

Circuit Analysis I (DC)

University Electrical Engineering

What this course is, and how to use it

This is the first course in University Electrical Engineering: **direct-current circuit analysis**. A DC circuit is one where currents and voltages settle to steady values – batteries, sensors, logic supplies, and the bias networks inside every amplifier.

The course is a ladder, and every rung is used by the next:

- Units 1.1 and 1.2 need only algebra: quantities, sign conventions, Ohm's law and Kirchhoff's laws.
- Units 1.3 to 1.5 build systematic methods – node and mesh analysis, network theorems, and the ideal op-amp – that turn long circuits into short calculations.
- Units 1.6 and 1.7 add the energy-storing elements and the first **transient response**, which is where calculus begins to matter.

How to read each unit: a short story says why it matters; a picture or a rule of thumb comes before any formula; then the precise statement; then worked examples that get harder, one step at a time, each with a check at the end; then the mistakes people actually make; then a three-line summary. Every number in the examples is something you can verify a second way – a different loop, an unused node, a limiting case – and you should build that habit here: in later courses the check is the difference between a caught error and a burnt resistor.

Quantities, units and measurement

A phone charger labelled “20 W” and a kettle labelled “2000 W” both move charge. Why is one a hundred times more powerful? To answer, you need the six quantities every circuit question is asked in – and one convention that keeps their signs honest.

Learn them as a chain, not a list: charge moves; its rate is current; current through a difference in potential transfers energy; the rate of that transfer is power.

- **Charge** q is measured in **coulombs** (C). One coulomb is about 6.24×10^{18} elementary charges. Charge is **conserved**: it is never created or consumed in a circuit, only moved.
- **Current** i is the rate of charge flow, in **amperes** (A): $i = dq/dt$, so 1 A = 1 C/s.
- **Voltage** v is energy per unit charge, in **volts** (V): 1 V = 1 J/C. It is always measured **between** two points; “5 V at a node” means 5 V between that node and the reference (ground) node.


- **Power** p is energy per unit time, in **watts** (W): $p = vi$, and 1 W = 1 J/s.
- **Energy** w is power integrated over time, $w = \int p dt$, in **joules** (J). Electricity meters sell the **kilowatt-hour**: 1 kWh = 3.6 MJ.

Deriving the power formula. Move a charge dq through a potential difference v and the energy transferred is $dw = v \cdot dq$. Divide by dt : $p = dw/dt = v \cdot (dq/dt) = vi$. Power is voltage times current – nothing deeper.

The **passive sign convention** removes all ambiguity. Pick a direction for the current through an element and a polarity for its voltage; when the current arrow enters the + terminal, $p = vi$ is the power the element **absorbs**. A positive result means the element consumes power; a negative result means it supplies power. Getting this backwards is the most common source of wrong signs in the whole course.

Where each meter goes

12 V battery



ammeter: in series,
nearly a short circuit

voltmeter: across the element,
nearly an open circuit

Worked example 1 (power and energy). A 12 V battery drives a lamp that draws 0.25 A, with the current arrow drawn into the lamp's + terminal. For the lamp, $p = vi = 12 \times 0.25 = 3$ W absorbed. Over 2 minutes the lamp takes $w = pt = 3 \times 120 = 360$ J. With the same arrow drawn for the battery, the battery's current enters its – terminal, so its absorbed power is -3 W: it supplies exactly what the lamp takes, and the bookkeeping closes.

Worked example 2 (the meter reads the wrong number). A 12 V ideal source feeds two 10 MΩ resistors in series, so each should have 6 V across it. A real voltmeter with 10 MΩ input resistance is connected across the lower one to "check". The meter is now a second 10 MΩ resistor in parallel: the lower part combines to 5 MΩ, and the meter reads $\frac{5}{10+5} \times 12 = 4$ V instead of 6 V. The check itself changed the circuit. (Across the ideal source itself the meter would read 12 V correctly: an ideal source holds its voltage whatever is connected.) Rule of thumb: the meter's resistance must be at least a hundred times the resistance it measures. Escalate further: an **ammeter** in current mode is the mirror case – a series meter with non-negligible resistance steals voltage and reads low current.

Common mistakes.

- Reporting “the voltage at point A” without saying what the other point is. Voltage is always between two points; name them.
- Mixing units: minutes with watts, mA with kΩ. Convert first; then the answer’s unit falls out by itself (V/Ω = A).
- Trusting a meter without asking what resistance it adds.

Summary: six quantities, one chain (charge to current to energy to power); power is vi under the passive sign convention; meters must be nearly invisible to the circuit (series ammeter = near short, parallel voltmeter = near open).

Ohm’s law and Kirchhoff’s laws

Two resistors arrive at a node and one leaves. How much current takes each exit? The answer needs the two laws named after Kirchhoff – and they are not new physics. One is “charge does not pile up”; the other is “energy books balance”. Both were hiding in unit 1.1.


Ohm’s law links a **resistor** : $v = Ri$, with R in **ohms** (Ω).
Conductance $G = 1/R$ (siemens) is sometimes handier. Power in a resistor is always absorbed, and combining $v = Ri$ with $p = vi$ gives the three equivalent forms $p = vi = i^2R = v^2/R$ – use whichever variable you already know.

For **series** elements the resistances add and one current is shared; for **parallel** elements the conductances add and one voltage is shared. Two resistors in parallel combine as the product over the sum, $R_1 \parallel R_2 = \frac{R_1R_2}{R_1+R_2}$ – and the result is always smaller than either.

One node, four signed branch currents

KCL at A (signed sum out = 0): $-2 - 0.5 + 1 + 1.5 = 0$

Choose every arrow before computing; a negative answer is information, not a mistake.



Deriving the divider. Two resistors in series across a source carry one current $i = v_s/(R_1 + R_2)$. The voltage on R_2 is that current times R_2 , so $v_2 = v_s \cdot R_2/(R_1 + R_2)$ – the **voltage divider**. The mirror argument on two parallel resistors fed by a current source gives the **current divider** : $i_1 = i_s \cdot R_2/(R_1 + R_2)$. Notice the swap: each branch takes the share proportional to the **other** resistor.

Kirchhoff’s current law (KCL) restates conservation of charge: the signed currents leaving any closed surface sum to zero. A **node** is a junction of two or more elements; every connection of the same wire is one node, however many branches meet there.

Kirchhoff’s voltage law (KVL) restates energy conservation: the signed voltages around any closed **loop** sum to zero. Walk a loop and tally the rises and falls: a label that goes up in one traversal is the same label that goes down when you return, so the total must be zero – you cannot gain energy by walking in a circle.

Worked example 1 (divider, checked two ways). A 9 V source feeds $R_1 = 6$ kΩ in series with $R_2 = 3$ kΩ. The current is $i = 9 \text{ V}/9 \text{ kΩ} = 1 \text{ mA}$, so $v_2 = 1 \text{ mA} \times 3 \text{ kΩ} = 3 \text{ V}$ – the same answer as the divider formula $9 \times 3/9$. Check by KVL: $-9 + 6 + 3 = 0$. Check by power: source $9 \times 1 \text{ mA} = 9 \text{ mW}$; resistors $6 \text{ mW} + 3 \text{ mW} = 9 \text{ mW}$.

Worked example 2 (current divider, then the node it feeds). A 2 A source feeds 6 Ω in parallel with 3 Ω. The parallel pair is $6 \times 3/9 = 2$ Ω, so both see 4 V, and the currents are $4/6 = 0.667 \text{ A}$ and $4/3 = 1.333 \text{ A}$ – the divider formula gives the same shares, and $0.667 + 1.333 = 2 \text{ A}$ at the node, KCL closed. Now escalate: at that node a further branch of unknown current i_4 also leaves; if another 0.5 A flows **in** from elsewhere, KCL reads $-0.667 - 1.333 - 0.5 + i_4 = 0$, so $i_4 = 2.5 \text{ A}$ out. Charge did not vanish; it re-routed.

Common mistakes.

- Drawing KCL with some arrows in and some out, then adding magnitudes. Write "sum of currents leaving = 0" and let the signs speak.
- Treating two separate blobs of wire as one node (they are one node only if connected).
- Parallel resistors adding to something **bigger** – the combination is always smaller than the smallest.

Summary: Ohm's law plus two conservation laws solve every DC resistive circuit; dividers come from them, not from memory; label every arrow first, then trust the algebra – a negative answer is information, not a mistake.


Nodal and mesh analysis

A circuit with four nodes and six branches can be solved by clever guesses – once. The same circuit appears a hundred times in a design, and the guess-work has to become a recipe a computer could run. Two recipes exist; both end in a few simultaneous equations.

Nodal analysis uses KCL. Choose a **reference node** (ground), label the remaining $n - 1$ node voltages, and write one KCL equation per unknown node, with each branch current written as a voltage difference over a resistance. This method is the workhorse of circuit simulators (SPICE solves exactly these equations, millions of them).


Mesh analysis uses KVL. Assign a circulating **mesh current** to each window of a planar circuit and write one KVL equation per mesh; a resistor shared by two meshes carries the **difference** of the two mesh currents. Pick the method with fewer equations, and remember: a voltage source fixes a node-voltage difference (nodal gets easier), while a current source fixes a branch current (mesh gets easier).

Nodal: node voltages



one KCL equation per unknown node

Mesh: loop currents



shared branch carries $i_A - i_B$

Deriving the nodal equation. At a node A with resistors R_1 to a source node V_s and R_2 to ground, KCL (sum of currents leaving) gives $\frac{V_A - V_s}{R_1} + \frac{V_A - 0}{R_2} = 0$. Multiply through and collect: the equation is linear in V_A – always. Every nodal analysis is this one equation repeated per node.

Worked example 1 (one unknown node). A 12 V source connects through $2\ \Omega$ to node A ; from A , a $6\ \Omega$ resistor goes to ground and a 1 A source pushes current into A . KCL at A , currents out:

$$\frac{V_A - 12}{2} + \frac{V_A}{6} - 1 = 0$$

Multiply by 6: $3(V_A - 12) + V_A - 6 = 0$, so $4V_A = 42$ and $V_A = 10.5$ V. Check with the source branch: $(12 - 10.5)/2 = 0.75$ A flows from the source, plus 1 A from the current source; the $6\ \Omega$ carries $0.75 + 1 = 1.75$ A, and $10.5/6 = 1.75$ A – the unused route confirms the answer.

Worked example 2 (two meshes, shared resistor). Mesh A: a 12 V source (rising in the clockwise direction) with a $2\ \Omega$ resistor. Mesh B: a $6\ \Omega$ resistor with a 3 V source (opposing). The meshes share a $4\ \Omega$ resistor. KVL with clockwise mesh currents i_A, i_B :

$$-12 + 2i_A + 4(i_A - i_B) = 0 \quad 4(i_B - i_A) + 6i_B + 3 = 0$$

Collect: $6i_A - 4i_B = 12$ and $-4i_A + 10i_B = -3$. Solving, $i_A = 2.455$ A and $i_B = 0.682$ A, so the shared resistor carries $i_A - i_B = 1.773$ A downward in mesh A's sense. Check by substituting back into either original equation: both return zero to within rounding.

Common mistakes.

- Forgetting the shared resistor's current is a difference of mesh currents.
- Choosing a method without counting equations: current sources usually fit nodal analysis directly, while voltage sources usually fit mesh analysis directly.
- Assuming the reference node must be a source's negative terminal. Any node can be the reference; choosing a voltage-source endpoint may simplify the equations.

Summary: nodal = KCL per unknown node (simulators' choice); mesh = KVL per window; shared elements see current differences; choose by counting equations.

Source transformations and network theorems


You replace a sensor board and the lamp it feeds brightens slightly. What did the board look like **to the lamp**? Not its schematic – one source and one resistance. Three theorems make that replacement exact, and they all say one thing: a linear network is indistinguishable from its behaviour at two terminals.

- **Thévenin's theorem** : any linear two-terminal network of sources and resistors behaves, at its terminals, like one voltage source V_{th} in series with one

resistance R_{th} . V_{th} is the **open-circuit voltage** ; R_{th} is the resistance seen with sources zeroed (voltage sources shorted, current sources opened).


- **Norton's theorem** is the same statement as a current source $I_N = V_{th}/R_{th}$ in parallel with R_{th} . The two are the **source transformation** : swap between them whenever the arithmetic is friendlier.
- **Superposition** : in a linear circuit with several independent sources, the response is the sum of the responses to each source acting alone (others zeroed). It is a computing strategy, not a property of any one component.

Thévenin: series form



measure: V_{th} open-circuit,
 R_{th} with sources zeroed

Norton: parallel form



swap the pair freely:
 $I_N = V_{th}/R_{th}$

Deriving maximum power transfer. A Thévenin source V_{th} , resistance R_{th} , feeds R_L . The load power is $P = \left(\frac{V_{th}}{R_{th} + R_L} \right)^2 R_L$. Differentiate with respect to R_L and set the result to zero: the maximum sits at $R_L = R_{th}$, where $P_{max} = V_{th}^2 / (4R_{th})$ – and the source resistance then burns exactly as much as the load delivers: 50% efficiency, no more.

Maximum power transfer is a *signal* rule, not an *energy* rule: audio and RF systems match impedances to keep signals strong; power grids never do (50% efficiency would double every bill). A Thévenin equivalent does not preserve the internal losses of the original network; its efficiency must be calculated from the actual sources.

Worked example 1 (find the equivalent, then load it). A 10 V source and a 4 Ω resistor form one series branch across a 6 Ω resistor. Find the equivalent seen by an external load connected across that 6 Ω resistor. With the external load disconnected, the terminal voltage is $V_{th} = 10 \times 6/10 = 6$ V (a divider). Zeroing the source leaves 4 Ω parallel with 6 Ω : $R_{th} = 2.4$ Ω . With $R_L = R_{th} = 2.4$ Ω , the loop current is $6/(2.4 + 2.4) = 1.25$ A and the load power is $1.25^2 \times 2.4 = 3.75$ W. Check against the formula: $V_{th}^2 / (4R_{th}) = 36/9.6 = 3.75$ W.

Worked example 2 (superposition, checked head-on). A 12 V source through 2 Ω feeds node A; a 3 A source also injects into A; 4 Ω runs from A to ground. By superposition: the 12 V source alone (3 A opened) drives the 4 Ω through a 2 + 4 divider: $V_{A1} = 12 \times 4/6 = 8$ V. The 3 A source alone (12 V shorted) sees 2 $\Omega \parallel 4 \Omega = 1.333$ Ω : $V_{A2} = 3 \times 1.333 = 4$ V. Total $V_A = 8 + 4 = 12$ V. Direct check by one nodal equation: $(V_A - 12)/2 + V_A/4 - 3 = 0$ gives $V_A = 12$ V – the two methods agree exactly.

Common mistakes.

- Computing R_{th} with independent sources still active. Deactivate them: short ideal independent voltage sources and open ideal independent current sources. Retain dependent sources and use a test source when needed.
- Applying superposition to **power** (it is linear in voltage and current, not in their squares – compute responses first, then square).
- Believing the 50%-efficiency point is a design target for power delivery.


Summary: measure two things (V_{th} open-circuit, R_{th} zeroed) and any linear one-port becomes a battery and a resistor; swap series $\leftrightarrow$ parallel forms freely; add responses, never powers; match resistance for maximum signal power, not for economy.

Ideal operational amplifiers

A microphone produces 10 mV. A speaker wants 2 V. Something must multiply by 200 – with a dial to set exactly how much. That something is the **operational amplifier**, and the whole unit is about one trick: feedback makes an absurdly high-gain device behave like an exact, mild-gain one.

An **operational amplifier** is a differential voltage amplifier with very large gain, drawn as a triangle with + and – inputs and one output. Real op-amps have a **gain-bandwidth product** and finite input currents; the **ideal op-amp** model used in analysis assumes infinite open-loop gain, infinite input resistance and zero output resistance.

The symbol: five terminals **The inverting amplifier**



Why the golden rules follow from feedback. With **negative feedback**

(a path from output back to the $-$ input), the output does whatever is needed to make the input difference nearly zero – with gain above 10^5 , an output of volts needs an input difference of microvolts, which rounds to zero. Two **golden rules** hold:

1. No current flows into either input (infinite input resistance).
2. The voltage difference between the inputs is zero – a **virtual short**. The inputs track each other without being connected.

Deriving the inverting gain. The workhorse is the **inverting amplifier**

: input through R_{in} to the $-$ node, feedback through R_f , $+$ input grounded. Rule 2 pins the $-$ node at 0 V (a **virtual ground**). Rule 1 says the input current v_{in}/R_{in} cannot enter the chip, so it all flows through R_f :

$$\frac{v_{in}}{R_{in}} = -\frac{v_{out}}{R_f} \Rightarrow \frac{v_{out}}{v_{in}} = -\frac{R_f}{R_{in}}$$

The gain is set by two resistors – the chip's enormous gain cancels out of the answer. The **non-inverting amplifier** puts the signal at $+$ and a divider at $-$: $v_{out}/v_{in} = 1 + R_f/R_{in}$, always at least unity (at unity it is the **voltage follower**, a buffer).

Worked example 1 (inverting stage, fully accounted). $R_{in} = 10 \text{ k}\Omega$, $R_f = 100 \text{ k}\Omega$, $v_{in} = +0.2 \text{ V}$. Gain is $-100/10 = -10$, so $v_{out} = -2 \text{ V}$, comfortably inside $\pm 15 \text{ V}$ supplies. The feedback current is $0.2/10\text{k} = 20 \mu\text{A}$ and flows into the output; the ideal input carries no current. Check by the virtual ground: the $-$ node sits at 0 V, so the drop on R_f is $0 - (-2) = 2 \text{ V} = 20 \mu\text{A} \times 100 \text{ k}\Omega$.

Worked example 2 (a mixer, and a saturation trap). Extend the inverting stage to two inputs ($R_1 = R_2 = 10 \text{ k}\Omega$, same $R_f = 100 \text{ k}\Omega$): the virtual ground collects both currents, so $v_{out} = -\left(\frac{R_f}{R_1}v_1 + \frac{R_f}{R_2}v_2\right)$ – a **summing amplifier**

. With $v_1 = 0.2 \text{ V}$ and $v_2 = -0.05 \text{ V}$: $v_{out} = -(10 \times 0.2 + 10 \times (-0.05)) = -(2 - 0.5) = -1.5 \text{ V}$. Now the trap: swap R_f for $1 \text{ M}\Omega$ (gain -100 per input) and the same inputs ask for $-(100 \times 0.2 + 100 \times (-0.05)) = -15 \text{ V}$. The **saturation**

rule stops it: the output cannot leave its supply rails, and a $\pm 15 \text{ V}$ op-amp rests near -13.5 V . The golden rules have silently stopped applying – report the clamp, not the impossible voltage.

Common mistakes.

- Forgetting the inverting stage's minus sign (or “correcting” it away).
- Applying rule 2 when there is no negative feedback path – no feedback, no virtual short.
- Pushing the answer past the rails and reporting it anyway.

Summary: infinite gain tamed by feedback gives two golden rules; the inverting stage turns them into $-R_f/R_{in}$, the non-inverting into $1 + R_f/R_{in}$; virtual ground makes mixers; saturation is a real output state, not a rounding error.

Capacitors, inductors and energy

Every flash in a phone camera collects energy slowly and releases it in a millisecond. No arrangement of resistors can hold anything – resistors burn energy; the two new elements **store** it, and that single change is what makes circuits dynamic and worth a calculus course.

A **capacitor** stores charge separated by an insulator: two plates with electric field between them. Its **capacitance** C , in **farads** (F), is charge per volt, $q = Cv$. Differentiate: current is the rate of change of voltage,

$$i = C \frac{dv}{dt}$$


so a capacitor carries current only while its voltage changes: steady DC is an open circuit to it. Integrating the same law gives the stored energy $w = \frac{1}{2}Cv^2$ – a state memory that depends only on the present voltage, not on how it got there.

An **inductor** stores energy in a magnetic field around a coil. Its **inductance** L , in **henrys** (H), gives the mirror law,

$$v = L \frac{di}{dt}$$

so an inductor develops voltage only while its current changes: steady DC is a short circuit to it. Stored energy is $w = \frac{1}{2}Li^2$.

Burn it, or store it in a field



Why the combinations mirror the resistor rules. Parallel capacitors share one voltage, so charges add: $C_{eq} = C_1 + C_2$ (areas add). Series inductors share one current, so flux linkages add: $L_{eq} = L_1 + L_2$ (lengths add). Series capacitors and parallel inductors combine by the reciprocal rule – the exact mirror of resistor arithmetic, because the storage laws swap the roles of v and i .

Worked example 1 (two energies, two units). A 100 μF capacitor charged to 12 V holds $w = \frac{1}{2} \times 100 \times 10^{-6} \times 12^2 = 7.2 \text{ mJ}$. A 10 mH inductor carrying 2 A holds $w = \frac{1}{2} \times 0.01 \times 2^2 = 20 \text{ mJ}$. Both are small by kettle standards, but the capacitor's 7.2 mJ delivered in 1 ms is 7.2 W – and in 1 μs it would be 7.2 kW. Stored energy is modest; **rate** makes it dramatic, which is why large capacitors are discharged before you touch them.

Worked example 2 (slope law and combinations). A 10 μF capacitor's voltage rises at 100 V/s; the current is $i = C dv/dt = 10 \times 10^{-6} \times 100 = 1 \text{ mA}$ – constant, because the slope is constant. Now combine: a 2 μF and a 3 μF capacitor in series give $2 \times 3/5 = 1.2 \mu\text{F}$ (each holds the same charge, so the "capacity per volt" shrinks); in parallel they give 5 μF . A 10 mH and 20 mH inductor in series give 30 mH; the same two in parallel give $10 \times 20/30 = 6.67 \text{ mH}$.

△ An inductor's current is continuous: opening a switch in an inductor's path forces di/dt to be huge, and the induced voltage $L di/dt$ can arc across the opening contacts. This is real, not theoretical – it is how old ignition coils made their spark.

Common mistakes.

- Using $w = \frac{1}{2} C v^2$ with volts swapped for something else – check the units: farads times volts squared is joules.
- Adding parallel capacitances as reciprocals (that is the series rule).
- Believing a capacitor "blocks" DC instantly – it blocks **steady** DC; during changes it carries current gladly.

Summary: capacitors remember voltage, inductors remember current; $i = C dv/dt$ and $v = L di/dt$; energies $\frac{1}{2} C v^2$, $\frac{1}{2} L i^2$; combinations mirror resistor arithmetic with v and i swapped; never open an inductor's path suddenly.

First-order RC and RL response

You press a button and the light fades up over a second – or the sensor reading creeps to its new value over milliseconds. Switch a source, an energy-storing element and a resistor into a new arrangement, and the stored quantity cannot jump: it slides to its new steady value. A circuit with one storage element is **first order**, and every such circuit has the same shape of answer – learn the shape once.


Define the **natural response** as the decay from an initial state with the source removed, and the **step response** as the move from an initial state to a new forced value V_F . For a capacitor voltage,

$$v(t) = V_F + (V_0 - V_F) e^{-t/\tau}$$

Deriving it. After the switch, a capacitor C discharges through R : KCL gives $C dv/dt + v/R = 0$. Separate variables: $dv/v = -dt/(RC)$, integrate: $\ln v = -t/(RC) + \text{const}$, so $v = V_0 e^{-t/RC}$ – an exponential with **time constant**

$\tau = RC$. Add a source and the same equation gains a constant solution V_F ; the general answer is the constant plus a decaying difference: exponential to the new value, always. The RL mirror replaces $\tau = RC$ with $\tau = L/R$.

At $t = \tau$ the gap has closed to 37% of its starting size; after 5τ the circuit is within 1% and is treated as settled. The initial value V_0 comes from the stored energy (voltage continuity for C, current continuity for L), and V_F from the circuit the switch left behind, with the storage element treated as its steady-state open (C) or short (L).



Worked example 1 (RC discharge). A 100 μF capacitor charged to 12 V discharges through 10 $\text{k}\Omega$. $\tau = RC = 10\,000 \times 100 \times 10^{-6} = 1$ s, so $v(t) = 12e^{-t}$ volts. After 2 s, $v = 12 \times e^{-2} = 1.62$ V. After 5 s (5τ) it is $12e^{-5} = 0.08$ V – settled. Check the energy story: the initial 7.2 mJ has all gone into the resistor by then.

Worked example 2 (charging between two non-zero values). The same RC network, but the switch moves the resistor from a 2 V supply to a 9 V supply with the capacitor already at 2 V. $V_0 = 2$, $V_F = 9$, so $v(t) = 9 + (2 - 9)e^{-t} = 9 - 7e^{-t}$ volts. At $t = 1$ s ($= \tau$) it reaches $9 - 7 \times 0.368 = 6.42$ V – the gap (7 V) has closed by 63.2%, exactly the definition. Endpoint checks: $v(0) = 2$ V and $v(5) = 9 - 0.047 \approx 9$ V.

Worked example 3 (the RL mirror, fast). A 20 mH inductor in series with 100 Ω switches onto a 10 V supply. $\tau = L/R = 0.02/100 = 0.2$ ms, final current $I_F = 10/100 = 0.1$ A, so $i(t) = 0.1(1 - e^{-t/0.2\text{ms}})$. At $t = 0.6$ ms (three time constants) it has reached 95%: 0.095 A. The same shape as the capacitor voltage – one law, two elements.

Method for every switching problem: (1) find V_0 from the old circuit; (2) find V_F from the new circuit; (3) compute τ from the new circuit with the source zeroed; (4) write the one-line answer and check it against both endpoints.

Common mistakes.

- Using the old circuit for V_F or the new one for V_0 – the values come from different snapshots, deliberately.
- Computing τ with the source still present – zero the source to see the true resistance the storage element faces.
- Believing 5τ is a physical constant – it is a 1% engineering convention.

Summary: one storage element means one exponential; V_0 from continuity, V_F from the final circuit, τ from the zeroed-source resistance; $\tau = RC$ or L/R ; 5τ is settled for practical purposes, and every answer checks at both ends.