an energy–temperature curve. In ferritic steels the lower shelf has low energy, the upper shelf has higher energy, and the transition lies between them. An explicit energy criterion defines one transition temperature. For example, T_{41} is the temperature at which the test energy reaches 41 J. Other definitions need not give the same numerical temperature.

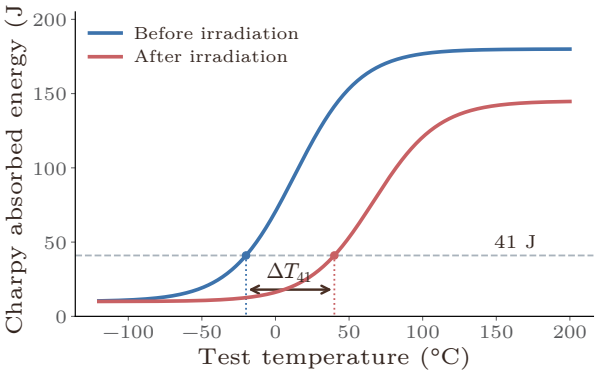


Figure 5.5: A higher transition temperature and a lower upper shelf are separate possible irradiation effects.

To read the transition figure, choose the energy criterion first, draw a horizontal line to each curve, and then read the corresponding temperatures. The horizontal separation gives the transition shift. The upper-shelf energy is read from the high-temperature plateau instead. A curve can move horizontally, change its upper shelf, or do both; one reported number cannot describe all three parts of the curve.

Compare reference and irradiated states using the same specimen geometry, loading procedure and energy criterion:

$$\Delta T_{41} = T_{41,\mathrm{irr}} - T_{41,\mathrm{ref}}.$$

A positive shift means that a warmer specimen is needed to attain the same energy. A temperature difference has the same numerical value in kelvin and degrees Celsius. The classic cleavage transition is especially relevant to ferritic steels; one should not assign the same response to every alloy family.

Hardening helps explain the shift. Cooling ferritic steel raises the resistance to plastic flow. Irradiation adds further dislocation resistance, reducing the opportunity for crack-tip plastic relaxation before cleavage. A simplified local model assumes that the fracture-stress criterion

is unchanged and that the reference yield strength decreases with temperature at m MPa/K, where $m > 0$. A rise ΔT then offsets an added hardening $\Delta\sigma_y$ when

$$m\Delta T \approx \Delta\sigma_y, \qquad \Delta T \approx \frac{\Delta\sigma_y}{m}.$$

The slope is specific to the chosen local model. This is a physical estimate, not a universal conversion from hardening to transition shift. Boundary changes and the crack-tip stress state can matter independently.

Key point

Keep the measured quantities separate: **Charpy energy** has units of joules, **transition temperature** has units of degrees Celsius or kelvin, and **fracture toughness** has units such as MPa$\sqrt{\text{m}}$. A shift in one test can motivate closer assessment with another, but it is not a unit conversion between them.

An existing crack concentrates the applied loading. For an appropriate linear-elastic opening-mode geometry,

$$K_I = Y\sigma\sqrt{\pi a}.$$

Here K_I is the applied stress intensity in MPa$\sqrt{\text{m}}$, σ is tensile stress in MPa, a is the geometry-defined crack size in metres and Y is a dimensionless geometry factor. At fixed Y, doubling stress doubles K_I, whereas quadrupling crack size doubles it. Millimetres must be converted to metres before using a toughness expressed per square root metre.

Compare the applied K_I with the relevant measured crack resistance at the temperature and material condition of interest. K_{IC} specifically requires plane-strain validity and small-scale yielding. Charpy energy is not K_{IC}: the notch, loading rate, geometry and reported dimensions differ. A calculation of K_I determines crack loading; it does not measure toughness or establish a component acceptance limit.

Example: Transition shift and crack loading

Given. The 41 J transition temperature changes from -30 to 15 °C. Separately, take crack size $a = 4$ mm, geometry factor $Y = 1$ and tensile stress $\sigma = 150$ MPa.

Model and substitution. Compare the same impact-energy criterion:

$$\Delta T_{41} = 15 - (-30) = 45 \text{ K}.$$

For the stated linear-elastic crack model, convert $a = 0.004$ m:

$$K_I = Y\sigma\sqrt{\pi a} = 1(150)\sqrt{\pi(0.004)} = 16.8 \text{ MPa}\sqrt{\text{m}}.$$

Meaning. The first result describes a change in an impact-test curve. The second describes the applied loading of a crack. At the same stress, a 16 mm crack would give twice this K_I, or 33.6 MPa$\sqrt{\text{m}}$. Assessing resistance requires a toughness measured for the relevant material condition and temperature; the transition shift supplies no such numerical value by itself.

5.4 Reactor pressure vessel steels

The reactor pressure vessel contains the core and coolant and must tolerate pressure, temperature changes and flaws throughout service. Low-alloy ferritic steels supply strength and toughness, but irradiation can change their crack resistance. The beltline is the vessel region near the height of the fuel, where neutron exposure is especially relevant. It is a location in the vessel, not a separate material layer.

Several microstructural changes contribute. Defect clusters provide obstacles to dislocations. Copper-rich features can cause substantial hardening when enough residual copper is present. Manganese–nickel–silicon precipitates can also matter, including in low-copper steels, depending on composition and exposure. Segregation at grain boundaries can alter resistance to intergranular fracture. A description based only on matrix strength misses that boundary contribution.

Reducing residual copper can reduce a particular source of embrittlement, but it does not remove every mechanism. The amount, composition and evolution of solute-rich features depend on the alloy and irradiation history. Two welds at the same fluence can therefore show different transition shifts. Initial microstructure, irradiation temperature and other solutes must be considered before assigning a difference entirely to one element. This is why measured surveillance data remain useful even for well-controlled compositions.

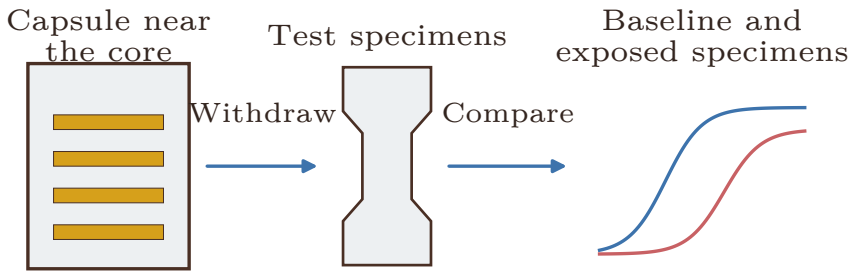


Figure 5.6: Representative specimens and dosimetry connect irradiation history to measured property changes.

A surveillance programme places representative vessel materials in irradiation capsules. Capsules are withdrawn, their exposure is assessed by dosimetry, and the specimens are tested. Tensile measurements, Charpy curves and fracture tests answer different questions about the evolving material. The programme samples representative materials and conditions; it does not directly test every point in the vessel. Differences between capsule and vessel exposure must therefore be considered when interpreting the results.

The surveillance sequence links three kinds of evidence. Dosimetry estimates the irradiation exposure. Mechanical tests describe the response of the exposed specimens. Comparison with the reference material identifies a property change. Keeping these links explicit helps distinguish an exposure difference from a material-composition difference when two capsules or welds behave differently.

Key point

A vessel assessment combines **material resistance** with **applied crack loading**. Surveillance informs the resistance side; pressure and temperature histories inform the loading side. Neither a fluence value nor a transition shift alone contains both.

A simple comparison can use $T - T_{41}$ at a stated operating or test temperature T. A decreasing value shows that the temperature has moved closer to the chosen transition criterion. This temperature difference is a comparison quantity, not a sufficient safety margin: it does not contain flaw size, geometry, applied stress or a measured toughness curve. Crack loading still requires those inputs. If a constant fluence rate $\dot{\Phi}$ is assumed, accumulated fluence is $\Phi = \dot{\Phi}t$, with fluence in $\mathrm{n/m^2}$ and rate in $\mathrm{n/(m^2\,s)}$. Fluence counts incident neutrons per area; it is not interchangeable with dpa without a damage calculation.

Pressurized thermal shock is a loading scenario in which rapid cooling of the inner vessel region occurs while pressure loading remains important. The warmer surrounding wall restrains contraction of the cooled region, producing thermal stresses. These act together with pressure stresses. At the same time, cooling can reduce ferritic-steel crack resistance. It is the combination of material condition, temperature history, stress and flaws that makes the transient important; pressure alone does not describe it.

For the simple linear-elastic crack model in Section 5.3, stress contributions acting in the same opening direction may be added before calculating K_I, when their stated geometry factors are the same. Real through-wall thermal stress distributions require an appropriate crack analysis; a uniform-stress exercise is only an illustration. A vessel assessment also includes uncertainty and the required evaluation margins.

Cooling makes the pressurized thermal-shock problem particularly coupled. The thermal gradient can increase crack-opening stress while the lower local temperature moves the steel toward the less tough part of its response. Thus the same transient changes both sides of the comparison between K_I and crack resistance. The important temperature is the temperature experienced by the material around the flaw during the transient, together with the corresponding stress field.

Annealing can allow some defects and solute structures to change, partly recovering hardening and transition temperature. The extent depends on composition, annealing temperature and hold time. Recovery must be measured; it is not automatically complete, and later neutron exposure can cause further embrittlement. A recovered tensile strength alone cannot demonstrate recovered crack resistance.

Example: Quantifying partial recovery

Given. A steel has reference $T_{41} = -35\,^{\circ}\mathrm{C}$, irradiated $T_{41} = 45\,^{\circ}\mathrm{C}$ and post-anneal $T_{41} = 5\,^{\circ}\mathrm{C}$. Use the same test criterion for all three states.

Model and substitution. Measure both changes relative to their stated starting points:

$$\Delta T_{\mathrm{irr}} = 45 - (-35) = 80\ \mathrm{K}, \qquad \Delta T_{\mathrm{recovered}} = 45 - 5 = 40\ \mathrm{K}.$$

The recovered fraction of the irradiation shift is $40/80 = 0.50$, or 50%. The remaining shift from the reference is $5 - (-35) = 40$ K.

Meaning. Half of the measured transition shift has recovered; this does not mean that half of every defect has disappeared. Different microstructural features contribute differently to the property. Subsequent exposure, temperature and loading history are needed before drawing a service-life conclusion.

5.5 Irradiation creep, swelling and growth

Creep is time-dependent plastic deformation under sustained stress. It can occur below the yield stress obtained in a short tensile test. During irradiation, a continuing supply of point defects can support additional dislocation motion. A test performed only after irradiation probes the retained damage structure; it cannot reproduce every effect of supplying defects while the load is applied. Beam-on and beam-off comparisons must use the same temperature and stress and account for evolving microstructure.

Glide moves a dislocation within its slip plane. Climb moves it normal to that plane by absorption or emission of point defects. Irradiation supplies vacancies and interstitials that can support climb. In stress-induced preferential absorption (SIPA), differently oriented dislocations capture defects differently under stress. Unequal climb produces a directional strain. This orientation comparison differs from the sink bias between interstitial and vacancy capture considered in swelling.

Stress-induced preferential nucleation (SIPN) means that stress changes the energy cost of forming differently oriented loops. Their unequal formation can contribute to a transient strain, but SIPN alone generally does not explain sustained steady creep. In climb-controlled glide, stress first drives glide until a dislocation is blocked. Defect absorption permits climb past the obstacle, and glide resumes. Climb can control the waiting time even if glide carries much of the strain. Several mechanisms may contribute to one measured rate; a linear stress dependence does not uniquely identify one of them.

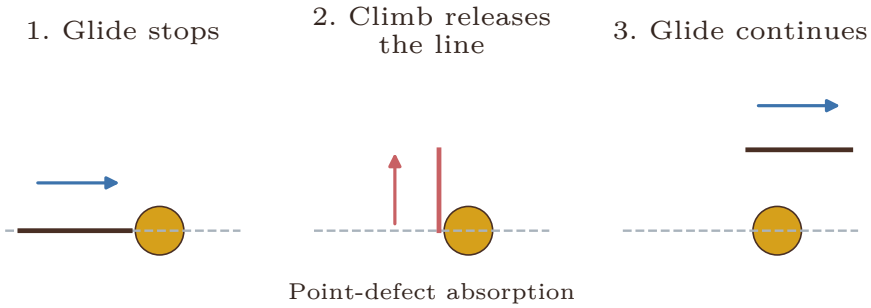


Figure 5.7: Climb can release a blocked dislocation so that stress-driven glide continues.

In the climb-controlled glide sequence, the obstacle stops motion within the slip plane. Point-defect absorption then changes the dislocation position normal to that plane. Once the obstacle is bypassed, the applied stress drives further glide. A small amount of climb can therefore unlock a much larger glide distance; the rate-limiting motion need not carry most of the final strain.

A simple steady uniaxial irradiation-creep law is

$$\dot{\varepsilon}_c = B_0 \sigma \dot{\mathcal{D}}, \qquad \varepsilon_c = B_0 \sigma \mathcal{D} \quad \text{for constant } B_0, \sigma \text{ and zero initial creep strain.}$$

Dose $\mathcal{D}$ is in dpa, stress σ in MPa and compliance B_0 in $\mathrm{MPa}^{-1}\,\mathrm{dpa}^{-1}$. The time rate is in s^{-1} if $\dot{\mathcal{D}}$ is in dpa/s. At fixed state, twice the dose rate doubles strain accumulated in the same time, but gives the same strain at the same total dose. Actual dose-rate changes can alter microstructure and B_0.

Swelling is a volume increase. For a small isotropic volumetric swelling fraction $S = \Delta V/V$, the axial expansion is approximately $S/3$. Irradiation growth instead changes shape through directional strains, potentially with little volume change and without applied stress. HCP zirconium alloys are an important example; crystal anisotropy and grain texture control the component response. Swelling, growth and stress-driven creep can coexist.

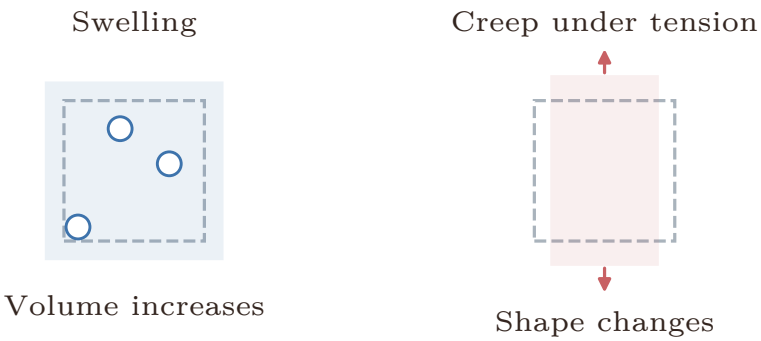


Figure 5.8: Volume expansion and stress-driven shape change contribute differently to dimensions.

Definition: Three dimensional changes

Creep is strain accumulated under sustained stress. **Swelling** is a volume increase. **Irradiation growth** is a directional change of dimensions that can occur without applied stress and with little volume change. A measured length change may contain all three, so a mechanism cannot be identified from one length measurement alone.

For small principal strains, the volume strain is approximately their sum: $\Delta V/V \simeq \varepsilon_x + \varepsilon_y + \varepsilon_z$. Equal positive strains therefore give isotropic swelling, whereas one positive and two compensating negative strains can change shape with little volume change. This explains why the two shapes in the dimensional-change figure cannot be compared using axial extension alone.

A swelling-dependent empirical creep law is

$$\dot{\varepsilon}_c = \sigma(B_0\dot{\mathcal{D}} + C_s\dot{S}), \qquad \frac{d\varepsilon_c}{d\mathcal{D}} = \sigma\left(B_0 + C_s\frac{dS}{d\mathcal{D}}\right),$$

where C_s has units MPa^{-1}. Stress multiplies both terms, and S is a fraction, not a percentage number. Both terms describe creep. Direct swelling expansion $S/3$ is added separately when calculating axial strain; do not add the creep coupling twice.

Example: Adding creep and swelling strains

Given. Over 10 dpa, assume constant $\sigma = 80$ MPa, $B_0 = 10^{-6}$ MPa^{-1} dpa^{-1}, $C_s = 0.005$ MPa^{-1} and $dS/d\mathcal{D} = 2 \times 10^{-4}$ dpa^{-1}. Other strains are absent.

Model and substitution. The creep strain per unit dose is

$$\frac{d\varepsilon_c}{d\mathcal{D}} = 80\left[10^{-6} + 0.005(2 \times 10^{-4})\right] = 1.6 \times 10^{-4}\ \text{dpa}^{-1}.$$

Multiplying by 10 dpa gives $\Delta\varepsilon_c = 0.0016$. The volumetric swelling is $\Delta S = (2\times10^{-4})(10) = 0.002$. Small isotropic swelling adds axial strain $S/3 = 0.000667$, so

$$\Delta\varepsilon_{\text{axial}} = 0.0016 + 0.000667 \simeq 0.00227 = 0.227\%.$$

Meaning. The C_s term has already increased the stress-dependent creep. Direct swelling expansion is a separate geometric contribution. For a 0.5 m initial length, this total strain gives extension $0.5(0.00227) \simeq 1.13$ mm under the stated constant-stress conditions.

If total extension is held fixed, creep can relax stress. With constant Young's modulus E and no changing thermal, swelling or growth strain, $\varepsilon_{\text{total}} = \sigma/E + \varepsilon_c$ gives $\Delta\sigma = -E\Delta\varepsilon_c$. The creep strain replaces elastic strain. This is a fixed-extension condition; it must not be confused with the constant-stress integration above.

5.6 Helium effects and materials selection

Nuclear reactions can generate helium inside a metal. For example, neutron absorption by boron-10 can produce lithium-7 and an alpha particle. A nickel route first converts nickel-58 to

nickel-59 by neutron capture; a subsequent (n,α) reaction produces iron-56. An alpha particle becomes a helium atom after slowing and capturing electrons. These are examples rather than a complete reaction list; gas production depends on composition, isotope inventory, neutron spectrum and exposure history.

$$^{10}_{5}\mathrm{B} + {}^{1}_{0}\mathrm{n} \longrightarrow {}^{7}_{3}\mathrm{Li} + {}^{4}_{2}\mathrm{He}, \qquad {}^{58}\mathrm{Ni}(n,\gamma)^{59}\mathrm{Ni}(n,\alpha)^{56}\mathrm{Fe}.$$

The notation omits branch-dependent excitation and energy terms. For an assumed constant helium production ratio r_{He} in appm/dpa,

$$c_{\mathrm{He}} = r_{\mathrm{He}}\mathcal{D}.$$

Atomic parts per million (appm) count helium atoms per million reference material atoms in this dilute convention. Multiply appm by 10^{-6} to obtain an atomic fraction, or divide appm by 10^{4} to obtain at.%. A constant He/dpa ratio is a convenient assumption for an exercise, not a universal property of fission, fusion or a particular alloy.

Helium has very low solubility in most metals. Interstitial helium in α-iron has a low migration barrier, but freely moving helium and helium trapped in vacancy complexes do not have the same mobility. Vacancies, dislocations, particle interfaces and grain boundaries can collect it. Where helium is trapped affects bubble nucleation and the eventual damage distribution.

For a spherical interface, a capillarity estimate of bubble pressure difference is

$$\Delta p = p_{\mathrm{in}} - p_{\mathrm{out}} \approx \frac{2\gamma}{r},$$

where surface energy γ is in $\mathrm{J/m^2}$ and bubble radius r is in metres. The result is in pascals. A smaller radius gives a larger estimated pressure difference at fixed surface energy. Solid constraint, the gas equation of state, atomistic effects and nonequilibrium growth are neglected; this expression is not a complete nanoscale bubble model.

Example: Helium concentration and bubble pressure

Given. Assume helium production ratio $r_{\mathrm{He}} = 6$ appm/dpa over 15 dpa. Separately consider a spherical bubble with $\gamma = 1.5\ \mathrm{J/m^2}$ and radius $r = 6$ nm.

Model and substitution. The dilute concentration convention gives

$$c_{\mathrm{He}} = 6(15) = 90\ \text{appm} = 90 \times 10^{-6}\ \text{atomic fraction} = 0.009\ \text{at.\%}.$$

The capillarity estimate is

$$\Delta p \simeq \frac{2\gamma}{r} = \frac{2(1.5)}{6 \times 10^{-9}} = 5 \times 10^{8}\ \mathrm{Pa} = 500\ \mathrm{MPa}.$$

Meaning. The pressure estimate concerns one assumed bubble radius. The total helium concentration does not determine that radius: the number of bubbles and the partition of helium among traps are also needed. Halving the assumed radius doubles this pressure difference within the same capillarity model.

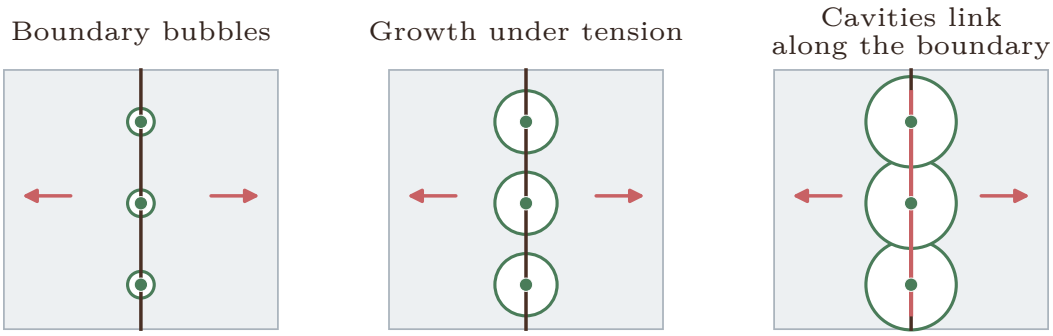


Figure 5.9: Helium-stabilized boundary cavities can grow and join, reducing load-bearing area.

At suitable temperature and loading, helium can stabilize grain-boundary cavities that grow and link under tension. This reduces the load-bearing boundary area and can promote intergranular failure with low ductility. Vacancy transport, stress and boundary structure matter along with gas pressure. Heating may remove matrix obstacles and reduce yield strength while helium remains at boundaries or redistributes into larger cavities. Recovery of strength therefore does not prove recovery of ductility or crack resistance.

The boundary-cavity figure should be read as a possible mechanism under suitable conditions. Isolated cavities first reduce the intact area locally; growth and joining can create a connected fracture path. The spatial distribution matters as much as the specimen-average helium concentration: concentrating the same inventory at a small fraction of boundaries can change where deformation and separation occur.

Key point

Removing matrix obstacles can lower yield strength while boundary cavities remain. Evaluate recovery using the property needed for the application: a tensile strength measurement cannot by itself reveal how an existing crack or a weakened boundary will behave.

Stable intragranular interfaces, as in oxide-dispersion-strengthened (ODS) alloys, can trap helium away from major boundaries. Trapping redistributes helium; it does not remove it. Capacity, stability, joining, anisotropy and toughness still matter. This feature alone cannot establish radiation resistance.

Material choice follows the failure mode. Austenitic steels offer fabrication, toughness and corrosion benefits but face swelling, creep and helium-assisted boundary damage. Ferritic/martensitic steels often resist swelling better in comparable regimes; hardening, cleavage transitions and high-temperature strength remain concerns. Nickel-base alloys offer strength and corrosion resistance but require spectrum-dependent helium assessment. Tungsten-based plasma-facing materials face brittle fracture, hardening and gas damage.

Compare temperature, dose rate, recoil spectrum, helium, microstructure, loading and test conditions alongside dpa; for ion–neutron comparisons, include damage depth and sampled volume. Measure crack resistance for cold vessel transients, swelling and creep for duct clearance, and combined gas and displacement effects for fusion structures.

References. G. S. Was, *Fundamentals of Radiation Materials Science: Metals and Alloys* (Springer, 2007), Chs. 12–14. K. L. Murty and I. Charit, *An Introduction to Nuclear Materials: Fundamentals and Applications* (Wiley-VCH, 2013), Ch. 5 and Sec. 6.2. US NRC, *Backgrounder on Reactor Pressure Vessel Issues* and *Reactor Vessel Integrity*. H. Ke et al., *Thermodynamic and kinetic modeling of Mn–Ni–Si precipitates in low-Cu reactor pressure vessel steels* (2017), arXiv:1707.08072. C.-C. Fu and F. Willaime, *Physical Review B* **72**, 064117 (2005). P. D. Edmondson et al., *Scripta Materialia* **65**, 731–734 (2011).

Unit 5 questions

Conceptual questions

5.C1 Why can higher yield strength accompany lower uniform elongation?

5.C2 Why does reducing obstacle spacing increase the stress needed for slip?

5.C3 Why can irradiation hardening favour cleavage at a higher temperature?

5.C4 Distinguish creep, swelling and irradiation growth.

5.C5 Why does recovery of yield strength not demonstrate recovery from boundary helium damage?

Numerical problems

5.N1 Find engineering stress for $F = 5$ kN and $A_0 = 20$ mm^2.

5.N2 Find the transition shift if T_{41} changes from -45 to $25\,^\circ$C.

5.N3 Find helium concentration in appm and at.% for 8 appm/dpa over 25 dpa.

Combined problems

5.P1 Initially $N = 4 \times 10^{22}$ m^{-3} and $d = 4$ nm. Coarsening doubles d and changes N to $N/8$. Calculate both spacings and the final strength if the initial hardening increment was 240 MPa and the reference strength is 300 MPa. Keep other factors fixed.

5.P2 Over 20 dpa, take $\sigma = 100$ MPa, $B_0 = 10^{-6}$ MPa^{-1} dpa^{-1}, $C_s = 0.004$ MPa^{-1} and $dS/d\mathcal{D} = 5 \times 10^{-4}$ dpa^{-1}, all constant. Find creep strain, swelling fraction and total extension of a 0.5 m length. Ignore other strains and use small isotropic swelling.

Unit 6: Nuclear Fuel Materials and Performance

6.1 Fuel systems and oxide fuel 75

6.2 Heat transfer and mechanical response 77

6.3 Burnup, fission gas, and fuel restructuring 79

6.4 Comparing ceramic fuel families 82

6.5 Metallic fuels and fuel–cladding interaction 84

6.6 Coated particles and accident-tolerant concepts 87

Questions 90

Learning Objectives

1. Connect fuel structure, porosity, and composition to uranium inventory.
2. Calculate temperature rises along the fuel-to-coolant heat path.
3. Relate burnup, fission-gas inventory, pressure, and dimensional change.
4. Compare oxide, carbide, nitride, metallic, and coated-particle fuel systems.

Key quantities

Quantity	Symbol	Common units
Linear power	q'	W/m
Thermal conductivity	k	W/(m K)
Burnup	B	GWd/tHM
Gas release fraction	F_{rel}	dimensionless
Gas pressure	P	Pa

6.1 Fuel systems and oxide fuel

A nuclear fuel must contain the required actinide inventory, produce heat, and transfer that heat to the coolant while its dimensions and chemistry evolve. Fuel selection therefore concerns a system: fuel, interfaces, cladding, coolant, and operating history. A material with high conductivity can still be unsuitable if it reacts with its surroundings or cannot retain an acceptable shape during irradiation. Reproducible fabrication and evidence from relevant tests are part of the materials decision.

Key Concepts

- **Inventory:** heavy-metal mass contained in the fuel volume.
- **Heat removal:** temperature needed to carry the generated power.
- **Integrity:** barrier loads caused by evolving dimensions, gas and chemistry.

A *pellet* is a small fuel body, commonly a sintered ceramic cylinder. A *rod*, or pin, contains fuel inside a sealed cladding tube. An *assembly* supports a group of rods and provides coolant passages. In a conventional oxide rod, a stack of pellets occupies the heated length. The initial pellet–cladding gap permits fabrication tolerances and dimensional change. A gas plenum provides free volume, and a spring holds the stack. Cladding separates the fuel from coolant and provides a barrier to fission products. These functions explain why the rod cannot be represented by the pellet alone.

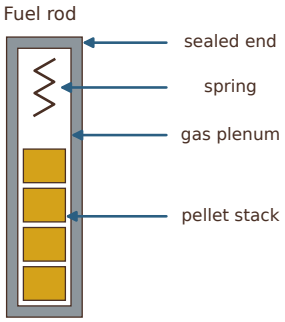


Figure 6.1: Components of a generic oxide-fuel rod; dimensions are schematic.

UO_2 is an important reference fuel because it combines high melting temperature, established ceramic processing, and extensive operating experience. Its limitations include relatively low thermal conductivity and brittle fracture. Defects, pores, gas bubbles, and cracks develop or change under irradiation, so the initial material properties do not fully describe its later behavior. The useful question is how those properties change over the intended temperature and irradiation range.

6.1.1 Crystal structure, density, and defects

UO_2 has the fluorite structure. Uranium occupies an FCC sublattice, and oxygen occupies all eight tetrahedral sites in the conventional cell. Counting shared uranium atoms gives four U

atoms per cell; with eight O atoms, the cell contains four UO_2 formula units. Each U has eight nearest O neighbors, while each O has four nearest U neighbors. The two coordination numbers refer to different central species and need not be equal.

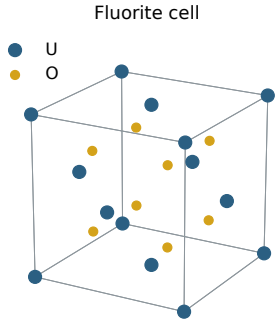


Figure 6.2: Fluorite conventional cell: four uranium atoms and eight oxygen atoms after sharing is counted.

Read the cell drawing by separating **atom counting** from **coordination**. Corner and face uranium atoms are shared with adjacent cells when calculating Z. Coordination instead counts the immediate neighbors around one chosen atom, including neighbors in adjacent cells. The density equation uses the shared-cell count, not the coordination number.

For a fully dense ideal crystal,

$$\rho_{\mathrm{TD}} = \frac{ZM}{N_A a^3}, \qquad \rho_{\mathrm{pellet}} = d\rho_{\mathrm{TD}}, \qquad \rho_U = \rho_{\mathrm{pellet}} w_U.$$

Here Z is formula units per cell, M is compound molar mass, a is the lattice parameter, N_A is Avogadro's constant, and d is the fraction of theoretical density. Uranium density ρ_U is uranium mass per total pellet volume, not the density of the compound. For UO_2, $w_U \simeq 238/270$. If pores contain negligible mass, porosity is approximately $1 - d$. Use consistent length units: 1 Å is 10^{-8} cm.

Example: From pellet density to uranium inventory

Given. A pellet occupies 2.0 cm^3 and has density fraction $d = 0.94$. Use $\rho_{\mathrm{TD}} = 10.96$ $\mathrm{g/cm}^3$ and $w_U = 238/270$.

Model. First find compound mass per volume, then retain only its uranium mass fraction:

$$\rho_{\mathrm{pellet}} = d\rho_{\mathrm{TD}}, \qquad m_U = V d \rho_{\mathrm{TD}} w_U.$$

Substitution. $\rho_{\mathrm{pellet}} = 0.94(10.96) = 10.30$ $\mathrm{g/cm}^3$; $\rho_U = 10.3024(238/270) = 9.08$ $\mathrm{g/cm}^3$. Therefore $m_U = 2.0(9.0814) = 18.2$ g.

Meaning. The pellet contains about 20.6 g of compound, of which 18.2 g is uranium. Its approximate 6% porosity reduces both masses per external volume. Multiplying the compound density by volume alone would count oxygen as uranium.

Oxygen excess or deficiency is written UO_{2+x} or UO_{2-x}. Charge compensation can involve changes in uranium valence, defects, and defect clusters. Temperature and oxygen chemical potential determine which states are accessible. These changes matter because defects scatter

heat-carrying excitations and because oxygen exchange can affect phase stability and fuel–cladding reactions. Excess oxygen does not generally improve UO_2 conductivity.

Pores reduce heavy-metal inventory and interrupt conduction paths. Their size and connectivity also influence densification and gas behavior. Fabrication therefore controls more than bulk density: powder preparation, pressing, sintering, finishing, inspection, loading, and sealing establish dimensions, grain size, pore distribution, composition, and rod integrity. A conductivity value must be accompanied by temperature, composition, microstructure, and irradiation state. Phonon scattering strongly affects the lower-temperature contribution in UO_2; additional high-temperature contributions can alter its trend. A schematic curve is insufficient as a property correlation.

6.2 Heat transfer and mechanical response

Fission energy is deposited inside the fuel, whereas cooling occurs outside the cladding. The temperature at the fuel center therefore depends on the complete path: fuel conduction, gap transfer, cladding conduction, and the coolant film. Confusing fuel surface temperature with coolant temperature omits several resistances and can substantially underestimate the center temperature.

6.2.1 Radial conduction

Consider a long, solid cylindrical pellet of radius R, with uniform volumetric heating q''' in $\mathrm{W/m^3}$. Assume steady radial conduction, no axial heat flow, and constant thermal conductivity k in W/(m K). Linear power is $q' = \pi R^2 q'''$ in W/m. A balance on a cylinder of radius r gives $2\pi r q''(r) = \pi r^2 q'''$, so

$$q''(r) = \frac{q'''r}{2}, \qquad T(r) = T_s + \frac{q'}{4\pi k}\left(1 - \frac{r^2}{R^2}\right), \qquad T_c - T_s = \frac{q'}{4\pi k}.$$

The radial flux is zero at the center by symmetry, even though the center is hottest. Flux grows with radius because each larger cylindrical surface must carry the heat generated inside it. The temperature gradient is zero at the center and largest in magnitude at the outer surface.

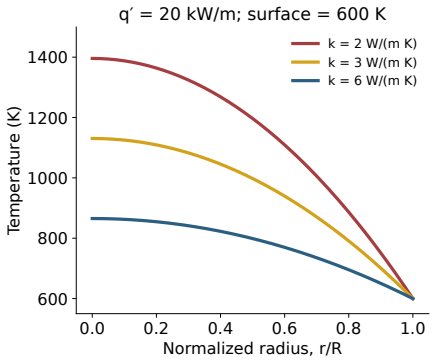


Figure 6.3: At fixed linear power and surface temperature, lower conductivity increases the internal temperature rise.

Radius disappears from the center-rise expression only because linear power is fixed. At fixed q''', the rise is $q'''R^2/(4k)$ and increases with radius squared. Doubling k halves the center-to-surface rise; it does not halve the absolute temperature. When conductivity depends on temperature, the corresponding relation is $\int_{T_s}^{T_c} k(T)\,dT = q'/(4\pi)$ under the same geometry and source assumptions. A valid property model is needed to evaluate that integral.

Key point: what is held fixed?

At fixed **linear power**, $\Delta T_f = q'/(4\pi k)$ is independent of radius in this model. At fixed **volumetric heating**, $\Delta T_f = q'''R^2/(4k)$. A larger pellet therefore has a larger rise in the second comparison because it generates more power per unit length.

6.2.2 Outer resistances and dimensional change

For a thin gap with effective conductance h_g in W/(m^2 K),

$$\Delta T_g \simeq \frac{q'}{2\pi R h_g}, \quad \Delta T_c^{\text{clad}} = \frac{q'}{2\pi k_c}\ln\frac{r_o}{r_i}, \quad \Delta T_{\text{film}} = \frac{q'}{2\pi r_o h_c}.$$

Here r_i, r_o are cladding radii, k_c is cladding conductivity, and h_c is coolant-film conductance. Gap size, gas composition, pressure, radiation, and solid contact affect h_g. The center temperature is the bulk coolant temperature plus the film, cladding, gap, and fuel drops. Thermal margin must compare that operating temperature with the relevant degradation or phase limit, including property changes and uncertainty.

Example: Adding the temperature drops

Given. Let $q' = 18$ kW/m, $k = 3$ W/(m K), $R = 4$ mm and $h_g = 6000$ W/(m^2 K). Bulk coolant is 570 K; film and cladding drops together are 40 K.

Model. Use constant-property cylindrical conduction and add successive temperature rises from the coolant toward the center.

$$\Delta T_f = \frac{18000}{4\pi(3)} = 477.5\ \text{K}, \qquad \Delta T_g = \frac{18000}{2\pi(0.004)(6000)} = 119.4\ \text{K}.$$

The fuel surface is $T_s = 570 + 40 + 119.4 = 729.4$ K. Hence

$$\boxed{T_c = 729.4 + 477.5 \simeq 1207\ \text{K}.}$$

Meaning. The fuel supplies the largest drop in this assumed case. If h_g doubles with all other inputs fixed, only the gap drop halves: T_c falls by about 59.7 K. If k doubles instead, the fuel rise halves and T_c falls by about 238.7 K. These comparisons identify which property change matters most for this set of conditions.

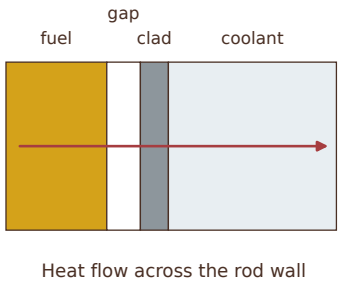


Figure 6.4: Heat crosses several interfaces between the heat-generating fuel and bulk coolant.

Temperature gradients also create mechanical mismatch. Free linear thermal strain is approximately $\epsilon_{\text{th}} = \alpha \Delta T$, where α is the expansion coefficient in K^{-1}. A freely expanding radius changes by $\Delta R = R\epsilon_{\text{th}}$. This is a strain estimate, not a stress solution: neighboring material and cladding constrain expansion. Brittle pellets can crack and redistribute stress while the cladding remains intact. Fragment relocation may narrow the external gap and improve its conductance, but internal cracks can impede conduction and produce uneven contact.

Densification removes or changes fabrication pores and can shrink the pellet. Swelling from solid fission products and retained gas bubbles increases volume. Expansion, relocation, swelling, and cladding deformation together determine contact. Pellet–cladding mechanical interaction (PCMI) concerns the resulting loads; pellet–cladding interaction (PCI) can also include chemistry. A susceptible cladding state, tensile stress, and aggressive species such as iodine can act together. Pellet edge geometry, liners, suitable dimensions, and power-history control address different parts of this problem. Neither a contact sketch nor a simple expansion calculation predicts cladding failure.

6.3 Burnup, fission gas, and fuel restructuring

Burnup measures thermal energy extracted per initial heavy-metal mass:

$$B = \frac{E_{\text{thermal}}}{m_{\text{HM},0}}, \qquad B\,[\text{GWd/tHM}] = \frac{P\,[\text{MW}]\,t\,[\text{d}]}{1000 m_0\,[\text{tHM}]}.$$

The second form assumes constant thermal power. At variable power, integrate power over time; equivalent full-power days give the same energy at a stated reference power. Heavy-metal mass excludes oxygen, carbon, or nitrogen in the compound. Electrical output is not the required energy basis. The useful unit identity is 1 MWd/kgHM = 1 GWd/tHM. Local burnup, rod-average burnup, and an inventory average describe different spatial scales.

FIMA means *fissions per initial heavy-metal atom*. Assuming 238 g/mol initial heavy metal and 200 MeV thermal energy per fission gives approximately 0.107% FIMA per GWd/tHM. This is a conversion under stated assumptions, not an exact identity for every isotope mixture. It follows by dividing energy per kg by energy per fission and then by initial atoms per kg. Burnup records integrated exposure but does not specify temperature, power ramps, or the resulting

microstructure by itself.

Example: From operating history to burnup

Given. A 0.50 tHM inventory produces 20 MW thermal for 100 days, then 10 MW for 40 days.

Model. Add the thermal energy in each interval before dividing by the initial heavy-metal mass:

$$E = 20(100) + 10(40) = 2400 \text{ MWd}, \qquad B = \frac{2400}{1000(0.50)} = 4.8 \text{ GWd/tHM}.$$

The same energy corresponds to $2400/20 = 120$ equivalent full-power days at 20 MW, although elapsed operation is 140 days. With the stated conversion, FIMA is approximately $4.8(0.107\%) = 0.514\%$.

Meaning. Equal burnup means equal extracted energy per initial heavy-metal mass. It does not imply equal instantaneous power or fuel temperature. The two operating intervals contribute different energy per day and may drive different rates of gas transport.

6.3.1 Gas transport, release, and pressure

Xenon and krypton produced by fission can remain in the lattice, form bubbles within grains, or reach grain boundaries. A boundary bubble is not automatically a released bubble: a connected route to a free surface is needed. Connected boundary pores and cracks can carry gas into the gap and plenum. For diffusion over distance L with effective diffusivity D_g,

$$t_{\text{diff}} \sim \frac{L^2}{D_g}, \qquad F_{\text{rel}} = \frac{N_{\text{released}}}{N_{\text{produced}}}.$$

The diffusion expression is a scale estimate. Larger grains can increase the distance to a boundary, but trapping, irradiation, cracking, and changing diffusivity prevent a universal inverse relation between grain size and measured release. For stable-gas accounting with no other losses, produced equals retained plus released. State the species, interval, and release boundary; radioactive inventories may also require decay and transmutation terms.

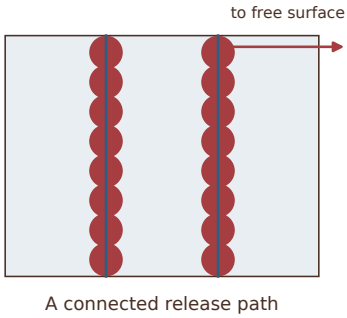


Figure 6.5: Release requires connected paths from gas-containing regions to a free surface.

Gas retained in bubbles contributes to swelling and local stress. Gas released into an intact rod changes the free-volume inventory, not the environmental release. Xe and Kr mixed into

helium-rich gap gas can reduce gap conductance, which can raise fuel temperature and accelerate transport. Gap closure may counteract this effect. These couplings make temperature and gas release part of the same performance problem.

An ideal-gas estimate for a uniform gas temperature is

$$P = \frac{nR_gT_g}{V_g}, \qquad R_g = 8.314\ \mathrm{J\,mol^{-1}\,K^{-1}}.$$

Use total gas, including initial fill gas, absolute temperature, and the current free volume. One cm^3 is 10^{-6} m^3. Actual rods have nonuniform gas temperature and changing volume, so this estimate does not replace a full pressure calculation.

Example: Gas release and rod pressure

Given. Of 2.0 mmol stable fission gas produced, 1.8 mmol remains in the fuel. The rod initially contains 0.8 mmol fill gas. Current free volume is 20 cm^3 at 500 K.

Model. Apply the gas inventory balance before the ideal-gas law:

$$n_{\mathrm{rel}} = n_{\mathrm{produced}} - n_{\mathrm{retained}}, \quad F_{\mathrm{rel}} = \frac{n_{\mathrm{rel}}}{n_{\mathrm{produced}}}, \quad P = \frac{(n_{\mathrm{fill}} + n_{\mathrm{rel}})R_gT_g}{V_g}.$$

Substitution. Released gas is 0.2 mmol, so $F_{\mathrm{rel}} = 0.2/2.0 = 0.10$. Total free gas is 1.0 mmol = 0.001 mol:

$$P = \frac{0.001(8.314)(500)}{20 \times 10^{-6}} = 2.08 \times 10^5\ \mathrm{Pa} = 0.208\ \mathrm{MPa}.$$

Meaning. Using all produced gas would overestimate the free-gas inventory; ignoring fill gas would underestimate it. This is absolute internal pressure. A cladding load calculation must also include external pressure and the current tube geometry.

6.3.2 High-burnup structure

The outer pellet can develop a high-burnup structure: original grains subdivide, and gas redistributes into a changed pore population. High local burnup together with relatively low local temperature favors this rim structure; there is no single threshold applicable to every fuel history. Closed pores can retain gas, whereas connected pores provide possible release paths. Heating, stress, and cracking can change that connectivity. Fine subdivisions and pores also alter heat transport, swelling, and mechanical response.

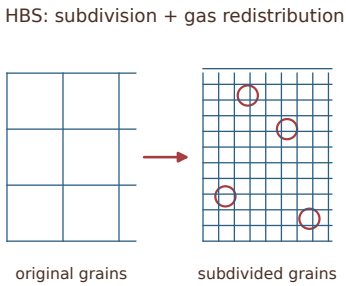


Figure 6.6: High-burnup restructuring changes grain subdivision and pore populations.

The rim is therefore not uniformly beneficial or harmful. Its effect depends on the property and transient under consideration. Fragmentation means breaking into smaller pieces; dispersal means those pieces leave their original containment region through a path. Fragmentation alone does not determine the mass that escapes. Cladding condition, available openings, and flow also matter. A useful performance assessment keeps gas location, pore connectivity, and the containment boundary distinct.

Definition: Retained and released fission gas

Retained gas remains within the fuel material, including its bubbles and closed pores. **Released** gas has crossed the chosen fuel boundary into free volume. Release from fuel inside an intact rod is therefore different from release beyond the cladding.

6.4 Comparing ceramic fuel families

A fuel comparison must specify the quantities being held fixed. Conductivity at one temperature cannot rank complete fuel systems. Density, porosity, composition, irradiation state, power, and boundary conditions must be matched or explicitly stated. High uranium density can increase inventory per volume, while high conductivity can reduce the internal temperature rise. Neither benefit establishes chemical compatibility, acceptable swelling, or reliable fabrication.

Key point: use a common comparison basis

Compare **uranium density** to assess inventory in a fixed volume. Compare **temperature rise** at matched power and boundary conditions to assess heat transport. Compare **margin** using each material's relevant phase or degradation limit. These three comparisons answer different questions.

Mixed oxide fuel (MOX) commonly combines UO_2 and PuO_2 in a fluorite-related oxide matrix. Using recovered plutonium is one motivation, but fuel-cycle benefits depend on the full accounting basis. Plutonium fraction and distribution affect local fission rates and gas generation. Pu-rich regions are a reminder that average composition need not describe local behavior. Fabrication controls this distribution; a schematic heterogeneous pellet does not by itself demonstrate an isolated thermal spike. Comparing MOX with UO_2 requires matched temperature, density,

oxygen stoichiometry, and power and burnup history. There is no universal MOX conductivity or release factor.

ThO_2 is also fluorite and has high thermal stability. Thorium-232 is fertile: neutron capture forms ^{233}Th, followed by beta decays through ^{233}Pa to fissile ^{233}U. The route requires an initial fissile driver; an inventory of fertile thorium alone does not supply the required initial chain reaction. Mixed composition, fabrication, handling, and fuel-cycle infrastructure affect practicality. Thorium-based fuel is not automatically waste-free or free of proliferation concerns.

6.4.1 Carbide and nitride fuels

UC and UN have the rock-salt structure. The uranium sublattice is FCC, and C or N occupies the octahedral sublattice. Both species have six nearest neighbors of the other species, unlike the 8:4 coordination of fluorite. High uranium density and useful thermal conductivity motivate these compounds. Their carbon or nitrogen balance, phase constitution, irradiation behavior, and reactions with cladding or coolant must also be assessed.

For UC, carbon balance affects phase constitution and chemical compatibility. For UN, nitrogen stoichiometry and high-temperature stability matter, and nitrogen isotope selection adds another consideration:

$$^{14}\mathrm{N}(n,p)^{14}\mathrm{C}.$$

This reaction absorbs a neutron and produces carbon-14. Nitrogen enriched in ^{15}N can reduce that contribution, but isotope supply and cost enter the fuel-cycle decision. The specification depends on the reactor and fuel design. Water contact and oxidation of alternative fuels require specific evidence, especially for conditions after cladding breach.

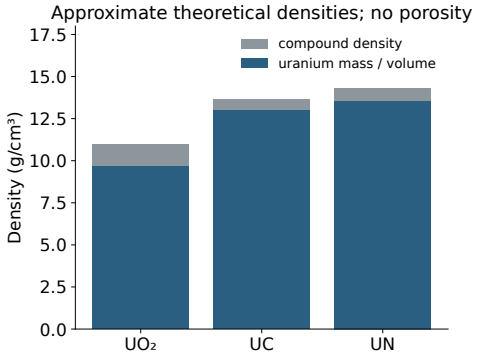


Figure 6.7: Approximate fully dense compound and uranium densities; pores lower actual pellet values.

The reference densities for dense UO_2, UC, and UN are approximately 10.96, 13.63, and 14.32 g/cm^3, respectively. Uranium mass fractions are approximately 238/270, 238/250, and 238/252. Thus compound density and uranium density give different numerical comparisons. Actual pellets need their density fraction d, with $\rho_U = d\rho_{\mathrm{TD}} w_U$. The approximation uses a uranium molar mass of 238 g/mol and neglects pores in the theoretical reference values.

Example: Comparing uranium inventories

Given. UN is 93% of its reference density, 14.32 g/cm^3. Take $w_U = 238/252$ and a fuel volume of 5 cm^3.

Model and substitution. The same density balance used for oxide fuel gives

$$\rho_U = d\rho_{\mathrm{TD}} w_U = 0.93(14.32)\frac{238}{252} = 12.58 \text{ g/cm}^3, \qquad m_U = 5(12.578) = 62.9 \text{ g}.$$

For comparison, the 94%-dense UO_2 calculated earlier contains $5(9.0814) = 45.4$ g U in this volume. The UN case therefore holds about $62.9/45.4 = 1.39$ times the uranium inventory.
Meaning. This comparison holds external fuel volume fixed and uses the two stated density fractions. It isolates inventory; conductivity, enrichment, neutron absorption and chemical compatibility require their own information. The nitrogen isotope fraction cannot be inferred from the uranium mass fraction.

A thermal comparison uses the same geometry and heat-transfer assumptions developed earlier. For example, at fixed q' and T_s, the fuel with twice the conductivity has half the internal rise. If it has a lower relevant phase limit or a worse external gap, the total margin may not improve. Margin must use the actual center temperature and the appropriate degradation limit, rather than a ranking based only on melting temperature.

Fuel	Main attraction	Important constraint
UO_2	Established oxide-fuel basis	Conductivity, cracks, and evolving gas
MOX	Actinide use in an oxide matrix	Composition distribution and history
ThO_2	Thermal stability	Fissile driver and fuel cycle
UC	Density and conductivity	Carbon balance and compatibility
UN	Density and conductivity	Nitrogen reactions and compatibility

The table identifies questions, not a universal order of merit. Irradiation tests must establish whether favorable initial properties persist as gas, defects, phases, and dimensions change. Measurements at relevant conditions and consistent fabrication are necessary before a material-level advantage becomes a demonstrated fuel-system advantage.

6.5 Metallic fuels and fuel–cladding interaction

Metallic fuel concepts include U–Zr and U–Pu–Zr alloys. Alloying changes phases and properties, so a metallic fuel is not described adequately by a single uranium property. High conductivity reduces the internal temperature gradient under matched conditions. However, an alloy may have a lower solidus than oxide fuel, and its local solidus can change with composition. Operating temperature and the appropriate phase limit must be considered together.

A suitable metallic-fuel pin can use a sodium bond between a dense fuel slug and its cladding. The bond transfers heat across the initial separation while free space allows dimensional change. This arrangement addresses two different requirements: an initially open region need not be a

poorly conducting gas gap, and good heat transfer need not require initial mechanical contact. Bond material, cladding, fuel, and coolant must all be compatible with the design. The geometry changes as fuel swells.

6.5.1 Smear density and swelling space

For a dense circular slug, define geometric area loading by

$$f_A = \frac{A_{\text{fuel}}}{A_{\text{clad,inside}}} = \left(\frac{R_f}{r_i}\right)^2, \qquad \frac{\Delta A}{A_0} = \frac{1}{f_A} - 1.$$

Here R_f is initial slug radius and r_i is cladding inner radius. The second expression gives relative area increase needed to fill the available cross section. At fixed length, area and volume increase fractions are the same. If the slug itself is porous, a density-based smear convention includes that porosity; always state the convention. Geometric loading is not automatically the actual material density.

The difference between available-area fraction and relative fuel expansion is important. If fuel occupies 80% of the available area, the remaining region is 20% of the available area, but it is 25% of the initial fuel area. These are different denominators. Neither percentage is a qualified swelling limit: axial expansion, changing cladding shape, contact, and constraints can modify the usable volume.

Example: Space for a swelling fuel slug

Given. A dense circular slug occupies $f_A = 0.64$ of the cross section inside cladding of inner radius 5 mm. Hold length and cladding dimensions fixed.

Model and substitution. Since area is proportional to radius squared,

$$R_f = r_i\sqrt{f_A} = 5\sqrt{0.64} = 4 \text{ mm}, \qquad \frac{\Delta A}{A_0} = \frac{1}{0.64} - 1 = 0.5625.$$

The initial radial separation is 1 mm. Filling it requires a 25% radius increase, $(5-4)/4$, and a 56.25% area increase. The initially unoccupied fraction is instead $1 - 0.64 = 36\%$ of the cladding's interior area.

Meaning. The 36% and 56.25% values describe the same space with different denominators. At fixed length, the latter also gives relative fuel-volume increase. Actual axial expansion, cladding deformation and mechanical constraint determine how much of this geometric capacity is usable.

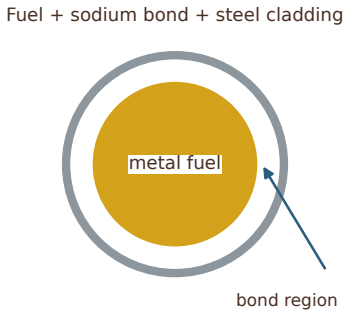


Figure 6.8: A metallic-fuel pin combines a heat-transfer bond, swelling space, cladding, and gas free volume.

Fission gas can form bubbles and contribute to swelling. As bubbles interconnect, gas can reach free volume. This can alter the later swelling response, but a single saturation value or universal plateau cannot represent every alloy, temperature, and history. After contact, axial motion, mechanical constraint, and gas pressure still matter. The initial geometry supplies useful capacity information but not the complete evolution law.

6.5.2 Chemical interaction and system response

Fuel–cladding chemical interaction (FCCI) changes the interface through reactions and interdiffusion. Fuel constituents and fission products, including important lanthanide contributions in metallic systems, can modify cladding composition and properties. Reaction layers can consume useful wall thickness or reduce local margin. Barriers and temperature limits must be evaluated for the actual material pair; conductivity alone does not remove this interface problem.

Temperature and chemical-potential gradients can also redistribute alloy constituents within the fuel. A uniform initial alloy need not remain uniform. Local phase constitution and solidus can therefore change even when average composition stays fixed. A performance model must track spatial composition along with temperature and burnup.

For a simple thin tube under uniform pressure, the pressure-only hoop stress is

$$\sigma_\theta \simeq \frac{(P_i - P_o) r_m}{t_c},$$

where r_m is mean radius and t_c is wall thickness. Consistent pressure units produce the same stress units. Positive internal-minus-external pressure produces tensile hoop stress. This thin-wall estimate omits thermal stress, pellet contact, creep, irradiation, and corrosion; it is not a failure criterion. A reduction in effective wall thickness raises the estimate at fixed pressure and radius.

For example, a 20% reduction in effective wall thickness multiplies the pressure-only stress by $1/0.80 = 1.25$. The interface reaction can therefore affect a mechanical calculation even if pressure stays unchanged. The metallic-pin drawing should be read as an evolving arrangement: swelling changes the bond region while chemical interaction can change the wall carrying the load.

The 1986 EBR-II loss-of-flow and loss-of-heat-sink tests without normal scram demonstrated a response of a particular reactor system. Negative reactivity feedback and heat removal limited that response; metallic fuel was one part of the tested configuration. Temperature-induced

feedback opposed the initiating temperature rise by reducing power. The result does not establish the same response for every reactor using metallic fuel. Geometry, feedback coefficients, coolant behavior, and heat removal remain essential.

6.6 Coated particles and accident-tolerant concepts

TRISO places several material functions around each small fuel kernel. The name means tristructural isotropic. The kernel generates heat and fission products; a porous carbon buffer accommodates dimensional change and provides space for gas. Dense inner pyrocarbon, silicon carbide, and outer pyrocarbon complete the coating system. SiC provides mechanical strength and an important transport barrier. Pyrocarbon layers affect stress as well as transport: irradiation-induced dimensional changes, anisotropy, defects, and interface debonding can change load transfer. A correct layer count alone does not establish coating quality.

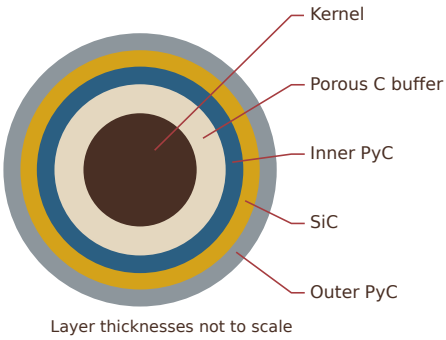


Figure 6.9: Kernel, buffer, pyrocarbon, and SiC have distinct functions in a TRISO particle; layers are not to scale.

Kernel composition also changes internal chemistry. UO_2 and UCO are examples; UCO contains oxide and carbide phases that alter oxygen and carbon chemistry. Gas chemistry and kernel evolution influence coating loads. Particles may be embedded in graphite-based pebbles or compacts. The matrix supports particles and conducts heat, while element geometry and coolant paths determine larger-scale removal. Fully ceramic microencapsulated concepts instead use a ceramic matrix such as SiC. Particle loading, matrix damage, heat flow, and manufacturing must then be evaluated together.

6.6.1 Failure, release, and qualification

Coating failure fraction is the number of particles meeting a defined structural-failure criterion divided by the number assessed. Species release fraction is the inventory of a named species transported beyond a stated boundary divided by its specified reference inventory. They are different measurements. Some species, including silver in studied TRISO systems, can migrate through nominally intact coatings. Conversely, gas leaving the kernel may remain inside the coating system. Temperature, time, species, and boundary must accompany any retention claim.

Qualification connects controlled kernel and coating specifications to irradiation tests, post-irradiation examination, and heating tests over relevant burnup, temperature, and fluence.

Models require comparison with independent measurements. One heating temperature or a small measured failure fraction does not characterize every particle design and history. The particle barrier and the reactor cooling system perform different jobs, and both must remain part of the assessment.

Key point: name the retention boundary

A kernel, coating system, compact and reactor each define a different boundary. A release fraction is meaningful only with a **named species**, **reference inventory**, **boundary** and **exposure history**. Coating failure and species release need not have the same numerical fraction.

6.6.2 Accident-tolerant fuel systems

Accident-tolerant fuel concepts aim to improve response under specified accident conditions while maintaining acceptable normal-operation behavior. They do not remove the need for cooling or make a reactor accident-proof. In zirconium-alloy cladding, high-temperature steam oxidation releases heat and hydrogen:

$$\mathrm{Zr} + 2\mathrm{H_2O} \longrightarrow \mathrm{ZrO_2} + 2\mathrm{H_2}.$$

The balanced equation gives amounts, not the reaction rate. Temperature, alloy condition, and oxide integrity affect the kinetics. In an ideal diffusion-controlled regime, scale thickness x follows $x^2 \simeq k_p t$, where k_p is a parabolic rate coefficient with units of length squared per time. Cracking, scale loss, or composition changes can invalidate that relation.

Cr-coated zirconium retains the Zr-alloy substrate while adding a protective surface. Adhesion, strain, wear, and interdiffusion determine whether protection persists. FeCrAl can form a protective alumina scale; aluminum supply and thermal history matter, alongside neutron absorption, irradiation behavior, strength, and tube design. SiC/SiC composites use fibers and interfaces to deflect or bridge cracks. Their direction-dependent properties, joining, coolant compatibility, and gas tightness need separate evaluation: improved toughness does not automatically make a hermetic tube.

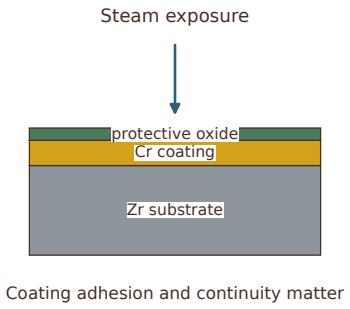


Figure 6.10: A surface coating changes oxidation behavior while the underlying cladding remains part of the system.

Modified fuel pellets address other mechanisms. Dopants can alter UO_2 grain growth, diffusion distances, creep, defects, and thermal properties. A larger grain can slow a simple diffusion process but does not guarantee a fixed reduction in gas release. UN and U_3Si_2 offer density and

thermal motivations, yet swelling, phases, and compatibility still require evidence. Water and steam reactions are particularly relevant for U_3Si_2; experience with a dispersion form does not automatically qualify a monolithic water-reactor pellet.

Example: Reading a parabolic oxidation law

Given. An initially bare surface follows $x^2 = k_p t$ with assumed $k_p = 0.25\ \mu\text{m}^2/\text{h}$. Find thickness at 16 h and the time to reach twice that thickness.

Model and substitution.

$$x = \sqrt{0.25(16)} = 2\ \mu\text{m}, \qquad t(4\ \mu\text{m}) = \frac{4^2}{0.25} = 64\ \text{h}.$$

Meaning. Doubling thickness requires four times the total exposure time in this regime. The additional time from 2 to 4 μm is $64 - 16 = 48$ h. The layer grows more slowly as its transport distance increases. A cracked or detached layer may no longer follow the same law, so these values describe the assumed regime rather than an alloy-specific accident prediction.

References. K. L. Murty and I. Charit, *An Introduction to Nuclear Materials: Fundamentals and Applications* (Wiley, 2013), Chapter 7; INL BISON, *UO2Thermal* and *Sifgrs* model documentation; Argonne National Laboratory, EBR-II Fuels Irradiation and Physics Database; P. A. Demkowicz, D. W. Marshall, and J. D. Hunn, *Advanced Gas Reactor Fuel Specification Technical Bases*, INL/RPT-23-71992 (2023); U.S. NRC, *Accident Tolerant Fuel Technologies*; E. S. Wood et al., *U_3Si_2 behavior in H_2O: Part I, flowing steam and the effect of hydrogen* (2018).

Unit 6 questions

Conceptual questions

6.C1 Distinguish a pellet, a rod, and an assembly, and state the purpose of the plenum.

6.C2 Distinguish PCMI from PCI, and explain why thermal strain alone does not predict failure.

6.C3 Distinguish the consequences of retained gas and gas released into an intact rod.

6.C4 Why must high conductivity and the alloy solidus be considered together?

6.C5 State the distinct functions of the kernel, porous buffer, pyrocarbon layers, and SiC.

Numerical problems

6.N1 A pellet is 92% of 10.96 g/cm^3 theoretical density. Find its compound density and approximate porosity.

6.N2 A 2 tHM inventory produces 40 MW thermal for 300 d. Find burnup.

6.N3 Find uranium density of fully dense UN using 14.32 g/cm^3 and 238/252 uranium mass fraction.

Combined problems

6.P1 A rod has $q' = 24$ kW/m, $k = 4$ W/(m K), $R = 4$ mm, and $h_g = 6000$ W/(m^2 K). Coolant is 580 K; film and cladding drops total 40 K. Calculate gap drop and center temperature.

6.P2 Of 4 mmol stable fission gas produced, 90% stays in fuel. Add released gas to 0.6 mmol fill gas, and find pressure at 600 K in 15 cm^3 using $R_g = 8.314$ J/(mol K). Assume a sealed rod and ideal gas.

Unit 7: Corrosion and Environmental Degradation

7.1 Electrochemical reactions and metal loss 92

7.2 Passive films and localized corrosion 94

7.3 Reactor-water chemistry and radiolysis 96

7.4 Stress corrosion cracking and irradiation effects 98

7.5 Zirconium oxidation and hydrogen 101

7.6 Compatibility with advanced coolants 103

Questions 106

Learning Objectives

1. Connect electrochemical reactions and current balance to metal loss.
2. Distinguish film stability, passivity, and localized corrosion.
3. Explain how water chemistry, irradiation, and tensile loading affect cracking.
4. Relate zirconium oxidation to hydrogen generation and pickup, and compare coolant-compatibility requirements.

Key quantities

Quantity	Symbol	Common units
Electrode potential	E	V versus a reference
Current density	j	$\mu\mathrm{A/cm^2}$
Oxide thickness	x	$\mu\mathrm{m}$
Parabolic rate coefficient	k_p	$\mu\mathrm{m^2/h}$
Hydrogen pickup fraction	f_H	dimensionless

7.1 Electrochemical reactions and metal loss

Corrosion connects a material to its chemical environment. Reaction may remove load-bearing metal, produce a surface film, or create a local site that later develops into a pit or crack. The first distinction is between driving force and rate. A reaction that lowers Gibbs free energy is thermodynamically permitted under the stated conditions; it need not proceed rapidly. Surface barriers, charge-transfer kinetics and transport through the solution or film can all slow it. Conversely, a small average loss does not establish that every location retains adequate thickness.

Key Concepts

- **Thermodynamics** identifies permitted reactions and stable states.
- **Kinetics and transport** determine how rapidly material changes.
- **Localization and loading** determine whether that change threatens a component.

A corrosion cell contains oxidation and reduction reactions connected by two current paths. Electrons move through the metal, while ions carry charge through the electrolyte. The anode is the site of oxidation; the cathode is the site of reduction. Both may occupy different regions of the same component. For the chosen iron dissolution reaction,

$$\mathrm{Fe} \rightarrow \mathrm{Fe}^{2+} + 2e^-, \qquad \mathrm{O_2} + 2\mathrm{H_2O} + 4e^- \rightarrow 4\mathrm{OH}^-.$$

Two iron atoms supply the four electrons consumed by one oxygen molecule. In an acidic environment, hydrogen evolution, $2\mathrm{H}^+ + 2e^- \rightarrow \mathrm{H_2}$, can supply another cathodic route. These half-reactions specify electron bookkeeping; the actual dissolved species depend on solution chemistry and surface films.

An electrode potential must be reported relative to a reference. The standard hydrogen electrode (SHE) defines the conventional zero. A standard reduction potential describes specified standard states; a galvanic series instead ranks measured behavior in a specified environment. Neither alone predicts the rate of a particular component. For a written reduction reaction, the equilibrium potential is

$$E = E^\circ - \frac{RT}{nF} \ln Q,$$

where E and E° are in volts, $R = 8.314$ J mol^{-1} K^{-1}, T is absolute temperature, n is the number of electrons and $F = 96485$ C mol^{-1}. The reaction quotient Q uses dimensionless activities. For $\mathrm{Fe}^{2+} + 2e^- \rightleftharpoons \mathrm{Fe}$, pure solid iron has unit activity. At 25 °C,

$$E = E^\circ + 0.02958 \log_{10} a_{\mathrm{Fe}^{2+}}.$$

A lower ion activity shifts this equilibrium potential downward. It does not by itself specify a corrosion current.

At open circuit the signed partial currents balance: $\sum j_a + \sum j_c = 0$, taking anodic current as positive. The external current can therefore be zero while iron dissolves continuously. The corrosion potential is the mixed potential at which all surface currents balance. Faster cathodic kinetics can move that balance toward higher dissolution on an active anodic branch; a passive

branch may respond differently. Faraday's law converts the dissolution current into reacted mass. For constant current I, exposure time t, molar mass M and unit dissolution efficiency,

$$m = \frac{MIt}{nF}, \qquad x = \frac{m}{\rho A} = \frac{Mjt}{nF\rho}, \qquad j = I/A.$$

Here ρ is metal density and A is exposed area. Using grams, centimetres, amperes and seconds gives m in grams and x in centimetres. Only the part of the current that dissolves metal belongs in this expression. A uniform penetration estimate spreads the loss over A and cannot locate a deep pit. For Fe^{2+} dissolution with $M = 55.85$ g/mol and $\rho = 7.87$ g/cm^3, 1 μA/cm^2 corresponds to 0.01160 mm per 365-day year.

Definition: Corrosion current and net current

The **corrosion current** represents the magnitude associated with metal oxidation under a stated reaction model. The **net external current** is the signed sum of all partial currents. Zero net current at open circuit does not imply zero metal dissolution.

Example: Separating a rate from an equilibrium potential

Given. Iron dissolves uniformly as Fe^{2+} at $j = 3$ μA/cm^2 for two 365-day years. Separately, compare Fe^{2+} activities of 10^{-2} and 10^{-4} at 25 °C.

Rate model. With unit dissolution efficiency, $x = Mjt/(nF\rho)$. Convert j to 3×10^{-6} A/cm^2 and t to $2(365)(86400) = 6.3072 \times 10^7$ s:

$$x = \frac{55.85(3 \times 10^{-6})(6.3072 \times 10^7)}{2(96485)(7.87)} = 0.006959 \text{ cm} = 0.0696 \text{ mm}.$$

Equilibrium model. The Nernst relation gives

$$\Delta E = 0.02958 \log_{10} \frac{10^{-4}}{10^{-2}} = -0.05916 \text{ V}.$$

Meaning. The first result is an accumulated uniform loss under a specified current. The second is a shift in one reaction's equilibrium potential. Predicting a new corrosion rate after changing chemistry also requires the competing electrode reactions and their kinetics.

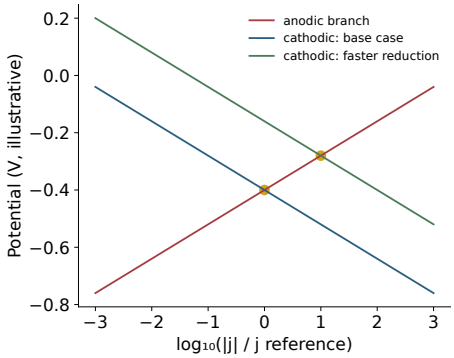


Figure 7.1: Changing cathodic kinetics shifts the mixed potential and current on the illustrated active anodic branch.

The intersection of the partial-current curves represents a **balance of reaction rates**. At that potential the positive anodic current and negative cathodic current cancel externally. Their individual magnitudes remain nonzero. Moving the cathodic curve changes the intersection because the active anodic branch must then supply the matching current. The drawing separates this mixed-potential condition from the equilibrium potential of either isolated half-reaction.

7.2 Passive films and localized corrosion

A passive film lowers reaction rates by separating the metal from the solution and restricting charge or species transport. A film can be thermodynamically stable yet offer little protection if it is porous, cracked, poorly attached or discontinuous. Real films may contain several oxides and hydroxides. Protection depends on their structure and on whether a damaged region can repassivate under the local conditions. Alloy composition supplies film-forming elements, but the environment helps determine which compounds form and remain stable.

An active–passive polarization curve records changes in surface behavior. Current can rise during active dissolution, fall as an effective film develops, and remain comparatively low over a passive range. At higher potential, current may rise again. Localized film breakdown, transpassive dissolution and oxygen evolution are distinct processes that can contribute to this rise. Current magnitude alone does not distinguish them. Surface examination and knowledge of the reactions are needed before assigning a mechanism.

A Pourbaix diagram maps thermodynamic states against electrode potential and pH for stated temperature, activities and selected species. A metal-stability region represents immunity; a dissolved-species region permits dissolution; a stable solid compound may support passivity. Such a map supplies neither a rate nor a test of film adhesion. In particular, a solid-stability field does not guarantee a protective film. Local chemistry inside a pit may also lie far from the measured bulk condition.

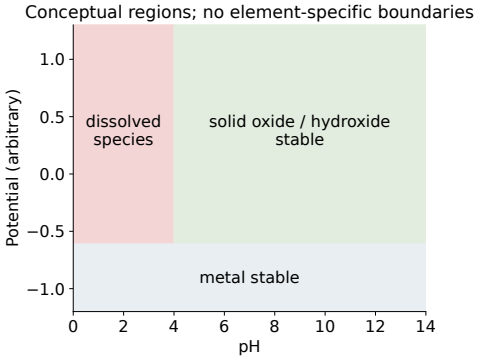


Figure 7.2: Potential–pH regions distinguish stable chemical states. The boundaries are schematic and do not specify an alloy's operating window.

To interpret a potential–pH point, first identify the selected temperature, activities and species used to construct the map. Then ask which state is stable there. A stable dissolved species suggests a dissolution route; a stable solid suggests a possible surface product. Whether that solid forms a continuous, adherent barrier remains a separate materials question. For water at

25 °C, unit gas fugacity relative to 1 bar and water activity near one, the equilibrium lines are

$$E_{\mathrm{H^+/H_2}} = -0.05916\,\mathrm{pH}, \qquad E_{\mathrm{O_2/H_2O}} = 1.229 - 0.05916\,\mathrm{pH}$$

in V versus SHE. Their separation is 1.229 V under these assumptions. Raising pH shifts both downward without changing that separation. These are equilibrium relations; overpotential and kinetics affect observed gas evolution, and the component's mixed potential need not shift by the same amount.

A pit begins at a locally active site in an otherwise less active surface. Dissolved metal ions can hydrolyze, changing local acidity. Ion migration maintains charge balance, and restricted exchange with bulk water can preserve aggressive conditions. Crevices create a related transport restriction geometrically: joints, deposits or narrow gaps isolate a small volume of solution. Oxygen consumption, hydrolysis and ion movement then alter that volume. Bulk-water measurements alone may miss the local environment driving attack.

Galvanic area balance explains why concentrated loss can be severe. If A_a and A_c are anodic and cathodic areas, average partial current densities satisfy

$$j_a A_a = |j_c| A_c.$$

A large cathode can support a high current density on a small anode. This equation balances charge; polarization and transport must still establish the actual current. It is not valid to hold a real cathodic current density fixed through arbitrary changes of geometry without checking those effects.

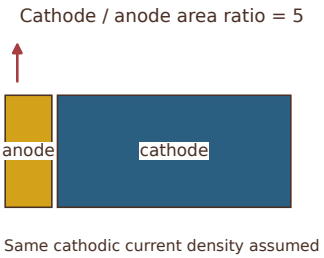


Figure 7.3: A small anodic region can carry the current supplied by a much larger cathodic surface.

Flow-accelerated corrosion (FAC) also couples chemistry and transport. Oxide dissolves into water, flow removes dissolved material, and continued metal reaction replaces the lost barrier. Chemistry, temperature, alloy content and flow geometry influence the thinning. FAC therefore differs from purely mechanical erosion by particles or cavitation. A thickness map, surface observations and water history provide more useful evidence than an average mass loss alone.

Key point: area averaging can hide damage

Uniform corrosion spreads metal loss over a broad area. **Localized corrosion** concentrates it in a pit, crevice or small anodic region. Equal total mass loss can leave very different minimum wall thicknesses, so remaining section should be assessed where attack is deepest.

Example: Concentrating current on a small anode

Given. A 2 cm^2 anode couples to a 12 cm^2 cathode with assumed $|j_c| = 1.5\ \mu\text{A/cm}^2$. The anode dissolves uniformly as Fe^{2+}.

Model. Charge conservation requires equal total anodic and cathodic current magnitudes, although their current densities differ.

$$I = |j_c|A_c = 1.5(12) = 18\ \mu\text{A}, \qquad j_a = \frac{18}{2} = 9\ \mu\text{A/cm}^2.$$

Using the iron conversion above gives an annual thickness loss of

$$x = 9(0.01160) = 0.1044\ \text{mm}.$$

Meaning. The area ratio is six, so the anodic current density is six times the cathodic magnitude. If the same total current were confined to 1 cm^2, the local average dissolution rate would double. In a real component, polarization and transport may change the total current as geometry changes; the fixed-current comparison isolates the area effect.

7.3 Reactor-water chemistry and radiolysis

Water in a reactor acts as a coolant and a reacting environment. Temperature, pressure, radiation, dissolved species, flow and boiling determine the history experienced by a surface. In a pressurized water reactor (PWR), the bulk coolant remains subcooled, although local nucleate boiling can occur at the cladding. In a boiling water reactor (BWR), water boils in the core. Phase separation changes the transport of species, so a concentration measured at one location cannot automatically be assigned to all wetted surfaces.

PWR chemistry control coordinates several functions. Dissolved hydrogen influences oxidizing radiolysis products and surface electrochemical conditions. Boron and alkalinity must support reactivity control and pH management while remaining compatible with materials. High purity limits aggressive impurities and unwanted deposits. Appropriate concentrations depend on the plant and operating state; the mechanisms here do not define plant chemistry specifications. Startup, steady operation and changing power can expose surfaces to different histories.

Radiolysis begins when ionizing radiation deposits energy in water. Selected products include the hydrated electron e^-_{aq}, hydrogen radicals $H^\bullet$, hydroxyl radicals $OH^\bullet$, molecular hydrogen and hydrogen peroxide. These products form a reaction network, not a single balanced decomposition equation. $OH^\bullet$ and H_2O_2 can supply oxidizing pathways, while e^-_{aq} and $H^\bullet$ participate in reducing pathways and rapid radical reactions. Production does not equal the concentration arriving at

a component: species react with each other, with dissolved additives and with surfaces while flowing through the system.

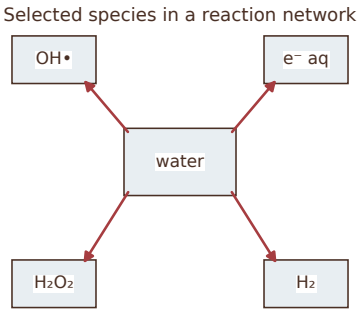


Figure 7.4: Radiation creates several reactive species whose subsequent reactions and transport determine local water chemistry.

A one-species balance makes this distinction explicit:

$$\frac{dc}{dt} = G - kc - \frac{c - c_{\text{in}}}{\tau}, \qquad c_{\text{ss}} = \frac{G + c_{\text{in}}/\tau}{k + 1/\tau}.$$

Here c is concentration, G is production per volume per second, k is a first-order removal coefficient in s^{-1}, c_{in} is inlet concentration, and τ is residence time in seconds. At steady state the derivative is zero. With concentration in μmol/L, G has units $\mu\text{mol L}^{-1}\ \text{s}^{-1}$. The reaction loss is kc and the net flow loss is $(c - c_{\text{in}})/\tau$. A negative net flow loss means the inlet supplies species. This balance omits the coupled reactions of a real radiolysis calculation; its purpose is to connect local inventory with production and removal.

For zero inlet concentration, shorter residence time lowers the steady inventory at fixed G and k. Increasing reaction removal does the same. When the inlet is nonzero, the sign and size of the flow contribution must be evaluated rather than assumed. A measured low concentration could reflect weak production, strong reaction loss or rapid transport; concentration by itself does not identify the controlling process.

The balance can also be written as $G + c_{\text{in}}/\tau = (k + 1/\tau)c_{\text{ss}}$. The left side contains two sources: local production and delivery by the inlet. The right side contains reaction removal and outflow. This form makes the units and signs easier to check. If $k = 0$ and $G = 0$, it reduces to $c_{\text{ss}} = c_{\text{in}}$, as expected for replacement by unchanged inlet water.

Electrochemical corrosion potential (ECP) reflects the balance of partial surface currents. Oxidants, hydrogen, catalytic surface condition, temperature and transport all contribute. Dissolved oxygen alone cannot determine ECP in irradiated water. Hydrogen water chemistry in a BWR can reduce oxidizing species and lower ECP in relevant regions, and surface catalysts can promote recombination. The effectiveness depends on flow, boiling and surface state. A low ECP is not a universal guarantee against cracking: susceptible nickel-base alloys can still undergo primary-water SCC under reducing PWR conditions.

Corrosion products generated on one surface can be transported and deposited elsewhere. Deposits can alter heat transfer and concentrate local chemistry; source and damage locations need not coincide. Assessment therefore combines bulk chemistry records with local surface

conditions, deposits, transport and temperature. Changes in chemistry should be evaluated for their system-wide consequences as well as their effect on one component (IAEA, 2011).

Key point: bulk chemistry and local chemistry

A sample describes the water at its sampling location and time. The surface of interest may experience different temperature, deposits, boiling or residence time. Connecting the measurement to corrosion therefore requires the **transport path** and **local surface condition**, as well as concentration.

Example: Production balanced by reaction and flow

Given. Let $G = 12$ μmol L^{-1} s^{-1}, $k = 2$ s^{-1}, $\tau = 1$ s and $c_{\text{in}} = 3$ μmol/L. Assume a well-mixed volume and steady conditions.

Model and substitution. Set $dc/dt = 0$ in the species balance:

$$c_{\text{ss}} = \frac{G + c_{\text{in}}/\tau}{k + 1/\tau} = \frac{12 + 3/1}{2 + 1/1} = 5 \ \mu\text{mol/L}.$$

Balance check. Reaction removes $kc = 2(5) = 10$ and net flow removes $(5 - 3)/1 = 2$, both in μmol L^{-1} s^{-1}. Their sum equals production, 12.
Meaning. Most removal occurs by reaction in this case. If residence time falls to 0.5 s, the same model gives $c = (12 + 3/0.5)/(2 + 1/0.5) = 4.5$ μmol/L. Reaction then removes 9 and net flow removes 3. Faster replacement lowers the inventory here because the incoming concentration is below the local steady concentration.

7.4 Stress corrosion cracking and irradiation effects

Stress corrosion cracking (SCC) requires a susceptible material state, a relevant environment and tensile crack-opening loading. Applied load is only one source of tension. Welding, machining and cold work can leave local residual stress and alter microstructure. Initiation depends on surface condition, local strain and exposure history; propagation additionally requires a crack and suitable crack-tip conditions. Final mechanical fracture depends on the remaining ligament and fracture resistance. A final fracture surface alone cannot reconstruct every preceding stage.

Intergranular cracks follow grain boundaries, while transgranular cracks cross grains. Morphology identifies the path but does not uniquely identify its cause. Several chemical and mechanical processes can produce an intergranular path. Records of processing, stress, water chemistry and irradiation must be considered alongside boundary chemistry and oxide observations. Small overall mass loss is consistent with a damaging crack because the affected volume can be small.

For a suitable linear-elastic geometry with limited crack-tip plasticity, mode-I loading is represented by

$$K_I = Y\sigma\sqrt{\pi a}, \qquad \Delta a = vt$$

when growth rate $v = da/dt$ is imposed as constant. Here Y is dimensionless, σ is nominal tensile stress in MPa and a is the crack dimension in metres for that geometry. Then K_I is in MPa$\sqrt{\text{m}}$.

Fixed Y and stress imply $K_I \propto \sqrt{a}$, not a. Environmental growth rates may vary with loading, chemistry and material state. A plateau or threshold observed in one test is conditional on its duration, detection limit and history; it is not a universal no-cracking guarantee.

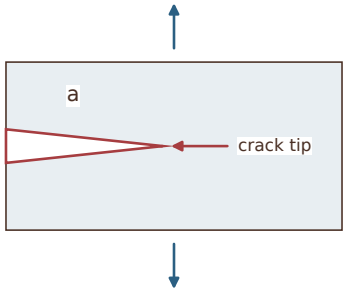


Figure 7.5: Crack geometry and nominal tensile stress jointly determine the mode-I loading estimate.

Thermal sensitization can occur when chromium-rich carbides precipitate during suitable exposure. Carbide growth draws chromium from adjacent matrix, leaving locally depleted zones even when bulk chromium remains high. A chromium-rich carbide and a depleted neighboring matrix are therefore consistent observations. Welding can place sensitization and residual tension in the same region. Lower carbon or suitable stabilization reduces this route under relevant processing conditions; neither substitutes for examining the actual material condition.

BWR intergranular SCC in susceptible stainless steel can combine sensitization, weld residual stress and oxidizing water. Primary-water SCC has affected Alloy 600 and related nickel-base weld metals in PWR service. Higher-chromium Alloy 690 and suitable weld metals generally improve resistance, with qualification still tied to processing and exposure. At Davis-Besse in 2002, leakage through a nozzle region exposed vessel-head steel to borated coolant and severe corrosion removed steel above the stainless cladding. Improving nozzle cracking resistance addresses one link; leakage detection, inspection and response to deposits address others (NRC, 2002).

Irradiation-assisted SCC (IASCC) introduces several coupled material changes. Radiation-induced segregation arises when vacancy and interstitial fluxes toward sinks couple differently to alloying elements. This differs from thermal sensitization, which consumes chromium in carbides. Irradiation also changes strength and deformation: slip may become concentrated in narrow dislocation channels. A channel terminating at a boundary can create high local loading, dependent on boundary orientation and slip transmission. Average applied stress does not fully describe that local condition.

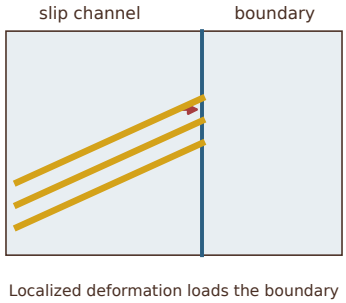


Figure 7.6: Localized deformation can concentrate loading where a dislocation channel meets a grain boundary.

Boundary oxidation and amplified local stress can work together to promote initiation in irradiated stainless steels. Thus a low-carbon stainless steel can still be susceptible: low carbon addresses thermal sensitization, while irradiation can change boundary chemistry and deformation through other routes. Dose alone cannot distinguish these mechanisms. Boundary chemical profiles, strain maps, oxide characterization and matched crack tests provide complementary evidence. A mitigation should target the demonstrated contribution rather than rely on morphology or dose alone (Was et al., 2023).

Key point: similar cracks can have different causes

Thermal sensitization changes boundary-adjacent chemistry through precipitation. **Irradiation-induced segregation** changes it through defect-mediated redistribution. **Localized deformation** changes boundary loading. An intergranular crack path can be consistent with more than one of these contributions; mechanism assessment needs chemical and mechanical evidence together.

Example: Crack loading and a specified growth rate

Given. Assume $Y = 1$, $\sigma = 150$ MPa and initial crack dimension $a = 2$ mm. A specified constant growth rate is $v = 2 \times 10^{-10}$ m/s for 500 h.

Model and substitution. Use metres in the stress-intensity expression:

$$K_I = 150\sqrt{\pi(0.002)} = 11.9 \text{ MPa}\sqrt{\text{m}}, \qquad \Delta a = 2 \times 10^{-10}(500)(3600) = 0.00036 \text{ m}.$$

Thus the extension is 0.36 mm and the final crack dimension is 2.36 mm. If stress and Y remain fixed, the final loading estimate is

$$K_{I,\text{final}} = 150\sqrt{\pi(0.00236)} = 12.9 \text{ MPa}\sqrt{\text{m}}.$$

Meaning. The crack grows 18% longer while K_I rises by about 8.6%, consistent with its square-root dependence. The imposed constant rate supplies the time dependence; the loading formula does not predict that rate or establish a fracture limit.

7.5 Zirconium oxidation and hydrogen

Zirconium-alloy cladding reacts with water to form oxide while generating hydrogen:

$$\mathrm{Zr} + 2\mathrm{H_2O} \rightarrow \mathrm{ZrO_2} + 2\mathrm{H_2}.$$

One reacted zirconium atom corresponds to two hydrogen molecules, or four hydrogen atoms. Some generated hydrogen enters the alloy. Oxide-growth rate and hydrogen-pickup fraction are separate quantities: the same reacted amount need not produce the same absorbed hydrogen inventory in different alloys or exposure histories. Metal converted to oxide no longer contributes to the metallic load-bearing section.

A compact oxide can restrict transport and slow further reaction. Protective zirconium oxide contains monoclinic and tetragonal phases; protection is not established by naming one phase alone. Stress, porosity, connected defects and cracking affect the available transport paths. As these features evolve, the oxide can undergo transitions from a slowing growth regime to renewed faster growth. Repeated transitions can produce an approximately linear long-time trend without implying that each growth interval was linear.

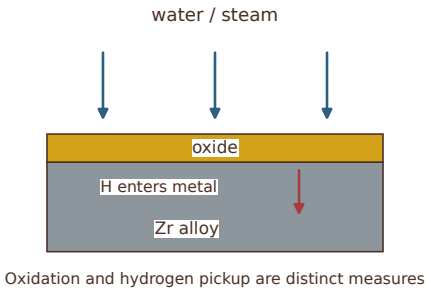


Figure 7.7: Oxidation consumes zirconium and produces hydrogen; only a condition-dependent fraction of that hydrogen enters the metal.

A useful description over a specified regime is $x^m = kt$, for an initially bare surface. Thickness x varies as t, $t^{1/2}$ and $t^{1/3}$ for linear, parabolic and cubic kinetics, respectively. The units of k depend on m. If thickness is in μm and time in hours, a parabolic constant k_p is in $\mu\mathrm{m}^2/\mathrm{h}$. For a layer of initial thickness x_0 that continues in the same parabolic regime,

$$x^2 - x_0^2 = k_p \Delta t.$$

An increasing transport distance can explain a decreasing rate: if flux and growth rate scale as $1/x$, integration gives $x^2 \propto t$. The exponent must still be interpreted alongside film structure; it is not a unique mechanism diagnosis.

Mass gain and metal consumption use different balances. If stoichiometric $\mathrm{ZrO_2}$ is completely retained and net specimen gain represents oxygen alone,

$$\frac{m_{\mathrm{Zr,consumed}}}{m_{\mathrm{O,gained}}} = \frac{91.22}{32.00} = 2.8506.$$

The oxide mass is the sum of consumed zirconium and gained oxygen. Oxide thickness additionally requires oxide density and area. Spallation, dissolution or additional mass changes invalidate a direct conversion from measured net gain to total reaction. A specimen can lose oxide while continuing to consume metal.

On a hydrogen-atom basis, define

$$f_H = \frac{n_{H,\text{absorbed}}}{n_{H,\text{generated}}}, \qquad n_{H,\text{absorbed}} = 4 f_H n_{\text{Zr,reacted}}.$$

The same relation works for numbers of atoms or moles, provided the basis is consistent. Pickup depends on alloy, oxide and exposure history; a selected percentage is an assumption, not a universal constant. Absorbed hydrogen can remain dissolved or precipitate as hydrides. Dissolved-hydrogen capacity generally decreases during cooling, and precipitation and dissolution can show hysteresis. A simple solubility trend therefore does not completely specify the precipitated amount after an arbitrary thermal history.

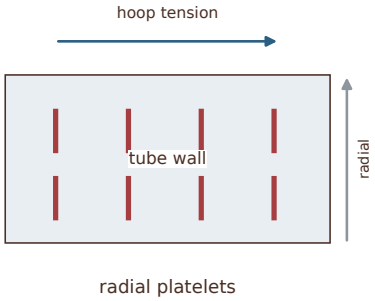


Figure 7.8: Radial hydride platelets can provide through-wall paths that open under hoop tension; connectivity and metal condition also matter.

In the unwrapped-wall drawing, the radial direction runs through the wall thickness and the hoop direction runs around the tube. Hoop tension opens a connected radial path across that thickness. This geometrical link explains why a hydrogen inventory alone cannot describe the fracture response: the same amount can be arranged in platelets with different orientations and connections. Hydride orientation, distribution and connectivity influence resistance to fracture along with total hydrogen. Texture, cooling history and stress during precipitation affect morphology. Two tubes with equal hydrogen inventory can have different resistance to hoop loading. Delayed hydride cracking (DHC) can proceed through hydrogen redistribution near a stressed crack tip, local hydride precipitation, hydride fracture and incremental advance. These events require time. Hydrides in a micrograph alone do not demonstrate active DHC; the loading, temperature history, available hydrogen and time-dependent crack advance must also be established (IAEA, 2010).

At elevated temperature, steam oxidation can accelerate sharply, release heat and generate hydrogen. Oxygen uptake also changes the metal condition. Normal-operation growth laws should therefore not simply be extrapolated through an accident transient with changing temperature, phases or transport paths.

Key point: keep three balances separate

Oxygen mass gain measures retained added oxygen under the stated assumptions. **Metal consumption** measures zirconium converted to oxide. **Hydrogen pickup** measures the fraction of generated hydrogen entering the alloy. Converting between them requires stoichiometry and, for pickup, a specified fraction.

Example: Oxide growth and hydrogen pickup

Given. An initially bare surface reaches 1.5 μm oxide at 50 h in a parabolic regime. Separately, 2 mmol Zr reacts with assumed hydrogen-pickup fraction $f_H = 0.10$.

Growth model. From $x^2 = k_p t$,

$$k_p = \frac{1.5^2}{50} = 0.045\ \mu\text{m}^2/\text{h}, \qquad x(200\ \text{h}) = \sqrt{0.045(200)} = 3.0\ \mu\text{m}.$$

Hydrogen balance. Each mole of reacted Zr generates four moles of H atoms, equivalent to two moles of H_2:

$$n_{H,\text{generated}} = 4(2) = 8\ \text{mmol}, \qquad n_{H,\text{absorbed}} = 0.10(8) = 0.80\ \text{mmol}.$$

Meaning. Four times the exposure produces twice the oxide thickness in this regime. The hydrogen result uses an atom basis; counting H_2 molecules as H atoms would introduce a factor-of-two error. Neither calculation specifies how the absorbed hydrogen partitions between solution and hydrides during later cooling.

7.6 Compatibility with advanced coolants

Coolant compatibility belongs to a complete material–environment system. Alloy composition, processing, temperature, stress, impurities, flow and exposure duration all matter. A material that forms a useful protective film in water may fail to form or retain that film in another coolant. Similarly, a static, isothermal coupon test cannot reproduce every mechanism in a circulating loop with temperature gradients. The test should reproduce the conditions that drive the degradation being assessed.

A nonisothermal loop can dissolve alloy constituents in one region and deposit them elsewhere. Temperature changes solubility and reaction equilibria, while circulation transports dissolved material. Deposition may restrict passages or alter surface chemistry. Local mass gain at a cold region can therefore coexist with structural metal loss elsewhere. The transport path links components that would appear independent in isolated tests. Mass change must be interpreted together with chemical analysis and spatial examination.

Sodium systems require attention to dissolution, mass transfer and oxygen or carbon impurities. Structural alloys can show useful compatibility under controlled conditions, but a statement of compatibility must include the alloy state and coolant chemistry. Lead and lead–bismuth eutectic (LBE) systems have an additional oxygen-control problem. Too little oxygen may leave an alloy

vulnerable to dissolution; a suitable range can support a protective oxide; too much can cause coolant-oxide precipitation. The boundaries depend on temperature and material chemistry. A qualitative oxygen-window drawing cannot supply numerical operating limits.

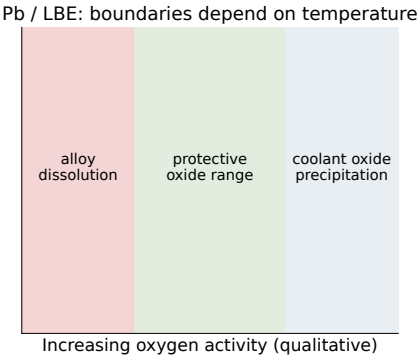


Figure 7.9: Pb/LBE oxygen control must avoid both inadequate alloy protection and coolant-oxide precipitation. The boundaries depend on the selected system.

The window expresses competing requirements. Increasing oxygen activity may help sustain an alloy oxide, but can eventually favor precipitation of coolant oxide. Because hot and cold regions have different equilibria, a single oxygen reading must be interpreted against the loop's temperature distribution. A useful control range must satisfy the relevant conditions throughout the system. Even an oxide that is chemically stable must adhere and withstand the local flow and thermal history. Oxygen control also does not resolve every liquid-metal embrittlement concern: wetting, local stress and the selected alloy–coolant pair affect susceptibility. Tests must examine whether a protective condition persists throughout the loop rather than merely whether an oxide can form at one temperature (OECD NEA, 2015).

Halide molten salts can selectively dissolve less noble alloy constituents, including chromium. Salt composition, redox state and impurities affect reaction driving force. Moisture can introduce additional reactive species. Container materials, thermal gradients and flow can change the transport and deposition of dissolved products. A chromium-bearing alloy's resistance in an aqueous passive state is therefore insufficient evidence for its performance in a salt. The relevant question is whether the alloy, reaction products and coolant remain compatible under the actual chemical conditions.

Helium itself is chemically inert, but trace water and carbon-containing gases can react with hot alloys. Their ratios influence oxygen and carbon chemical potentials, allowing oxidation, carburization or decarburization under different conditions. A protective scale also draws a continuing supply of its forming element from the metal. Depletion can limit repair, while thermal cycling can create mismatch stresses and spallation. Chemical stability and mechanical integrity must persist together. Supercritical water presents another combination of oxidation, cracking and changing transport properties rather than a simple extension of ordinary water behavior.

A rate constant is sometimes represented over one defensible mechanism range by

$$k = A \exp\left(-\frac{Q_a}{RT}\right), \qquad \ln\frac{k_2}{k_1} = \frac{Q_a}{R}\left(\frac{1}{T_1} - \frac{1}{T_2}\right).$$

Here T is in kelvin, Q_a is activation energy in J/mol, $R = 8.314$ J mol^{-1} K^{-1} and A has the

units of k. A higher temperature gives a larger k for positive Q_a within this model. Changes in phase, film integrity or rate-controlling mechanism can invalidate extrapolation. For parabolic growth, $x \propto \sqrt{k}$ at fixed time; doubling a rate constant does not double thickness.

Test selection follows the damage question. Mass loss or thickness tests need controlled chemistry, temperature and product removal. Initiation tests need the appropriate surface and stress history. Precracked tests measure propagation under specified geometry and loading. Flowing-loop tests assess chemistry transport and thermal gradients. A useful diagnosis combines morphology, composition and structure, with exposed and unexposed material states compared under controlled conditions. One observation, such as mass gain or an intergranular path, rarely establishes the full mechanism.

Example: Temperature sensitivity of a rate constant

Given. Assume a single rate-controlling regime with $Q_a/R = 6000$ K between $T_1 = 600$ K and $T_2 = 750$ K. Compare the rate constants and equal-time parabolic thicknesses.

Model and substitution.

$$\ln \frac{k_2}{k_1} = 6000 \left(\frac{1}{600} - \frac{1}{750} \right) = 2, \qquad \frac{k_2}{k_1} = e^2 = 7.39.$$

For initially bare surfaces following $x^2 = kt$,

$$\frac{x_2}{x_1} = \sqrt{\frac{k_2 t}{k_1 t}} = \sqrt{7.39} = 2.72.$$

Meaning. The rate constant increases more strongly than thickness at a fixed time. Reaching the same selected thickness at 750 K would take $1/7.39 = 0.135$ times the exposure at 600 K. This comparison requires the same kinetic law and intact film; a temperature-driven change of mechanism needs a different model.

References. G. S. Was, *Fundamentals of Radiation Materials Science* (2007), Ch. 15; K. L. Murty and I. Charit, *An Introduction to Nuclear Materials* (2013), Secs. 5.3 and 6.4. IAEA, *Stress Corrosion Cracking in Light Water Reactors: Good Practices and Lessons Learned* (2011); IAEA-TECDOC-1649, *Delayed Hydride Cracking of Zirconium Alloy Fuel Cladding* (2010). G. S. Was, A. Singlawi, S. Swaminathan and K. Sun, *The Mechanism of Irradiation Assisted Stress Corrosion Cracks in Stainless Steels*, DOE LWRS (2023). NRC Bulletin 2002-01 (2002). OECD NEA, *Handbook on Lead-bismuth Eutectic Alloy and Lead* (2015).

Unit 7 questions

Conceptual questions

7.C1 Why can metal dissolve when the measured external current is zero?

7.C2 Why can a stable oxide fail to protect the underlying metal?

7.C3 Why is dissolved oxygen alone insufficient to predict ECP in irradiated water?

7.C4 What three contributions are required for SCC, and how can residual stress supply one?

7.C5 Why does aqueous passivity not establish molten-salt compatibility?

Numerical problems

7.N1 A cathode is 8 cm^2 and an anode is 2 cm^2. Find j_a if $|j_c| = 3$ $\mu\text{A/cm}^2$.

7.N2 Calculate K_I for $Y = 1$, $\sigma = 100$ MPa and $a = 1$ mm.

7.N3 For an initially bare surface, $k_p = 0.09$ $\mu\text{m}^2/\text{h}$. Find oxide thickness at 100 h.

Combined problems

7.P1 A constant 0.020 A dissolves Fe as Fe^{2+} for 965 s. Find charge, mass loss and uniform thickness loss over 10 cm^2. Use $F = 96485$ C/mol, $M = 55.85$ g/mol and $\rho = 7.87$ g/cm^3.

7.P2 A specimen gains 16 mg oxygen, solely as retained ZrO_2. Use $M_{\text{O}} = 16.00$ and $M_{\text{Zr}} = 91.22$ g/mol to find reacted mmol Zr and consumed Zr mass. If $f_H = 0.125$, find absorbed mmol H atoms.