

Photon–Atom Coupling: Problem Class

Eight parts, one per chapter

Photonics Engineering

Each part matches one chapter of the lecture notes. The exercises are built around technologies you can find today in laboratories, data centres, hospitals or your pocket; every one of them reduces to the physics of the course. Constants are in Appendix B of the notes. Sketch every curve you derive: axes, limits, extrema, widths.

Constants used throughout

$\hbar = 1.055e-34$ J s, $h = 6.626e-34$ J s, $c = 3.00e8$ m/s, $\varepsilon_0 = 8.85e-12$ F/m, $e = 1.602e-19$ C, $m_e = 9.11e-31$ kg, $k_B = 1.381e-23$ J/K, $a_0 = 5.29e-11$ m, $ea_0 = 8.48e-30$ C m, $1 \text{ u} = 1.66e-27$ kg.

Part 1 — The classical oscillator (Chapter 1)

Exercise 1.1: Rayleigh scattering and the choice of the telecom wavelength. The intrinsic loss floor of a silica fibre is Rayleigh scattering off frozen-in density fluctuations, tiny Lorentz oscillators driven far below their UV resonances.

1. Using the Lorentz model in the regime $\omega \ll \omega_0$, show that the scattered power per oscillator scales as ω^4 , i.e. as $1/\lambda^4$. State the two physical ingredients.
2. The Rayleigh contribution at $\lambda = 1550$ nm is $\alpha_R \approx 0.15$ dB/km. Derive the general scaling law $\alpha_R(\lambda) = \alpha_R(1550) \times (1550/\lambda)^4$ and evaluate it at 1310 nm and 850 nm (the two historical telecom windows).
3. Sketch $\alpha_R(\lambda)$ on log–log axes (slope?). Silica also has an infrared absorption wing rising steeply beyond $1.6 \mu\text{m}$ (Si–O phonons, Chapter 8). Explain graphically why the global minimum sits near $1.55 \mu\text{m}$, and why we do not simply go to $3 \mu\text{m}$.

Exercise 1.2: amplitude and phase of the driven oscillator (THz molecular sensing). THz spectrometers used for security imaging and pharmaceutical quality control read both the *amplitude* and the *phase* of the field transmitted through a sample; both are dictated by the Lorentz response.

1. From the steady-state solution of the Lorentz equation, show that the displacement amplitude and its phase lag relative to the drive are

$$|r_0|(\omega) = \frac{eE_0/m_e}{\sqrt{(\omega_0^2 - \omega^2)^2 + \gamma_c^2\omega^2}}, \quad \tan \varphi = \frac{\gamma_c\omega}{\omega_0^2 - \omega^2}.$$

2. Show that $|r_0|$ peaks at $\omega_{\text{res}} = \sqrt{\omega_0^2 - \gamma_c^2/2}$, and that for $\gamma_c \ll \omega_0$ the full width at half maximum of $|r_0|^2$ is $\Delta\omega \simeq \gamma_c$. Sketch $|r_0|(\omega)$ and $\varphi(\omega)$; give φ at $\omega \ll \omega_0$, $\omega = \omega_0$, $\omega \gg \omega_0$, and show that the phase slope at resonance is $d\varphi/d\omega|_{\omega_0} = 2/\gamma_c$.
3. A crystalline mode of a common drug molecule sits at $\omega_0/2\pi = 1.0$ THz with $\gamma_c/2\pi = 50$ GHz. Give the quality factor $Q = \omega_0/\gamma_c$ and the phase slope in rad/THz; why does measuring the phase help identify thin samples?

Exercise 1.3: the quantum origin of the polarizability. The Lorentz model postulates a spring; quantum mechanics must *derive* the polarizability. Place an atom in a static field $E\hat{z}$, so that $H' = -\hat{d}_z E$ with $\hat{d}_z = -e\hat{z}$.

1. Show that for an eigenstate $|g\rangle$ of definite parity the first-order shift vanishes: $\langle g|\hat{d}_z|g\rangle = 0$. (Use the parity operator Π : $\Pi|g\rangle = \pm|g\rangle$ and $\Pi^\dagger\hat{z}\Pi = -\hat{z}$.)
2. Using second-order stationary perturbation theory (Appendix D of the notes), show that the ground state shifts by $\Delta E_g = -\frac{1}{2}\alpha_0 E^2$ with the static polarizability

$$\alpha_0 = 2 \sum_e \frac{|\langle e|\hat{d}_z|g\rangle|^2}{E_e - E_g}.$$

3. Truncate to two levels and show $\alpha_0 = 2d^2/\hbar\omega_0$. Evaluate for sodium ($d = 2.5ea_0$, $\omega_0 = 3.2e15$ rad/s), express the result as a polarizability volume $\alpha_0/4\pi\epsilon_0$ in $\text{\AA}^3$, and compare with the measured $24 \text{\AA}^3$.
4. Compare with the classical static Lorentz value $\alpha_{cl} = e^2/m_e\omega_0^2$ and show that the ratio is exactly the oscillator strength $f_{ge} = 2m_e\omega_0 d^2/\hbar e^2$. What does this exercise change conceptually about the optical tweezer of Exercise 2.2?

Part 2 — Coherent dynamics (Chapter 2)

Exercise 2.1: from the unitary transformation to the quantum gate (and its crosstalk).

In a superconducting quantum processor each qubit is driven by microwave Rabi pulses; neighbouring qubits, detuned in frequency, must *not* respond. We first *derive* the tools, then use them.

1. A state transformed as $|\tilde{\psi}\rangle = U|\psi\rangle$ with a time-dependent unitary $U(t)$ obeys a Schrödinger equation with the effective Hamiltonian $\tilde{H} = UHU^\dagger + i\hbar\dot{U}U^\dagger$: prove this from $i\hbar\partial_t|\psi\rangle = H|\psi\rangle$. Then take $U = \exp(i\omega t|e\rangle\langle e|)$ and $H = \hbar\omega_0|e\rangle\langle e| - \hbar\Omega\cos\omega t(\sigma + \sigma^\dagger)$, apply the RWA, and recover $\tilde{H} = -\hbar\delta|e\rangle\langle e| + \frac{\hbar\Omega}{2}(\sigma + \sigma^\dagger)$.
2. On resonance ($\delta = 0$), $\tilde{H} = \frac{\hbar\Omega}{2}\sigma_x$ with $\sigma_x = \sigma + \sigma^\dagger$. Using $\sigma_x^2 = \mathbb{I}$, sum the exponential series to prove

$$U(t) = e^{-i\Omega t\sigma_x/2} = \cos\left(\frac{\Omega t}{2}\right)\mathbb{I} - i\sin\left(\frac{\Omega t}{2}\right)\sigma_x,$$

and write the 2×2 matrices of the π - and $\pi/2$ -pulses. Which logic gates are these?

3. A qubit is driven on resonance with $\Omega/2\pi = 50$ MHz. How long is a π -pulse? With a coherence time $T_2 = 100 \mu\text{s}$, roughly how many gates fit within the coherence window?
4. A neighbouring qubit sees the same pulse but is detuned by δ . From the generalized Rabi formula, derive the *maximum* excitation probability $P_{\max}(\delta)$, identify the shape of this curve and its half width. Evaluate it for $\delta/2\pi = 300$ MHz: acceptable for fault tolerance (threshold $\sim 10^{-3}$ – 10^{-2})? Name two remedies used in real devices.

Exercise 2.2: designing an optical-tweezer array for a neutral-atom quantum computer. Neutral-atom machines (hundreds of qubits today) hold one ^{87}Rb atom per optical tweezer. Data: trapping laser $\lambda = 850$ nm; Rb resonance $\lambda_0 = 780$ nm, $\Gamma/2\pi = 6.07$ MHz; waist per tweezer $w_0 = 1 \mu\text{m}$; required depth $U_0/k_B = 1$ mK; mass 87 u.

1. Compute the detuning $\delta/2\pi$ and justify treating the trap with the far-detuned dipole potential of Chapter 2.
2. From $U_{\text{dip}} = \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\delta} I$, compute the peak intensity, then the power per tweezer ($I = 2P/\pi w_0^2$).
3. With 5 W of useful laser power, how many qubits can the array hold? Compare with current experiments.
4. Expand the Gaussian potential near its bottom and derive the radial trap frequency $\omega_r = \frac{2}{w_0} \sqrt{\frac{U_0}{m}}$; evaluate it. Why must $\omega_r \gg$ (photon-recoil heating rate) for a good qubit?

Exercise 2.3: Ramsey fringes and the stability of an optical clock.

1. For the Ramsey sequence (two $\pi/2$ -pulses separated by T), $P_e = \frac{1}{2}[1 + \cos \delta T]$. Derive the sensitivity $|dP_e/d\nu|$ ($\delta = 2\pi\nu$ -detuning) and show it is maximal on the fringe side ($\delta T = \pi/2$), where it equals πT .
2. A projective measurement of one atom has binomial uncertainty $\sigma_P = \sqrt{P(1-P)} = \frac{1}{2}$ on the side of the fringe. Show that the one-shot, one-atom frequency uncertainty is $\sigma_\nu = 1/(2\pi T)$, and that N atoms and M repeated cycles improve it to $\sigma_\nu = 1/(2\pi T\sqrt{NM})$.
3. Evaluate for a strontium lattice clock: $T = 1$ s, $N = 1000$ atoms, $\nu_0 = 429$ THz, one day of averaging ($M \approx 86\,400$). Compare the fractional stability with the state of the art ($\sim 10^{-18}$) and comment.

Part 3 — Dissipation and the Bloch equations (Chapter 3)

Exercise 3.1: optimising the repetition time of an MRI scan. In a T_1 -weighted MRI sequence, each voxel's longitudinal magnetisation recovers as $S(T_R) = S_0(1 - e^{-T_R/T_1})$ between excitations, where T_R is the repetition time chosen by the radiographer. At 1.5 T, healthy white matter has $T_1^{(a)} = 0.8$ s and a tumour $T_1^{(b)} = 1.4$ s.

1. Explain physically (Bloch equation for w) where $S(T_R)$ comes from, and why a tissue with longer T_1 appears *darker*.
2. The image contrast is $C(T_R) = S_a(T_R) - S_b(T_R)$. Show that C vanishes at both $T_R \rightarrow 0$ and $T_R \rightarrow \infty$, and derive the optimum

$$T_R^* = \frac{\ln(T_1^{(b)}/T_1^{(a)})}{1/T_1^{(a)} - 1/T_1^{(b)}}.$$

3. Evaluate T_R^* and $C(T_R^*)$ (in units of S_0). Sketch S_a , S_b and C versus T_R . Why does a shorter T_R also mean a faster scan, and what trade-off does the radiographer face?

Exercise 3.2: power broadening and the laser lock of a wavemeter. Compact laser systems (LiDAR, wavemeters, cold-atom gravimeters) stabilise their frequency on the saturated absorption of a Rb vapour cell, using the D₂ line ($\lambda = 780$ nm, $\Gamma/2\pi = 6.07$ MHz, $\tau = 26.2$ ns).

1. Compute $I_{\text{sat}} = \pi\hbar c/3\lambda^3\tau$ for this line.
2. From the steady-state Bloch solution, recall why the observed linewidth is $\Delta\omega = \Gamma\sqrt{1+s_0}$ (for $T_2 = 2T_1$) and evaluate the intensity at which the line is *twice* its natural width. Sketch $\Delta\nu(s_0)$, with its two asymptotes.
3. The error signal of the lock is the slope of the line centre. Explain why cranking up the probe power degrades the lock twice over, and give the good operating regime.

Exercise 3.3: T_1 and T_2 from the master equation (qubit datasheets). The Lindblad master equation for a two-level system decaying at rate Γ is

$$\dot{\rho} = -\frac{i}{\hbar}[\tilde{H}, \rho] + \Gamma\left(\sigma\rho\sigma^\dagger - \frac{1}{2}\{\sigma^\dagger\sigma, \rho\}\right), \quad \sigma = |g\rangle\langle e|.$$

1. With $\rho = \begin{pmatrix} \rho_{ee} & \rho_{eg} \\ \rho_{ge} & \rho_{gg} \end{pmatrix}$ and no drive ($\tilde{H} = 0$), compute the three operator products $\sigma\rho\sigma^\dagger$, $\sigma^\dagger\sigma\rho$, $\rho\sigma^\dagger\sigma$ as explicit 2×2 matrices, and deduce $\dot{\rho}_{ee}$, $\dot{\rho}_{gg}$ and $\dot{\rho}_{eg}$.
2. Identify T_1 and T_2 , and prove the fundamental inequality $T_2 \leq 2T_1$ for purely radiative decay.
3. Add a pure-dephasing dissipator $\frac{\Gamma_\varphi}{2}(\sigma_z\rho\sigma_z - \rho)$ and show that it leaves the populations untouched while adding $-\Gamma_\varphi\rho_{eg}$ to the coherence: $1/T_2 = 1/2T_1 + \Gamma_\varphi$.
4. A transmon datasheet reads $T_1 = 100$ μ s, $T_2 = 120$ μ s. Check consistency with question 2 and extract the pure-dephasing time $T_\varphi = 1/\Gamma_\varphi$.

Part 4 — Fermi’s golden rule and Einstein coefficients (Chapter 4)

Exercise 4.1: photon occupancy, or why lasers exist. The ratio of stimulated to spontaneous emission into a given mode equals the mode’s photon occupancy $\bar{n}$.

1. Prove first, from the commutator $[a, a^\dagger] = 1$ alone, that $a^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle$ (compute the norm of $a^\dagger |n\rangle$). Deduce that the emission matrix element into a mode holding n photons scales as $\sqrt{n+1}$, the rate as $n+1$, and identify the stimulated and spontaneous parts.
2. For thermal light, derive $\bar{n}(\omega, T) = 1/(e^{\hbar\omega/k_B T} - 1)$ (partition function of one mode). Sketch $\bar{n}$ versus $x = \hbar\omega/k_B T$ with its two asymptotes.
3. Evaluate $\bar{n}$ for sunlight ($T = 5778$ K) at $\lambda = 500$ nm, and for a room-temperature (300 K) mode at $\lambda = 1.55$ μm . Conclude: why can no thermal source lase, and why is the noise added by an erbium fibre amplifier quantum (spontaneous) rather than thermal?

Exercise 4.2: forbidden lines as a technology: the strontium clock transition. ^{87}Sr has an allowed transition $^1S_0 \rightarrow ^1P_1$ at 461 nm ($\tau = 5$ ns) used for laser cooling, and a doubly-forbidden transition $^1S_0 \rightarrow ^3P_0$ at 698 nm ($\tau \approx 118$ s) used as the clock.

1. Compute the natural linewidths $\Delta\nu = 1/2\pi\tau$ of both transitions and the quality factor $Q = \nu_0/\Delta\nu$ of the clock line.
2. Using $\Gamma \propto \omega^3 d^2$, estimate the ratio $d_{\text{clock}}/d_{\text{cool}}$ of the dipole moments.
3. Explain the design logic: why does a good clock *need* a nearly forbidden line, what is the price to pay (think Rabi frequency), and how does the aurora of Chapter 4 illustrate the same physics?

Exercise 4.3: can a single atom block a single photon? Quantum networks would love a node where one photon deterministically talks to one atom.

1. A resonant photon is focused to a Gaussian spot of waist w . Taking the photon’s effective area as $\pi w^2/2$ argue that the absorption probability is $p \approx \sigma_0/(\pi w^2/2)$ with $\sigma_0 = 3\lambda^2/2\pi$.
2. Evaluate p for $\lambda = 780$ nm and a strong objective, $w = \lambda$. Can free-space coupling be made deterministic?
3. What are the two standard ways to push $p \rightarrow 1$? (One of them is Chapter 7.)

Part 5 — Field quantization and spontaneous emission (Chapter 5)

Exercise 5.1: the efficiency and the colour problem of OLED pixels. An OLED pixel emits by spontaneous emission of molecular excitons. The measured decay of the excited population is e^{-t/τ_m} with $1/\tau_m = \Gamma_r + \Gamma_{nr}$ (radiative + non-radiative).

1. Define and derive the internal quantum efficiency $\eta = \Gamma_r/(\Gamma_r + \Gamma_{nr}) = \tau_m/\tau_r$. A green emitter has $\tau_r = 10$ ns and $\tau_m = 6$ ns: give η .
2. Assuming equal dipole moments, use $\Gamma_r \propto \omega^3$ to compare the radiative rates of a blue (460 nm) and a red (620 nm) emitter. Sketch $\Gamma_r(\lambda)$.
3. If Γ_{nr} (set by molecular chemistry) is the same for both, which pixel has the better efficiency? Real displays show the opposite: blue pixels age fastest. What does that tell you about Γ_{nr} (or the photochemistry) at high photon energy?

Exercise 5.2: how forbidden is erbium? (the transition that runs the internet) The $^4I_{13/2} \rightarrow ^4I_{15/2}$ transition of Er^{3+} in silica radiates at $\lambda = 1536$ nm with $\tau \approx 10$ ms.

1. Invert the spontaneous-emission formula to obtain the dipole moment d ; express it in units of ea_0 and compare with an allowed transition ($\sim 2.5 ea_0$ for Na).
2. Why is such a small dipole a *feature* for an amplifier? (Think population inversion, pump efficiency, and what a large Γ would do.)

3. The natural linewidth is $1/2\pi\tau$. Compare it with the actual gain bandwidth of an EDFA (~ 4 THz) and explain the discrepancy in the vocabulary of Chapter 3.

Part 6 — Maxwell–Bloch and propagation (Chapter 6)

Exercise 6.1: open-path methane sensing by laser absorption (TDLAS). Drones and fixed stations monitor methane leaks by sending a $1.65\text{ }\mu\text{m}$ diode laser across open air and reading the absorption (TDLAS: tunable diode laser absorption spectroscopy). Data: peak absorption cross section of the targeted CH_4 line $\sigma_0 \approx 2e-24\text{ m}^2$; path $L = 100\text{ m}$; the photodetector resolves relative intensity changes down to $\Delta I/I = 10^{-4}$; air density $n_{\text{air}} = 2.5e25\text{ m}^{-3}$.

1. Derive Beer’s law $I(L) = I_0 e^{-N\sigma L}$ from the loss of a slab dz , and linearise it for small absorption.
2. Compute the minimum detectable methane density and convert it to a mixing ratio (ppb). Compare with the atmospheric background (≈ 1.9 ppm) and with a leak.
3. The laser is swept across the line, whose shape is Lorentzian, $\sigma(\delta) = \sigma_0/(1 + 4\delta^2/\Gamma_{\text{tot}}^2)$. Sketch the transmitted signal versus laser frequency, and explain why sweeping (rather than sitting at the peak) makes the measurement immune to dust and power drifts.

Exercise 6.2: a slow-light delay line for a quantum repeater. In a $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ crystal a narrow spectral window (width $\Delta\nu = 100\text{ kHz}$) is prepared, inside which the group index reaches $n_g \approx 10^7$ over a length $L = 5\text{ mm}$.

1. Recall (three lines) why $v_g = c/n_g$ with $n_g = n + \omega dn/d\omega$, and compute the traversal (delay) time of the crystal.
2. What pulse duration τ_p can be delayed without distortion? (Its spectrum must fit inside $\Delta\nu$.) Deduce the delay–bandwidth product $\tau_{\text{delay}}\Delta\nu$: how many pulse lengths are stored?
3. Why is this the natural front-end of a quantum memory, and what does Chapter 3 say limits the ultimate storage time?

Part 7 — Cavity QED and the Purcell effect (Chapter 7)

Exercise 7.1: a telecom single-photon source for quantum key distribution. Quantum communication wants on-demand single photons at $\lambda = 1.55\text{ }\mu\text{m}$. Design brief: an InAs quantum dot ($\tau_0 = 1.5\text{ ns}$ in bulk) in an InP photonic cavity ($n = 3.17$), Purcell factor $F_P = 30$, mode volume $V = 2(\lambda/n)^3$.

1. What quality factor Q is required? Give the cavity linewidth $\kappa/2\pi = \nu_c/Q$.
2. Compute the Purcell-shortened lifetime and the fraction $\beta = F_P/(F_P + 1)$ of photons emitted into the cavity mode. What clock rate (photons per second) can the source sustain if one waits ~ 10 lifetimes between triggers?
3. The enhancement requires spectral matching. Show, from the Lorentzian mode density of Chapter 7, that

$$F_P(\delta_c) = \frac{F_P(0)}{1 + (2\delta_c/\kappa)^2},$$

sketch this curve, and give the detuning at which the enhancement is halved. Why does a larger Q make the source better but harder to fabricate?

Exercise 7.2: dispersive readout of a superconducting qubit. A transmon (ω_0) is coupled to a readout resonator (ω_c) with $g/2\pi = 100\text{ MHz}$, $\kappa/2\pi = 1\text{ MHz}$, $\gamma/2\pi = 10\text{ kHz}$, and operated far detuned: $\Delta_{ac}/2\pi = 1\text{ GHz}$.

1. Show by explicit commutator algebra that the excitation number $\hat{N} = \sigma^\dagger\sigma + a^\dagger a$ is conserved by the Jaynes–Cummings Hamiltonian: $[\hat{N}, H_{\text{JC}}] = 0$. (You will need $[\sigma^\dagger\sigma, \sigma^\dagger] = \sigma^\dagger$ and

- $[a^\dagger a, a] = -a$; prove both.) Then diagonalize the resonant n -excitation block $\{|e, n\rangle, |g, n+1\rangle\}$ and show the splitting is $2\hbar g\sqrt{n+1}$.
2. Compute the cooperativity $C = g^2/\kappa\gamma$ and the vacuum-Rabi period $2\pi/2g$ at resonance. In which regime is the system at $\Delta_{ac} = 0$?
 3. In the dispersive regime, the adiabatic-elimination result of Chapter 7 shows the cavity acquires a qubit-state-dependent shift $\chi = g^2/\Delta_{ac}$. Evaluate $\chi/2\pi$ and compare with $\kappa/2\pi$: can a probe tone distinguish the two qubit states in one cavity linewidth?
 4. Explain in words the full readout chain (probe, phase shift, amplifier), and why the same g must be *large* for readout but the qubit must stay *detuned* to protect it (Purcell decay $\Gamma_{\text{Purcell}} \approx (g/\Delta_{ac})^2\kappa$: evaluate it and compare with γ).

Part 8 — Optical response of solids (Chapter 8)

Exercise 8.1: the colour of a QLED television (2023 Nobel Prize in chemistry). The pixels of a QLED display are CdSe nanocrystals whose emission colour is set purely by their size. Model: electron and hole confined in a sphere of radius R ; infinite walls; lowest level of a particle in a sphere $= \hbar^2\pi^2/2mR^2$. Data: bulk gap $E_g = 1.74$ eV; reduced mass $\mu = 0.10 m_e$.

1. Prove the confinement formula first: solve the Schrödinger equation for a particle of mass m in an infinite spherical well of radius R . For the ground (s -wave) state, write $\psi(r) = u(r)/r$, show that u obeys the free 1D equation $-\frac{\hbar^2}{2m}u'' = Eu$ with $u(0) = 0$ and $u(R) = 0$, and deduce

$$\psi_1(r) \propto \frac{\sin(\pi r/R)}{r}, \quad E_1 = \frac{\hbar^2\pi^2}{2mR^2}.$$

2. Show that the emission energy is then $E(R) = E_g + \frac{\hbar^2\pi^2}{2\mu R^2}$ (neglecting the Coulomb term) and sketch $E(R)$. Which way does the colour shift as the dot shrinks, and why physically?
3. Compute the radii needed for red (2.0 eV), green (2.33 eV) and blue (2.64 eV) pixels. Comment on the fabrication tolerance: what does dE/dR at the green dot imply for the size dispersion allowed if the colour must stay within 0.02 eV?
4. Why is this *the same physics* as the exciton of Chapter 8, and in which limit (compare R with the bulk exciton Bohr radius $a_X \approx 5.6$ nm for CdSe)?

Exercise 8.2: how thick must a solar absorber be? Just above their gaps, typical absorption coefficients are: GaAs (direct) $\alpha \approx 10^4 \text{ cm}^{-1}$; silicon at 800 nm (indirect) $\alpha \approx 10^3 \text{ cm}^{-1}$ and near its gap (~ 1000 nm) $\alpha \approx 10^2 \text{ cm}^{-1}$; halide perovskite (direct) $\alpha \approx 10^5 \text{ cm}^{-1}$.

1. From Beer's law, derive the thickness L_{90} absorbing 90% of the light and compute it for the four cases.
2. Relate the direct/indirect difference to the golden rule of Chapter 8 (what extra ingredient does an indirect transition need, and what does it cost?).
3. Commercial Si cells are $\sim 180 \mu\text{m}$ thick and use textured light-trapping; perovskite films are $\sim 0.5 \mu\text{m}$. Justify both numbers, and give one advantage and one drawback of ultra-thin absorbers.

Exercise 8.3: transparent electrodes and low-emissivity windows (Drude engineering). Indium tin oxide (ITO), the transparent conductor of every touchscreen, is a degenerate semiconductor described by the Drude model with effective mass $m^* = 0.35 m_e$ and background permittivity $\varepsilon_\infty = 4$.

1. With damping neglected, show that $\varepsilon(\omega) = \varepsilon_\infty(1 - \tilde{\omega}_p^2/\omega^2)$ with the *screened* plasma frequency $\tilde{\omega}_p^2 = Ne^2/\varepsilon_0\varepsilon_\infty m^*$, and recall the optical behaviour on either side of $\tilde{\omega}_p$.
2. Specification: the film must stay transparent across the visible but reflect the thermal infrared (low-emissivity window), i.e. $\tilde{\lambda}_p = 1.2 \mu\text{m}$. Compute the required carrier density N and compare with a metal ($\sim 10^{29} \text{ m}^{-3}$).

3. Sketch reflectivity versus wavelength from the UV to $10\text{ }\mu\text{m}$ for this film, marking the visible window and the thermal-IR band ($\sim 10\text{ }\mu\text{m}$). Explain in one paragraph how the same layer serves the touchscreen (electrical) and the window (radiative) markets.

How to use this sheet

Every exercise is solvable with the boxed results of the lecture notes plus the constants of Appendix B. Work the analytics first (curve shapes, limits, extrema), plug numbers last, and always finish with the sanity check: does the order of magnitude match the device that actually exists?