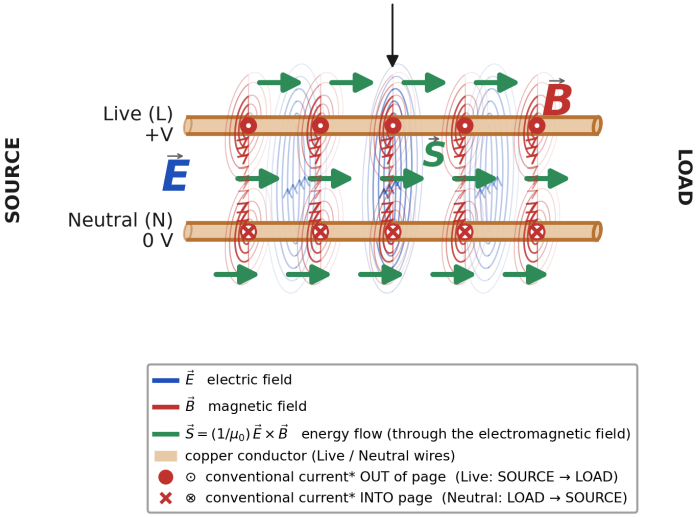
AC, DC, and the Electromagnetic Field

A Practical Quick Reference
from wall sockets to Maxwell

Abandoning the false water-in-a-pipe analogy —
in reality, electrons barely move inside the copper:
they drift slower than a snail.

Rough cross-section — $\vec{E}$ and $\vec{B}$ fill the space between the wires AND fringe around them;
the drawn rings and arcs are sampled slices only — the fields are everywhere, always.



** Conventional current is the slow drift of free electrons in the copper,*
written with the sign of $\vec{E}$ (+V → 0 V) — opposite to how the electrons actually drift.
Nothing flows; only the sign is flipped.

The thin grey arrow above $\vec{E}$, $\vec{B}$ and $\vec{S}$ marks them as vectors, by this pamphlet's convention.
A vector is nothing fancy: it just carries a direction and a magnitude (how strong the push is along that direction).

Energy flows through the electromagnetic field in the space around the conductors —
not inside them. The electrons themselves barely move.

Colophon

Written: April 2026.

Origin note. This pamphlet is the summary of a fruitful interactive discussion held with an AI assistant (via Cursor) in April 2026. Once the discussion had run its course, the AI was subsequently tasked with distilling that conversation into a self-contained, general-purpose quick-reference document. The text that follows is the result of that distillation.

All examples have been generalised. The pamphlet intentionally retains no domain-specific context from the original discussion.

License. This project uses two licenses, each matched to what it actually covers — see the last page or the accompanying `LICENSE.md`. Copyright (c) 2026 rf aabrewery@proton.me.

- The **pamphlet text and its figures** (this document and everything under `images/`) are released under **CC BY 4.0** — a standard license for creative and educational works; you may share and adapt them freely, including commercially, as long as you give appropriate attribution.
- The **build helpers** (`build_em_fields_pamphlet.py` and `make_cover_figure.py`) are released under the **MIT License**, since they are actual software.

Neither license extends to the third-party references cited in the Bibliography chapter, which retain their own copyright and licenses.

How to read this. The chapters are designed to be read in order on first pass, but each one stands on its own as a quick-reference once you are familiar with the material. Cross-references are explicit; the glossary (chapter 11) and bibliography (chapter 12) can be consulted independently.

Disclaimer

This pamphlet is not meant to teach you physics. Please do not cite it as a reference for your school studies — there are proper textbooks for that, a few of which are listed in the Bibliography at the end.

What this pamphlet *is* for is something more modest and, I think, more useful: to show you that **what you were probably taught about electricity is, on its most important point, not true at all**. Most school courses quietly build on the image of "electrons flowing through the wire like water through a pipe" — and then layer everything else (voltage, current, resistance, power, AC, DC, motors, transformers) on top of that single picture.

That picture is a **misconception**. It is not a harmless simplification; it is a wrong starting point, and it makes every later concept harder than it needs to be.

If your studies, like mine, were quietly built on top of that misconception, the good news is that the foundation can be replaced without tearing the house down. The right foundation is the one these pages describe: **energy does not flow inside the copper of a wire — it flows in the electromagnetic field in and around the wire**. Rebuilding on that foundation was, for me, a moment of unusually clean intellectual satisfaction. I hope the same will be true for you.

Hang on for the next few pages: the misconception goes first, the real picture comes right after.

Contents

Colophon 2

Disclaimer 3

Introduction — the field, not the pipe 8
The one idea this pamphlet is built on . . . 8
A word on the hands-on sections . . . 9
A word on the math . . . 9
What energy is, in simple terms . . . 10
What "the field" actually is — a plain-language picture . . . 11
The journey of one joule to a socket . . . 19

Chapter 1 — AC vs DC: the core difference 21
1.1 Two very different beasts . . . 21
1.2 The waveform equation . . . 23
1.3 Why the name "alternating" is literal . . . 24
1.4 Why we use AC at all . . . 24
1.5 What the word "push" actually means . . . 25
1.6 You have already felt the field — the car-door shock . . . 30
1.7 Take-aways . . . 33

Chapter 2 — Why conduction is chemistry, not plumbing 35
2.1 The formal counting definition . . . 35
2.2 The field-based one-liner . . . 36
2.3 Three consequences the old picture hid from you . . . 37
2.4 What each common instrument is actually telling you . . . 37
2.5 A footnote on other media: copper is not the only conductor . . . 39
2.6 The atom, in sixty seconds . . . 40
2.7 What an ion is . . . 40
2.8 Quantum mechanics, defused . . . 41
2.9 Copper specifically: 29 electrons, one of them loose . . . 42
2.10 When copper atoms pack into a solid: the lattice and the shared cloud . . . 43
2.11 What does a "gas" or "cloud" inside a solid actually look like? . . . 46
2.12 The two forces: bound electrons stay, free ones drift . . . 47
2.13 Why the plumbing model cannot explain any of this . . . 49
2.14 Take-aways . . . 50

Chapter 3 — How AC is actually generated 52

3.1 The real machine: a synchronous alternator 52
3.2 What spins the rotor? 54
3.3 Why the rotor is an electromagnet, not a permanent magnet 54
3.4 Why 50 Hz means exactly what you think 55
3.5 Three-phase in one page 55
3.6 Take-aways 56

Chapter 4 — Mains in the wall: Live, Neutral, Earth 57
4.1 Where the three wires come from 57
4.2 What actually happens over one cycle (20 ms at 50 Hz) 57
4.3 Why touching Live shocks you but touching Neutral (usually) doesn't 58
4.4 The Earth (PE) wire 58
4.5 Grounding architectures in one paragraph each 59
4.6 IEC colour codes 59
4.7 Transients on the mains: lightning, network surges, SPDs 60
4.8 Take-aways 64

Chapter 5 — RMS (Root Mean Square) vs Peak 65
5.1 Both numbers are true, they measure different things . 65
5.2 Why can't we just use the simple average? 65
5.3 A worked example 66
5.4 Why "equivalent DC" matters — two resistors 68
5.5 What would you measure in a home wall socket? . . . 69
5.6 Practical consequences — why this matters 69
5.7 The 220 V vs 230 V history 70
5.8 Quick cheat sheet for a pure sine 70
5.9 Take-aways 70

Chapter 6 — Voltage vs current: what they really are 71
6.1 The core confusion 71
6.2 Revisiting: why the water analogy is discarded here . . 71
6.3 What voltage actually is 72
6.4 What current actually is 72
6.5 Ohm's law, without the valve 73
6.6 Where the energy actually is 75
6.7 Worked example: the electric kettle 76
6.8 Where the ±325 V "went" 76
6.9 Current alternates together with voltage 77
6.10 Zero-crossings, 100 times a second 78
6.11 Switching an inductive load: the flyback spike and its taxonomy 80

6.12 Why this feels counter-intuitive . . . 85
6.13 The current transformer — galvanic isolation in working clothes . . . 88
6.14 Take-aways . . . 90

Chapter 7 — The electromagnetic field: the real carrier of energy 92
7.1 The textbook diagram is a useful lie . . . 92
7.2 What actually travels . . . 92
7.3 The Poynting vector: direction and rate of energy flow 93
7.4 What the electrons are doing . . . 94
7.5 Where the electrons come from (and where they go) . 95
7.6 Experimental proof . . . 95
7.7 Better analogies . . . 96
7.8 The one-line truth . . . 97
7.9 Take-aways . . . 97

Chapter 8 — Maxwell in one chapter 98
8.1 What a “field” even is . . . 98
8.2 The electric field E . . . 98
8.3 The magnetic field B . . . 99
8.4 The deep magic: E and B create each other . . . 99
8.5 Why light travels at the speed of light . . . 100
8.6 Maxwell’s four equations — the source code of electromagnetism . . . 100
8.7 Fields in insulators (dielectrics) . . . 104
8.8 What the field carries . . . 105
8.9 Take-aways . . . 105

Chapter 9 — Fields in everyday wiring: loop area, coupling, EMI (electromagnetic interference) 106
9.1 Analog and digital are the same field . . . 106
9.2 The guiding principle of cable design . . . 109
9.3 Three scenarios . . . 110
9.4 Cable design in one table . . . 111
9.5 Capacitive vs inductive coupling . . . 111
9.6 Worked example: a corridor lamp switched from three locations . . . 112
9.7 Five practical rules . . . 113
9.8 Signal-loop example: a 4-20 mA transmitter . . . 114
9.9 Jumpers, headers, DIP switches: stubs in miniature . . 114
9.10 Why specific high-frequency tricks exist . . . 118
9.11 A final beautiful truth . . . 119

9.12 Take-aways . . . 119
9.13 Footnote — on “instantaneously” . . . 120

Chapter 10 — How to actually see and measure fields 122
10.1 Why “invisible” fields are not really invisible . . . 122
10.2 Classic demonstrations . . . 122
10.3 Instruments . . . 122
10.4 Interactive simulators . . . 123
10.5 Looking at the field without measuring it: the “visual audit” . . . 123
10.6 Take-aways . . . 124

A closing note on the math 125

Chapter 11 — Glossary and symbols 127
11.1 Voltage, current, and waveform quantities . . . 127
11.2 Field quantities . . . 127
11.3 Mains wiring . . . 128
11.4 Power, loads, and signals . . . 129
11.5 Electromagnetic interference and compatibility . . . 131
11.6 Standards bodies and common documents . . . 133
11.7 Take-aways . . . 133

Chapter 12 — Bibliography and further reading 134
12.1 Videos and interactive simulators . . . 134
12.2 Canonical textbooks . . . 134
12.3 Practical electrical and electronic engineering . . . 135
12.4 EMC, cable design, and grounding . . . 135
12.5 Historical and narrative . . . 136
12.6 Standards . . . 136
12.7 A suggested “watch next” sequence for Maxwell . . . 137

License 138
Pamphlet text and figures — CC BY 4.0 . . . 138
Build helpers — MIT License . . . 138
Scope . . . 139
Why split? . . . 139

Introduction — the field, not the pipe

The one idea this pamphlet is built on

Almost every introductory account of electricity starts with the same picture: voltage is pressure in a pipe, current is flow rate through the pipe, resistance is a valve, and "electricity" is stuff flowing inside the wire.

This pamphlet refuses to use that picture. We do not use it, we do not refine it, we do not "improve" it. We set it aside here on page one and never pick it up again, and there are **two reasons** to do so — one of which is the real reason.

- **The main reason: the analogy is simply not true.** It is not a harmless simplification, a useful approximation, or a "lie-to-children" that you later trade in for the real thing. It is a **misconception** — a wrong mental model that contradicts what actually happens in a live conductor. Its single most damaging element is the picture of *electrons flowing through the wire like water through a pipe*: in reality electrons barely move at all (net drift of the order of **0.1 mm/s** in DC, a small back-and-forth oscillation in AC), they do not "flow" to the load, and they do not carry the energy. The water analogy misrepresents what is physically happening, and everything built on top of it inherits that error.
- **The practical reason: every downstream concept is easier without it.** Ohm's law, transformers, motors, antennas, VFDs, EMC, and Maxwell's equations are all shorter, sharper, and more accurate when you start from the electromagnetic field rather than from imaginary water in an imaginary pipe.

What replaces it is a single idea, used everywhere from here on:

> **The electromagnetic field is the primary object. The wires are guides, not pipes.**

- The **field** — the electric field $\vec{E}$ between the conductors and the magnetic field $\vec{B}$ wrapping around each of them (the conductors) — is where the action is. It is what carries the energy, what does work on electrons, what couples the primary and secondary of a transformer, what radiates from an antenna.
- The **wires** shape and guide the field. They do not carry the energy *inside* them. They hold the field in the right places and

route it from where it is generated to where it is consumed.
- The **electrons** in the wires are the medium on which the field acts. They barely move — roughly **0.1 mm/s** of net drift in DC, a fraction of a millimetre of back-and-forth oscillation in mains AC. They are not the message.

No pipes. No fluid. No flow inside the wire. Only a field, some guides for it, and a lot of electrons sitting where they have always been, being pushed around gently by that field.

A word on the hands-on sections

This pamphlet is billed as a field-first rethink of AC/DC electricity, and the bulk of it — Chapters 1 through 5, 7, 8, and most of 9 and 10 — stays at that level. Three sections break pace and dive, with unusual specificity, into hardware:

- **§4.7** on mains transients and surge-protective devices;
- **§6.11** on switching an inductive load — the flyback spike and its protection taxonomy;
- **§9.9** on jumpers, headers and DIP switches as transmission-line stubs.

These are not detours and they are not a change of genre. **They are the theory earning its keep.** Each of them is a problem whose conventional "plumbing" treatment (fit a diode because the electrons back up; fit a varistor because the manual says so; ignore the jumper because it's just a bit of wire) is opaque, while the field-first picture makes the reason immediate. We include them precisely because this is where the alternative view returns the most concrete value — the right component, the right place to fit it, and the right reason it works. Each of the three is flagged inline at its opening, so you will never wonder whether you have wandered out of the book's main thread.

A word on the math

Some chapters carry formulas — sometimes trivial (Ohm's law), sometimes a handful of derivatives and integrals, occasionally Maxwell's four equations in a single block. **None of it is required reading.** The formulas are **signposts, not gatekeepers**: each one carries, in symbols, the same meaning as the words next to it; neither side of that pairing is more correct than the other. The

math is there because the fields are among the most precisely *representable* objects in physics, and we wanted to show the range — from $V = IR$ at one end to a glimpse of Heisenberg's matrices and Schrödinger's wave function for the electron at the other. If a formula slows you down, skip to the words beside it; a longer reflection on why we bothered appears at the very end of the pamphlet, after Chapter 10.

What energy is, in simple terms

Before we trace how a socket gets its energy, it is worth naming what energy is.

Energy is the capacity to do work — where "work" has a precise physical meaning: changing the state of something, typically by applying a force over a distance, or by heating it.

The unit is the **joule (J)**. For intuition:

- Lifting a 1 kg object by 1 m takes about **10 J**.
- Heating 1 g of water by 1 °C takes about **4.2 J**.
- Running a laptop for 1 s at 50 W takes **50 J**.

Power is simply how many joules are transferred per unit of time — not a different kind of quantity, just energy divided by seconds:

$$\text{power} = \frac{\text{energy (J)}}{\text{time (s)}}, \quad P = \frac{\Delta E}{\Delta t}$$

Its unit is the **watt (W)**: $1\ \text{W} = 1\ \text{J/s}$. So if an appliance is rated 2 kW it is transferring **2 000 joules every second**, no more, no less. Multiply by how long it runs and you recover the total energy:

$$\text{energy} = \text{power} \times \text{time} \quad \rightarrow \quad E = P \cdot t$$

Run that 2 kW kettle for 3 minutes (180 s): it delivers $2000 \times 180 = 360\,000\ \text{J} = 360\ \text{kJ}$.

Practical unit conversion for the field electrician

Electricity bills, nameplates, and datasheets almost never talk in joules; they talk in **watts** (instantaneous power) and **kilowatt-hours** (energy consumed over time). The bridge is:

- **1 Wh** = 1 W running for 1 hour = `1 × 3 600 s = 3 600 J = 3.6 kJ`
- **1 kWh** = 1 000 W running for 1 hour = `3 600 000 J = 3.6 MJ`

So on a utility bill, **1 kWh is just a bigger joule** — 3.6 million of them, bundled into a unit that survives a month of consumption without scientific notation.

A few shop-floor reference points to calibrate intuition:

Appliance / load	Power (W)	Runs for	Energy (J)	Energy (kWh)
LED bulb	10 W	1 hour	36 000 J	0.010 kWh
Laptop	50 W	1 hour	180 000 J	0.050 kWh
Electric kettle	2 000 W	3 min	360 000 J	0.100 kWh
Electric oven	3 000 W	1 hour	10 800 000 J	3.000 kWh
Typical house (EU)	~400 W	1 day	34 560 000 J	9.600 kWh
Small workshop pump	750 W	8 hour	21 600 000 J	6.000 kWh

Read the table either direction: **watts are the rate**, **kWh are the total**, and **joules are the physics-textbook form of the same total**. They are one quantity viewed at three zoom levels.

Everything else in this pamphlet is, at some level, a description of how the electromagnetic field moves joules from wherever they were generated to whatever is consuming them.

What “the field” actually is — a plain-language picture

Later in the pamphlet we will say things like *“the field arrives at the wall socket”*, *“energy flows in the field”*, or *“a wave travels down the cable”*. It is worth, once and for all, making clear what kind of object *the field* is — and what kind of thing is actually *travelling* in it.

The field is not a substance

> **This single idea — "the field is not a substance" — is the pillar of this pamphlet. Every chapter that follows is an application of it. If you remember one thing from these pages, remember this.**

An **electromagnetic field is not matter.** It is **not** a fluid, not a gas, not a cloud of tiny particles, not any kind of stuff filling the space between the wires.

The field is a **property of space itself** — more precisely, a pair of properties that space acquires whenever **charges** and **currents** are around. By **charges** we mean the free electrons inside a conductor; by **currents**, the collective drift of those electrons once the field acts on them (§§2.1–2.2 give the formal definition).

At every point in the space around a cable, two invisible quantities are defined:

- The **electric field** $\vec{E}$: a direction + a strength. It tells you *how hard* a **still** electric charge — e.g., an electron inside the copper — sitting at that point would be pushed, and *which way*: in the **same direction** $\vec{E}$ is pointing for a positive charge, the **opposite** direction for a negative one like the electron (§2.2). *(The blue arcs between the conductors on the cover.)*
- The **magnetic field** $\vec{B}$: another direction + strength. It tells you *which way* a **moving** electric charge — again, think of a drifting electron — at that point would be shoved sideways. *(The red rings along each conductor on the cover.)*

That is the field. It **is** those two directed quantities — $\vec{E}$ and $\vec{B}$ — defined at every point in space.

*(So the field is **not** a couple of arrows drawn somewhere; it is a **condition that space has taken on** because of the charges and currents around it — a condition that, at every point, carries the information "how hard and which way would a charge here be pushed". A single cubic metre of space already carries infinitely many such* $\vec{E}$ / $\vec{B}$ *pairs, one per point inside it.)*

Remove all the air, remove all the matter — the field is still there. The field does **not** need a medium to exist.

Is the field real, or "pure math"?

The field **is real**. It is not a mere representation or an accounting device; it is a physical entity, on the same footing as matter itself. Modern physics stopped treating it as "just a mathematical convenience" more than a century ago, for three concrete reasons:

- **It carries energy.** Every point of space where the field exists stores a measurable amount of energy per cubic metre — the *electromagnetic energy density* u, given by the formula boxed below. It is measurable, conserved, and convertible into heat or motion. Something that carries energy is not a representation.
- **It carries momentum.** Light pushes on things (radiation pressure, solar sails). While the light is in flight, that momentum has to live *somewhere* — it lives in the field.
- **It outlives its source.** A star a million light-years away emits a pulse, then dies. The pulse still reaches Earth a million years later, with no charges involved along the way — just the field, propagating through empty space, carrying the energy and the information.

What is "pure math" is only the **language** used to describe the field, just as mathematics is the language used to describe gravity or fluid flow. The equations are not the field — they describe it. The field has its own independent existence; the equations only tell us its shape, its dynamics, and how it couples to charges and currents.

The difficulty of *visualising* the field is real, but it is not evidence that the field is unreal. Human intuition is trained on *matter* — things made of atoms, with surfaces and weights — and the field has none of those attributes. It is not a substance. This is why every attempt to picture it as a fluid, a gas, an ether, or "stuff flowing in the wire" fails: those analogies impose a matter-ontology on something that is simply a different kind of physical object. The honest answer is that the field is invisible *by its nature* and is best understood through what it *does* (energy, momentum, forces, waves), described with vector calculus, rather than through what it "looks like".

Electromagnetic energy density u [J per m^3]

$$u = \frac{\varepsilon_0}{2} |\vec{E}|^2 + \frac{1}{2\mu_0} |\vec{B}|^2$$

Where:

- u — electromagnetic energy density at the point of interest, in joules per cubic metre (J/m^3). A scalar.
- $\vec{E}$ — electric-field vector at that point, in volts per metre (V/m).
- $\vec{B}$ — magnetic-field vector at that point, in teslas (T).
- $|\vec{E}|^2$, $|\vec{B}|^2$ — squared *magnitudes* (lengths) of the E- and B-vectors; each is a scalar, units V^2/m^2 and T^2 respectively. See the forthcoming "math of vectors" primer.
- ε_0 — electric permittivity of vacuum, $\approx 8.854 \times 10^{-12}$ F/m. A scalar constant of nature.
- μ_0 — magnetic permeability of vacuum, $\approx 4\pi \times 10^{-7}$ H/m. A scalar constant of nature.

What it says. Every cubic metre of space that has an E-field or a B-field in it stores a measurable amount of energy, in joules. Half of that energy lives in the electric field, half in the magnetic field. This is why the field can push a solar sail, heat a resistor, or light a bulb — it literally carries energy, everywhere it exists. It is also why claims of the form "the field is just bookkeeping" cannot be right: bookkeeping does not boil water.

What "flows" in the field, then?

When the voltage at the source changes (and in AC it flips 50 times a second — for 50 Hz European mains; 60 in the US — more on Hertz in §3.4), the **field itself** everywhere must re-arrange to follow it *(most rigorously: the condition of space at every point must re-adjust)* — and that re-arrangement cannot travel faster than light *(more on why in Chapter 8, where c is derived from Maxwell's equations alone)*.

A **rippling, coordinated change of $\vec{E}$ and $\vec{B}$ as it spreads outward from the source is what we call an electromagnetic wave** *(a ripple in the condition of space itself)*.

It **is** a wave — but not a wave *of matter*. **The field itself is what**

oscillates (most rigorously it could be said as *"it is a wave in the condition of space itself"*): $\vec{E}$ and $\vec{B}$ change in value and direction at every point, each point pulling its neighbours into the same pattern, the whole pattern racing outward at (nearly) the speed of light.

This moving pattern carries two things that are very real and measurable:

- **Energy** — stored in the field, in an amount proportional to $|\vec{E}|^2 + |\vec{B}|^2$ at each point (here $|\vec{E}|$ is the *electric-field* magnitude (strength) and $|\vec{B}|$ the *magnetic-field* magnitude (strength) — not to be confused with the energy they carry).
- **Momentum** — which is why radio waves can, in principle, push a sail in space.

So when we later say *"the field arrives at the socket"*, what physically arrives is:

- a specific pattern of $\vec{E}$ and $\vec{B}$ values now present between Live and Neutral, and in the space around the cable;
- carrying a specific amount of energy in that pattern;
- ready to act on any charge that closes the circuit.

Nothing *material* flowed from the power plant to your socket. A **configuration of the field**, with energy inside it, propagated along the cables at nearly the speed of light.

*("Energy" here is not a substance, not a fluid, not a tiny packet. It is the field's **capacity to cause change** — to push a charge, heat a wire, spin a motor. The boxed energy-density formula just above — essentially $|\vec{E}|^2 + |\vec{B}|^2$ weighted by two constants of nature — counts **how much** of that capacity the field is currently carrying there.)*

Energy is an abstract thing.

Richard Feynman, in his *Lectures on Physics*, put it bluntly:
"It is important to realize that in physics today, we have no knowledge of what energy is. We do not have a picture that energy comes in little blobs of a definite amount. [...] It is an abstract thing."
That is the honest answer, and it is worth keeping as a quiet caution behind every further use of the word **energy**. The rest

is **what we do know with certainty**:

- **Energy is not a substance.** It is not a fluid, not a particle, not a little packet, not a stuff of any kind. Looking inside a charged battery for "the energy" the way one looks inside a bottle for the water is a search with nothing to find.
- **Energy is a capacity** — specifically, a system's *capacity to cause measurable change*: a charge pushed, atoms made to jiggle harder (heat), a bond broken (chemistry), a weight lifted, light emitted.
- **Energy is conserved.** It never appears or disappears, only transforms between forms, and the totals always balance. That conservation is what makes it *feel* like a substance (we speak of it flowing, being stored, being carried) — but the substance is a metaphor; the underlying reality is an abstract number the universe always balances.

The different **forms of energy** are just different kinds of change a system is currently set up to cause:

- **Kinetic** — a moving thing's ability to slow down and move something else.
- **Thermal** — atoms jiggling randomly; hot transfers its jiggle to cold.
- **Chemical** — electrons arranged in bonds that, if rearranged, release heat or light (fuel, food, battery).
- **Gravitational potential** — a raised weight's ability to fall and drive something.
- **Field (EM) energy** — a configuration of $\vec{E}$ and $\vec{B}$ in space that, if it meets a charge, pushes it; if it meets a resistive wire, heats it; if it meets a motor coil, spins it.

Each form is the same one abstract thing (capacity-to-cause-change) stored in a different physical arrangement. When a kettle boils, **chemical**-form energy in the fuel at the power plant becomes **field**-form energy travelling down the cable becomes **thermal**-form energy in the water. The total is always the same number; only the form changes.

So whenever this pamphlet says "energy in the field", read it as: the field, in its current $\vec{E}$ / $\vec{B}$ configuration, is in a position to make things happen the moment it meets matter. The field *is* that capacity; *energy* is the name of *how much* of it is parked at each point.

One visualization you already know

If you have ever sprinkled iron filings on a sheet of paper over a magnet, you have seen a field made visible. The curved lines along which the filings line up are **not** the field — the field was already there, invisible, everywhere around the magnet. The filings just reveal its direction.

The electromagnetic field around a mains cable is the same idea, with one addition: the source (generator, transformer) keeps flipping its polarity 50 times per second (for 50 Hz European mains; 60 in the US), so the field flips with it *(a wave in the condition of space itself)*. The flipping pattern races down the length of the cable at almost the speed of light, and **the energy is carried in that travelling pattern of $\vec{E}$ and $\vec{B}$** — not in the copper, not in the electrons.

Wires give the pattern a place to live. Electrons let it do work on the load (the field pushes them and that push, despite them barely flowing/moving, becomes heat, motion, or light inside the load). But the thing that actually moves from the plant to your socket is a **configuration of $\vec{E}$ and $\vec{B}$ in space** — a field, not a fluid.

A better ladder of analogies

Keeping the concept strictly, but trying to give a picture of it. See in next box how chemistry did exactly this, elegantly, with quantum mechanics.

1. The gravitational field is created by the planet — the grammar all three use.

- Mass (m) $\rightarrow$ g-field $\rightarrow$ asteroid feels $\mathrm{F} = \mathrm{m} \cdot \mathrm{g}$.
- Charge (q) $\rightarrow$ E-field $\rightarrow$ electron feels $\mathrm{F} = \mathrm{q} \cdot \mathrm{E}$.
- Current (I) $\rightarrow$ B-field $\rightarrow$ moving charge feels a sideways shove.
- Note: the landmine is English (“gravity” = source / field / force all collapsed into one word), not the physics.

2. The deepest point: the arrows are already an analogy.

- The field is *not made of arrows* — arrows are how we *draw* it on paper.
- The field itself is a **condition of space**: at every point, space now “knows” how it would push any charge that happens to be there.

- Mathematically: a function that assigns a direction-and-magnitude to every point.
- Physically: a real property of space, not a fluid, not a gas, not an ether.
- Every diagram in this pamphlet is a map; the thing mapped is that condition of space — a pattern the space has taken on because of what is elsewhere.

3. A cleaner three-step ladder.

- Step 1 — *scalar field*: a temperature map in a room. Every point has a number; no direction.
- Step 2 — *vector field*: the TV weather / wind map. Every point has a number **and** a direction.
- Step 3 — *EM field*: same idea as the wind map, but the value+direction at each point encodes *"what push any charge placed here would feel"*.

How chemistry solved the analogy problem

The same trap we are trying to avoid here — inventing a new substance as the analogy — is one that chemistry and quantum physics faced a century ago and solved without cheating. Bohr, Schrödinger, and Heisenberg never replaced electrons with buoys or the field with water. They kept the actual physical objects intact and competed only on how to **picture** them.

Historical accuracy note: when Bohr built his model in 1913 and when Heisenberg and Schrödinger invented matrix and wave mechanics in 1925–26, the neutron had not yet been discovered (Chadwick, 1932). The atom was understood as a positive nucleus plus orbiting electrons. None of that changed under the "visualization debate" — only the way of picturing the electrons did.

1925 — Heisenberg's matrix mechanics. Built from what is *measurable* (transition frequencies between atomic states). No pictures, no orbits, no orbitals — just matrices acting on abstract state vectors. Fully correct, mathematically elegant and beautiful, for high-school students utterly unvisualizable.

1926 — Schrödinger's wave mechanics. An electron becomes a wavefunction $\psi(x,t)$ you can *plot*. Orbitals have a *shape*: s, p, d, f. It looks like something. Schrödinger explicitly rejected the matrix approach because he wanted to *see* the

atom.
1926 (later) — equivalence proof. Schrödinger proves the two formulations are mathematically equivalent. Two languages, one physics.
Lesson for this pamphlet. We want to build upon the Schrödinger-Bohr way — pictures, orbitals, field-lines, arrows, iron filings. We **embrace** that, because a visual-but- imprecise picture is a better entry point than a correct-but- unpicturable one. We just add the discipline of naming pictures as pictures. No electron becomes a buoy; no field becomes water. The picture decorates the physics; it does not replace it.

The journey of one joule to a socket

Follow a single joule of energy from its origin all the way to a kettle plugged into a wall. At no point in the journey is it "inside the wire".

1. **Source.** In a hydro plant the joule starts as gravitational potential energy of water held high in a reservoir. In a gas or coal plant, chemical potential energy locked in the fuel. In a photovoltaic panel, radiant energy from the sun. In a wind turbine, kinetic energy of moving air.
2. **Mechanical stage.** Falling water, expanding steam, or wind spins a rotor. The joule is now kinetic energy of a rotating mass.
3. **Generator.** The rotor drags magnets past coils of wire. The changing magnetic field induces an electric field in the coils (Faraday's law). The joule leaves the rotor as **energy in the electromagnetic field around the generator windings** — not as anything "flowing" inside the copper.
4. **Transmission.** The field is guided along the transmission lines — conductors running parallel for hundreds of kilometres. The energy travels in the space *between and around* those conductors at nearly the speed of light. Transformers step the voltage up to hundreds of kilovolts for the long haul (low current → low $\mathrm{I}^2\mathrm{R}$ wire losses) and step it back down for distribution.
5. **Distribution.** The field is routed through the substation, the street cable, the service drop, the consumer unit, and the house wiring. The conductors are just rails that shape where the field is allowed to go.
6. **Socket.** The field arrives at the wall socket, standing ready at 230 V RMS between Live and Neutral (50 Hz in Europe). It is *waiting*. Nothing flows yet, because there is nowhere for it to

go.

7. **Load.** Plug in the kettle. Its heating element closes the loop. The field in the space around the supply cable and the element now drives the element's electrons — does **work** on them *(it pushes each electron into collisions with the metal atoms of the element; those countless tiny collisions make the atoms jiggle harder, and "atoms jiggling harder" is what "hot metal" physically **is** — the electrons themselves barely travel: in AC they just oscillate back-and-forth in place; §6.6 unpacks the mechanism)* — and heats the metal. The joule is delivered. Meanwhile the electrons in the cable itself have barely moved.

Not once in that chain does anything flow inside a pipe. A joule of electrical energy spends its entire electrical life as a feature of a field in the space *around* the conductors, with the conductors acting as shaping guides.

Every chapter that follows is built on top of this picture. The rest of the pamphlet's job is to make it precise, give it units, and show it at work in real apparatus — from wall sockets to Maxwell's equations.

Chapter 1 — AC vs DC: the core difference

1.1 Two very different beasts

DC (Direct Current).

- Voltage polarity is **fixed**: the positive terminal stays positive, the negative terminal stays negative, for as long as the source is on.
- Equivalently, the electric field between the conductors points in one steady direction.
- As a (slow) consequence, the electrons in the wire acquire a small net drift in that direction — about **0.1 mm/s** at 10 A in household copper. They do not race through the wire; they barely creep. The energy is not carried by them (chapter 7).
- Voltage as a function of time = a **flat horizontal line**.
- Examples: batteries, solar panels, USB power, almost every internal rail inside a piece of electronics.

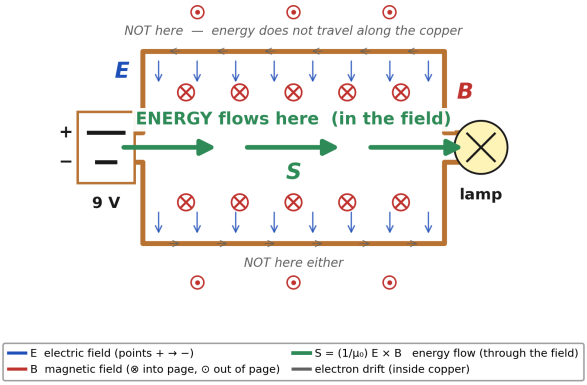
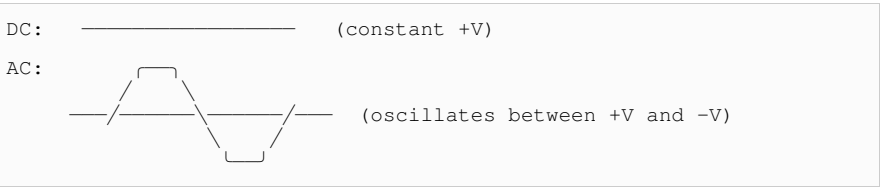


Figure 1.1 — A 9 V battery (west) powers a lamp (east) through a deliberately long rectangular wire path. The thick green arrows are the Poynting vectors $\vec{S} = (1/\mu_0)\,\vec{E} \times \vec{B}$: **energy flows directly across the field**, from battery to lamp, through the space enclosed by the loop — *not* along the copper. The blue arrows are the E-field (top rail → bottom rail); the red ⊗ markers mark $\vec{B}$ into the page inside the loop (and ⊙ out of the page outside it). The copper wires only *guide* the field; the "NOT here" labels mark where the energy is **not** travelling.

AC (Alternating Current).

- Voltage polarity **reverses periodically**: what was "+" becomes "−", then "+" again, at frequency $\mathtt{f}$ (50 Hz in Europe, 60 Hz in the US).
- Equivalently, the electric field between the conductors flips direction $\mathtt{f}$ times per second.
- As a consequence, the electrons in the wire oscillate back and forth about a fixed mean position — by a fraction of a millimetre, 50 or 60 times a second. They do not travel anywhere. Again, the energy is carried by the field around the wires, not by them (chapter 7).
- Voltage as a function of time = a **sine wave**.
- Examples: mains power, generators, transformers.



Where does the field come from, and why does it "wave"?

The single sentence that ties every piece of this section together, read in the field-based way:

> *The household circuit carries energy because of the potential difference between Hot and Neutral; and in AC the **time-variation** of that potential difference is what turns the field from a static configuration into an electromagnetic wave.*

Unpacking it.

1. Why the field exists at all. If the two conductors sat at the same potential, there would be no $\vec{E}$-field between them ($\vec{E}$ is the gradient of potential, so no ΔV means no $\vec{E}$). No voltage would push current through the load, so the wires would carry no current either — which means no $\vec{B}$-field. With no $\vec{E}$ and no $\vec{B}$, the Poynting vector $\vec{S} = (1/\mu_0)\,\vec{E} \times \vec{B}$ is zero, and **no energy flows anywhere**. The sustained voltage difference between the source's two terminals is, therefore, the single ingredient

that *cocks* the whole field configuration — cover picture, green Poynting arrows and all.

2. Why DC is a steady river and AC is a wave. A potential difference by itself only gives you a **static** field. A 9 V battery across a resistor has constant $\vec{E}$ between the leads, constant $\vec{B}$ around each wire, and a constant Poynting vector pointing battery $\rightarrow$ load. Energy flows, but nothing oscillates; there is no "wave" in the usual sense, just a steady configuration of the field.

What turns the field into a **propagating wave** is the source's voltage *changing in time*. The moment $\vec{E}$ changes, Maxwell's equations force a changing $\vec{B}$ to appear; a changing $\vec{B}$ forces a changing $\vec{E}$; that mutual forcing *is* what an electromagnetic wave mechanically is. In AC mains, the source drives $V_{\text{Hot}}(t) - V_{\text{Neutral}}(t)$ to oscillate sinusoidally at 50 Hz (or 60 Hz), so the whole field pattern around the wires is forced to re-adjust fifty times a second — and that re-adjustment propagates outward at the speed of light.

3. Why AC mains does not *look* like a wave. At 50 Hz the wavelength is $\lambda = c/f = 3\times10^8\,\text{m/s}/50\,\text{Hz} = 6\,000\,\text{km}$. Over the length of a house wire (tens of metres) you are sitting **very deep inside a single wavelength**, so the field looks almost static to the eye. It is still genuinely a wave; you simply never see a full cycle laid out in space.

In one line: the *existence* of ΔV creates the field; the *time-variation* of ΔV (AC) turns that field into a propagating wave; in DC, ΔV is fixed and the field stays static while energy still flows steadily through it.

1.2 The waveform equation

For a pure AC sine wave:

$$v(t) = V_{peak} \cdot \sin(2\pi \cdot f \cdot t)$$

Where:

- $v(t)$ — instantaneous voltage at time t
- V_{peak} — peak amplitude of the waveform
- f — frequency in hertz (cycles per second)

- $2\pi \cdot f$ — angular frequency in radians per second

1.3 Why the name “alternating” is literal

In AC, **both voltage and current alternate together**. During the positive half-cycle, current — the free electrons’ collective drift through the copper, a tiny forward push at a peak of roughly **0.1 mm/s** (an instantaneous peak, **not** a net travel rate — it reverses every 10 ms, so in AC no electron ever actually goes anywhere) — flows one way through the circuit; during the negative half-cycle the drift **reverses**: the very same electrons inch at the same peak **0.1 mm/s** the other way, so across one full AC period each electron just wiggles a fraction of a millimetre back and forth in place. That direction-reversal is what the name says, and it is the single most important fact to hold onto when the rest of the pamphlet starts building on it.

Those electrons, however, **don’t go anywhere in the sense the old plumbing analogy implies** — their net displacement per cycle is a fraction of a millimetre. They **oscillate** in step with the driving voltage, and the *energy* is carried not by their slow drift but by the electromagnetic field around the wire (chapter 7).

*(A clarification on the word “current”. In a copper wire the free electrons **are** the charge carriers, so their collective drift is literally what we measure and call “the current”. Saying “the current” and “the electrons” refers here to **the same population of particles**, just counted in different units — amperes on one side, electrons-per-second on the other. It is **not** that current is some separate fluid which the electrons carry; in metals, current is the electron drift. Energy is what is separate: that lives in the field around the wire, not in those electrons.)*

1.4 Why we use AC at all

A quick historical note — the famous “War of Currents”:

1. **Transformers only work with AC.** They rely on a *changing* magnetic field to induce voltage in the secondary coil. DC produces a constant field → no induction → no voltage transfer.
2. **Long-distance transmission efficiency.** Transformers let power companies step voltage way up (400 kV or higher) for transmission, then step it back down near the consumer. Since

power loss in wires is $P_{loss} = I^2 \cdot R$, **high voltage = low current = low losses**.

3. **AC motors** (especially squirrel-cage induction motors) are simpler, cheaper, and more reliable than DC motors — no brushes or commutators to wear out.

This is exactly why Tesla & Westinghouse's AC won over Edison's DC for grid distribution in the late 1800s.

1.5 What the word "push" actually means

Throughout this pamphlet you will see words like "push", "drive", and "force" when we talk about how the field acts on the electrons in a wire. That vocabulary is not a leftover from the plumbing picture we threw away in the Introduction. It names a real, physical, empirically measurable force. This short section says exactly *which* force, and shows the experimental evidence that it exists.

The push is the Coulomb force F = q$\vec{E}$

Whenever an electric field $\vec{E}$ (volts per metre) exists somewhere, any charge q sitting in it feels a force.

Notice the payoff already. This is the first working equation the pamphlet writes down, and it is only writable because the previous chapters put the field on solid ground. The idea we built — that space itself carries a physical condition called $\vec{E}$ — is already bearing fruit: in a single line we can now name the exact force felt by every electron in every wire everywhere. From here on, most physical claims in this pamphlet will have the same shape as this one: *"the field is doing X at this point in space, so the charge sitting there feels Y."*

Coulomb force on a charge in an E-field [N]

$$\vec{F} = q\vec{E}$$

Where:

- $\vec{F}$ — force on the charge, in newtons (N). Its direction is the same as $\vec{E}$ for a positive charge, and opposite for a negative charge (like an electron).

- q — amount of charge on the particle, in coulombs (C). A scalar. For a single electron, $q = -e \approx -1.6 \times 10^{-19}$ C.
- $\vec{E}$ — electric-field vector at the location of the charge, in volts per metre (V/m). A vector: it has a **magnitude** (the field's strength — what ultimately sets how hard each free electron gets pushed, and therefore the peak drift speed we met in §1.3, e.g. the 0.1 mm/s in mains AC) and a **direction** (which way along the wire the field points — when $\vec{E}$ reverses, the electron drift reverses with it, which is exactly what "alternating" means in AC).

What it says. Wherever there is an electric field, every charge sitting in that field feels a real, physical force proportional to the field's strength and the particle's own charge. It is the force a balloon feels toward a wall after you rub it on your hair; it is also the force every free electron in a copper wire feels when a battery is connected. The entire mechanism of electrical conduction rests on this one equation.

This is the Coulomb force — the same force that makes the balloon stick to the wall. It exists everywhere the $\vec{E}$-field exists, on every charge, all the time.

When a battery or transformer establishes a voltage V across a copper wire of length L, the $\vec{E}$-field inside the copper is roughly $|\vec{E}| \approx V/L$ volts per metre. Every one of the $\sim 10^{28}$ free electrons per cubic metre of copper therefore feels a real force of magnitude $|\vec{F}| = e \cdot |\vec{E}| = eV/L$ newtons, where $e \approx 1.6 \times 10^{-19}$ C is the electron charge.

That is "the push". Nothing fluid, nothing mechanical, nothing hydraulic. A real physical force, in newtons, with a sign, a direction, and a magnitude you can compute from the voltage and the geometry.

Four independent empirical proofs that the push is real

$\vec{F} = q\,\vec{E}$ is one of the best-tested equations in physics. Four different kinds of experiment all agree that it is the literal, real force at work in every wire:

- **Cathode ray tubes (J. J. Thomson, 1897).** In a vacuum tube, electrons are liberated from a hot filament and fly through empty space. A voltage across two deflection plates creates an E-field; the electrons deflect in flight along a parabolic trajectory that matches $\vec{F} = q\,\vec{E}$ exactly. Thomson used the numbers to measure the electron's charge-to-mass ratio — the experiment that *proved* electrons existed as discrete particles.
- **Ohm's law.** Double the voltage across a resistor and the current doubles, because the E-field inside the wire doubles and the Coulomb force on every electron doubles with it. Ohm's law is empirical — it was measured before it was understood — and it is exactly what $\vec{F} = q\,\vec{E}$ on a sea of scattering electrons predicts (the Drude model, 1900).
- **Particle accelerators.** Every linear accelerator from a CRT to CERN works by placing charged particles in an E-field and measuring the resulting acceleration. The observed $\vec{a} = q\,\vec{E}/m$ matches theory to many decimal places. If the "push" were a metaphor, these machines would not work.
- **Hall effect.** Run a current through a conductor placed in a magnetic field, and a transverse voltage appears across the conductor. This only makes sense if the current is carried by real, discrete charges being pushed sideways by the Lorentz force $\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$. The sign of the Hall voltage even tells you whether the carriers are positive or negative.

In short: the "push" is not a figure of speech. It is the direct, continuously-acting, experimentally-verified Coulomb force on every free electron in the wire.

The accelerator and Hall-effect items above each name a law of nature of their own. Here they are, boxed in the pamphlet's standard format.

Electron (or charge) acceleration in an E-field [m/s²]

$$\vec{a} = \frac{q\,\vec{E}}{m}$$

Where:

- $\vec{a}$ — acceleration of the charged particle, in metres per sec-

ond squared (m/s^2). A vector: it points the same way as the force.
- q — amount of charge on the particle, in coulombs (C). A scalar.
- $\vec{E}$ — electric-field vector at the particle's location, in volts per metre (V/m).
- m — mass of the particle, in kilograms (kg). A scalar. For an electron, $m_e \approx 9.11 \times 10^{-31}$ kg.

What it says. Newton's second law applied to a charged particle in an electric field. The E-field creates a force $\vec{F} = q\vec{E}$, and that force accelerates the particle according to $\vec{F} = m\vec{a}$. Combine the two and you get how fast a given field can speed a given charge up — the working equation of every particle accelerator ever built.

Lorentz force — the full force on a moving charge [N]

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$$

Where:
- $\vec{F}$ — total electromagnetic force on the charge, in newtons (N).
- q — charge, in coulombs (C). A scalar.
- $\vec{E}$ — electric-field vector at the particle's location, in volts per metre (V/m).
- $\vec{v}$ — velocity vector of the particle, in metres per second (m/s).
- $\vec{B}$ — magnetic-field vector at the particle's location, in teslas (T).
- $\vec{v} \times \vec{B}$ — cross product; a vector perpendicular to both $\vec{v}$ and $\vec{B}$, with magnitude $|\vec{v}|\,|\vec{B}| \sin\theta$. See the forthcoming "math of vectors" primer.

What it says. This is the complete, general electromagnetic force on any charge, anywhere. The first term, $q\vec{E}$, is the Coulomb force from the previous box — it acts on every charge regardless of motion. The second term, $q\,\vec{v} \times \vec{B}$, is the magnetic force — it only acts on *moving* charges, always sideways to their

motion. Together they account for everything the field can do to a charge: speed it up, slow it down, or deflect it.

So why don't the electrons accelerate to the speed of light?

If there is a real force on every electron all the time, why don't they keep picking up speed indefinitely? Because copper is not empty space. Every $\sim 3 \times 10^{-14}$ seconds a free electron collides with a copper atom (more precisely, scatters off a lattice vibration or an impurity), loses most of the kinetic energy it has just picked up from the field, and starts over in a random direction. The field immediately begins pushing it again; it collides again; and so on, about 10^{14} times per second.

What survives this relentless sequence, averaged over billions of collisions per second, is a tiny **net** velocity in the direction of the push — the **drift velocity**, roughly $0.1\ \mathrm{mm/s}$ at 10 A in household copper. The kinetic energy dumped into the lattice at every collision is what you feel as heat: the famous I^2R losses in a wire are literally this.

The paradigm change in one sentence

Old (wrong) picture:

> *A source pushes current through a wire the way a pump pushes water through a pipe.*

Correct picture:

> *A source establishes an electric field both **inside** the wire (a tiny field along the wire's length, whose only job is to keep the free electrons drifting) and **around** the wire (a much larger field, between the conductors, in which the **energy** then flows — chapter 7). The **inside** field exerts a real Coulomb force $\vec{F} = q\vec{E}$ on every free electron sitting there; the electrons respond by drifting slowly, bouncing off copper atoms as they go.*

(So the field is not just one thing in one place: it lives wherever the source's voltage can reach. Inside the copper it is feeble but non-zero — enough to drive the drift; outside, in the space between the conductors, it is much stronger and carries all the energy. §1.5

is zoomed in on the inside piece because that is where the "push" story lives; chapter 7 zooms back out to the outside piece, where the energy-flow story lives.)

From this one substitution everything else in the pamphlet follows:

- **"Voltage"** is the name for how hard the field pushes per coulomb of charge ($\mathrm{V/m} \times \mathrm{m} = \mathrm{V}$).
- **"Current"** is the name for the slow drift of electrons that results.
- **"Resistance"** is the name for how much the lattice gets in the way of that drift (more collisions → more scattering → less drift per volt).
- **"Energy"** is not delivered by the drifting electrons at all. It is delivered by the field itself, in the space around the wires (chapter 7).

No fluid. No pipe. Just a field, a force, a lot of electrons getting nudged along, and a lot of atoms getting warm from the collisions.

1.6 You have already felt the field — the car-door shock

Everything §1.5 just argued — that the field is a real, physical condition of space, and that it **carries energy** — rests on laboratory evidence (cathode-ray tubes, Ohm's law, particle accelerators, the Hall effect). You don't need a laboratory to test it. You just need a dry day and a car.

Step out of a car on a cold, dry winter morning, walk around to the boot, and reach for the door handle. Somewhere between your fingertip and the metal — a millimetre before you actually touch — a small blue spark fires. You feel it. You may hear it. Then it is gone. Three billion people have felt that shock. Nobody had to explain *current*, *Ohm's law*, or *Maxwell's equations* to them to do so.

Here is what actually happened, in three acts:

1. **Charging up (while driving).** Friction between your clothes and the seat shuffled a few billionths of a coulomb of charge onto your body (this is called *triboelectric* charging). It is far too little for you to feel; your body potential climbs silently from zero to a few thousand volts.
2. **Setting up the field (as you reach for the handle).** You stepped out, still charged, and approached the grounded metal

of the car body with a finger. The air gap between fingertip and handle was now sitting in an **electric field**. *Nothing was flowing yet.* But the space itself had acquired a condition — the same condition the pamphlet has been pointing at all along — and that condition was already **carrying energy**, exactly at the density $(\varepsilon_0/2)|\vec{E}|^2$ we named back on page 12.

3. **Release (the moment you touch).** At roughly a millimetre of gap, the field strength exceeded the breakdown voltage of air (~3 MV/m). A narrow, briefly conductive channel of ionised air formed between finger and handle. The field in the whole gap — not just along the spark — dumped its stored energy through that channel in under a microsecond. You felt the shock.

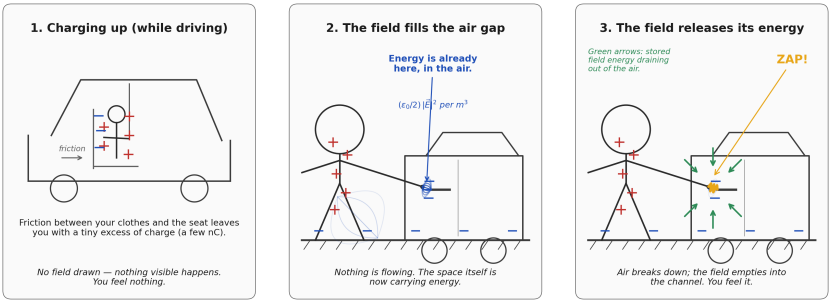


Figure 1.2 — The car-door shock, in three acts. **Left:** during the drive, friction with the seat leaves you with a tiny excess of charge; you feel nothing. **Centre:** you step out and reach for the handle. Blue arcs mark the E-field that has appeared in the air gap between fingertip and handle (and, fainter, between body and ground). Nothing is flowing yet — but the space itself is now carrying energy, at $(\varepsilon_0/2)|\vec{E}|^2$ per cubic metre. **Right:** at about a millimetre of gap the air breaks down; a narrow conductive channel opens and the green arrows show the stored field energy draining out of the surrounding air into that channel. You feel the shock.

The pedagogical punchline is the part that makes the car-door shock the pamphlet's cleanest everyday exhibit:

> **No wire was ever involved.** No circuit from the car to you, no battery, no source, no copper path anywhere. And yet energy was delivered — enough to sting a fingertip. **That energy had been sitting, quietly, in the air between your fingertip and the handle.** The spark is

just the moment the air let go of it.

If the field-carries-energy story is hard to swallow for a copper wire (where you can still wrongly imagine that the *electrons* bring the energy along with them, a picture Chapter 7 will dismantle in detail), try swallowing it for the car-door shock. There are no wires. There are no electrons travelling from source to load. There is only you, the ground, the car, and the air in between — and the shock is the direct demonstration that **the air was carrying energy**, which is to say that **the field was**. Same story as every live wire in your wall; just a cleaner exhibit.

An everyday experiment a child can run

Pick a cold, dry afternoon — humidity under roughly 40% is best. (In damp air the charge bleeds off slowly through the moisture and the shock never builds up; the field is still there every step of the way, it just discharges smoothly instead of suddenly.)
Walk across a synthetic-fibre carpet in rubber-soled shoes for a few seconds while shuffling your feet; or slide out of a car seat. Now reach a finger toward a metal door handle, a radiator, or a tap. You will see a small blue spark jump the last millimetre, and feel a pinch.
A five-year-old can hold the whole pamphlet's core claim in two sentences after doing this:

- **Before the spark:** "The air was holding something."
- **The spark:** "The air just let it out."

That is the field. Nothing more mystical is needed.

Two side notes close §1.6.

- **A number check.** A typical dry-day body potential before a discharge is a few kilovolts; the body acts as a capacitor of roughly 100 pF (picofarads, 10^{-12} F) to ground; the stored energy is therefore $\frac{1}{2} C V^2$ (with C the capacitance in farads and V the voltage in volts) — tens of microjoules to a few millijoules. Startling to a fingertip, utterly harmless to a healthy human — and **instantly fatal to an unprotected MOSFET** (Metal-Oxide-Semiconductor Field-Effect Transistor) gate, whose oxide is a few atoms thick. That is why §§4.7 and 6.11 spend as many pages as they do on transient protection: the same air-gap field that pinches a fingertip can punch a sub-nanometre hole through the gate oxide of a CMOS (Complementary Metal-

Oxide-Semiconductor) input if it finds one. This is the physical origin of **ESD** (Electro-Static Discharge) damage to electronics.

- **Cross-reference forward.** The car-door shock is the simplest case of what the cover diagram and Chapter 7 show for a live wire: energy travels in the field, *in the space around (or between) the conductors*, not along the copper. In the car-door case the simplification goes all the way — there is no copper at all. Same story, fewer moving parts.

1.7 Take-aways

- DC = a steady electric field in one direction; AC = an electric field that reverses direction f times per second (50 Hz in Europe, 60 Hz in the US).
- The "push" on the electrons is the Coulomb force $\vec{F} = q\vec{E}$ — a real, empirically measurable physical force, confirmed by cathode-ray deflection, Ohm's law, particle accelerators, and the Hall effect. It is **not** a fluid-flow metaphor.
- Electrons do not "flow" through the wire the way water flows through a pipe. They drift slowly (DC) or oscillate in place (AC), colliding with copper atoms $\sim 10^{14}$ times per second. Those collisions are what heats the wire ($\mathrm{I}^2\mathrm{R}$ losses).
- AC is described by a sine wave; its angular frequency is $2\pi \cdot \mathrm{f}$.

What is "angular frequency", and what is a "radian"?

The last bullet says that AC has *angular frequency* $\omega = 2\pi f$. Two definitions are hiding in those four symbols. Both are worth spelling out once, because every transformer, motor, alternator, current-transformer, antenna and resonant LC circuit later in the pamphlet quietly relies on them.

1. Angular frequency ω. A sine wave is the vertical projection of a point spinning at constant speed around a circle. Two ways to describe how fast that point spins:

- how many *full turns* it makes per second — the **frequency** f, measured in hertz (Hz, "cycles per second"). Mains in Europe is $f = 50\,\mathrm{Hz}$: 50 full turns of the spinning point each second.
- how many *radians* the point sweeps per second — the **angular frequency** ω, measured in radians per second

(rad/s).

Because one full turn is 2π radians, the two are forced by geometry to be related by $\omega = 2\pi f$. For 50 Hz that is $\omega \approx 314.16$ rad/s. Same physical fact, scaled by 2π: f is the right number to put on a nameplate, ω is the right number to put inside $\sin(\cdot)$.

2. The radian itself. A **radian** is a unit of angle defined geometrically rather than by an arbitrary slicing of the circle into 360 degrees. Its definition in one sentence:

> *One radian is the angle at the centre of a circle for which the arc length along the rim equals the radius of the circle.*

The recipe, with no diagram needed: take any circle of radius r, mark its centre O and a point A on the rim, then walk along the rim until the path you have walked has length exactly r and call the arrival point B. The wedge $A\,O\,B$ you have just constructed is, by definition, one radian. Because the full circumference is $2\pi r$, one full lap around the rim consumes $2\pi r / r = 2\pi$ of those radian-wedges — so a full turn is 2π rad, a half-turn is π rad, and a quarter-turn (90°) is $\pi/2$ rad.

3. Why radians, not degrees, in field formulas. Two reasons, both small but cumulative across the pamphlet's later chapters:

- Calculus is clean only in radians. The familiar $\frac{d}{dt}\sin(t) = \cos(t)$ is true *only* when t is in radians; if you wrote sin of a degree count, every derivative and integral would carry an extra $\pi/180$ conversion factor. Maxwell's equations, Faraday's law EMF $= -N\,d\Phi/dt$, and the alternator / CT / RLC formulas in later chapters *all assume* the angle inside the sine is in radians — which is why the pamphlet writes $\sin(\omega t)$ with ω in rad/s.
- Arc length and angle are interchangeable. A wheel of radius r turning by an angle θ radians moves its rim by exactly $s = r\,\theta$, no fudge factor. That is the natural language of every rotating machine in the pamphlet, from the alternator in §3.1 to the CT in §6.13.

So whenever a formula in this pamphlet contains ωt, $2\pi f$, sin, cos or $d\Phi/dt$, the angle inside is in radians, and the rate ω is in radians per second — that is all "angular frequency" is saying.

Chapter 2 — Why conduction is chemistry, not plumbing

In this chapter "current" gets a field-based definition, and we open a piece of copper up to see what is actually drifting inside it.

The pamphlet has now introduced the field-based picture in pieces — the "clarification on the word 'current' " after §1.3, the paradigm-change bullets in §1.5, and the one-line glossary entry in §11.1. This chapter gathers those pieces into a single canonical recap, so the word **current** can be used for the rest of the document without any of the water-pipe baggage still clinging to it — and then it opens the copper up to show what those *"free charges"* actually are, where they live, and why copper has them at all.

The chapter has three movements:

- **Movement A** (§§2.1–2.5) recaps what "current" means in the field picture. It is the content you may remember from the earlier drafts as §1.6 "Current, redefined", lifted here because it is really a deep-dive the rest of the pamphlet draws on, not a sub-topic of AC vs DC.
- **Movement B** (§§2.6–2.12) opens the copper up: atoms, ions, the electron gas, and the two forces that decide which electrons move and which ones stay. This is new material, and it is written assuming no prior chemistry.
- **Movement C** (§§2.13–2.14) collects the consequences and states, once and for all, why the plumbing analogy cannot be patched up.

2.1 The formal counting definition

Current is the rate at which electric charge crosses a chosen surface of a conductor:

Current — the formal counting definition [A]

$$I = \frac{dQ}{dt}$$

Where:

- I — current across the chosen surface — i.e., the collective **drift** of free charge carriers (in a copper wire, the free

electrons) past that surface — in amperes (A). A scalar. One ampere is one coulomb per second: $1\ \text{A} = 1\ \text{C/s} \approx 6.24 \times 10^{18}$ elementary charges per second.
- Q — net electric charge that has crossed the surface so far, in coulombs (C). A scalar.
- t — time, in seconds (s).
- dQ/dt — rate of change of that charge; i.e., how many coulombs cross the surface each second at this instant.

What it says. An "ampere" is nothing more than a count of charge per unit time crossing a surface you chose. There is no fluid and no flow in the hydraulic sense — only discrete charges going past a mathematical slice through the conductor, tallied up per second. This is the definition every ammeter in the world implements.

That is all the ampere is: a **count of charge per unit time** past a surface of your choosing. Nothing about "flow of electricity", nothing about a fluid. A number of coulombs per second.

2.2 The field-based one-liner

> **Current is the field's footprint on the free charges of a conductor.** It is what you see — and what you measure — where the field happens to overlap with a population of mobile charge carriers.

The order of causation is always the same:

```
source ΔV
    │
    ▼
E-field in/around the wire
    │
    ▼
Coulomb force on each free electron   (see §1.5)
    │
    ▼
electron drift through the copper
    │
    ▼
[ what we count and call "current" ]
```

The "Coulomb force" step is the formula $\vec{F} = q\vec{E}$ from the `formulabox` in §1.5: $\vec{E}$ is the electric-field vector at the electron's location (V/m), $\vec{F}$ is the force on that electron (N), and

q is the electron's own charge, $q = -e \approx -1.6 \times 10^{-19}$ C — negative, which is why $\vec{F}$ points *opposite* to $\vec{E}$ and the electron drifts "upstream" against the field.

Current is at the end of that chain. It is never the beginning.

The second half of this chapter (Movement B, §§2.6–2.12) is dedicated to the phrase *"free charges of a conductor"*: what those free charges physically are, where in the copper they live, and why copper has them in the first place.

2.3 Three consequences the old picture hid from you

1. **Current is the effect, the field is the cause.** The arrow of causation is field → electrons → current, never the other way round. The field does not "ride on" the current; the current exists **because** the field is acting. Switch the source off — the field collapses, the force on each electron vanishes, drift goes back to thermal random motion, and the ammeter reads zero.
2. **Current is local to conductors; the field is global.** Where there are no free charges — in the air between Live and Neutral, in the PVC insulation, in vacuum — there is **no current**, yet the field and its energy are still there. This is exactly why the cover figure and Figure 1.1 draw the big green Poynting arrows in the **empty space** between the conductors: that is where the energy is travelling, **and there is no current to be measured there because there is nothing mobile to push**.
3. **Current does not transport the energy.** This is the pamphlet's pillar restated in current-language. The 4 A of current in the supply cable of a 1 kW kettle is not "carrying" the kilowatt; it is the field's local effect on the cable's electrons — useful for sizing a fuse, useful for computing $\mathrm{I}^2\mathrm{R}$ losses, but the kilowatt itself travels by the field route **around** the cable and only meets the kettle's electrons inside the heating element, where the field does its heating work (chapter 7).

2.4 What each common instrument is actually telling you

Translating each ordinary meter reading into field language:

Instrument reads…	…and in field language that means
Voltmeter: 230 V between Hot and Neutral	The source is currently maintaining an $\vec{E}$-field of that strength in the space between those two conductors. The field exists right now, whether or not anything is plugged in.
Ammeter: 4 A through this cable	The field is succeeding in pushing electrons through this conductor at a rate of 4 coulombs per second. The cable is a conduit for the **drift**, not for the energy.
Wattmeter: 920 W into the load	The field around the cable is delivering 920 J/s to the load via the Poynting vector. The familiar $P = V \cdot I$ is the **outcome** of that field energy flow, not its mechanism.
Wire reads 0 A but is “live”	The source ΔV is there, and the $\vec{E}$-field around the conductor is set up — but there is no closed path for charges to drift along, so no current. The field is still there. Touching the wire completes the circuit through **you**, and the field instantly starts pushing charge through your body. See also the earlier discussion of why you can be shocked by a wire that an ammeter says is doing “nothing”.

The last row is the safety-critical one: a current reading of zero on a conductor does **not** mean the conductor is inert. It means only that **right now no path exists for the field to push charge along** — a condition you can turn into a lethal one by simply offering the field that path.

The crispest restatement

Old picture: *"Current is the flow of electricity through a wire."*

Field-based picture:

> **Current is the local rate at which the E-field is succeeding in pushing free charges across a chosen surface of a conductor. It is measured in amperes (coulombs per second), it exists only where the field meets mobile charge carriers, and it is the *symptom* of the field at work — never the carrier of the energy that work is delivering.**

Every formula you already know — Ohm's law, Kirchhoff's laws, $P = VI$, I^2R losses, fuse and breaker sizing — still works exactly as before. What has changed is only what the word **means** underneath: it stopped being the protagonist of the circuit and became a witness to what the field is doing.

2.5 A footnote on other media: copper is not the only conductor

In a copper wire, "current" and "electron drift" are literally the same phenomenon measured in different units (amperes vs electrons per second). In other media the charge carriers are different, but the definition of current is not:

- **Salt water / electrolysis / the wet inside of a human body** — carriers are positive and negative **ions** (Na^+, Cl^-, K^+, H^+, ...) drifting in opposite directions. The ammeter still reads dQ/dt across a surface.
- **Semiconductors** — carriers can be free electrons (n-type) or "holes" (p-type, effectively a moving absence of an electron), often both at once. Same definition.
- **Vacuum** (CRT, electron gun, particle accelerator) — carriers are free electrons travelling through empty space. The surface an ammeter integrates over is still a surface; the dQ/dt across it is still the current.

In every case the pattern is the same: **the field pushes whatever free charge is available, wherever that charge happens to be; and we call the resulting rate of charge flow across a chosen surface "current".**

Different media, different carriers — but the word **“ion”** has just shown up twice in the same sentence, and the pamphlet has not yet said what one is. Before the chapter goes any further, we have to open a copper atom up and actually look inside.

2.6 The atom, in sixty seconds

Everything the rest of this chapter says about copper rests on two prior facts about ordinary atoms. Both are gentle-grade chemistry; neither needs mathematics.

1. **An atom has a nucleus in the middle and electrons around it.** The nucleus is a tight cluster of **protons** (positive charge $+1$ each) and **neutrons** (no charge). The electrons (negative charge -1 each, roughly 1 836 times lighter than a proton) occupy the space around the nucleus. A neutral atom has *exactly as many electrons as protons*, so the charges cancel and the outside world sees a net charge of zero.
2. **The electrons are arranged in shells, not at random.** The Bohr picture draws those shells as concentric rings, labelled K, L, M, N, … from the innermost outward (we will meet a specific example — the copper atom — in Figure 2.1 of §2.9). Each shell can hold only a fixed number of electrons — K up to 2, L up to 8, M up to 18, N up to 32, and so on. Atoms fill the inner shells first, and the **outermost** occupied shell is, chemically, the one that does the interesting work.

For this pamphlet, only the outermost electron(s) matter. Everything interesting about electricity happens because **a handful of the outermost electrons in certain atoms are barely held on at all**.

2.7 What an ion is

An **ion** is an atom that is no longer electrically neutral — it has either lost or gained one or more electrons, so the number of electrons no longer matches the number of protons, and the atom ends up with a net charge.

- An atom that has **lost** one or more electrons has too few negative charges to balance the nucleus. It ends up with a net positive charge and is called a **cation**. Example: a sodium atom Na has 11 protons and 11 electrons (neutral); when it donates

its one outermost electron, it becomes Na^+ — a sodium cation with 11 protons but only 10 electrons. Still the same nucleus, still the same element, just missing one electron, and now carrying a net $+1$.

- An atom that has **gained** one or more electrons has too many negative charges. It ends up with a net negative charge and is called an **anion**. Example: a chlorine atom Cl that has grabbed one extra electron becomes Cl^-.

The two names sound similar but the thing to remember is simple: **positive ion (cation) = lost an electron; negative ion (anion) = gained one.** The rest of this chapter talks almost exclusively about cations, because the atoms inside a copper wire are cations — every one of them has donated its lone outer electron into shared space, leaving the nucleus plus the 28 tightly-held inner electrons behind as a Cu^+ cation.

2.8 Quantum mechanics, defused

The phrase *quantum mechanics* tends to make readers brace themselves for something alien. It should not. For the purposes of this pamphlet, a one-line definition suffices:

> **Quantum mechanics is just the mechanics of the atom — the same way automotive mechanics is just the mechanics of the car.**

It is not a parallel universe, not a philosophical movement, not a religion. It is the body of rules that describe how things behave once you zoom in small enough that ordinary billiard-ball mechanics stops giving the right answers. Atoms are that small. Electrons in particular are that small.

For the rest of this chapter, only **one** quantum fact is actually used:

> **Inside an atom, electrons can only sit at certain discrete energy levels — the shells from §2.6. The outermost occupied shell is the loosest; its electrons are the easiest to disturb, to pull out, or to share with neighbours. Inner-shell electrons are enormously more tightly bound and effectively inert.**

That single rule is what separates “metal” from “insulator” in the sections that follow. Everything else in quantum mechanics — wave-

functions, uncertainty, spin, entanglement — is irrelevant to the story of how a 230 V wall socket lights a lamp. It is fine to leave it unread.

That said — if you've got this far, odds are you're the curious kind. For you, §12.5 of the bibliography cites Pietro Greco's *Quanti. La straordinaria storia della meccanica quantistica* (Carocci, 2020): a warmly-written Italian-language narrative of where quantum mechanics actually came from, from Planck and Einstein's light-quantum through the Bohr–Einstein debates. It picks up more or less exactly where this pamphlet stops, and is a genuinely enlightening read about one of the great chapters in the history of science.

2.9 Copper specifically: 29 electrons, one of them loose

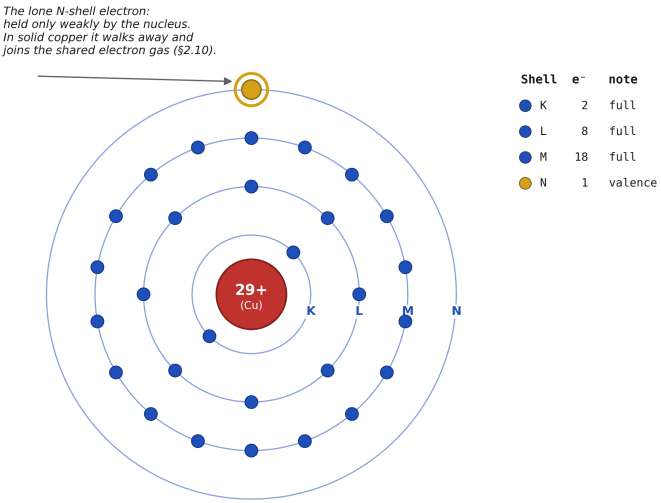


Figure 2.1 — Bohr picture of a copper atom. The nucleus (red) holds 29 protons and therefore attracts 29 electrons. They are arranged in four shells: $K = 2$, $L = 8$, $M = 18$, and $N = 1$. The lone electron in the outermost N shell is highlighted: it is held so weakly that, in a solid piece of copper, it walks away from its nucleus and joins the shared electron gas (§2.10).

A copper atom has **29 electrons** because its nucleus has 29 protons. The shells fill from the inside out, following the pattern intro-

duced in §2.6:

- **K shell:** 2 electrons (full).
- **L shell:** 8 electrons (full).
- **M shell:** 18 electrons (full).
- **N shell:** 1 electron. *(This is the loose one.)*

Those first three shells account for 28 of the 29 electrons, and they are held tightly: they sit close to the nucleus, where the full $+29$ nuclear pull is not yet diluted by the screening of other electrons. The **29th** electron — alone in the outermost N shell — is a different story. It is far from the nucleus, the inner 28 electrons partly hide that nucleus from it, and it is held so feebly that ordinary room-temperature thermal vibration is already enough to jostle it loose.

This is a **chemistry** fact, not a matter of electrical theory. It is why copper conducts and, say, glass does not. Glass's outermost electrons are tightly bound in covalent molecular arrangements; copper's lone outer electron is almost not bound at all.

2.10 When copper atoms pack into a solid: the lattice and the shared cloud

Pack roughly 10^{28} copper atoms per cubic metre together in a regular three-dimensional grid (this is what a piece of copper *is*: a very large crystal lattice, in wires usually assembled from many small grains fused at the boundaries), and something remarkable happens to those lone N-shell electrons.

Each atom's outer electron is already barely attached to its own nucleus. The instant the atom has neighbours — other copper atoms sitting only fractions of a nanometre away — the lone electron no longer has any strong reason to stay with any particular one. It walks off into the space *between* the atoms, joining the lone electrons from every other atom in the lattice.

What is left behind, at each atomic site, is a **Cu^{+} cation** — the nucleus plus the 28 tight inner electrons — bolted into place by the lattice's ionic forces. And sloshing through the gaps between those cations are roughly **10^{28} loose electrons per cubic metre**, none of them belonging to any particular atom, all of them shared by the whole metal.

This collective of loose electrons has a name. Chemists call it the **electron gas**; solid-state physicists more often call it the **electron**

sea or the **Fermi sea**; the everyday phrase *"shared cloud of electrons"* means exactly the same thing. In this pamphlet the three terms are interchangeable, and §2.11 is dedicated to saying what that gas/sea/cloud looks like up close.

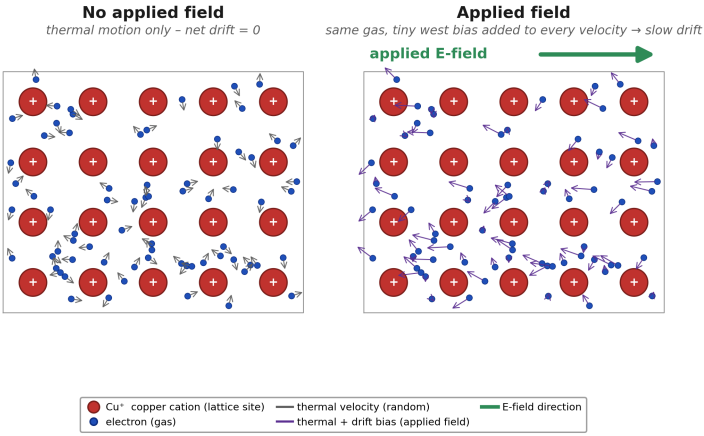


Figure 2.2 — A slice through solid copper. **Left panel:** no applied field. The red circles are the **Cu**$^+$ cations held in place by the crystal lattice; the blue dots are the electron gas — the 29th electron of every atom, walked away from its nucleus and shared with the whole metal. The grey arrows are random thermal velocities: net motion of the gas is zero. **Right panel:** a source has set up an external E-field pointing to the right. Every electron still zips around at thermal speed ($\sim 10^6$ m/s), but a tiny systematic bias is added to each velocity — and the whole gas drifts slowly in the opposite direction to $\vec{E}$ (electrons are negative, so they drift against the field). That slow collective drift is the current.

Reading Figure 2.2 the right way: current is a statistical bias, not a march

The two panels look more similar than different — and that is the point. What is *physically identical* between them is striking:

- The lattice is the same — the Cu^+ ions sit at the same sites in both panels.
- The electrons themselves occupy the same positions (the figure is drawn from a fixed random seed).
- Each individual electron is still zipping around at the same

thermal speed, roughly 10^6 m/s. *Thousands of kilometres per second*, every one of them, in both panels.

What is *different* is only the small vector added to every electron's velocity:

- **Left panel, no field.** Each arrow is a **velocity vector**: length = speed, direction = where the electron is currently moving. The arrows point every which way at random (isotropic distribution), so if you added them all head-to-tail, the sum would be **zero**. No net motion. No current.
- **Right panel, field on.** The thermal arrows are still random, but a tiny **westward vector** (same size, same direction) has been added to every single one of them. Arrows that were already pointing west got a bit longer; arrows that were pointing east got a bit shorter, or even flipped; arrows that were pointing north or south tilted slightly westward. The distribution is no longer isotropic — it leans west. Add all the arrows head-to-tail now and the sum is a small but non-zero **westward vector**.

That non-zero vector sum, scaled by charge and density, is exactly what the ammeter reads. It is the current. Nothing else in the figure changed.

Two things this picture is trying to burn into your mental model:

1. **Current is not a march.** The electrons are not lined up and moving together like a file of soldiers. Each one is still chaotic and fast. The "current" is a faint statistical bias — a fraction of a millimetre per second — riding on top of a hurricane of thermal motion at a *million* metres per second.
2. **The figure is grossly exaggerated.** In the real copper of a mains wire the drift bias is about 10^{-10} of the thermal speed (0.1 mm/s vs. 10^6 m/s). To scale, the two panels would look *pixel-identical*: you could not tell them apart even with a microscope trained on the electron gas itself. The figure turns the bias up by a factor of roughly a billion just so the eye can catch what the ammeter measures.

If all of that is held in mind, then **"the current flows westward at 0.1 mm/s"** stops being a hydraulic statement ("something is flowing in a pipe") and becomes what it really is: a statement about the **vector sum of the velocities of 10^{28} electrons per cubic metre** being slightly non-zero, in a particular direction, at this instant.

2.11 What does a “gas” or “cloud” inside a solid actually look like?

Saying “there is an electron gas inside copper” invites a wrong mental picture: invisible mist seeping through a pipe, or a fog of something vaguely fluid. The right picture is quite different, and the pamphlet owes the reader a concrete one before using the phrase any further.

The solid you see when you look at copper is the ion lattice. That is what gives copper its shape, its hardness, its colour, its weight, and the fact that a bar of it does not flow through your fingers. The Cu^{+} cations are bolted to their lattice sites and do not leave them; they only jiggle thermally, like billiard balls chained to their positions. Everything a bystander calls “the solidness of copper” is that ionic scaffolding.

The electron gas is not a literal gas — no wispy medium, no fluid. If you could shrink to atomic size and walk into the metal, what you would actually see is:

- A regular three-dimensional grid of fixed positive cores (the Cu^{+} ions), each about 0.1 nm across, spaced roughly 0.3 nm apart.
- An enormous number of tiny negative particles (the free electrons) zipping around in the empty space between those cores at thermal speeds of roughly 10^{6} m/s — *thousands of kilometres per second* — with no fixed home, bouncing off the ions about 10^{14} times per second.
- No coordinated direction. Each electron picks a new direction after each bounce; averaged over any macroscopic region, their net velocity is zero.

The word *cloud* or *gas* is used because the **statistical density** of those particles is spatially uniform — any two cubic nanometres deep inside the metal contain essentially the same count of electrons — which is exactly what a uniform gas looks like when sampled at scales above the particle spacing. Zoom in to individual-electron scale, and the “cloud” resolves into individual particles; zoom out, and it is a smooth electron density, just like a gas of any other kind.

How can a gas live inside something solid? Because the lattice is not space-filling. Atoms are overwhelmingly empty: a Cu^{+} ion is about 10^{-10} m across, but the lattice spacing is about 2.5×10^{-10} m, so the ions themselves occupy only a small fraction of

the total volume. The rest — the majority — is exactly the room the free electrons need to roam. The solid-ness of copper lives in the **ionic scaffolding**; the electron gas lives in the **interstitial space** between the scaffolding. Neither picture is whole without the other.

What does copper look like to an outside observer, then? Exactly what you already see: a uniform reddish solid, shiny when polished, heavy, malleable, and an outstanding electrical and thermal conductor. All four of those properties — colour, shine, heat conduction, electrical conduction — come from the *same* electron gas:

- **Colour and shine.** Visible light striking the surface interacts with the free electrons; they absorb and re-emit it almost elastically, and copper's particular electron density gives the wavelength-dependent response we call *"reddish metallic"*.
- **Heat conduction.** Drop one end of a copper rod into a flame and the free electrons carry thermal energy rapidly to the other end by bouncing off jiggling ions and passing their kinetic energy along. Insulators cannot do this — they have no electron gas to do the carrying.
- **Electrical conduction** — the whole subject of this pamphlet — is what happens when an external E-field adds a tiny systematic bias to the otherwise random motion of the same gas. §2.12 handles that last step in detail.

Nothing about this picture is exotic. It is just what happens when you pack 10^{28} atoms per cubic metre of an element whose outer electron is barely bound in the first place.

2.12 The two forces: bound electrons stay, free ones drift

At this point one question remains: why, exactly, does the field in §2.2 push the *free* electrons but leave the inner ones alone? The answer is that **every electron in copper feels two forces at all times**, and the outcome of the tug-of-war decides whether the electron drifts or stays put.

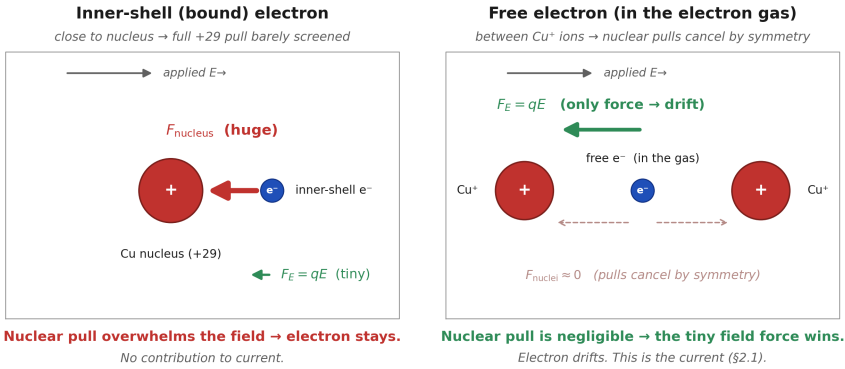


Figure 2.3 — The two forces on an electron in copper. **Left:** an inner-shell electron. The red arrow is the nuclear Coulomb attraction: enormous, because the electron sits close to the nucleus and very few other electrons stand between it and the full $+29$ charge. The green arrow is the force from the applied external E-field, $\vec{F} = q\vec{E}$: present, but smaller by many orders of magnitude. The inner electron stays home. **Right:** a free electron in the gas, sitting between two **Cu**$^{+}$ ions. Its original atom's nucleus is now far away, and whatever attraction remains from surrounding ions cancels out by symmetry. The green $\vec{F} = q\vec{E}$ arrow is the *only* unbalanced force left. The electron drifts. That drift is the current.

One reminder worth pinning down before we list the two forces: the Fermi sea already exists inside solid copper thanks to lattice formation (§2.10) — no battery required. The $\sim 10^{28}$ free electrons per cubic metre are already there, already zipping around at $\sim 10^{6}$ m/s, with or without anything plugged in. What follows is **not** a story about how a battery creates free electrons. It is a story about what happens to each electron — bound or free — once a battery *then* imposes an extra E-field force on top of the lattice that is already there.

For any electron in copper, the two forces are:

- **(i) The nuclear Coulomb pull.** The electrostatic attraction from the positive nucleus that, in the isolated atom, was responsible for holding the electron in its shell. Its strength depends sharply on how close the electron is to the nucleus and how effectively other electrons between it and the nucleus

screen (i.e. partly cancel out) that pull.

- **(ii) The applied E-field force $\vec{F} = q\vec{E}$.** The force on a charged particle q sitting in an external electric field $\vec{E}$, as introduced in §1.5. This is the force that the 9-V battery, the 230-V mains, or any other source imposes on every charge inside the wire.

Both forces act on every electron. Their *ratio* is what decides the electron's fate.

- For an **inner-shell electron** (the 28 tight ones per copper atom), force (i) is enormous — the electron is close to the full $+29$ nucleus, and the screening by the even tinier number of inner electrons between it and the nucleus is ineffective. Force (ii) from ordinary laboratory fields is smaller than (i) by *many* orders of magnitude. The inner electron barely notices the applied field. **It stays home, and it does not contribute to current at all.**
- For the **free electron** — the 29th electron, now living in the electron gas — force (i) is essentially zero. The electron sits outside any particular atom; it is approximately equidistant from several Cu^{+} ions on all sides, so their pulls cancel out by symmetry. Force (ii), although tiny in absolute terms, is now the *only* unbalanced force on the particle. **The free electron drifts. Collectively, those drifts are the current of §2.1.**

This is the last piece of the chapter's argument: **current happens, in a copper wire, because there is exactly one electron per atom that feels essentially no nuclear force and can therefore respond to the field, and 28 electrons per atom that feel far too much nuclear force to care about the field at all.** Neither half of that statement is something the plumbing analogy can represent.

2.13 Why the plumbing model cannot explain any of this

The plumbing picture — water through a pipe — survives, at best, at the level of "draw a closed loop and label the direction of flow". Everything in §§2.6–2.12 lies below that cartoon in a place the cartoon cannot reach:

- The plumbing picture has **no concept of bound vs free carriers.** Water molecules in a pipe are all alike; none of them is "bound to the pipe" while others are "free to flow". The distinction does not exist in a fluid. In a metal it is the *whole story*.
- The plumbing picture has **no ions and no electron donation.** Nothing in water "gives up" part of itself when the fluid enters the pipe. In copper, *every atom does exactly that* — it donates its outer electron into the shared cloud. Without that donation there is no electron gas, no conduction, no current.
- The plumbing picture has **no lattice and no shared cloud.** Pipe walls are passive; the fluid does not care about them. In a metal, the lattice of Cu^{+} cations is the very thing that makes the electron gas possible, and the cloud is what the E-field actually pushes.
- The plumbing picture **cannot say why copper conducts and glass does not.** Both are, in plumbing language, just "pipes", maybe pipes of different diameters. It has no handle on the fact that glass simply has no free electrons to push around.

Every chemistry fact in §§2.6–2.12 is something the water analogy is structurally incapable of representing. That is why the pamphlet discards it — not as a didactic preference, but because the analogy does not carry the information the rest of the story needs.

Which is why the chapter title is what it is.

2.14 Take-aways

- **Current is the field's footprint on the free charges of a conductor** — measured in amperes (coulombs per second), it exists only where the field meets mobile charge carriers, and it is the *symptom* of the field at work, never the carrier of the energy.
- **In copper, the "free charges" are a genuine electron gas** — one donated electron per atom, roughly 10^{28} of them per cubic metre, drifting in the empty space between the Cu^{+} ions of the lattice.
- **The solidness of copper is the ion lattice; the conductivity is the electron gas.** Both live inside the same cubic centimetre of metal at the same time — the lattice holds its shape, the gas carries the current.
- **An ion is an atom that has lost or gained one or more electrons.** Copper in a wire is a lattice of Cu^{+} cations (each

has lost one electron). Salt water conducts because it is full of Na^+ and Cl^- ions. Same definition, different medium.

- **Quantum mechanics is the mechanics of the atom** — nothing mystical. Only one quantum fact was used here: electrons sit in shells, and the outermost shell is the loosest.
- **Two forces decide every electron's fate in copper:** the nuclear Coulomb pull from its home nucleus, and the applied E-field force $\vec{F} = q\vec{E}$. Inner electrons: nuclear pull wins by many orders of magnitude → they stay home. Free electrons: nuclear pull ≈ 0 → the tiny field force wins → they drift.
- **The plumbing analogy cannot represent any of this** — it has no atoms, no shells, no ions, no lattice, no electron gas. That, not pedagogical preference, is why the pamphlet does not use it.

Chapter 3 — How AC is actually generated

The textbook “rotating coil in a magnet” is a teaching abstraction. Let’s look at what is physically sitting inside a real power plant.

3.1 The real machine: a synchronous alternator

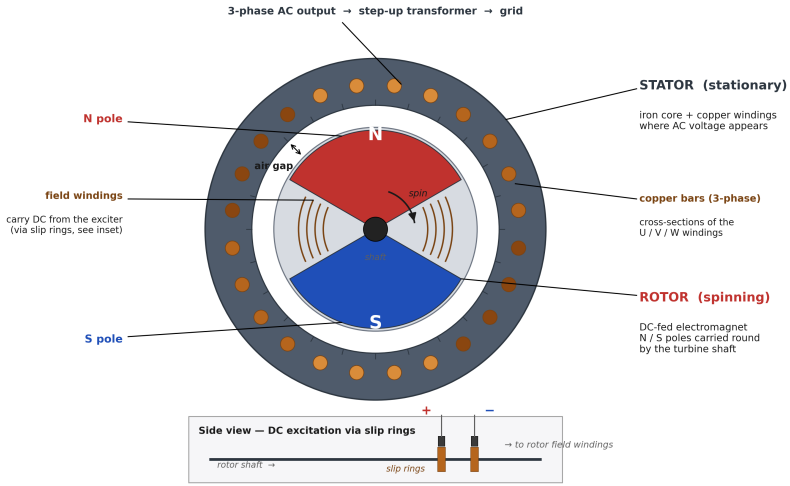


Figure 3.1 — End-on cross-section of a synchronous alternator. The outer grey annulus is the **stator**: a stationary iron core whose inside wall carries thick copper bars (small circles) arranged as three-phase (U/V/W) windings — these are the wires where AC voltage actually appears and leaves for the grid. The inner disc is the **rotor**: a DC-fed electromagnet with a North pole (red wedge) and a South pole (blue wedge), draped with field windings. The rotor is mechanically bolted to a turbine and spins inside the stator across a narrow *air gap*; as its N–S magnetic field sweeps past the stationary stator windings, it induces AC voltage in them. The inset at the bottom is a side view showing how DC excitation is fed onto the spinning rotor through a pair of *slip rings* riding on the shaft.

Every power plant — whether it burns coal, splits uranium, dams a river, or catches wind — ultimately spins a big rotating machine called a **synchronous alternator** (or just “the generator”). This is what physically produces the AC that eventually reaches a wall socket.

It has **two main parts**:

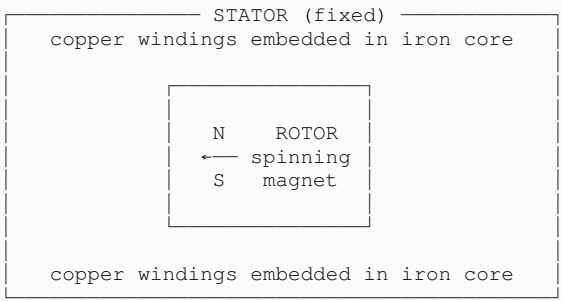
3.1.1 The rotor (the spinning part)

- A massive cylindrical shaft with **copper windings wrapped around it** — NOT a permanent magnet. *(A **winding** is simply a length of insulated wire wound many times into a coil — used wherever a magnetic field has to be built up by a current or picked up by one. From here on, every "motor winding", "transformer primary", or "stator winding" means the same thing: a coil of wire.)*
- A small DC current is fed into those rotor windings through **slip rings** (or via a brushless exciter). That DC current turns the rotor into a powerful **electromagnet** — a spinning electromagnet with a North pole on one side and a South pole on the other.
- The rotor is mechanically coupled to a **turbine** that forces it to rotate at a very precise speed.

3.1.2 The stator (the stationary part surrounding the rotor)

- A heavy iron cylinder with **thick copper bars/wires** laid into slots along its inside wall.
- These are the coils where AC voltage actually **appears** at the output terminals. Their wires run (through step-up transformers) all the way out to the transmission grid.
- The stator windings **do not move**. They just sit there while the magnet spins past them.

3.1.3 How they work together



As the N-S magnetic rotor spins, its magnetic field sweeps past the fixed stator windings → induces AC voltage in those fixed wires.

So the “rotating magnet inside a stationary coil” from the textbook is **literally what happens**, just at industrial scale:

- Rotor: typically **1–1.5 m** in diameter, **5–10 m** long, **50–200 tonnes**.
- Rotational speed: **3000 RPM** (for 50 Hz) or **3600 RPM** (for 60 Hz) in thermal plants.
- Hydroelectric generators spin more slowly (100–500 RPM) but have many more magnetic poles on the rotor to still produce 50/60 Hz.
- Output per unit: **100 MW to 1.8 GW**.

3.2 What spins the rotor?

The shaft of that big electromagnet-rotor is **mechanically bolted to a turbine**. Different power plants simply use different ways to spin that turbine:

Plant type	What spins the turbine
Coal / gas / nuclear	High-pressure **steam** (from boiling water) blasts through blades
Hydroelectric	Falling / flowing **water** pushes blades (a giant water wheel)
Wind	**Wind** turns the blades directly; a gearbox steps up the RPM
Gas turbine peaker	Hot **combustion gases** (same principle as a jet engine)

In every case: mechanical rotation → spinning electromagnet → AC voltage induced in stationary copper windings → transmission lines → a wall socket somewhere.

3.3 Why the rotor is an electromagnet, not a permanent magnet

Three practical reasons:

1. **Excitation control.** By varying the DC current fed into the rotor windings, the grid operator can vary the strength of the magnet, and therefore the voltage the alternator produces.

This is how generators follow load and share reactive power across the grid.

2. **Scale.** No permanent magnet would be strong enough, durable enough, or cheap enough in the sizes required.
3. **Fault response.** An electromagnet can be de-energised quickly in an emergency — a permanent magnet cannot.

3.4 Why 50 Hz means exactly what you think

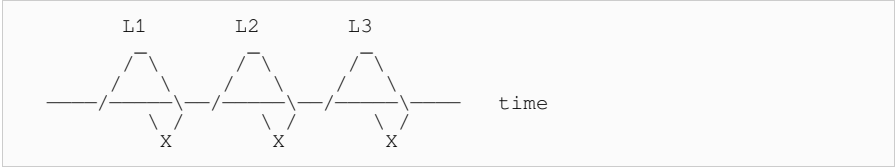
Hertz (Hz) = cycles per second. At 50 Hz the voltage completes 50 full cycles every second. Each cycle lasts **1 / 50 = 20 ms**. Polarity reverses **100 times per second** (two reversals per cycle), and zero-crossings occur at the same rate.

Region	Frequency	Cycle period	Polarity reversals/sec
Europe, most of Asia, Africa, Australia	**50 Hz**	20 ms	100
North America, western Japan, parts of South America	**60 Hz**	16.67 ms	120

This frequency is **rigidly** locked to how fast the generators at the power plant spin; the whole national grid stays in sync. A 2-pole generator must spin at exactly 3000 RPM to produce 50 Hz (3000 / 60 = 50 rev/s); a 4-pole generator at 1500 RPM; and so on.

3.5 Three-phase in one page

Large alternators have **three** sets of stator windings, spaced 120° apart. As the rotor spins, each set sees the magnetic field peak at a different moment, producing three sine waves offset by a third of a cycle:



Three-phase matters because:

- **Rotating magnetic field.** Three-phase current in three windings at 120° produces a smoothly rotating magnetic field — which is exactly what makes induction motors work, without brushes or permanent magnets.
- **Higher power density** per unit of copper and iron.
- **Lower pulsation** in power delivery (sum of three sine waves = constant average power).
- **No need for a return conductor** in balanced loads (the three phases cancel out at the star point).

The “three wires plus neutral” you see on large buildings and industrial feeds is the same three-phase system, split down to single-phase (one of L1, L2, L3 versus N) for domestic circuits.

3.6 Take-aways

- Real alternators have a spinning **electromagnet** (DC-excited coil), not a permanent magnet.
- The stator is where AC actually appears; the rotor just provides the moving magnetic field.
- Frequency is locked to rotational speed × pole pairs. That is why the grid frequency is so stable.
- Three-phase is the natural output of a three-winding stator and is what makes industrial motors simple.

Chapter 4 — Mains in the wall: Live, Neutral, Earth

4.1 Where the three wires come from

At the neighbourhood distribution transformer, the high-voltage three-phase supply is stepped down to the domestic level. On the low-voltage side, one end of each secondary winding is bonded to earth; that bonded point becomes **Neutral**. The other end of each winding is a **Live** (phase) conductor.

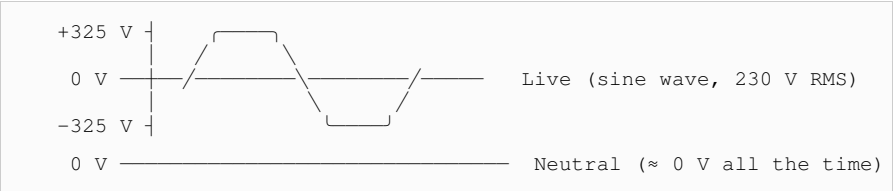
For a single-phase socket, two conductors run to the consumer:

- **Live (L / Hot / Phase)** — oscillates between about **+325 V** and **−325 V** relative to Neutral (230 V RMS in Europe).
- **Neutral (N)** — held at ~0 V, bonded to earth at the transformer.
- **Earth / Ground (PE)** — safety conductor, normally carries no current, tied to the building's grounding system.

4.2 What actually happens over one cycle (20 ms at 50 Hz)

Neutral — the boring one. Sits at roughly 0 V relative to the building's earth for the entire cycle. In a healthy installation it never oscillates far from 0 V even when heavy current flows through it.

Live — the dramatic one. Swings from 0 V up to +325 V, back down through 0 V, all the way to −325 V, and back to 0 V, continuously, 50 times per second. During each half-cycle the **current** reverses direction in the circuit.



4.3 Why touching Live shocks you but touching Neutral (usually) doesn't

You are always at earth potential — you stand on the floor, which is bonded to earth. When you touch a conductor, what matters is the potential difference between that conductor and earth.

- **Touching Neutral.** Neutral sits at ~0 V. The potential difference between your hand and earth is ~0 V. No shock.
- **Touching Live.** Live swings between ±325 V relative to earth. Enough potential difference to drive current through your body — anything above ~50 V RMS is dangerous.

4.3.1 Important caveats about "safe" Neutral

- **Neutral is not exactly 0 V.** Under load it develops a few volts of "neutral-to-earth" offset because of wiring resistance. That is normal but not trivial — the offset grows with the current through the Neutral conductor.
- **A broken Neutral is extremely dangerous.** If the Neutral conductor breaks upstream of the load, the "Neutral" wire inside the appliance can float up to close to Live potential. The loss of the safety assumption that Neutral ≈ 0 V is the single most hazardous fault in typical low-voltage installations.
- **Wrong wiring.** If a socket is wired with L and N swapped, the device's switch and fuse end up on the wrong wire. A lamp can *look* off yet have its "neutral" side still live to earth.
- **IT (Isolated Terra) installations.** In some industrial and medical installations there is no Neutral bonded to earth at all; both poles sit symmetrically around earth. Everything written above needs rethinking in those contexts.

4.4 The Earth (PE) wire

Earth (Protective Earth, PE) is the safety conductor. Normally it carries no current. Its job is to:

1. Provide a low-impedance return path if Live ever touches an accessible metal part, so that the **overcurrent protection** (breaker or fuse) trips.
2. Keep all accessible metal surfaces at the same potential (equipotential bonding), so nobody can get a shock by touching two different earthed surfaces at once.

3. Give Residual Current Devices (RCDs / GFCIs) something to detect — a fault current leaking away from L/N into PE triggers an RCD in tens of milliseconds.

4.4.1 Summary of voltages relative to earth (healthy installation)

Conductor	Typical potential vs earth
Live (L)	−325 V … +325 V, sinusoidal
Neutral (N)	~0 V (may have a small offset under load)
Earth (PE)	0 V by definition (the reference)

4.5 Grounding architectures in one paragraph each

- **TT.** Consumer installation has its own local earth electrode; Neutral is earthed at the transformer only. Common in older rural installations. Always requires an RCD.
- **TN-C / TN-S / TN-C-S.** The transformer's earthed neutral is carried along with the supply. In **TN-S**, Neutral and PE are separate conductors all the way. In **TN-C**, they share a single combined "PEN" conductor. **TN-C-S** combines PEN on the supply side and splits into separate N and PE at the consumer.
- **IT.** No point of the supply is solidly earthed. A single ground fault does not drive fault current; the system keeps running, and a monitoring device raises an alarm so the fault can be cleared before a second one occurs. Used where continuity of supply is critical (hospitals, some industrial processes).

4.6 IEC colour codes

Modern IEC conductor colours (IEC 60446 / 60757 family):

Function	Colour
Live (L1)	**Brown**
Live (L2)	**Black**
Live (L3)	**Grey**
Neutral (N)	**Blue**
Protective Earth	**Green/Yellow striped**

Older installations (and regional variants) use different colours — always verify with a tester rather than trusting insulation colour alone.

4.7 Transients on the mains: lightning, network surges, SPDs

(One of the three hands-on sections flagged in the Introduction — a place where the field picture earns its keep.)

Everything §§4.1–4.6 said is a **steady-state** description: L/N/PE sit at their nominal potentials, the sine wave repeats, life is predictable. That picture is accurate roughly 99.9 % of the time. The remaining 0.1 % is what this section is about: **analog transients on the mains** (§9.1 vocabulary), short-lived but often violent voltage excursions that the ±325 V peak of §4.2 entirely fails to bound.

The word "transient" here means the same thing as in §9.1 and §6.10 — a short, non-steady voltage event — but applied to the mains it stops being a neat ns-scale CMOS (Complementary Metal-Oxide-Semiconductor) curiosity and becomes a several-kilovolt, several-kiloampere hazard that destroys electronics and, occasionally, buildings.

4.7.1 Where transients come from

Three families dominate:

- **Lightning surges.** A direct strike is rare; a strike to a nearby pole, line, tree, or grounded structure is not. The strike's current (tens of kA, rise time in the µs range) couples into the low-voltage network by direct injection, by magnetic induction into overhead lines, and by the earth-potential rise of the local ground. At the consumer's entry panel the surge is typically already attenuated to a few kV peak, but still carries hundreds to thousands of joules.
- **Network switching surges.** Events elsewhere on the distribution network: a utility capacitor bank switching in or out, a large contactor operating at a neighbouring factory, a line autoreclose after a fault. Amplitudes are smaller than lightning (hundreds of V to a few kV), but the events are far more frequent.
- **Local switching surges.** Large loads inside the same installation switching off — motors, transformers, contactors — cou-

ple their own inductive kicks (same mechanism as §6.11 below, different scale) onto the shared L/N. The very same event that blows a relay coil open at the load side radiates a transient upstream onto the mains.

4.7.2 How surges travel: differential-mode and common-mode

A mains transient rarely appears as a clean "voltage between L and N". It appears in two simultaneous forms, because §4.4's PE conductor is a third player:

- **Differential-mode (L–N).** The surge shows up as an extra voltage between Live and Neutral. The load sees it directly.
- **Common-mode (L+N against PE).** L and N both rise together against the protective earth. Any equipment whose signal ground is bonded to PE sees this as an apparent ground shift, sometimes by kilovolts.

Surge-protection strategies have to address **both** modes, which is why proper protection is never a single component.

4.7.3 The standard test waveform: IEC 61000-4-5

Equipment immunity to mains surges is tested to **IEC 61000-4-5** (International Electrotechnical Commission standard 61000-4-5; see §11.6). Its canonical waveform is the *combination wave*:

- **1.2 / 50 µs** open-circuit voltage pulse: rises in 1.2 µs, decays to half amplitude in 50 µs.
- **8 / 20 µs** short-circuit current pulse: rises in 8 µs, decays to half in 20 µs.

Typical installation-class numbers for consumer mains entry are **6 kV / 3 kA** (Installation Category III in the IEC classification); subpanels and end-equipment see scaled-down versions (4 kV / 2 kA, 2 kV / 1 kA, down to a few hundred volts at a socket behind several metres of wiring and one SPD).

4.7.4 SPDs (Surge Protective Devices)

A **Surge Protective Device (SPD)** is a component whose job is to **clamp the voltage during a surge by dumping the surge energy through itself**. In the field-first language of this pamphlet:

during the surge an enormous energy pulse arrives along the line in its E- and B-field; if it has nowhere to deposit that energy, it deposits it in whatever reaches its dielectric strength first (the insulation of a motor winding, the gate oxide of an MCU, the air gap of a relay contact). An SPD is a sacrificial low-impedance path that *volunteers* to swallow the pulse instead.

Real SPDs combine three technologies because no single technology is simultaneously fast, high-energy, and low-leakage:

- **MOV (Metal-Oxide Varistor).** A ceramic disc (typically zinc-oxide based) whose bulk resistance collapses sharply above a threshold voltage. Physics: the ceramic is a sintered aggregate of tiny conducting grains with insulating grain boundaries; above the threshold those boundaries break down in parallel and the disc conducts. Good energy absorption, moderate speed (ns–tens of ns), *wears out* with every surge (the grain boundaries degrade) and so must be paired with a thermal disconnector / fuse.
- **GDT (Gas-Discharge Tube).** A sealed glass or ceramic capsule containing a low-pressure inert gas between two electrodes. Above the strike voltage the gas ionises (plasma, §2.5) and the tube conducts as a near-short-circuit. Very high energy capacity, very low leakage when idle, but *slow to strike* (hundreds of ns) and prone to *follow current*: after the surge, the ionised gas continues to conduct the 50 Hz mains until the next zero-crossing (§6.10) breaks the arc.
- **Mains-class TVS (Transient-Voltage-Suppression) diode.** A silicon avalanche diode engineered for short, high-current pulses (same family as the smaller §6.11 TVS, rated for mains duty). Clamps in *picoseconds*, but limited total joules per device.

A well-designed SPD uses them in combination: **GDT** to dump the gross energy, **MOV** to absorb the middle range, **TVS** to clamp the fast leading edge before the MOV has even started conducting. The three sit in parallel across L–N and between each line and PE, so both differential-mode and common-mode surges are clamped.

4.7.5 Installation classes (IEC 61643-11)

SPDs are classified by *where* in the installation they sit:

- **Type 1 (Class I).** Designed to withstand a direct-lightning

residual. Installed at the service entrance, immediately after the main breaker. Uses a heavy GDT or spark-gap as its backbone.
- **Type 2 (Class II).** For switching surges and residual lightning past a Type 1. Installed at the distribution board. Predominantly MOV-based.
- **Type 3 (Class III).** For fine protection of sensitive end equipment: a desktop, a PLC, a CNC controller. Installed at the socket or inside the equipment. Fast MOV + TVS combination.

The three classes **must be coordinated** — a Type 3 alone cannot handle a direct lightning hit; a Type 1 alone leaves residual voltages that still destroy electronics.

4.7.6 Why §4.3's "Neutral ≈ 0 V" is a statistical statement

In §4.3 we said touching Neutral is usually safe because N sits near earth potential. That statement holds on time-averages, not during a surge. During a nearby lightning event, **the whole building's N and PE can rise by several kV** against remote earth for a few tens of microseconds. Locally (hand-on-N, feet-on-PE-bonded-floor) you still see a small difference, because N and PE move together in common-mode. That same common-mode rise is, however, fatal to any device whose reference is held at *remote* earth via a data cable (network cable to a distant switch, long sensor cable to a different building) — which is exactly the scenario the §9 EMI chapter and §9.8's 4–20 mA loop example prepare the reader for.

4.7.7 The load-side cousin of a mains surge

A mains surge is one class of analog transient. The *other* class, smaller in amplitude but far more common inside a control cabinet, is the one you **create yourself** every time you switch off an inductive load (relay coil, solenoid valve, contactor, motor winding). That is §6.11's territory. The underlying physics is identical — energy stored in a field has nowhere to go and drives the voltage across the breaking switch up until something clamps it — but the scale is different: tens of kV and kJ on the mains; tens to hundreds of V and mJ at the load side.

Both deserve a clamp. Both are tamed by giving the stored energy a deliberate, non-destructive place to land. The taxonomy in §6.11 and the SPD stack in this section are two sides of the same coin.

4.8 Take-aways

- Live is the dangerous one because it is the only conductor whose potential swings far from earth.
- Neutral is "cold" *only* because it is bonded to earth at the transformer; break that bond and it becomes just as dangerous as Live.
- Earth is a safety conductor, not a working conductor — current in PE during normal operation is a fault.
- **The mains is a battlefield for transients.** Lightning surges, network-switching surges, and local inductive kicks (§6.11) routinely push voltages on L, N, and PE well beyond the ±325 V steady-state peak. SPDs (MOV + GDT + TVS, in coordinated Type 1 / 2 / 3 stages) exist because the "Neutral ≈ 0 V" of §4.3 is a steady-state statement, not a transient-free one.

Chapter 5 — RMS (Root Mean Square) vs Peak

5.1 Both numbers are true, they measure different things

When a multimeter shows “230 V AC”, that is the **RMS (Root Mean Square)** value — the DC-equivalent voltage that would deliver the same average power to a resistive load. The actual waveform is a sine swinging much higher at the peaks.

The relationship for a pure sine wave is:

$$V_{peak} = V_{rms} \cdot \sqrt{2} \approx V_{rms} \cdot 1.414$$
$$V_{rms} = V_{peak} / \sqrt{2} \approx V_{peak} \cdot 0.707$$

So:

- Europe: $230\ V \times 1.414 \approx 325\ V\ peak$ → swings between **+325 V and −325 V**.
- United States: $120\ V \times 1.414 \approx 170\ V\ peak$ → swings between **+170 V and −170 V**.

Both numbers describe the same waveform; they are just different ways of characterising it.

5.2 Why can’t we just use the simple average?

Take the plain mathematical average of a sine wave over one full cycle and you get **zero**, because the positive and negative halves cancel. That is obviously useless for describing how “strong” the AC is.

5.2.1 Attempt: average of the absolute value

Rectifying first (taking |v(t)|) and then averaging gives a non-zero number, but it does not correspond to the actual heating power of the waveform.

5.2.2 The RMS (Root Mean Square) procedure

1. **Square** the instantaneous voltage at every moment (squaring makes negative values positive *and* weights larger values more heavily, because power goes like V^2).

2. Take the **Mean** (average) of those squared values over one full cycle.
3. Take the square **Root** of that mean.

For a pure sine wave, this conveniently works out to `Vpeak / √2`. The **result represents the equivalent DC voltage that would dissipate the same heat in a resistor** — which is what actually matters for energy and power calculations.

5.3 A worked example

Take a pure sine with peak `Vpeak = 10 V`, angular frequency $\omega = 2\pi/T$, and instantaneous voltage

$$V(t) \;=\; V_{\text{peak}}\,\sin(\omega t)$$

For each of the three candidate "averages" from §5.2 we will do the integral in full, then plug in `Vpeak = 10 V` at the end. The point is not the numbers — it is seeing *why* plain averaging fails, why the absolute-value average misses the physics, and why squaring is the step that finally lands on the heat-dissipation-equivalent DC voltage.

5.3.1 Plain average over one cycle

The arithmetic mean of $V(t)$ over one full period T:

$$\begin{aligned}
\langle V\rangle \;&=\; \frac{1}{T}\int_0^T V_{\text{peak}}\sin(\omega t)\,dt \\
&=\; \frac{V_{\text{peak}}}{T}\left[-\frac{\cos(\omega t)}{\omega}\right]_0^T \\
&=\; \frac{V_{\text{peak}}}{\omega T}\bigl(1-\cos(\omega T)\bigr) \\
&=\; \frac{V_{\text{peak}}}{2\pi}\bigl(1-\cos(2\pi)\bigr) \;=\; \frac{V_{\text{peak}}}{2\pi}\cdot 0 \;=\; 0\ \text{V}.
\end{aligned}$$

The positive and negative half-cycles cancel exactly. For our `Vpeak = 10 V` sine: `⟨V⟩ = 0 V`. Mathematically true, physically useless

— it tells us nothing about how “strong” the waveform is.

5.3.2 Average of the absolute value

Rectify first (take $|V(t)|$), then average. Because $|\sin(\omega t)|$ has half the period of $\sin(\omega t)$, the integral over $[0, T]$ is just twice the integral over $[0, T/2]$:

$$\begin{aligned}\langle |V| \rangle &= \frac{1}{T}\int_0^T \left|V_{\text{peak}} \sin(\omega t)\right| dt = \frac{2}{T}\int_0^{T/2} V_{\text{peak}} \sin(\omega t)\, dt \\ &= \frac{2V_{\text{peak}}}{T}\left[-\frac{\cos(\omega t)}{\omega}\right]_0^{T/2} \\ &= \frac{2V_{\text{peak}}}{\omega T}(1 - \cos\pi) = \frac{2V_{\text{peak}}}{2\pi}\cdot 2 = \frac{2\,V_{\text{peak}}}{\pi}.\end{aligned}$$

Plugging in `Vpeak = 10 V`: $\langle |V| \rangle = 20/\pi \approx 6.37$ V. Non-zero — so at least the number reacts to the waveform’s amplitude — but this is *not* the DC-equivalent for power. Power goes like V^2, not like V, and that quadratic weighting is what the next average finally captures.

5.3.3 RMS (Root Mean Square)

Square first, then average, then take the square root. The squaring step is what maps each instantaneous voltage onto the instantaneous heating power of a resistor ($P \propto V^2$), so the mean of the squares is proportional to the mean *power*.

Using the identity $\sin^2 x = (1 - \cos 2x)/2$:

$$V_{\text{rms}}^2 \;=\; \frac{1}{T}\int_0^T V_{\text{peak}}^2 \sin^2(\omega t)\,dt \;=\; \frac{V_{\text{peak}}^2}{T}\int_0^T \frac{1-\cos(2\omega t)}{2}\,dt$$

$$= \frac{V_{\text{peak}}^2}{2T}\left[\, t \;-\; \frac{\sin(2\omega t)}{2\omega}\,\right]_0^T$$

$$= \frac{V_{\text{peak}}^2}{2T}\left(T \;-\; \tfrac{\sin(4\pi)}{2\omega}\right) \;=\; \frac{V_{\text{peak}}^2}{2T}\cdot T \;=\; \frac{V_{\text{peak}}^2}{2}.$$

Taking the square root:

$$V_{\text{rms}} \;=\; \frac{V_{\text{peak}}}{\sqrt{2}}.$$

For `Vpeak = 10 V`: $V_{\text{rms}} = 10/\sqrt{2} \approx 7.07$ V.

That 7.07 V is the only one of the three numbers with a direct physical meaning: a DC source of 7.07 V feeding the same resistor would dissipate *exactly* the same average power as this 10 V-peak sine. That is what RMS (Root Mean Square) delivers and what the other two averages cannot.

5.4 Why “equivalent DC” matters — two resistors

Compare two resistors, each with the same resistance `R`:

- **Resistor A:** connected to 230 V DC. Power dissipated = `V²/R = 230²/R`.
- **Resistor B:** connected to 230 V RMS AC (which is 325 V peak). The instantaneous power swings between 0 and `325²/R`, but the **average** power over a cycle comes out to exactly `230²/R` — the same as resistor A.

That is the entire point of RMS: it gives you a number that lets you use the DC power formulas `P = V · I` and `P = V² / R` directly, without worrying about the waveform.

5.5 What would you measure in a home wall socket?

- **Multimeter set to “AC Voltage”** → 230 V (RMS).
- **Oscilloscope, properly isolated** → a sine wave with peaks around +325 V and −325 V, period 20 ms.

5.6 Practical consequences — why this matters

5.6.1 Capacitor voltage ratings

Capacitors must be rated for **peak**, not RMS. A 230 V RMS line will apply **325 V peak** across a capacitor; a “300 V” cap would be marginal.

5.6.2 Rectifier output voltage

After a bridge rectifier and smoothing capacitor, the DC output sits close to the AC **peak** (minus a diode drop or two). A 230 V AC input produces roughly 320 V DC across the smoothing capacitor, not 230 V.

5.6.3 Insulation and creepage

Insulation must withstand the peak voltage, plus transient over-voltages. Creepage distances on PCBs (Printed Circuit Boards — the gap measured along a surface) and clearance distances (the gap measured through air) are specified for peak, with additional margin for pollution degree.

5.6.4 Shock severity

Electric shock danger depends on RMS current through the body for steady-state AC — the body’s impedance acts more or less like a resistor, and the heating and muscular effects scale with RMS. Peak-limited transients (e.g., electrostatic discharge) behave differently.

5.6.5 “400 V three-phase” really means

On a three-phase supply, “400 V” is the **line-to-line RMS** voltage. Each phase is 230 V RMS to neutral; two phases 120° apart give a

line-to-line voltage of $\sqrt{3} \cdot 230 \approx 400$ V RMS, peaking at $400 \cdot \sqrt{2} \approx 566$ V.

5.7 The 220 V vs 230 V history

The EU harmonised on **230 V RMS +10 % / −6 %** in 1988, which means anything from roughly 216 V to 253 V RMS is compliant. Old Italian / European nameplates saying “220 V” refer to the pre-harmonisation value; new equipment is still designed to work across the wider band.

5.8 Quick cheat sheet for a pure sine

Quantity	Relation
Peak	V_{peak}
Peak-to-peak	$2 \cdot V_{peak}$
RMS	$V_{peak} / \sqrt{2} \approx 0.707 \cdot V_{peak}$
Average of |v(t)|	$2 \cdot V_{peak} / \pi \approx 0.637 \cdot V_{peak}$
Form factor (RMS / avg)	$\pi / (2 \cdot \sqrt{2}) \approx 1.11$
Crest factor (peak / RMS)	$\sqrt{2} \approx 1.414$

5.9 Take-aways

- RMS is the DC-equivalent voltage for power; that is literally why we use it instead of peak or average.
- $V_{peak} = V_{rms} \cdot \sqrt{2}$ for a pure sine; a 230 V outlet really swings to ±325 V.
- Capacitor voltage ratings, insulation clearances, rectifier outputs, and three-phase line-to-line figures all hinge on peak, not RMS.

Chapter 6 — Voltage vs current: what they really are

6.1 The core confusion

The single most common source of confusion in elementary electricity is conflating **voltage** and **current**. They are very different things, and the fact that a 230 V outlet can deliver anything from 0 A to 16 A (at the same voltage!) is surprising only until the distinction is clear.

This chapter makes the distinction in physical terms. Before we do, we have to clear away a piece of classroom furniture that actively makes it harder: the water-in-a-pipe analogy.

6.2 Revisiting: why the water analogy is discarded here

Every introductory textbook hands you this table:

Electrical quantity	Water analogue
Voltage (V)	Pressure in a pipe
Current (I)	Flow rate through the pipe
Resistance (R)	Restriction / valve narrowing

We are not going to use it. The analogy gets the Ohm's-law arithmetic right by accident, at the price of cementing the **single deepest misconception in electricity**: that something is flowing *inside the wire* the way water flows inside a pipe. Nothing of the sort is happening. The wires contain slowly-jiggling electrons that barely move at all, and the energy is carried by the **electromagnetic field in the space around the wires** (chapter 7). The pipe picture gets this backwards, and every downstream concept — inductance, transformers, antennas, shielding, EMI, VFD behaviour — becomes harder to understand while the pipe is still in your head.

So we set the analogy aside here and not pick it up again. The rest of this chapter describes voltage, current, and resistance as what they physically are, and recovers Ohm's law from those.

6.3 What voltage actually is

Voltage between two points is a potential difference set up by an electric field. It is the work, per unit charge, that the field would do in moving a test charge from one point to the other. It is a property of the field configuration, not of anything that is flowing.

The mains transformer maintains, via its windings and iron core, an electric field configuration between Live and Neutral such that the work per coulomb between them is 230 V (RMS). This exists **whether or not anything is plugged in**. An open socket is sitting at 230 V with no load connected; nothing is "pent up behind a closed valve", because there is no valve and there is no reservoir — there is a field.

In AC the field oscillates at 50 Hz, so the voltage is a sign-alternating function of time rather than a constant level.

6.4 What current actually is

Current is just a count: the number of charge carriers crossing a given cross-section of conductor per second. In copper carrying 10 A DC, the free electrons drift at roughly **0.1 mm/s** — slower than a snail. They do not race from the socket to the appliance. A single electron, at that drift speed, would take **hours** to travel one metre.

In AC at 50 Hz the electrons do not even drift from one end of the wire to the other. They **oscillate back and forth** by a fraction of a millimetre, 50 times per second, around a fixed mean position. Overlaid on this tiny orderly motion is a much larger random thermal motion (on the order of kilometres per second in arbitrary directions). The "AC current" you read on a clamp meter is this slow collective oscillation — a tiny perturbation on the random thermal sea of electrons already present in the copper.

Crucially: current does **not** carry energy along with it. The electrons are the medium, not the message (chapter 7). They are already there, everywhere in the wire, before the circuit is even closed.

6.5 Ohm's law, without the valve

Given a potential difference V across a component of resistance R, the steady-state drift rate of the electrons in that component is:

Ohm's law [A]

$$I = \frac{V}{R}$$

Where:

- I — current through the component, in amperes (A). A scalar.
- V — potential difference across the component, in volts (V). A scalar. ("Potential difference" is just the line integral of the E-field from one terminal to the other.)
- R — resistance of the component, in ohms (Ω). A scalar set by the material and geometry: longer, thinner, dirtier or hotter conductors have higher R.

What it says. For a well-behaved ("ohmic") component at a steady temperature, the current that flows when you apply a voltage is simply voltage divided by resistance. It is empirical — Georg Ohm measured it before anyone understood why it held — and the drift-in-copper picture of §§1.5 and 6.4 reproduces it exactly. Ohm's law is *not* the supply "pushing I amperes through the wire": it is a relation between three physical quantities — the source's field, the material's resistance, and the electron gas's response.

This is Ohm's law. It is a relation between three real physical quantities — a field-induced potential difference, a material-defined resistance, and the equilibrium drift rate of the charge carriers. It is **not** the supply "pushing I amperes through the wire".

The consequence most often mis-read as "the supply delivers the current" is that V and I are decoupled at the socket:

- 10 W LED on 230 V → about **43 mA** drifts through its leads.
- 2 kW kettle on 230 V → about **8.7 A** drifts through its leads.

Same socket. Same 230 V. Different loads → different I. The socket is not choosing what current to provide; the load (plus the voltage) is what determines what drift rate the electrons settle at.

6.5.1 Resistance vs reactance — the two ways AC pushes back

Ohm's law as written above is the DC story. As soon as the source alternates, a load can push back against the current in **two physically different** ways, and the pamphlet has been freely using the vocabulary of both — power factor (§5.5), inductive vs capacitive loads (§6.9), flyback (§6.11), capacitive and inductive coupling (§9.4–9.5) — without ever naming the umbrella concept. Naming it once, here, is free leverage for everything that follows.

The umbrella concept is **impedance**. Impedance has two components, both measured in ohms, but utterly different physically:

Resistance R	Converts electrical energy into *heat* and dissipates it into the environment. Defined for DC and AC alike. Set by material, geometry and temperature. *Energy out* of the circuit, irreversibly. The "I^2R" of §6.6 is exactly this.
Reactance X	Stores energy in a field and returns it the next half-cycle. *Dissipates nothing.* Only meaningful in AC, because it relies on the source alternating. Two flavours: $X_L = 2\pi f L$ — inductive: energy parked in the *magnetic* field of a coil, as $\frac{1}{2}LI^2$. Larger at higher frequency. $X_C = \dfrac{1}{2\pi f C}$ — capacitive: energy parked in the *electric* field of a capacitor, as $\frac{1}{2}CV^2$. Smaller at higher frequency.

The total opposition the AC sees is the **impedance**

$$Z = \sqrt{R^2 + X^2} \qquad \text{(in ohms; vector sum at the mains frequency } f \text{ in question).}$$

— the word "impedance" the pamphlet has been using since §4 without ever defining now has a one-line definition. The phase angle φ between voltage and current is just $\tan\varphi = X/R$, and the **power factor** $\cos\varphi$ already used in §5.5 (and in the $P = VI\cos\varphi$ form of §6.6) is exactly the cosine of that angle. The "three kinds of load" taxonomy in §6.9 is the same idea seen from a different angle: a purely resistive load has $X = 0$ and $\cos\varphi = 1$; a purely inductive

or capacitive load has $R = 0$ and $\cos\varphi = 0$, dissipating no real power even while the current sloshes through it.

A clean place to see *the absence* of added reactance be the whole point of a device — rather than a parasitic — is the **current transformer**, dissected in §6.13.

6.6 Where the energy actually is

Neither voltage nor current *is* the energy. The electrical power delivered to a load is given by any of the four equivalent forms below — choose whichever two of V, I, R you happen to know:

Electrical power — four equivalent forms [W]

$$P = V \cdot I = \frac{V^2}{R} = I^2 R = V \cdot I \cdot \cos\varphi$$

Where:

- P — electrical power delivered to (or dissipated in) the component, in watts (W). A scalar. One watt is one joule per second.
- V — potential difference across the component, in volts (V). For AC this is the RMS voltage (§5).
- I — current through the component, in amperes (A). For AC this is the RMS current.
- R — resistance of the component, in ohms (Ω). Used when the load is purely resistive.
- $\cos\varphi$ — *power factor* (§5.5); the cosine of the phase angle φ between the AC voltage and current. 1.0 for a purely resistive load; < 1 when inductance or capacitance shifts current out of phase with voltage.

What it says. Power is the rate at which the field delivers energy to the load, in joules per second. The first form, $P = V\,I$, is the general DC / in-phase-AC expression. The two substitutions $P = V^2/R$ and $P = I^2R$ come from Ohm's law and are handy when computing wire losses (I^2R) or heating-element ratings (V^2/R). The fourth form, $P = V\,I\,\cos\varphi$, is the real ("active") power drawn by an AC load whose current is phase-shifted from its voltage: only the in-phase component of the current actually delivers energy. All four forms are the *outcome* of the field's energy flow — the mechanism itself is the Poynting vector of

chapter 7.

That power does **not** travel inside the wire with the electrons. It travels through the electromagnetic field in the space *around* the two conductors — direction and density given by the Poynting vector $\vec{S} = (1/\mu_0)\,\vec{E} \times \vec{B}$ (chapter 7). The wires are guides that shape the field; the field is what carries the energy.

Heat dissipated in a resistive element comes from the field doing work on the jiggling electrons inside the element, which then collide with the crystal lattice. The energy arrived at the element via the field around the wires, not via electrons carrying joules in their pockets.

6.7 Worked example: the electric kettle

- Rated power: 2000 W at 230 V RMS.
- Current (drift rate): `I = P / V = 2000 / 230 ≈ 8.7 A` RMS.
- Effective resistance of the heating element: `R = V / I ≈ 26.4 Ω`.

During each half-cycle the voltage peaks at ±325 V, so the peak current through the element is `Ipeak = 325 / 26.4 ≈ 12.3 A`, and the peak instantaneous power is `325 × 12.3 ≈ 4000 W`. Averaged over the full cycle, the RMS of that current is `12.3 / √2 ≈ 8.7 A`, and the mean power is `230 × 8.7 = 2000 W` — consistent.

Every one of those numbers is what the old pipe picture would give you too. The difference is **what they mean physically**: an oscillating electric field between Live and Neutral is depositing 2000 joules per second into the heating element by doing work on its electrons — while the electrons themselves jiggle in place by a fraction of a millimetre.

6.8 Where the ±325 V “went”

The ±325 V is always present at the socket, regardless of what is plugged in. Plug nothing in → the 325 V peaks are still on the Live conductor; no current flows because there is no conductive path

between Live and Neutral. Plug in the kettle → the 325 V peaks drive 12.3 A peaks through the 26.4 Ω element, heating it.

Nothing is “released” or “let through”. The voltage is always there; the current is a response that appears the instant a conductive path is provided between Live and Neutral.

6.9 Current alternates together with voltage

A **resistive load** is a device whose current is proportional to the voltage across it and follows that voltage *instantly*, with no delay — a heating element, an incandescent bulb, or a fixed resistor. Inductive loads (motors, transformer windings) and capacitive loads (the input filter of a switched-mode power supply) don’t have that property; their current is shifted in time relative to the voltage. We take the resistive case first because it is the simplest, and we come back to the other two in the callout below.

Three kinds of load — a one-page taxonomy

Every load downstream will be classified as one of three idealised types. Real devices are usually a mix of two or three, but the classification is what every power-factor discussion (§5.5), every EMI analysis (Chapter 9), and every AC/DC contrast below silently assumes.

- **Resistive load.** Current is proportional to voltage and follows it *instantly* — no time delay. All the energy the field delivers is dissipated as heat or light. Examples: heating elements, incandescent bulbs, and **fixed resistors** (the plain two-terminal components, each with a single designed resistance value, found by the thousand on every circuit board). In a resistive load the current and voltage waveforms cross zero together and peak together — they are said to be **in phase**.
- **Inductive load.** Anything whose dominant part is a **winding** (see §3.1) — a coil of wire, usually around an iron core. The magnetic field the current builds up inside that coil stores energy and *resists changes* in the current, so the current arrives late: the current waveform is shifted in time *behind* the voltage waveform. That time shift is called **phase lag**. Examples: motor windings, transformer primaries, solenoids, inductive sensors, relay coils.

- **Capacitive load.** Anything whose dominant part is two conductors separated by a thin insulator. The electric field across that insulator stores energy, and as a result the *current* is ahead of the voltage — the current rises before the voltage has caught up (the mirror image of the inductive case). Examples: the input filter capacitors of an **SMPS** front end (every laptop charger, LED driver, PC PSU, control-cabinet 24 V supply), and — to a lesser extent — long shielded cables seen from the source.

Power factor (§5.5) is simply the accounting of how much of a real load is resistive versus the two others.

AC current is called "alternating" because in each half-cycle the electron drift reverses direction. In-phase on a purely resistive load:

```
Positive half (0 – 10 ms):
  Live: +0 → +325 → +0 V
  electron drift:  Live side → element → Neutral side

Negative half (10 – 20 ms):
  Live: +0 → −325 → +0 V
  electron drift:  Neutral side → element → Live side
```

The kettle element does not care which way the drift goes — it heats up either way. That is why resistive loads are "direction-blind" and why AC has no problem powering them.

6.10 Zero-crossings, 100 times a second

Because the sine wave passes through zero twice per cycle, at 50 Hz there are **100 zero-crossings per second**. At those instants the instantaneous voltage *and* current are momentarily zero (for a resistive load). This is why:

- Thyristors and triacs can self-commutate on AC (they turn off at the zero-crossing automatically).
- Zero-crossing optocouplers / relays switch at those instants to minimise inrush and EMI (electromagnetic interference).
- Incandescent bulbs flicker at 100 Hz (perceptually usually invisible, but measurable).

(Switching at the zero crossing is still an analog transient — just a very small one. See §9.1 for why the analog/digital distinction is an interpretation layer, not a physics layer.)

The five terms packed into those three bullets — *thyristor*, *triac*, *optocoupler / relay*, *inrush*, *EMI* — all recur through Chapters 9 and 10, so the next box gives a one-page unpacking while the reader is already thinking about the AC waveform crossing zero.

Zero-crossing switching glossary — what is behind the three bullets above

- **Thyristor (SCR).** A latching semiconductor switch. A brief gate pulse turns it on; once conducting, it stays on until the current through it falls naturally to zero. On a DC supply that never happens, so thyristors are mostly an AC component: they *self-commutate* off at the next zero-crossing, and the control circuit only has to decide *when* in the next half-cycle to re-trigger them. Classic uses: phase-controlled heater power, soft-starters, crowbar overvoltage protection.
- **Triac.** Electrically a pair of back-to-back thyristors in one package, so it conducts on *both* halves of the AC waveform. Still self-commutates at each zero-crossing. Classic uses: domestic dimmers, small-motor speed controls, low-cost solid-state relays for resistive AC loads.
- **Optocoupler / solid-state relay (SSR).** A LED shining across a tiny air gap onto a photodetector that drives a semiconductor switch (triac, MOSFET, IGBT). The galvanic isolation is done by light, not by magnetic coupling, so the control side is electrically independent of the power side. A *zero-crossing* variant adds a small detector that waits for the AC line to cross zero before closing the output switch, trimming the switching transient at the cost of a few milliseconds of latency.
- **Inrush current.** The short, high-amplitude current surge a load draws the instant it is energised, before the steady-state current settles. Cold tungsten filaments, transformer magnetising current, SMPS input capacitors (§6.9, capacitive load) and motor windings are the usual culprits; peaks of *5 to 20 times* the nominal current over a few ms to tens of ms are routine. Switching at the zero-crossing helps the resistive cases because the applied voltage starts from zero and ramps up; it helps less with transformer inrush, which depends on the residual-flux state at turn-on rather than on the instantaneous line voltage.

- **EMI** (**electromagnetic interference**). Unwanted coupling of one circuit's switching transients into another. A hard switch at non-zero voltage creates fast dV/dt and dI/dt edges that radiate both as a near-field magnetic disturbance (inductive coupling into nearby signal loops) and as a capacitively-coupled displacement current on the ground reference. Zero-crossing switching keeps both dV and dI small at the instant of transition, so less energy goes into radiated modes. Full treatment in Chapter 9.

In short: zero-crossing switching does *not* reduce steady-state current or power — it only removes the worst transient of the *turn-on* instant, which is exactly where inrush and EMI are born.

6.11 Switching an inductive load: the flyback spike and its taxonomy

(One of the three hands-on sections flagged in the Introduction — a place where the field picture earns its keep.)

§6.9 defined the inductive load: a coil of wire whose magnetic field stores energy and *resists changes* in the current through it. §6.10 dealt with the AC case, where current and voltage simply slide past each other with a phase lag. This section deals with the far harsher event every control cabinet lives through — a DC inductive load being **switched off**.

For any low-voltage control cabinet (Arduino / OpenPLC outputs driving relay coils, solenoid valves, motor contactors) this is the single most common way hardware dies. The fix is cheap and universal; understanding *why* it works is what this section is for.

6.11.1 The question §6.9 stopped short of

An inductor is energy stored in the magnetic field *B* wrapping around its winding. The energy is

$$W = \tfrac{1}{2} L I^2$$

where L is the inductance (henries) and I the current. For a typical 24 V relay coil, $L \approx 100\,\text{mH}$ and $I \approx 40\,\text{mA}$, so $W \approx \tfrac{1}{2} \cdot$

$0.1 \cdot (0.04)^2 \approx 80\,\mu$J — not much as energies go, but it has to go *somewhere* in a very short time.

Now open the switch. The circuit is broken, so I must fall from 40 mA to 0. The inductor's defining law,

$$V_L \;=\; L\,\frac{dI}{dt},$$

says: *however fast you force dI/dt, that is how big V_L will have to be*. A MOSFET switches off in roughly 200 ns. If the full 40 mA had to collapse in that time,

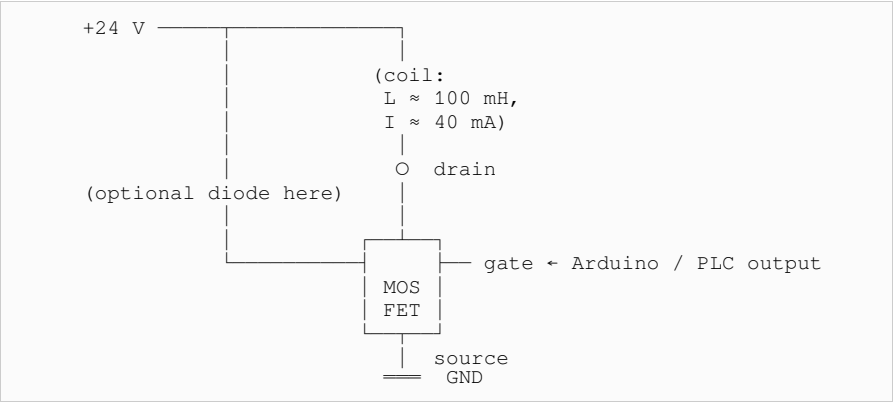
$$V_L \;\approx\; 0.1\,\text{H} \cdot \frac{0.040\,\text{A}}{200 \times 10^{-9}\,\text{s}} \;=\; 20{,}000\,\text{V}.$$

Of course 20 kV does not appear on a 24 V rail, because *something breaks down first*. That "something" is whatever part of the circuit has the lowest breakdown voltage: the MOSFET's drain-source junction (typically 40–60 V for a logic-level part) avalanches, the relay contacts arc, the insulation of a nearby winding punctures. In every case, **the coil's stored energy ends up dissipated inside a semiconductor or across an arc**, where it was never meant to go.

Field-first restatement: the collapsing magnetic flux has to induce a voltage large enough to keep the current circulating until the energy is bled off. If you give it no legal circuit to circulate in, it will carve out an illegal one — through the device's dielectric (§8.7), through an arc, through your gate oxide.

6.11.2 Worked example — 24 V relay coil on a MOSFET

Concrete scenario: an Arduino or OpenPLC digital output drives the gate of a logic-level N-channel **MOSFET (Metal-Oxide-Semiconductor Field-Effect Transistor)** whose drain sinks a 24 V relay coil to ground.



Case A — no clamp. The gate falls, the MOSFET turns off. The coil's 80 µJ has no legal path. The drain voltage shoots up until the MOSFET's avalanche breakdown conducts (≈ 60 V for a typical logic-level part), at which point *all 80 µJ is dissipated inside the MOSFET silicon*. A few thousand cycles of that and the gate oxide fatigues, the threshold voltage drifts, and the device fails — often intermittently, which is the worst kind of failure.

Case B — flyback diode. A plain silicon diode is wired in parallel with the coil, cathode to +24 V, anode to the MOSFET drain. During normal operation the diode is reverse-biased and does nothing. At turn-off, the collapsing flux drives the current through the closed loop {coil → diode → coil}, because that is now the lowest-voltage path. The diode drop is $V_f \approx 0.7\,\text{V}$, so the drain rises only to $24 + 0.7 \approx 25\,\text{V}$. The coil current decays through its own DC resistance (plus the diode's) over a few ms, dissipating the 80 µJ harmlessly as heat in the copper.

The field-first punchline: **the flyback diode is a return path for the collapsing magnetic flux**. It does not "absorb" the surge in any mysterious way; it simply gives the field's energy a quiet loop to circulate in until the coil's own resistance drains it.

6.11.3 Taxonomy — the full inductive-load-protection family

Different scenarios (DC vs AC, fast vs slow release, bidirectional vs polarised, cabinet vs mains) call for different clamps. The entire family is below; they are all variations on the same principle of giving stored energy a legal path.

Clamps for inductive loads — pick the right one for the right load

- **Flyback (freewheeling) diode.** A standard silicon diode across a *DC* inductive load, reverse-polarised in normal operation. Clamps the coil voltage to about $V_f \approx 0.7\,\text{V}$ during turn-off; the coil current decays through the closed coil + diode loop in a few ms. *Cheap and universal; the default choice for relay coils, solenoid valves, small DC motors driven non-PWM.* Downside: the slow decay means the relay armature releases slowly (typically 2–10 ms extra drop-out delay), which matters for safety-interlock relays where release time is a certified specification.
- **Zener-clamped diode.** A Zener diode (named after Clarence Zener, 1934, who first described the reverse-bias avalanche mechanism) in series with a regular diode, again across the coil but with the Zener facing the collapsing-flux direction. The clamp voltage is now the Zener voltage plus V_f — e.g., 39 V for a 24 V coil — so the current decays far faster than with a plain diode, at the cost of a bigger voltage spike on the switching device. *Used where relay release time matters (safety interlocks) and the MOSFET can tolerate the higher V_{DS}.*
- **TVS (Transient-Voltage-Suppression) diode.** A silicon avalanche diode engineered specifically for short, high-current transients. Electrically it is similar to a Zener but with a much higher peak-pulse-power rating and a sub-nanosecond clamp time. Available as:
  - *Unidirectional* (one-way, like a Zener) for DC clamping;
  - *Bidirectional* (two TVS back-to-back, one package) for AC or bipolar signals.

  Typical homes: across MOSFET drain-source for fast clamping, across communication-bus lines (RS-485, CAN, USB) for ESD (electrostatic discharge) protection, across relay-coil terminals when a flyback diode is not fast enough.
- **RC snubber.** A resistor R in series with a capacitor C, placed either across the switching contacts or across the load. Physics: the capacitor *absorbs* the $L\,dI/dt$ pulse because its own voltage cannot change instantaneously (§8.7), so it ramps up instead of the switching node; the resistor dissipates the stored capacitor energy on each

switching cycle instead of returning it to the circuit as a ring. Typical values: $R \approx 100\,\Omega$, $C \approx 100$ nF for a 230 V contactor. *Used on AC inductive loads (contactor coils on 230 V, triac-switched motors) where a unidirectional diode would short-circuit half the AC cycle, and on* contact *sides to protect the contacts themselves from arc-erosion.*

- **MOV (Metal-Oxide Varistor).** Same component family as the mains SPD of §4.7, in smaller disc sizes (S07, S10, S14). Clamps bidirectionally at a threshold voltage; good energy absorption; wears out. *Used across AC contactor coils, as a cheap bidirectional alternative to a TVS or RC snubber, and as a first-stage clamp on input terminals exposed to field wiring.*
- **Back-to-back Zeners.** Two Zener diodes in series, anodes or cathodes tied together, across a bipolar line. Clamps symmetrically at $V_{\text{Zener}} + V_f$ in either polarity. A poor-man's bidirectional TVS; still used where the exact Zener voltage is hard to find as a TVS part number.
- **GDT (Gas-Discharge Tube).** Rarely used at cabinet-load scale (slow to strike, follow-current prone) but noted for completeness: the same component used at the mains entry in §4.7 exists in miniature for telecom-line protection.

6.11.4 Choosing one — a short decision guide

A practical decision procedure for a control-cabinet designer:

- **DC load, release time does not matter** (general-purpose relay, solenoid valve, pilot lamp with inrush) → **plain flyback diode**. Cheapest, most universal.
- **DC load, release time matters** (safety interlock relay, fast solenoid, high-cycle valve) → **Zener-clamped diode**, sized so that $V_{\text{Zener}} + V_f$ is below the switching device's safe V_{DS} rating.
- **AC load** (230 V contactor coil, AC solenoid) → **RC snubber** (resistor–capacitor) or **bidirectional MOV** (metal-oxide varistor) across the coil; optionally another RC snubber across the switching contacts to save them from arc-erosion.
- **Signal-line ESD** (electrostatic discharge — differential bus, sensor input, header pin exposed to field wiring) → **bidirectional TVS** (transient-voltage-suppression diode).
- **Mains entry** → the full coordinated SPD stack of §4.7 (Type 1

– Type 2 + Type 3).

6.11.5 One honest caveat

A flyback diode protects only the **drain-source** path of the MOSFET against the coil's stored magnetic energy. It does not protect:

- the **gate** (the MOSFET's control pin) against electrostatic discharge or ringing on long gate-drive leads (needs a gate resistor, sometimes a gate-to-source TVS);
- the gate-driver circuit against **Miller-effect** (parasitic drain-to-gate capacitance couples a fast drain transient back into the gate, risking spurious turn-on) coupling during fast drain transients (needs a gate-drive resistor sized for the specific part);
- the coil itself against **inrush voltage** on *turn-on* (the flyback diode is reverse-biased then and does nothing).

Full MOSFET protection in a noisy cabinet is therefore always a *combination*: flyback diode on the coil, gate resistor, local decoupling, and (in hostile environments) a TVS across gate-to-source and drain-to-source. Cross-reference: §9 on cabinet wiring, §9.9 on jumpers and configuration headers (which are themselves a protection weak-spot).

6.12 Why this feels counter-intuitive

The elementary picture of electrons as little marbles shooting through a pipe from + to −, with voltage as "pressure", is almost universally taught. It lets students crank through Ohm's-law problems without the physics. The cost is high: once the pipe picture is in place, very little else about electricity makes sense.

- Why does a coil induce a voltage when current changes? (Because the field around it is changing — not because water is sloshing.)
- How does a transformer transfer power with no galvanic connection between primary and secondary? (Because the field links them — not because water jumps across the gap.)
- Why can a VFD's switching radiate interference far from the cable? (Because its fields fluctuate in the space around the cable — not because water leaks from the pipe.)

Three concepts drop in there in quick succession — *galvanic connection*, *primary and secondary*, *VFD switching radiated interference* — and each of them is load-bearing for Chapters 7 through 10, so the rest of this section is a figure and a grey box unpacking them.

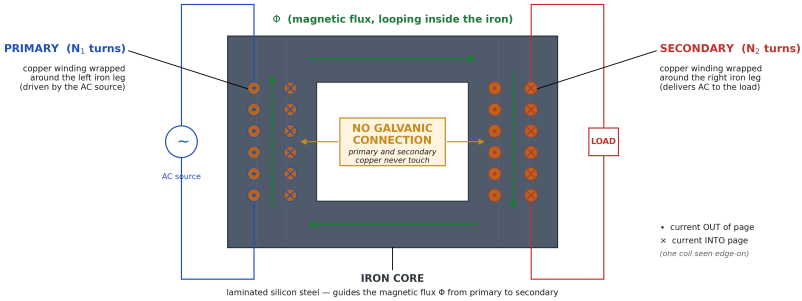


Figure 6.1 — Schematic cross-section of a single-window *core-type transformer*, the textbook example of energy transfer *without* a galvanic connection. The grey frame is the laminated silicon-steel **iron core**: a closed loop with a rectangular window cut out of it. Around the left vertical leg is wrapped the **primary winding** (N_1 turns, shown in blue): each coil turn is drawn edge-on as a pair of copper cross-sections, one on each side of the leg, with a dot / cross telling you which way the current is flowing at that instant. Around the right vertical leg is wrapped the **secondary winding** (N_2 turns, in red). The two windings are electrically **insulated from each other and from the core** — the enamel on the copper, plus the air gap in the window, means the primary copper and the secondary copper *never touch*, directly or through the iron. The green arrows trace the magnetic flux Φ that the primary's alternating current drives around the closed iron path; as Φ rises and falls it threads through the secondary coil and, by Faraday's law, induces an AC voltage there that drives current through the load. The only thing crossing from primary to secondary is the *field inside the iron* — never a wire.

Galvanic, primary / secondary, VFD interference — what those three bullets are saying

- **Galvanic connection.** Engineering shorthand for "direct conductive metal-to-metal path between two terminals, so a steady current can flow from one to the other without passing through anything but copper." A plain wire is a galvanic connection; a soldered joint is one; a screw ter-

minal is one. A transformer winding-to-winding is *not* one — see Figure 6.1: the energy crosses the gap as a magnetic flux through the iron core, not as a current through a wire. *Galvanic isolation* is the deliberate absence of such a path, and it is the backbone of mains safety (isolation transformers), ground-loop suppression (signal transformers, opto-isolated inputs, optocouplers from §6.10), and high-voltage measurement (current transformers — see §6.13 for a worked example). Capacitive and magnetic coupling both move *energy* across a galvanic gap; only a wire (or a short-circuit fault) re-introduces a galvanic path.

- **Primary and secondary.** The two windings of any magnetically-coupled two-port device. The **primary** is the side driven by the source; the **secondary** is the side that delivers power to the load. For an ideal transformer with turns counts N_1 and N_2, the voltage ratio follows the turns ratio ($V_2/V_1 = N_2/N_1$) and, because power is conserved, the current ratio is the inverse ($I_2/I_1 = N_1/N_2$): a step-up in voltage is a step-down in current, and vice versa. The same primary / secondary language applies to *current transformers* (CTs) used for metering and protection, to *ignition coils* (huge turns ratio, secondary delivers the spark), to *wireless chargers* (primary in the pad, secondary in the phone), and to every mains *SMPS* (primary on the rectified line, secondary on the galvanically-isolated low-voltage rail). In every case the two windings share only the magnetic field, never a wire.
- **VFD switching, radiated interference.** A **VFD** (variable-frequency drive, sometimes called an *inverter drive*) is the electronic box that feeds a three-phase motor from a DC bus through fast MOSFET / IGBT switching at a carrier frequency of typically 4–16 kHz. Each switching edge is a step of several hundred volts completed in tens of nanoseconds, so the slew rate is on the order of kV$/\mu$s. A rectangular voltage edge that fast has Fourier content extending well into the MHz band, and the long motor cable between the VFD and the motor becomes an *antenna* for it — the fluctuating near-fields around the cable couple both inductively into nearby signal loops (§9.3) and capacitively onto the ground reference (§9.4), radiating EMI far beyond the physical cable route. That is why VFD installations man-

date *shielded motor cable with 360-degree shield termination at both ends*, a dedicated grounding conductor routed alongside the three phase conductors, and — for the worst cases — an output *dV/dt* or sine-wave filter that slows the edges before they ever reach the cable. Chapter 10 returns to this as a worked near-field measurement example.

All three concepts have the same root: *energy is carried by the field in the space around the wires*, not by the wires themselves. A transformer exploits that to move power across an air gap; a VFD's cable does the same thing accidentally, and the installer's job is to contain it.

Replace the pipe with the picture developed in chapter 7 — field carries the energy, electrons are the medium, wires are guides — and all of these become obvious instead of puzzling.

6.13 The current transformer — galvanic isolation in working clothes

§6.12 named current transformers (CTs) as one of three textbook examples of galvanic isolation, alongside isolation transformers and optocouplers. This subsection zooms into that one example and shows what *"energy is carried by the field, not by the wires"* looks like in working clothes. The running example is deliberately a banal domestic one: an ordinary civil installation with a single-phase 230 V mains feeding a ~6.6 kW resistive load — say, an electric water heater on its own line — drawing about 29 A continuous on its phase conductor whenever the heater is on. Whatever the motivation for caring about that current — a high-priority / low-priority load-shedding controller in the main panel, an energy meter, an over-current trip, a simple display — the installation needs to *read* the current without breaking the conductor that carries it, and without putting the measuring electronics at the conductor's potential. The CT is the device that does exactly that.

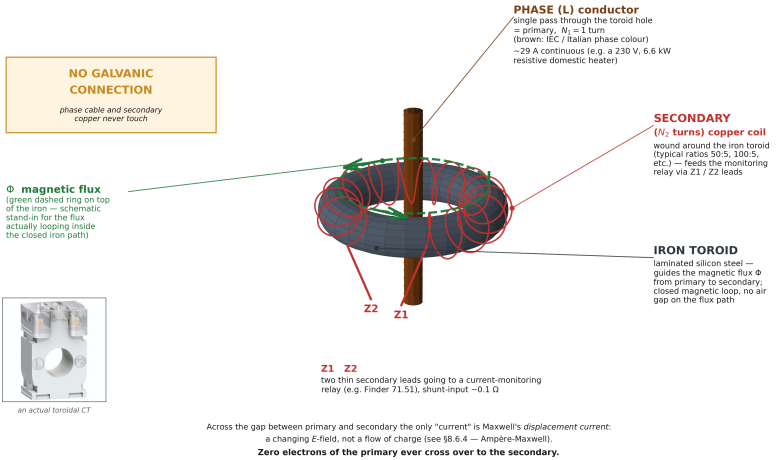


Figure 6.2 — 3D perspective view of a *toroidal current transformer* (CT). The grey donut is the laminated silicon-steel **iron toroid**: a closed magnetic loop with no air gap on the flux path, the standard geometry for CTs. The thick brown vertical line is the **primary** — in this case a single pass of the load's phase (L) conductor straight through the donut hole, so $N_1 = 1$ turn, no winding to make: the cable to be measured *is* the primary. Brown is chosen deliberately — in IEC / Italian wiring it is the colour reserved for phase, while blue is reserved for neutral, so the diagram already tells the installer which conductor of the 230 V line goes through the toroid. Around the toroid is wound the **secondary** (red helix), with two leads labelled **Z1 / Z2** that exit the front of the toroid and go to a small current-monitoring relay (e.g. a Finder 71.51) sitting in front of the assembly — not drawn here, to keep the figure focused on the magnetics. The dashed green ring with arrowheads, hovering above the iron, is a schematic stand-in for the magnetic flux Φ that the primary's alternating current drives *inside* the closed iron path; as Φ rises and falls it threads through the secondary turns and, by Faraday's law, induces a small secondary current proportional to I_1/N_2 (typical CT ratios are 50:5 or 100:5). The small inset bottom-left is a photo of an actual toroidal CT (Schneider METSECT5CC004 family), so the reader can mentally map the schematic onto the real product on a DIN-rail. The primary cable and the secondary copper *never touch* — not directly, not through the iron, not through the relay, not through ground.

The whole subsection condenses into a single sentence:

> *In a current transformer, primary and secondary share no electron and no copper — only the iron core's flux Φ and, in the air-and-insulation gap surrounding it, a tiny displacement current. It is the cleanest possible demonstration that energy is carried by the field, not the wires.*

Why the CT does not "load" the circuit it measures. A common question on first reading: if the CT couples to the primary, isn't it *adding a load* on top of the heater? The reactance vocabulary of §6.5.1 settles it cleanly. The primary is a single pass of cable through the toroid — exactly one turn — so its self-inductance through the core is on the order of microhenries, and at 50 Hz its inductive reactance $X_L = 2\pi f L$ is on the order of micro-ohms. The secondary is effectively short-circuited by the relay's input shunt (low-impedance current input, $\sim 0.1\,\Omega$ at full scale), so the impedance reflected back through the $1 : N_2$ ratio is shrunk by another factor of N_2^2 — pushing the impedance that the load on the line "sees" back from the CT into the *nano*-ohm range. The heater does not even know it is being measured; the CT contributes neither a real (R) drop nor a meaningful reactive (X) drop. In the language of §6.5.1, the device works *because* the reactance it adds to the circuit it measures is negligible, by construction.

Why the "no electron crosses" claim is exactly right. Across the gap between primary cable and secondary turns there is no copper, no solder, no carbon, no electrolyte — nothing for an electron to drift through. What links the two sides is the magnetic flux Φ inside the iron (Faraday, §8.6.3) and, in the air outside it, the *displacement current* term Maxwell added to Ampère's law precisely because there are no charges crossing (§8.6.4): a changing electric field that behaves like a current for the purposes of the field equations, but is not a flow of charge. The CT is the cleanest possible illustration of why Maxwell needed that term — without it, a real device that visibly transfers power across an empty gap would violate the original Ampère law every cycle, twice.

6.14 Take-aways

- **Discard the water-in-a-pipe analogy.** It is the single biggest source of misconception in basic electricity.
- **Voltage** is a potential difference set up by an electric field between two points. It exists whether or not current flows.
- **Current** is the slow collective drift (DC) or tiny oscillation (AC) of the electrons that are already everywhere in the conductor.
- **Neither voltage nor current *is* the energy.** Energy is carried by the electromagnetic field in the space around the wires; the wires are guides (chapter 7).
- **Ohm's law `I = V / R`** is a relation between three physical

quantities — not a valve equation for an imaginary fluid.

- A socket offers 230 V; the appliance, not the socket, sets how much current actually flows.
- Current and voltage alternate together in AC; for a resistive load they are in phase.
- **Switching off an inductive load is an analog-transient event.** The coil's stored magnetic energy has to go somewhere; a flyback diode (DC), Zener-clamp (fast-release DC), TVS (signal ESD), RC snubber (AC), or MOV (AC, bidirectional) is what keeps it from going through the switching device's dielectric (§6.11). The mains-side cousin is the SPD stack of §4.7.

Chapter 7 — The electromagnetic field: the real carrier of energy

7.1 The textbook diagram is a useful lie

That classic picture of electrons flying down a wire from the battery to the light bulb:

```
(Battery +)──▶──▶──▶──▶──▶── [bulb] ──▶──▶──▶──▶──▶──(Battery −)
         Fast-moving electrons racing through the wire?
```

This is **not** what actually happens. It is a pedagogical crutch that works for basic circuit calculations but totally misrepresents the physics.

Here is the reality:

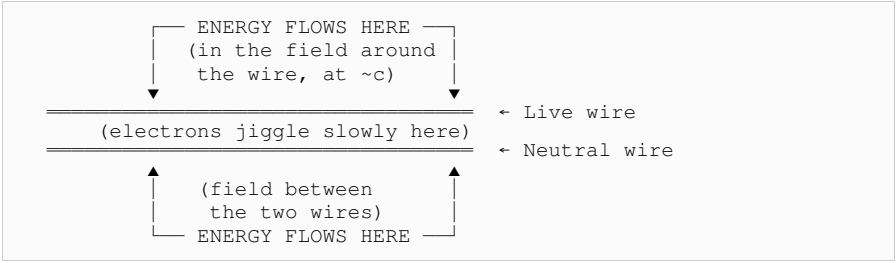
- **Individual electrons in a copper wire move incredibly slowly.** Typical drift velocity for a normal household current is about **0.1 mm/s** — a tenth of a millimetre per second, slower than a snail.
- **In an AC wire, electrons do not travel at all.** They oscillate back and forth by a fraction of a millimetre, 50 times a second, forever staying in roughly the same spot.
- **Yet the bulb turns on instantly** (within nanoseconds) when the switch closes.

If the bulb had to wait for electrons to travel from the battery to the filament at 0.1 mm/s through 2 m of wire, it would take about **5½ hours** to turn on. That is obviously not what happens, so something else is going on.

7.2 What actually travels

> **The energy travels through the electromagnetic field in the space around the wires — not through the wires themselves.**

Counter-intuitive when first heard, but rigorously true and provable. The wires act as **guides** for the field, not as pipes for the energy.



The electromagnetic field — an oscillating combination of an electric field (E) and a magnetic field (B) — fills the space between and around the Live and Neutral wires. It is that field that carries the energy from source to load, travelling at roughly the speed of light. The cover figure of this pamphlet is the same picture, labelled.

7.3 The Poynting vector: direction and rate of energy flow

There is a quantity in physics called the **Poynting vector** (named after John Henry Poynting, 1884). It represents the direction and magnitude of electromagnetic energy flow at a point in space.

Poynting vector — direction and rate of EM energy flow [W per m^2]

$$\vec{S} = \frac{1}{\mu_0}\vec{E} \times \vec{B}$$

Where:

- $\vec{S}$ — Poynting vector, in watts per square metre (W/m^2). Its direction is the direction in which electromagnetic energy is flowing at that point; its magnitude is how much power crosses a 1 m^2 surface held perpendicular to the flow.
- $\vec{E}$ — electric-field vector at the point, in volts per metre (V/m).
- $\vec{B}$ — magnetic-field vector at the point, in teslas (T).
- $\vec{E} \times \vec{B}$ — cross product; a vector perpendicular to both $\vec{E}$ and $\vec{B}$, with magnitude $|\vec{E}|\,|\vec{B}| \sin\theta$.
- μ_0 — magnetic permeability of vacuum, $\approx 4\pi \times 10^{-7}$ H/m.

What it says. At every point in space the Poynting vector tells

you two things about the electromagnetic field: *which way* its energy is moving, and *how fast* it is moving (power per unit area). In the cable between a battery and a bulb, $\vec{S}$ points along the cable, in the space *around* the copper — that is the picture on this pamphlet's cover. The field, not the electrons, is what carries the kilowatts.

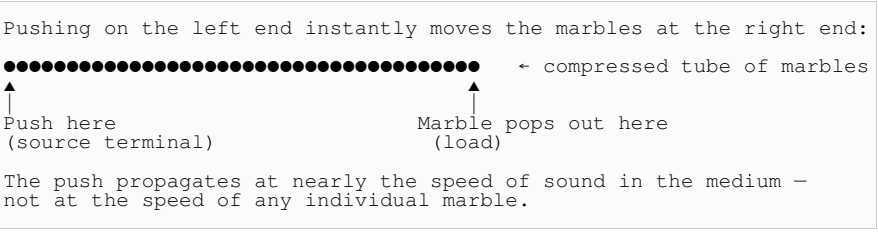
For a simple DC circuit with a battery and a lamp connected by two wires:

- **Inside the copper**, $\vec{S}$ points sideways (toward the centre of the wire), carrying the tiny amount of ohmic heating into the conductor.
- **In the space between and around the wires**, $\vec{S}$ points **from the battery toward the load** — this is the main energy flow.
- **The wires themselves do not carry the energy flow**; they **guide** the field that does.

7.4 What the electrons are doing

If electrons do not carry the energy, what are they doing?

Think of the electrons as **marbles already filling a long sealed tube**. When you push on one marble at one end, the push propagates down the line of marbles nearly instantly, and a marble pops out at the far end — but **no individual marble travelled the length of the tube**. Each marble just nudged its neighbour.



In a wire, the same idea applies with **electromagnetic push** rather than mechanical contact: the driving field at one end propagates along and between the wires at nearly the speed of light, and electrons everywhere in the wire start moving nearly simultaneously. The electrons themselves barely move; the **push**

that moves them travels at nearly c.

7.5 Where the electrons come from (and where they go)

They come from the wire itself. They were already there.

A copper wire is already packed with free electrons — about **10^{28} electrons per cubic metre**. Each copper atom contributes one "free" electron that can wander within the metal's crystal lattice, whether or not any source is connected.

When a circuit is closed:

- No electrons are **injected** from the source.
- No electrons are **used up** in the load.
- The source simply makes the electrons already in the wire **start jiggling collectively**, like starting a crowd dancing.

The electrons in the filament of a bulb were there before it was switched on and will be there after it is switched off. For every electron entering the filament from one side, another electron leaves from the other side at the same instant, because they are all being pushed by the same field.

7.6 Experimental proof

Four concrete pieces of evidence that energy flows through the field, not the conductor:

7.6.1 Bulbs turn on faster than electron drift allows

If energy travelled at the electron drift velocity (~0.1 mm/s), a bulb 1 m from a battery would take hours to light. Measured reality: a few nanoseconds. The transfer speed matches the speed of electromagnetic waves in the surrounding medium, not the speed of electrons.

7.6.2 The Veritasium "big circuit" video (2021)

A circuit is imagined with wires extending hundreds of thousands of kilometres out from a battery and back. The question: when the switch closes, how long before a bulb 1 m from the battery lights?

The intuitive (but wrong) answer is about 1 second — time for a signal to travel down the long wires and back at the speed of light. The **correct** answer is a few nanoseconds: the time for the field to cross the 1 m gap between the parallel wires, not the time to travel along them. See the bibliography for the video link.

7.6.3 Wireless power transfer

If energy were literally "in the electrons of the wire," wireless charging pads could not work, microwave ovens could not cook food, and radio waves could not reach an antenna. These all function because electromagnetic energy can exist and propagate **without any wire at all**. The wire is just a convenient guide when one is available.

7.6.4 Waveguides and coaxial cables at high frequencies

In radar, satellite, and telecom engineering, "transmission lines" at GHz frequencies become hollow metal tubes called **waveguides** — with **no central conductor at all**. Energy travels through the space inside the tube, guided by the walls. No electrons flow end to end; the energy was never in the electrons in the first place.

7.7 Better analogies

Mental models to replace "electrons flying through pipes" — none of them involves water in a pipe:

- **Sound wave in a tube of air.** Speak into one end; sound travels to the other end at ~340 m/s. Air molecules barely move — they vibrate in place. The pressure wave propagates, not the medium.
- **A long row of falling dominoes.** Push the first; the fall propagates to the last. No individual domino travels the length of the row.
- **Stadium wave ("La Ola").** Each fan stands up and sits down in place; nobody moves around the stadium. Yet the wave travels all the way around. This is exactly the AC electron jiggle.
- **A bicycle chain on sprockets.** The chain is already there all the way around the loop. Pedal, and the whole chain moves simultaneously; any given link moves only a short distance, but force on the rear sprocket is felt instantly.

7.8 The one-line truth

Electrons are the medium; the electromagnetic field is the actual message.

7.9 Take-aways

- Individual electrons move astonishingly slowly. The field does all the real travelling.
- The wires guide the field; the energy flows through the space around and between them.
- The Poynting vector $\vec{S} = (1/\mu_0)\,\vec{E} \times \vec{B}$ (§7.3) is the formal quantity that captures this; it points from source to load in the space around the conductors.
- Transformers, wireless charging, waveguides, and antennas are practical manifestations of this fact.

Chapter 8 — Maxwell in one chapter

8.1 What a “field” even is

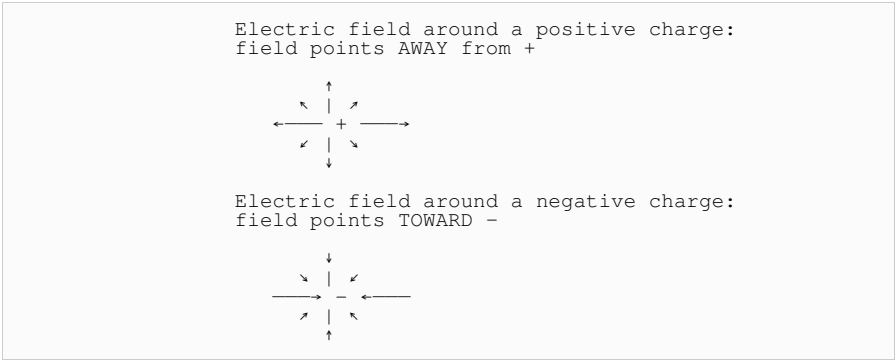
A **field** is a physical quantity that has a value at every point in space and time. That is the whole concept. Analogues you already know:

- **Temperature field.** Every point in a room has a temperature. Warm by the radiator, cool by the window. A **scalar field** — one number per point.
- **Wind velocity field.** Every point in the atmosphere has a wind speed *and* direction. A **vector field** — a value plus a direction per point.
- **Gravitational field.** Every point near Earth has a “downward pull strength” of about 9.81 m/s². A vector field — invisible but constantly felt.

The **electromagnetic field** is made of two vector fields that always come together: $\vec{E}$ (electric) and $\vec{B}$ (magnetic). Space is never “empty”; it is always filled with fields.

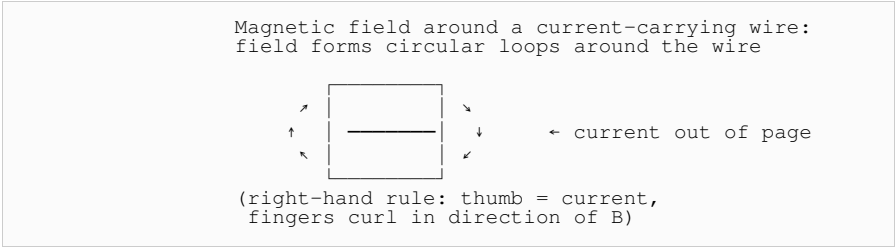
8.2 The electric field E

- **What it does:** pushes on electric charges.
- **Units:** volts per metre (V/m), or equivalently newtons per coulomb.
- **Direction:** from positive charge toward negative charge (by convention).
- **Sources:** static electric charges (batteries, charged balloons) and **changing magnetic fields**.



8.3 The magnetic field B

- **What it does:** pushes on *moving* charges and on magnets.
- **Units:** tesla (T), or equivalently V·s/m².
- **Direction:** always forms closed loops — there are no "magnetic charges" (no monopoles).
- **Sources:** magnets, steady currents, and **changing electric fields**.

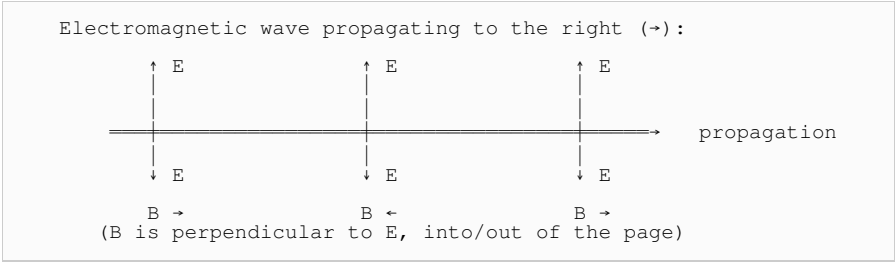


8.4 The deep magic: E and B create each other

Maxwell's insight (1860s):

> A changing electric field creates a magnetic field. A changing magnetic field creates an electric field.

This is a **self-sustaining loop**. Once an oscillation starts, E and B keep regenerating each other forever, propagating through space without needing anything else — not wires, not charges, not any medium.



- E oscillates in one plane (say, vertical).
- B oscillates in the perpendicular plane.
- The wave propagates in the third direction, perpendicular to both.
- All three are mutually perpendicular.

This is light. This is radio. This is how energy travels through space.

WiFi, 50 Hz mains power, sunlight, microwaves, X-rays — all the same phenomenon, differing only in frequency.

8.5 Why light travels at the speed of light

Maxwell derived the speed of the wave from pure theory.

Speed of light from Maxwell's equations [m/s]

$$c = \frac{1}{\sqrt{\varepsilon_0 \mu_0}}$$

Where:

- c — speed of electromagnetic waves in vacuum, in metres per second (m/s). A scalar. Numerical value: $\approx 2.998 \times 10^8$ m/s.
- ε_0 — electric permittivity of vacuum, $\approx 8.854 \times 10^{-12}$ F/m.
- μ_0 — magnetic permeability of vacuum, $\approx 4\pi \times 10^{-7}$ H/m.

What it says. Maxwell noticed that his four equations (§8.6), taken together, predict a wave in the electromagnetic field travelling at a speed fixed *entirely* by two constants of nature, ε_0 and μ_0, measurable in a lab with no mention of light at all. Plug them in — and the answer is the measured speed of light. That was the moment electricity, magnetism, and optics collapsed into a single theory: **light is simply an electromagnetic wave**, and the grid frequency, WiFi, sunlight, and X-rays are all the same phenomenon at different frequencies.

Plug in the numbers: $c \approx 3 \times 10^8$ m/s — exactly the speed of light measured experimentally. Maxwell had **predicted the speed of light from electromagnetic theory alone**, proving that light is an electromagnetic wave.

8.6 Maxwell's four equations — the source code of electromagnetism

All of classical electromagnetism — every phenomenon from home wiring to radio to light — follows from just four equations.

8.6.1 Gauss's law (for electricity)

Electric charges create electric fields that spread outward from them.

Gauss's law for electricity [V m]

$$\oint \vec{E} \cdot d\vec{A} = \frac{Q}{\varepsilon_0}$$

Where:

- $\vec{E}$ — electric-field vector at each point on the closed surface, V/m.
- $d\vec{A}$ — infinitesimal area vector on the surface; its direction is the outward normal, its magnitude is the element's area (m^2).
- $\oint \vec{E} \cdot d\vec{A}$ — *electric flux* through the closed surface: the net amount of E-field "poking out" of it.
- Q — total electric charge enclosed by the surface, in coulombs (C). A scalar.
- ε_0 — electric permittivity of vacuum, $\approx 8.854 \times 10^{-12}$ F/m.

What it says. Put a charge Q inside any closed balloon you can imagine, and the total amount of E-field poking outward through the balloon's skin is fixed — set by Q alone, not by the balloon's shape. Twice the charge, twice the flux. No charge inside, no net flux. Electric charges are the *sources* of electric field, and the field spreads outward from them.

Put a charge in space, and it fills the surrounding space with an E-field that weakens with distance.

8.6.2 Gauss's law for magnetism

Magnetic field lines always form closed loops — no magnetic monopoles exist.

Gauss's law for magnetism [T m^2]

$$\oint \vec{B} \cdot d\vec{A} = 0$$

Where:

- $\vec{B}$ — magnetic-field vector at each point on the closed surface, in teslas (T).
- $d\vec{A}$ — outward-pointing area element on the surface, m^2.
- $\oint \vec{B} \cdot d\vec{A}$ — net magnetic flux through the closed surface.

What it says. The net magnetic field poking out of any closed surface is always exactly zero. Whatever field lines enter the surface somewhere must leave it somewhere else: magnetic field lines form closed loops. In physical terms, magnetic monopoles do not exist — there is no "isolated north pole" the way there are isolated electric charges. Cut a bar magnet in half and you get two smaller, still-two-pole magnets, forever.

8.6.3 Faraday's law (of induction)

A changing magnetic field creates an electric field.

Faraday's law of induction [V]

$$\oint \vec{E} \cdot d\vec{l} \;=\; -\frac{d\Phi_B}{dt}$$

Where:

- $\vec{E}$ — electric-field vector at each point of the closed loop, V/m.
- $d\vec{l}$ — infinitesimal displacement vector along the loop, m.
- $\oint \vec{E} \cdot d\vec{l}$ — "electromotive force" (EMF) around the loop, in volts (V): the work the field would do moving one unit of charge once around the loop.
- Φ_B — magnetic flux through any surface bounded by the loop, in webers (Wb = T·m^2): $\Phi_B = \int \vec{B} \cdot d\vec{A}$.
- dt — time element, s.
- The minus sign — *Lenz's law*: the induced EMF opposes the change in flux.

What it says. Whenever the magnetic flux through a loop changes, an electric field is created that circulates around the loop — exactly strong enough to push charges around it with a total "EMF" equal (up to sign) to the rate of flux change. This

is the principle behind every generator, transformer, induction motor, inductive sensor, and induced voltage that plagues EMI-sensitive cables. Wiggle a magnet near a coil, and a voltage appears at the coil's terminals.

8.6.4 Ampère-Maxwell law

An electric current **or** a changing electric field creates a magnetic field.

Ampère–Maxwell law [T m]

$$\oint \vec{B} \cdot d\vec{l} \;=\; \mu_0 \left(I \;+\; \varepsilon_0 \frac{d\Phi_E}{dt} \right)$$

Where:

- $\vec{B}$ — magnetic-field vector around the closed loop, T.
- $d\vec{l}$ — displacement vector along the loop, m.
- $\oint \vec{B} \cdot d\vec{l}$ — line integral of $\vec{B}$ around the loop (how much the magnetic field "circulates").
- I — conduction current passing through any surface bounded by the loop, in amperes (A).
- Φ_E — electric flux through the same surface, in V·m.
- $\varepsilon_0 \, d\Phi_E/dt$ — Maxwell's *displacement current*: a current-like contribution from a *changing* electric field, even in vacuum with no charges present.
- μ_0 — magnetic permeability of vacuum, H/m.

What it says. Two things make a magnetic field: real currents (the I term — electromagnets, solenoids, motors) and *changing* electric fields (the $\varepsilon_0 \, d\Phi_E/dt$ term — Maxwell's great addition of 1864). The second term is what lets electromagnetic waves propagate through empty space: a changing E-field makes a B-field, which changes and makes an E-field again, forever. No charges required. **That self-regenerating loop — combined with Faraday's law above — is light.**

Current in a wire creates a B-field around it (electromagnet, solenoid, motor). The second term — Maxwell's key addition — says even a changing E-field, without any current, creates a B-field. This is what lets EM waves propagate through vacuum: no charges

needed, just oscillating E and B regenerating each other.

8.6.5 Together, they say…

- Electric charges → E-fields.
- Currents → B-fields.
- Changing B-fields → E-fields.
- Changing E-fields → B-fields.
- Once started, E and B regenerate each other and propagate forever through empty space at the speed of light.

8.7 Fields in insulators (dielectrics)

A **dielectric** is simply an insulator viewed from the field's point of view: a material that does not conduct DC current, but that the E-field can live inside, polarise, and store energy in. Air, plastic, glass, ceramic, and even pure vacuum are all dielectrics in this sense.

The field propagates through all of them — which is why it can reach across a transformer's air gap, between the plates of a capacitor, or between a wireless-charging pad and a phone.

Inside such an insulator (dielectric), three things happen:

1. The field propagates through it, carrying energy.
2. The atoms polarise slightly: electron clouds shift one way, nuclei the other, producing tiny dipoles that partially oppose the applied field. This is captured by the material's **relative permittivity** ε_r:

Material	ε_r (approx.)
Vacuum	1
Air	1.0006
PVC, polyethylene	2–4
Glass	4–10
Water	≈ 80
Ceramic dielectrics	1000+

3. Energy is stored in the polarisation. That is literally how capacitors work — energy lives in the E-field in the dielectric, not in the metal plates.

Above a critical field strength (the material's **dielectric strength**) the insulator's atoms ionise and current arcs through it — a **dielectric breakdown**. That is why insulation thicknesses and creepage distances are specified for peak voltage (see chapter 5).

8.8 What the field carries

When people say "the field is the message", the message contains four measurable things:

- **Energy.** Every point where the field exists carries an amount of energy per cubic metre, given by the electromagnetic energy-density formula introduced in chapter 1 ($u = \frac{\varepsilon_0}{2}|\vec{E}|^2 + \frac{1}{2\mu_0}|\vec{B}|^2$). Wherever the field is, there is energy — literally stored in that region of space.
- **Momentum.** Light can push things (solar sails; the radiation pressure that keeps stars from collapsing under their own gravity — in a star's core, the flood of photons streaming outward from fusion pushes on the surrounding gas hard enough to balance the inward pull of gravity, so the star neither collapses nor flies apart).
- **Information.** Patterns of oscillation encode data: radio and WiFi modulate the amplitude or frequency of an EM wave, and the receiver demodulates those patterns back into audio or bits.
- **Angular momentum.** Circularly polarised light carries spin — used in advanced optical communications.

8.9 Take-aways

- A field is just a value at every point in space.
- $\vec{E}$ and $\vec{B}$ are two vector fields that create each other in motion and propagate through space at c.
- Maxwell's four equations describe every classical electromagnetic phenomenon. They fit on a t-shirt.
- Fields exist in air, in insulation, and in vacuum. They carry energy, momentum, and information.

Chapter 9 — Fields in everyday wiring: loop area, coupling, EMI (electromagnetic interference)

9.1 Analog and digital are the same field

Before the rest of this chapter uses the phrase *"digital edge"* half a dozen times, it is worth retiring a misconception the industry quietly carries around: **there is no such thing as a digital signal at the physical layer.** There is only an analog voltage — a continuously varying potential between two conductors, and the continuously varying current (the drift of the free electrons, §2.10) that the accompanying E-field elicits in the copper — that the receiver agrees to round to `0` or `1` against a threshold. Everything this pamphlet has said so far about fields, edges, coupling and loop area applies *verbatim* to what most schematics label "the digital side" of a circuit.

An **analog signal**, in the plainest sense, is a voltage or current continuously proportional to something real — a temperature, a pressure, a microphone membrane, a 0–100 % valve opening. The industrial standards are familiar: **4–20 mA** (see §9.8) for process instruments, **0–10 V** for valves, a few millivolts for a thermocouple, a few tens of millivolts for a strain gauge. None of those are "a different kind of electricity"; they are all just a voltage or a current of the sort §2.4 taught us how to read.

A **digital signal** is the same thing, deliberately restricted. The transmitter drives only two voltage levels ("low", "high"), spending as little time as possible in between; the receiver accepts anything below one threshold as `0` and anything above another as `1`. Every rising edge is an *analog transient* with a finite rise time (`dV/dt`), a finite bandwidth, reflections on mismatched lines, and ringing on reactive loads — exactly the behaviour Chapter 8 told us to expect when an electric field has to rearrange itself quickly.

Analog transient vocabulary — decoding the sentence above

The paragraph above packs five phrases you will keep hearing across this chapter and the next. None of them has been defined yet, so here they are, anchored in physics already on the table.

- **Digital edge.** The brief transition moment when a digital line flips from "low" to "high" (*rising edge*) or from "high" to "low" (*falling edge*). Every clock tick, every data bit, every control pulse is built out of these edges. Physically, the edge *is* an analog transient (next bullet) — the voltage is on its way from one steady level to the other, and for that small time window it is neither "0" nor "1" in any strict sense.
- **Analog transient.** Any short-lived, non-steady voltage or current — a ramp, a spike, a ring, a bounce. The word "analog" is literal: it is just a voltage that happens to be in a hurry. A digital edge is the commonest example; so is the kick at the moment a relay closes, or the spike on an inductor when its current is interrupted.
- **Finite rise time** ($\mathrm{dV/dt}$). The time the voltage takes to climb from its "low" level to its "high" level. A 3.3 V CMOS gate (CMOS = *Complementary Metal-Oxide-Semiconductor*, the logic family inside virtually every modern digital chip; a "gate" is one of its output stages — an AND, OR, NOT, buffer, flip-flop, etc.) with a 2 ns rise time ramps from 0 to 3.3 V at roughly 1.6 V/ns $= 1.6 \times 10^9$ V/s. Nothing in the physical world can change instantaneously (Chapter 8), so $\mathrm{dV/dt}$ is always finite.

  Aside. Even the celebrated "quantum leap" is not truly instantaneous in modern quantum mechanics — see the footnote at the end of this chapter (§9.13) and Pietro Greco's *Quanti* in the bibliography (§12.5).
- **Finite bandwidth.** A fast edge does not contain a single frequency; it contains the fundamental of the signal *and* a wide spread of *higher harmonics* — extra sine-wave components, at integer multiples of the signal's base rate, that Fourier analysis adds together to produce the sharp edge shape (a pure sine wave alone has rounded peaks; the faster and squarer the edge, the more of these higher-frequency sines the sum must include) — reaching roughly $0.35/t_{\text{rise}}$. A 2 ns edge therefore has significant content up to $\sim$175 MHz — comfortably into the radio spectrum. That is why a "DC" control cabinet can radiate in the FM band.
- **Reflections on mismatched lines.** A sharp edge *is* a rapid change in the E- and B-field configuration around a cable or PCB trace (*PCB* = Printed Circuit Board; a *trace* is

a flat copper conductor etched on the board — the small-scale equivalent of a wire, running between components), and that field change propagates along the line at close to the speed of light — *in the dielectric between the "go" and "return" conductors* (*dielectric* = the insulating material separating those two conductors — air in open wiring, plastic in cables, fibreglass / FR-4 in PCBs; formally defined in §8.7), *not in the copper*: the electrons themselves barely stir (§2.10), the wire merely guides the wave. When the travelling edge meets a point where the line's geometry changes (a connector, a stub, an unterminated end), the line's characteristic impedance changes too, and part of the energy bounces back — just like sound echoing when a tunnel widens. The reflected edge returns to the transmitter, distorts the next edge, and keeps bouncing until losses damp it out. This is the origin of the "staircase" waveforms on scope traces of poorly routed PCB tracks (same object as "trace" above — "track" is the British usage, "trace" the American one).

- **Ringing on reactive loads.** The inductance of every real wire plus any capacitance at the receiver forms a small LC circuit (an *L–C resonator* — an inductor L and a capacitor C connected together, trading energy back and forth between the magnetic field of L and the electric field of C at their shared natural frequency $f_0 = 1/(2\pi\sqrt{LC})$). A sharp edge "kicks" it and it oscillates briefly at its own natural frequency before settling — exactly like striking a bell and hearing it ring. Reactive loads (see §6.9) are any inductive or capacitive load, including the parasitic L and C of the wiring itself.

Between the two extremes sit the deliberately hybrid schemes the glossary (§11.4) already lists:

- **PWM (Pulse-Width Modulation)** — digital in time, analog on average: a square wave whose duty cycle carries the analog quantity, which an LC filter (an L–C *low-pass filter*: the inductor blocks fast changes, the capacitor smooths them out; together they average a chopped square wave back into a near-DC level) or a motor winding's own inductance averages back into a smooth current.
- **SMPS (Switched-Mode Power Supply)** — chopped DC at

tens of kHz that looks like clean DC once the LC output filter (same L–C low-pass arrangement as above, sized for the SMPS switching frequency) has done its job.

- **Sigma-delta (ΣΔ) ADC (Analog-to-Digital Converter)** — the reverse of PWM: an analog voltage is chopped into a one-bit pulse stream whose *density* encodes the value, and a digital filter averages it back into a high-resolution number.

All three are analog physics wearing a digital costume (or vice versa). None of them violates Maxwell.

This is the bridge into the rest of this chapter. Fast digital edges are the worst EMI sources *not because they are "digital"* — the receiver's threshold has nothing to do with the field the wire emits — *but because their rise times are short*. A short rise time means wide analog bandwidth (roughly `f_BW ≈ 0.35 / t_rise`), and wide analog bandwidth means strong coupling through the capacitive and inductive mechanisms that §9.5 is about to describe. Analog signal lines, with their slower rates of change, are mostly *victims* rather than offenders — which is why the practical advice that follows spends most of its energy on keeping the aggressor's loop small, and keeping the victim's loop away from it.

9.2 The guiding principle of cable design

Everything in this chapter follows from one observation:

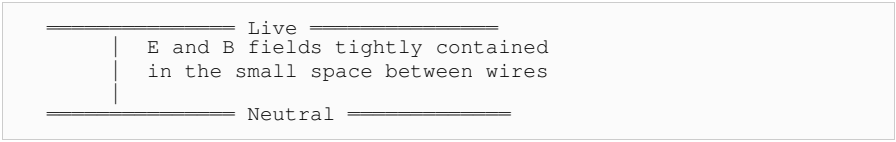
> **The field exists in the space between the "go" and "return" conductors. The geometry of those conductors decides how much space the field fills. Small space → contained field → clean transfer. Large space → spread-out field → radiation loss and interference.**

Every cable ever designed implements some version of that idea.

The key geometric quantity that makes this observation quantitative is the **loop area**: the two-dimensional area enclosed by a circuit's "go" and "return" conductors. It is the area through which an external magnetic field can thread to induce a voltage (Faraday's law), and, by reciprocity, the area from which the circuit itself radiates when current changes. Almost every EMI rule in this chapter boils down to *make that area small*. See also the glossary entry in §11.5.

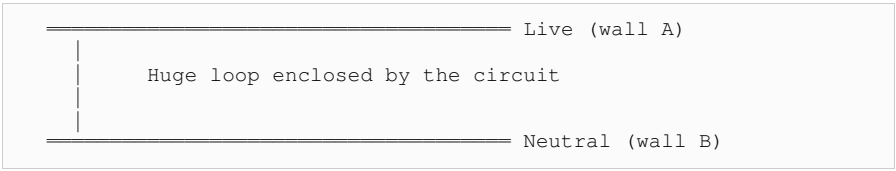
9.3 Three scenarios

9.3.1 Wires close together (normal installation)



- Field is tightly contained in the narrow gap.
- Almost all energy flows directly from source to load via the contained field.
- Tiny stray radiation, tiny loop inductance, tiny EMI.
- This is what good cables look like: twisted pairs, paired lines, coaxial cables.

9.3.2 Wires far apart (Live in one wall, Neutral in another)



- The field now fills a huge volume of room.
- The **loop area** enclosed by the circuit is enormous.
- Consequences:
  - Large loop inductance — the circuit resists current changes.
  - Strong magnetic field throughout the room — anything metallic in the loop experiences induced voltages and eddy currents.
  - The loop acts like an antenna, radiating energy as EM waves into the surroundings.
  - The same loop also *receives* interference from nearby sources.
- Energy still reaches the load, but some fraction is lost to radiation and to parasitic coupling.

9.3.3 Wires deliberately spread wide at high frequencies

At sufficient separation and sufficient frequency, a circuit stops trying to deliver energy to a load and instead radiates it. That is what

an antenna is. A 50 Hz widely-separated pair is a very inefficient antenna; a 2.4 GHz widely-separated pair is a deliberately engineered one.

9.4 Cable design in one table

Cable type	How it contains the field
Twisted pair	Conductors spiral around each other tightly. External B-fields induce equal-and-opposite voltages on alternating twists → they cancel in the differential signal. Loop area is essentially zero.
Coaxial cable	Inner conductor surrounded by a cylindrical outer shield with insulator in between. Field is 100 % contained inside the shield. Zero radiation, zero pickup.
Paired mains cable	L and N run parallel inside the same outer sheath. Not perfect, but loop area is small.
Shielded twisted pair	Twisted pair inside a metal braid / foil shield. Best of both worlds: twist cancels differential pickup, shield blocks field leakage.
Ribbon with alternating grounds	Every signal wire has a ground wire right next to it → tiny local return path, tiny loop area.

9.5 Capacitive vs inductive coupling

Two cables running parallel to each other exchange energy through both **capacitive** and **inductive** coupling:

- **Capacitive coupling.** The voltage on one cable pushes and pulls the charges on the other through the air / insulation acting as a dielectric. Proportional to $\mathrm{dV/dt}$ and the shared surface area. Dominant at high voltages and high frequencies with modest currents.
- **Inductive coupling.** The current in one cable creates a B-field that threads through any loop enclosed by the other, inducing a voltage. Proportional to $\mathrm{dI/dt}$ and the enclosed loop area. Dominant at high currents and high frequencies with modest voltages.

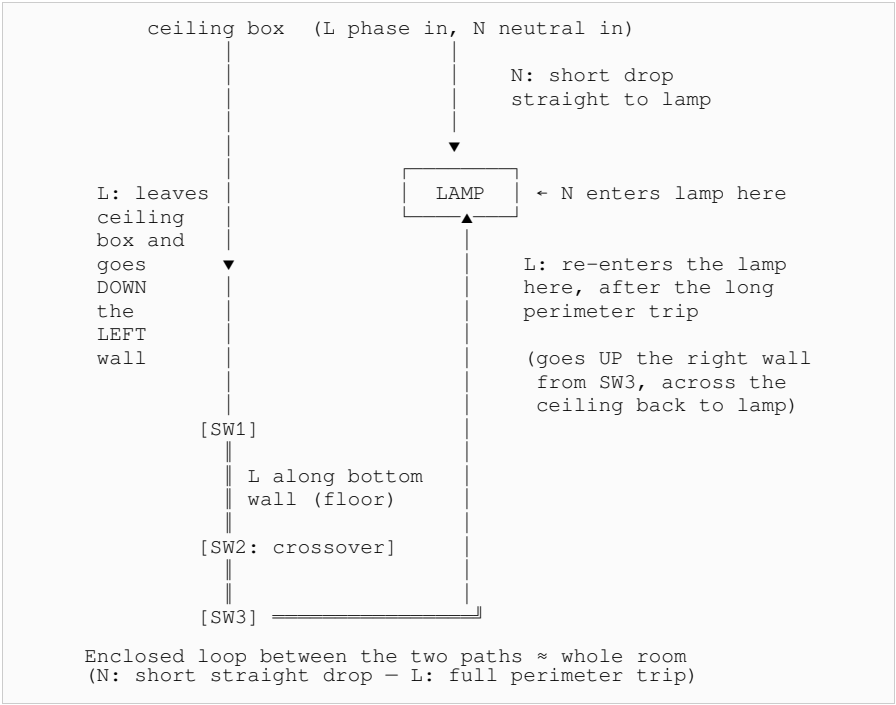
The two mechanisms are why **loop area** and **voltage slew rate**

together determine the EMI signature of a circuit.

9.6 Worked example: a corridor lamp switched from three locations

A classic residential installation uses two SPDT switches and one crossover (4-way) switch to control one lamp from three locations — the Italian “deviatore + invertitore + deviatore” configuration, or the US 3-way + 4-way. Physically, the hot-side wiring runs around most of the room’s perimeter (up one wall, across the ceiling, down the other wall, along the bottom, back up to the lamp), while Neutral takes the shortest possible path from the ceiling box straight down to the lamp.

Drawn together, these two paths enclose a closed loop of roughly **the whole room**:



- The enclosed loop area is **square metres**, not square centimetres.
- An AM radio held inside that room hums loudly at 50 / 100 Hz.

- Nearby metal objects pick up small induced voltages from the oscillating B-field.
- The lamp still works — the radiation loss at 50 Hz is negligible — but the wiring is a surprisingly efficient pickup antenna for any high-slew transient (lightning, nearby contactor operation, arcing appliance).

Why it is tolerated: 50 Hz is absurdly low frequency by EMI standards, the lamp itself is immune to its own field, and most other circuits in the room reject common-mode pickup. Why the same topology would be disastrous for a signal cable: see §9.8.

9.7 Five practical rules

Consolidating §§9.2–9.6 into installation-level rules that apply to any control installation:

1. **Keep “go” and “return” conductors in the same cable sheath.** Never use the building’s chassis or ground as an implicit return for a signal. By *building chassis* we mean the **metal structural frame of the installation** — steel girders, metal conduits, cable trays, cabinet and machine enclosures, pipework — all bonded together and to the earth / PE (protective earth) system. It is tempting to let a signal current “come back through the metalwork” because it is always available for free, but the return path then meanders through an enormous uncontrolled loop whose area you neither know nor control.
2. **Separate power cables from signal cables by at least 20 cm** when they run parallel. At any greater separation, coupling drops rapidly.
3. **Cross at right angles.** Where power and signal cables must cross, they should cross perpendicular to each other. Mutual inductance between two loops drops to zero when they are orthogonal.
4. **Use shielded twisted pair for low-level analog and communication signals.** Ground the shield at one end (typically the controller end) to avoid ground-loop circulating currents while still giving electrostatic protection.
5. **Do not share cable glands or conduit sections between power and signal.** Even a short shared run negates the separation rule’s benefit.

These rules are restatements of IEC EMC guidance and national

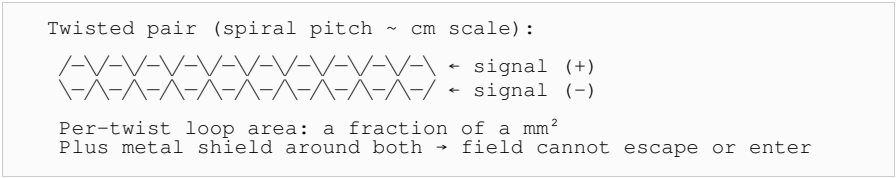
codes such as CEI 64-8 — but they are equally **direct consequences of Maxwell's equations** applied to cable geometry.

9.8 Signal-loop example: a 4-20 mA transmitter

A 4-20 mA current-loop transmitter sending a process variable to a controller over 15 m of cable.

Bad: single-ended wiring, signal return via building chassis. Loop area is the whole run, several m². Every contactor kick, every motor-drive commutation, every nearby 50 Hz power cable induces voltage into the loop. The controller sees chaos; the actual signal is buried in noise.

Good: shielded twisted pair, both conductors inside the same sheath, shield grounded at the controller only.



- Per-twist loop area: well under 1 mm².
- Successive twists alternate orientation, so any induced EMF on one twist is cancelled by the opposite induced EMF on the next.
- The shield intercepts external E-fields and routes them to ground (Gauss's law applied to a closed conductor).
- The controller sees the clean 4-20 mA signal with millivolt-level noise.

The difference between "works" and "does not work" in industrial sensor systems is overwhelmingly a consequence of loop area and shielding — pure Maxwell, applied.

9.9 Jumpers, headers, DIP switches: stubs in miniature

(One of the three hands-on sections flagged in the Introduction — a place where the field picture earns its keep.)

§9.1 introduced the idea of a **stub** — a piece of conductor that ends in air, off the main signal path, and **whose unused length acts as a reflection site for fast edges**. Every Arduino and OpenPLC

board is covered in hardware that is a stub by definition: jumpers, pin headers, DIP switches, solder bridges. The point of this section is to cash in the §9.1 vocabulary for that whole class of hardware and say, simply: **treat every jumper as a mid-bus stub** (not at either physical end of the bus) **and you will get it right**.

9.9.1 What each of these objects physically is

- **Jumper (shunt).** A small removable U-shaped connector — the black plastic cap you pull off an old motherboard. Mechanically it bridges two adjacent pins of a pin header; electrically it is a wire, typically a few mm long. Closed = signal continuous; open = signal broken. Binary configuration.
- **Pin header.** A row of through-hole or surface-mount pins that exposes internal signals. The header itself is the population of pins; a jumper is what you put *on* it. Headers of 2 × N pins are the standard for jumpers, DIP connectors and most probe points on a board.
- **DIP (Dual In-line Package) switch.** A small rectangular module carrying N independent mechanical slide switches, each with its own ON/OFF position. Electrically it is an N-bit jumper without loose parts — the same function, integrated so it cannot be lost.
- **Solder bridge.** Two copper pads side by side on the PCB (§9.1), normally separated by a sliver of silkscreen. The technician closes the bridge with a blob of solder at assembly time. Non-removable after the fact; used for "set once and forget" options that should not be user-adjustable.
- **0-Ω resistor.** A chip resistor whose nominal value is literally zero ohms, fitted in the place of a higher-value resistor. Electrically a solid piece of wire; functionally a factory-installable, robot-friendly version of the solder bridge.

9.9.2 Why they matter at fast edge rates

In DC and at mains-frequency AC, a 5 mm jumper and an 8 mm via stub are invisible: their series inductance is fractions of a nanohenry and their shunt capacitance is fractions of a picofarad. At the rise times of §9.1 (2 ns, ~175 MHz of analog bandwidth), those tiny parasitics start behaving like real circuit elements. Specifically:

- Every **unpopulated** pin of a header left dangling past the selected position acts as an **open-circuit stub**. The fast

edge reaches the branch point, splits, travels down the stub, bounces off the open end, and returns to the main line out of phase with the next edge. That is exactly the reflection mechanism of §9.1.

- Every **contact pair** inside a jumper or DIP switch has a small contact resistance and a real inductance from its internal geometry (tiny spring fingers, a plunger, a slider). That inductance adds to any dI/dt event and can ring with a nearby capacitance (§9.1 Ringing bullet).
- Every **transition** in PCB geometry — a thick trace narrowing down to a header pin, a jumper lead going through a via — changes the characteristic impedance. Each change is a small reflection site of its own.

A 2 mm dangling stub on a typical FR-4 PCB (**FR-4** = flame-retardant class 4, the glass-reinforced epoxy laminate of almost every PCB) presents an appreciable reactance at 500 MHz. On a 2 ns digital edge whose analog bandwidth reaches 175 MHz, that is not a theoretical concern; it is what shows up on the scope as ringing on the first edge after a jumper pin.

9.9.3 Where they show up on Arduino / OpenPLC hardware

The most common uses and the field-first reason each one works the way it does:

- **I²C (Inter-Integrated Circuit, pronounced "I-squared-C") address selection.** Many I²C peripherals (port expanders, RTCs, EEPROMs, ADCs) have two or three address-select pins that you strap high or low with jumpers. The jumper sits next to the chip, so the pin lead is short; still, the unused pin on the other side of the jumper header is an open stub. Route the jumper so the strap pin is *closest* to the chip; put the other pin (the "elsewhere" side) on the outside of the header.
- **Bus-termination jumpers for RS-485 and CAN.** A differential bus is terminated with a resistor (typically 120 Ω) matched to its characteristic impedance, **only at the two physical endpoints of the bus**. A terminator placed in the middle of the bus *is* a mid-bus stub in exactly the §9.1 sense, and spoils the reflections management the bus depends on. Boards expose this via a jumper labelled "TERM" or "120 Ω"; put it in only on the two physical ends.

- **ADC reference selection.** Many MCU (Microcontroller Unit) ADCs can take their reference from an internal bandgap, from VCC, or from an external precision reference (LM4040, REF3xxx, etc.) via a jumper. The jumper lead between the MCU's V_{REF} pin and the selected source is an **antenna** for nearby switching noise: keep it short, star-route it back to the source, and put a local 100 nF + 10 µF decoupling pair next to the pin.
- **MCU boot-mode pins.** On STM32 boards the BOOT0 pin, and on ESP boards the strapping pins GPIO0 / GPIO2, are sampled once at reset to choose between running the main program, entering a bootloader, or other firmware modes. Physically these are usually slow DC signals, but the pin is often adjacent to fast signals and is brought out on a jumper header that the user can touch. Leaving it as a long stub is an invitation for nearby switching edges to flip the boot mode spontaneously; keep the lead short, add a small capacitor to ground.
- **Voltage-selection jumpers (5 V / 3.3 V).** These pick which rail powers the logic level on an I/O pin or which supply drives a daughterboard. They carry real current, not just a signalling level, so the jumper must be treated like a miniature power rail: short lead, local bypass capacitor after the jumper, and a reminder that the jumper's contact resistance puts a small series resistor between the source and the load.
- **Interrupt / signal-routing DIP switches.** On many OpenPLC shields and industrial carrier boards, a bank of DIP switches routes sensor inputs to different MCU pins, connects or disconnects the onboard pull-ups, or enables features like watchdog reset. Each position is an internal stub whose length is fixed by the module geometry.

9.9.4 The one rule of thumb

When in doubt, **assume every jumper is a mid-bus stub and route it like one**:

1. Keep the lead from the live pin to the target as short as geometry allows.
2. Provide a clean reference return *immediately* underneath the jumper (a ground plane beneath, or at worst a return trace parallel and close).
3. Do not leave pins dangling past the active shunt if you can avoid it; where you must, accept that those pins are antennas

and keep them out of the signal-integrity critical path.
4. Terminate only at the ends, never in the middle.
5. For ESD (electrostatic discharge) protection on user-accessible jumpers, consider a small TVS (Transient-Voltage-Suppression, §6.11) from each exposed pin to ground.

The same sentence in Maxwell's language: every jumper is a local discontinuity in the geometry that guides the field, and the field responds to discontinuities by radiating, reflecting, and ringing. Short, well-referenced, terminated-only-at-the-ends is how you keep the discontinuity small enough to ignore.

9.10 Why specific high-frequency tricks exist

The same principles explain advanced cable and PCB (printed circuit board) tricks:

- **Ethernet (1 Gb/s and up) uses twisted pairs** because at hundreds of MHz, any appreciable loop area radiates massively and picks up interference. Twisting drives effective loop area to nearly zero.
- **High-frequency RF PCBs have ground planes** because a continuous ground plane under a signal trace provides an "image" return current flowing directly beneath it. The effective loop area collapses into a thin sliver between trace and plane.
- **CAN bus** (the two-wire serial bus every modern car uses internally — whenever a mechanic plugs a diagnostic tool into your OBD-II port, they are talking to the car's CAN bus) **and RS-485** (another two-wire industrial serial bus; it is the physical layer behind Modbus in solar inverters and building-automation gear, and behind DMX-512 in stage and concert lighting — so most people use RS-485 daily without knowing it) **are differential** for common-mode rejection: any external field induces equal voltages on both conductors, the two shift together, but the difference (the actual signal) is untouched.
- **Transformers use laminated iron cores** because iron's high permeability channels the B-field through the secondary winding, lifting coupling from ~20 % (air-core) to ~99 % (laminated iron). Laminations break up eddy-current paths that would otherwise dissipate energy inside the core.
- **Variable-frequency drives need EMC (electromagnetic compatibility) filters** because switching at kHz with very fast dV/dt radiates broadband EM fields from the motor cables.

The filter catches the HF noise and redirects it back to ground before it leaves the drive.
- **Shielded motor cables** exist for the same reason: contain the fast- switching field geometrically, so it does not broadcast across the building.

9.11 A final beautiful truth

Every cable is a compromise between containment and transmission.

- A **coaxial cable** perfectly contains the field inside its metallic shield → zero leakage, zero pickup, but no "reach" beyond the annular space between inner and outer conductors.
- A **dipole antenna** perfectly radiates its field outward → huge reach across kilometres, but total loss of energy containment.
- A **transformer primary** contains most of the field inside the iron core (reaching the secondary) but leaks a small fraction as stray inductance; engineers tune the core geometry to maximise containment.
- **Household wiring** falls in between — deliberately paired for containment, but not shielded because EMI is manageable at 50 Hz.

Every decision an electrical engineer makes about cable type, grounding strategy, component placement, PCB layout, or enclosure design is fundamentally **a decision about where to let the field go and where to block it**. That is the entire practice of electromagnetic engineering, distilled.

9.12 Take-aways

- Loop area drives inductive EMI. Shared dielectric volume drives capacitive EMI. Keep both small.
- Twisted pairs reduce loop area to nearly zero; shields block external E. Shielded twisted pair does both.
- Separate power and signal by ≥ 20 cm parallel, cross at right angles, no shared conduits.
- Engineering choices that look like unrelated conventions (twisted pair, ground plane, differential signalling, laminated core, EMC filter, shielded motor cable) are all the same idea applied to different geometries.

- **Every jumper, pin header, and DIP switch is a transmission-line stub** in the §9.1 sense: invisible at DC, visible as reflections and ringing at fast edge rates. Short lead, referenced return, terminate only at the ends (§9.9).

9.13 Footnote — on "instantaneously"

In §9.1 we underlined the sentence *"nothing in the physical world can change instantaneously"*. That claim deserves a careful footnote, because the one corner of physics where a lay reader might suspect an exception is quantum mechanics, and the popular imagination has the direction of that exception backwards.

The image of the **quantum "leap"** — an electron snapping from one Bohr orbit to another in zero time — is a **1913 cartoon** from the earliest version of atomic theory. It was a useful first pass: it told us that atoms emit light at discrete frequencies and that energy is exchanged in chunks of $h\nu$. But modern quantum mechanics, from Schrödinger and Heisenberg onwards, has long since retired that picture. A real transition between two quantum states is a **probabilistic, time-extended process**: the system's wavefunction is, for a while, a superposition of the two states with time-varying weights, and the transition has a *duration*, not an *instant*.

That duration is bounded below by the **time–energy uncertainty relation** $\Delta E \cdot \Delta t \gtrsim \hbar/2$. Equivalently: the narrower the energy spread of the transition (the longer the excited state's natural lifetime, which sets the spectral *linewidth*), the longer the transition takes. And since the Nobel Prize in Physics 2023 (Agostini, Krausz, L'Huillier), **attosecond spectroscopy** has been measuring these durations directly — tens to hundreds of attoseconds (10^{-17}–10^{-16} s). Short, but emphatically not zero.

The conclusion is the opposite of the popular one. Planck's constant h sets the *scale* of the quantum world; it does **not** license anything to happen in zero time. On the contrary, the same uncertainty relation that ties ΔE to Δt tells us that a *truly* instantaneous change would require infinite energy and infinite bandwidth — the universe does not supply either. So when this pamphlet writes that "$\mathrm{dV/dt}$ is always finite", the statement holds all the way down: at the classical level (inertia of charges, finite capacitance and inductance, causal propagation at c), and at the quantum level too (finite transition times set by $\hbar$).

Pietro Greco's *Quanti — La straordinaria storia della meccanica quantistica* (bibliography, §12.5) tells that story end to end: from Planck's 1900 quantum hypothesis, through Bohr's 1913 atom, to the 1925–27 revolution that replaced "leaps" with wavefunctions and probability amplitudes. It is an accessible read for anyone who wants the history to match the physics they now carry around.

Chapter 10 — How to actually see and measure fields

10.1 Why “invisible” fields are not really invisible

Fields do not emit light, but they interact with matter in predictable ways. Every one of the instruments and tricks below converts some aspect of that interaction into something a human can see, hear, or read.

10.2 Classic demonstrations

10.2.1 Iron filings and a magnet

Sprinkle iron filings on a sheet of paper placed over a bar magnet and tap the paper. The filings line up with the local magnetic field direction and visibly trace the field lines. Cheap, silent, stunning.

10.2.2 Fluorescent tube near a high-voltage line

Hold an unconnected fluorescent tube beneath an overhead HV line. The tube’s gas ionises from the ambient 50 Hz E-field and glows — a dramatic demonstration that fields carry real energy through the air with no wire contact at all. (Do this only from a safe distance; HV lines are not a playground.)

10.2.3 Compass near a wire

Place a magnetic compass near a straight wire carrying DC. The compass needle deflects, proving the current generates a B-field (Oersted, 1820). Reverse the current and the needle flips.

10.3 Instruments

10.3.1 Gaussmeter / magnetic field probe

A Hall-effect sensor that reads magnetic flux density in millitesla or gauss. Used for industrial EMC assessments, for sizing magnetic shielding, and for checking stray fields near transformers, motors, and bus bars.

10.3.2 E-field probe

Measures electric field strength in V/m. Used for RF exposure assessments, for verifying insulation clearances under HV, and for locating capacitive coupling sources.

10.3.3 Near-field EMI probes

Small sensing loops (magnetic) or capacitive plates (electric) used to sniff out interference sources on a PCB or inside an enclosure. Hold the probe near a suspect component, watch the oscilloscope: the field of that specific source becomes a visible time-domain trace.

10.3.4 Spectrum analyser

Reveals the frequency content of the ambient EM field. In a typical room you can literally see the WiFi, Bluetooth, 4G / 5G, and FM / AM broadcast carriers as separate peaks on the screen. When hunting EMC problems, a spectrum analyser plus a near-field probe narrows down "which component on this board is radiating at this frequency" in minutes.

10.3.5 A simple AM radio

Free, portable, surprisingly effective. Tune an AM radio between stations and walk around a room; near any current-carrying cable it will hum at 50 / 100 Hz. Near a SMPS or a VFD, it will buzz and hiss with the switching harmonics. The field is literally audible.

10.4 Interactive simulators

Browser-based field simulators let you play with Maxwell's equations without any hardware. See the bibliography for one canonical example (Falstad).

10.5 Looking at the field without measuring it: the "visual audit"

When assessing an installation for EMI risk, you can get very far with nothing more than a floor plan, a marker, and a flashlight:

1. **Sketch the floor plan.**
2. **Draw every cable as a line**, labelled with voltage and function (power, signal, motor, control, etc.).
3. **Look for "go and return" pairs that take different physical routes.** Those are your big-loop-area culprits.
4. **Identify long parallel runs (≥ 1 m) of power and signal cables.** Flag any with < 20 cm separation.
5. **Identify "field-noisy" devices:** VFDs, contactors, SMPS power supplies, radio transmitters, welding equipment. They are sources.
6. **Check every crossing.** Parallel crossings are bad; perpendicular ones are fine.

This procedure catches 90 % of real-world EMI problems before any measurement is taken. The remaining 10 % is where a spectrum analyser and a near-field probe earn their keep.

10.6 Take-aways

- Iron filings, compasses, and unconnected fluorescent tubes make fields directly visible — and they are cheap, fast demonstrations worth keeping on a bench.
- A gaussmeter, an E-field probe, a near-field EMI probe, and a spectrum analyser are the core field-engineering instruments. An AM radio is a free-but-surprisingly-useful fifth.
- Most EMI problems can be diagnosed by a visual cable-audit long before any instrument is plugged in.

A closing note on the math

You can close this pamphlet without resolving a single formula. The equations you met along the way were **signposts, not gatekeepers** — each one carries, in symbols, the same meaning as the words beside it; neither side of that pairing is more correct than the other.

They were added for a different reason. The fields — which are conditions of space, not arrows on paper — happen to be among the most precisely *representable* objects in physics: we can write them in symbols, and those symbols *are* the thing itself, not a picture of the thing. It is a claim that took two centuries to sharpen; it felt fair to show the range of what it looks like in practice.

You have therefore met math at three levels along the way:

- **Arithmetic algebra — the easy end.** Ohm's law ($V = IR$), the counting definition of current ($I = \mathrm{d}Q/\mathrm{d}t$), RMS for a pure sine wave ($V_{\text{RMS}} = V_{\text{peak}}/\sqrt{2}$), the power dissipated in a resistor ($P = VI$). The whole pamphlet can be read numerically using just these.
- **A light touch of calculus — the middle of the ladder.** A handful of derivatives and integrals where time matters: $V_L = L\,\mathrm{d}I/\mathrm{d}t$ for the inductor (§6.11), $W = \frac{1}{2}LI^2$ for the energy stored in a coil's magnetic field (§6.11), the RMS integral itself (§5.3), the Poynting flux $\vec{S} = (1/\mu_0)\,\vec{E} \times \vec{B}$ (§7.3). None of these require a calculus course to believe; they only require believing that rates of change, and sums of small pieces, are useful.
- **Maxwell's four equations — the hard end.** Chapter 8 shows you $\nabla \cdot \vec{E}$, $\nabla \times \vec{B}$ and their cousins. You will almost certainly not *work* with those on the job; but seeing them in a single block is the point: four lines that contain every radio wave, every lightning strike, every photon of visible light, and every current in every copper wire in your cabinet.

And then there is a fourth, deeper register of math — one this pamphlet only **names**, and never tries to teach — because it reaches where pictures and plain language cannot. Two sharp examples:

- **Schrödinger's wave function for the electron, $\psi(\vec{r}, t)$.** Inside the copper of your Live wire, each "free" electron is *not* a tiny billiard ball sitting at a definite spot. It is a delo-

calised, standing pattern of probability amplitude — a cloud of *"chance of being here"* spread across many atoms. $\psi(\vec{r}, t)$ is the complex-valued mathematical object that encodes that cloud: a function of position $\vec{r}$ and time t, whose squared magnitude $|\psi(\vec{r}, t)|^2$ gives the probability density of finding the electron at $\vec{r}$ at instant t, and whose evolution under the Schrödinger equation tells you how that probability shifts when an electric field — the same field this pamphlet is about — is applied. The slow *drift* you have been reading about since Chapter 2 is what that shifting probability looks like once you average over the $\sim 10^{22}$ free electrons in every cubic centimetre of copper. No plumbing analogy, no cartoon atom, no arrow-diagram reaches that description. The math does, exactly.

- **Heisenberg's matrix mechanics.** The same quantum physics, in a different dialect: observables (position, momentum, energy) are not plain numbers but **matrices**, and "measuring" them means applying those matrices to a state. Because certain pairs of matrices **do not commute** ($\hat{p}\,\hat{x} \neq \hat{x}\,\hat{p}$), the uncertainty principle $\Delta E \cdot \Delta t \gtrsim \hbar/2$ that we met in §9.13 is a *theorem*, not a slogan — it falls out of the algebra itself.

Both formalisms describe the same electron — the one drifting, barely, through the copper in your wall — and they agree exactly where they overlap. We hint at their existence in §2.8 and §9.13, and we point to Pietro Greco's *Quanti* in §12.5 for a readable, narrative history of how physicists arrived at those ideas. Neither is needed to operate a pump or wire a relay; both are needed if you want to know *what the electron is really doing*, at the deepest level available to human knowledge, inside that wire.

The practical takeaway, then: **if a formula slows you down, trust the words beside it** — they carry the same meaning in a different register. **If the words feel incomplete, come back to the formula** — that is often where the last 10% of the picture lives. And if you ever catch yourself wondering whether the math is *describing* the field or *being* the field, good: physicists have been arguing about that for a hundred years, and you are now in the room.

Chapter 11 — Glossary and symbols

A reference for the terms, acronyms, and symbols used throughout this pamphlet. Entries within each group are ordered roughly in the sequence you are likely to encounter them.

11.1 Voltage, current, and waveform quantities

- **V** — voltage, measured in volts (V).
- **I** — current, measured in amperes (A).
- **R** — resistance, measured in ohms (Ω).
- **P** — electrical power, measured in watts (W). `P = V · I` for DC or in-phase AC; `P = V · I · cos ψ` for AC with a non-zero power factor.
- **f** — frequency in hertz (Hz); 1 Hz = 1 cycle per second.
- **T** — period; `T = 1 / f`. At 50 Hz, `T = 20 ms`.
- **ω** — angular frequency in radians per second; `ω = 2π · f`.
- **Vpeak** — peak amplitude of a sinusoidal waveform.
- **Vpp** — peak-to-peak amplitude; for a sine, `Vpp = 2 · Vpeak`.
- **Vrms** — root-mean-square voltage; for a sine, `Vrms = Vpeak / √2`.
- **EMF** — electromotive force (historical term for voltage generated by a source — a battery, a generator, or Faraday-law induction).

11.2 Field quantities

Vector-valued quantities are marked with a grey half-arrow above the symbol ($\vec{E}$, $\vec{B}$, $\vec{S}$, $\vec{H}$); plain letters without the arrow denote scalars.

- $\vec{E}$ — electric-field vector, V/m (volts per metre).
- $\vec{B}$ — magnetic flux density (a vector field), T (tesla). Also called the B-field.
- $\vec{H}$ — magnetic-field-strength vector, A/m. $\vec{B} = \mu \vec{H}$.
- **Φ (phi)** — magnetic flux (a scalar), Wb (weber). $\Phi = \int \vec{B} \cdot d\vec{A}$, which for a uniform $\vec{B}$ perpendicular to area A simplifies to $\Phi = |\vec{B}|\, A$.
- **ε_0** — electric permittivity of vacuum ≈ 8.854 × 10^{-12} F/m.

- **μ_0** — magnetic permeability of vacuum = 4π × 10^{-7} H/m.
- **ε_r** — relative permittivity of a material ($\varepsilon_{\text{total}} = \varepsilon_r \, \varepsilon_0$).
- **μ_r** — relative permeability of a material ($\mu_{\text{total}} = \mu_r \, \mu_0$).
- $\vec{S}$ — Poynting vector, W/m²; direction and magnitude of EM energy flow. $\vec{S} = (1/\mu_0) \, \vec{E} \times \vec{B}$ (§7.3).
- **c** — speed of light (a scalar) ≈ 3 × 10^8 m/s, with $c = 1/\sqrt{\varepsilon_0 \, \mu_0}$ (§8.5).

11.3 Mains wiring

- **L (Live, Hot, Phase)** — the conductor whose potential swings with the sine wave relative to earth.
- **N (Neutral)** — the conductor bonded to earth at the transformer, normally near 0 V relative to earth.
- **PE (Protective Earth)** — safety conductor, tied to the building's earth electrode.
- **PEN** — a single combined conductor carrying both neutral and protective-earth functions (used in TN-C systems).
- **TT / TN / TN-C / TN-S / TN-C-S / IT** — grounding architectures, see §4.5.
- **L1, L2, L3** — the three line conductors of a three-phase supply.
- **RCD** — Residual Current Device (also called GFCI, ELCB, "salvavita" in Italian). Trips on differential current between L and N.
- **MCB** — Miniature Circuit Breaker. Overcurrent protection.
- **MCCB** — Moulded-Case Circuit Breaker.
- **RCBO** — Residual Current Breaker with Overcurrent (combined RCD + MCB).
- **Surge** — a short-lived, high-amplitude voltage transient on the mains, from lightning or network switching. Test waveform: IEC 61000-4-5 combination wave (1.2 / 50 µs voltage, 8 / 20 µs current). See §4.7.
- **SPD (Surge Protective Device)** — a component that clamps a mains surge by dumping its energy to earth or between L–N. Usually a coordinated stack of GDT, MOV, and TVS. See §4.7.
- **MOV (Metal-Oxide Varistor)** — a zinc-oxide ceramic disc whose resistance collapses sharply above a threshold voltage; the bulk of an SPD's energy capacity. Wears out with every surge. See §4.7 for mains-scale devices and §6.11 for cabinet-scale bidirectional clamps on AC inductive loads.
- **GDT (Gas-Discharge Tube)** — a sealed gas gap that ionises

at its strike voltage; very high energy capacity, slow strike time (hundreds of ns), prone to follow-current on AC. See §4.7.
- **Type 1 / Type 2 / Type 3 SPD** — installation classes from IEC 61643-11: Type 1 at the service entrance (direct-lightning capable), Type 2 at the distribution board, Type 3 at the equipment. Must be coordinated. See §4.7.

11.4 Power, loads, and signals

- **VA** — volt-amperes, unit of **apparent power**. For resistive loads, VA equals watts; for reactive loads, VA > W.
- **var** — volt-amperes reactive, unit of **reactive power**.
- **PF (Power Factor)** — $\mathrm{W} / \mathrm{VA} = \cos \varphi$. 1.0 for a purely resistive load; lower for inductive or capacitive loads. See §6.9 for the load-type definitions.
- **Resistive load** — a device whose current is proportional to the voltage across it and follows it instantly, with no delay; all delivered energy becomes heat or light (heater, incandescent bulb, fixed resistor). See §6.9.
- **Inductive load** — a device whose dominant part is a **winding**; current lags voltage in time (phase lag). Motors, transformer primaries, solenoids, relay coils. See §6.9.
- **Capacitive load** — a device whose dominant part is two conductors separated by an insulator; current *leads* voltage in time. SMPS input filters, long shielded cables seen from the source. See §6.9.
- **Phase lag / phase lead** — the time shift between the current and voltage waveforms of an AC load, expressed as an angle φ. Inductive loads lag (current behind voltage); capacitive loads lead (current ahead of voltage); resistive loads do neither. See §5.5, §6.9.
- **Winding** — a length of insulated wire wound into many turns to form a coil; the building block of every inductor, solenoid, motor stator/rotor, and transformer primary/secondary. First defined in §3.1.
- **Fixed resistor** — a two-terminal passive component whose resistance is a single designed value (not a control input); obeys Ohm's law inside its rated range. See §6.9.
- **THD** — Total Harmonic Distortion, a measure of how much a waveform deviates from a pure sine.
- **VFD** — Variable Frequency Drive; an AC motor speed controller that synthesises a variable-frequency output by switching a DC

bus at kHz rates.

- **SMPS** — Switched-Mode Power Supply. See §9.1.
- **PWM** — Pulse-Width Modulation; a common way of synthesising an average voltage by switching between full on and full off at high frequency. See §9.1.
- **4–20 mA loop** — industrial analog signal: the sensor modulates the current between 4 mA (low reading) and 20 mA (high reading). Robust against voltage drop and broken-wire detection.
- **RS-485, CAN** — differential serial buses used in industrial control.
- **PT100, PT1000** — platinum RTD (Resistance Temperature Detector); resistance increases linearly with temperature.
- **Analog signal** — a voltage or current whose magnitude is continuously proportional to something real (temperature, pressure, flow, valve position). Industrial examples: 4–20 mA, 0–10 V, thermocouple mV, strain-gauge mV. See §9.1.
- **Digital signal** — a voltage (or current) deliberately restricted to two designed levels, with the receiver deciding 0 or 1 against a threshold. Physically it is still an analog voltage; the digital/analog distinction is an interpretation layer, not a separate kind of electricity. See §9.1.
- **Sigma-delta (ΣΔ) ADC** — an analog-to-digital converter that encodes its input as a one-bit pulse stream whose *density* is proportional to the analog value, then digitally low-pass-filters it to recover a high-resolution number; the reverse trick of PWM.
- **MOSFET (Metal-Oxide-Semiconductor Field-Effect Transistor)** — a voltage-controlled semiconductor switch; the workhorse output stage of Arduino and PLC digital outputs driving DC loads. See §6.11 for the flyback case and the parts most likely to blow up.
- **Flyback / freewheeling diode** — a diode wired in parallel with a DC inductive load, reverse-polarised in normal operation; during turn-off it gives the coil's collapsing magnetic flux a closed circulation path, clamping the switching device to $V_f \approx 0.7\,\text{V}$. See §6.11.
- **Zener diode** — a silicon diode designed to conduct in *reverse* above a defined breakdown voltage (the *Zener voltage*), used as a voltage reference, a clamp, or a faster-release variant of the flyback diode. See §6.11.
- **TVS (Transient-Voltage-Suppression) diode** — a silicon

avalanche diode engineered for short, high-current pulses; picosecond clamp time. Unidirectional for DC, bidirectional (two back-to-back in one package) for AC or bipolar signals. Typical home: across MOSFET drain-source, across RS-485 / CAN / USB lines for ESD protection. See §6.11.

- **RC snubber** — a resistor in series with a capacitor, placed across an AC switching contact or load, to absorb the $L\,dI/dt$ pulse and dissipate its energy as heat instead of as contact arcing or ringing. See §6.11.
- **LC circuit (L–C resonator)** — an inductor L and a capacitor C connected together; energy shuttles between the magnetic field of L and the electric field of C at the natural frequency $f_0 = 1/(2\pi\sqrt{LC})$. A sharp edge "kicks" the resonator and it rings briefly before damping out. See §9.1.
- **LC filter** — an L–C low-pass filter: the inductor blocks fast changes in current, the capacitor smooths fast changes in voltage; together they average a chopped square wave back into a near-DC level. The output stage of every PWM drive and SMPS. See §9.1.
- **PCB (Printed Circuit Board)** — a rigid (or flexible) insulating board carrying copper traces, pads, and vias that electrically connect surface-mounted and through-hole components. The standard substrate is FR-4 (flame-retardant class 4 glass-reinforced epoxy).
- **PCB trace / PCB track** — a flat copper conductor etched on a PCB, running between components. "Trace" is American usage, "track" is British; same object. See §9.1.
- **FR-4** — flame-retardant class 4 glass-reinforced epoxy laminate; the standard PCB dielectric ($\varepsilon r \approx 4.3$, dielectric strength ≈ 20 kV/mm).
- **Dielectric** — an insulating material viewed from the field's point of view. The EM wave of a signal on a cable or PCB trace actually propagates in the dielectric between the "go" and "return" conductors, not in the copper. Formally defined in §8.7.

11.5 Electromagnetic interference and compatibility

- **EMI** — Electromagnetic Interference; unwanted EM coupling between circuits.
- **EMC** — Electromagnetic Compatibility; the engineering discipline of designing systems so they do not interfere with each

other.

- **Common-mode** — signal (or noise) that appears equally on both conductors of a differential pair, measured against ground.
- **Differential-mode** — signal (or noise) that appears as the difference between the two conductors of a pair.
- **CMRR** — Common-Mode Rejection Ratio; how well a differential receiver rejects common-mode noise.
- **STP** — Shielded Twisted Pair.
- **UTP** — Unshielded Twisted Pair.
- **Faraday cage** — a closed conducting enclosure that excludes external E-fields from its interior.
- **Ground loop** — an unintended closed loop of ground conductor through which circulating current can flow, coupling noise into sensitive signals.
- **Loop area** — the two-dimensional area enclosed by a circuit's "go" and "return" conductors. By Faraday's law, the voltage induced in the circuit by an external magnetic field is proportional to the rate of change of flux through this area; by reciprocity, the same area sets how efficiently the circuit radiates when its own current changes. The central geometric quantity of Chapter 9 — minimising it is the single most effective EMI defence (see §9.2–§9.4).
- **ESD (Electrostatic Discharge)** — a short, high-voltage, low-energy pulse from an electrostatically-charged body (human finger, synthetic-clothing brush) discharging into a grounded circuit. Classical "human body model": 2–8 kV, tens of A peak for tens of ns. Defeats most semiconductor inputs unless a TVS or equivalent clamp is present. Test standard: IEC 61000-4-2.
- **Stub** — a length of conductor branching off a main signal path and ending in open air, a jumper, or a via. Invisible at DC; acts as an open-circuit transmission-line stub at fast edge rates, causing reflections and ringing. See §9.1, §9.9.
- **Jumper (shunt)** — a small removable U-shaped connector that bridges two adjacent pins of a pin header. Binary configuration element; a stub in disguise. See §9.9.
- **Pin header** — a row of through-hole or SMD pins that exposes internal signals for jumpers, probes, or plug-in cables. See §9.9.
- **DIP (Dual In-line Package) switch** — a small rectangular module of N mechanical ON/OFF slide switches in a DIP footprint; an N-bit integrated jumper. See §9.9.

- **Solder bridge** — two copper pads on a PCB, normally separated, deliberately closed with solder at assembly time for "set once and forget" configuration. See §9.9.
- **0-Ω resistor** — a chip resistor whose nominal value is literally zero ohms; a robot-friendly solder bridge. See §9.9.
- **Bus termination** — a resistor matched to a transmission line's characteristic impedance, placed at the physical endpoints of the line to absorb incoming edges and prevent reflections. On differential buses (RS-485, CAN) a 120 Ω terminator goes at each of the two physical endpoints, never in the middle. See §9.9.

11.6 Standards bodies and common documents

- **IEC** — International Electrotechnical Commission. The main international electrical-engineering standards body.
- **IEC 60364** — low-voltage electrical installations (the international parent of most national wiring codes).
- **CEI** — Comitato Elettrotecnico Italiano. Italy's national standards body, aligned with IEC.
- **CEI 64-8** — the Italian implementation of IEC 60364.
- **IEC 60757** — IEC standard for conductor colour codes.
- **IEC 61000 series** — EMC; particularly relevant:
  - IEC 61000-4-x (immunity testing, e.g. ESD, surge, burst)
  - IEC 61000-6-x (generic emission / immunity standards)
- **IEEE** — Institute of Electrical and Electronics Engineers (the US-rooted, globally-used professional body).
- **CENELEC / EN** — European electrotechnical standards organisation; most IEC standards become EN standards after regional adoption.
- **UL** — Underwriters Laboratories (US product-safety standards).

11.7 Take-aways

Every symbol and acronym in this pamphlet can be decoded from this chapter. When in doubt, work back to the units — they almost always reveal what a quantity physically means.

Chapter 12 — Bibliography and further reading

The references below are organised from "most fun and most accessible" to "most rigorous". Every one of them is worth its time.

12.1 Videos and interactive simulators

- **Veritasium (Derek Muller), *"The Big Misconception About Electricity"* (2021).** https://www.youtube.com/watch?v=bHIhgxav9LY The single most cited demonstration in this pamphlet. Builds the long-wire thought experiment that proves energy travels in the field around the wires, not in the electrons inside them. Watch this first if you watch nothing else.
- **Veritasium, *"How Electricity Actually Works"* (2022).** Follow-up to the above, addressing the community's objections and carefully tightening the argument.
- **3Blue1Brown, *"Divergence and curl: the language of Maxwell's equations"*.** Animated, geometric explanations of the vector-calculus ideas behind Maxwell's four equations. The "divergence" and "curl" videos in particular make Gauss's and Faraday's laws visibly intuitive.
- **The Feynman Lectures on Physics — free online.** https://www.feynmanlectures.caltech.edu/ Volume II covers electromagnetism at undergraduate level, in Feynman's famously conversational style. The chapters on the Poynting vector and on "the vector potential" are especially rewarding.
- **Paul Falstad, *Electromagnetic field simulator*.** https://www.falstad.com/vector/ A browser-based interactive playground: drag charges and currents around and watch fields update in real time. Free, runs in the browser, no installation.

12.2 Canonical textbooks

- **Feynman, Leighton & Sands — *The Feynman Lectures on Physics, Vol. II: Mainly Electromagnetism and Matter* (Addison-Wesley).** The warmest introduction to fields at a physics-undergraduate level.

- **Purcell & Morin — *Electricity and Magnetism*, 3rd edition (Cambridge University Press, 2013).** The most elegant undergraduate text on classical E&M. Derives most classical results from relativistic arguments, which is unusual and illuminating.
- **Griffiths — *Introduction to Electrodynamics*, 5th edition (Cambridge University Press, 2023).** The standard modern physics-department textbook. Rigorous but readable.
- **Jackson — *Classical Electrodynamics*, 3rd edition (Wiley, 1998).** The graduate-level bible. Harder than the others; referenced for when the practical books are not detailed enough.

12.3 Practical electrical and electronic engineering

- **Horowitz & Hill — *The Art of Electronics*, 3rd edition (Cambridge University Press, 2015).** The indispensable practical reference for analog and digital circuits. The chapters on power, AC, rectifiers, and transformers are directly applicable to everything in chapters 3–5 of this pamphlet.
- **Ulaby, Michielssen & Ravaioli — *Fundamentals of Applied Electromagnetics*.** A textbook aimed at engineers rather than physicists; strong on transmission-line theory and antennas.

12.4 EMC, cable design, and grounding

- **Henry Ott — *Electromagnetic Compatibility Engineering* (Wiley, 2009).** The most-cited single reference on EMC. Comprehensive coverage of grounding, shielding, cable design, and filtering.
- **Clayton R. Paul — *Introduction to Electromagnetic Compatibility*, 2nd edition (Wiley, 2006).** More mathematical than Ott; excellent for understanding the field theory that underlies EMC rules.
- **Howard Johnson & Martin Graham — *High-Speed Digital Design* (Prentice Hall, 1993).** Classic “Black Magic” book

on high-speed PCB design. Explains why ground planes, controlled impedance, and short return paths matter, all in field-theoretic terms.

- **Tim Williams — *EMC for Product Designers*, 5th edition (Newnes / Elsevier, 2016).** More accessible EMC engineering with a product-focus.

12.5 Historical and narrative

- **Forbes & Mahon — *Faraday, Maxwell, and the Electromagnetic Field* (Prometheus Books, 2014).** Readable history of how the field concept was developed through the 19th century.
- **Bruce J. Hunt — *The Maxwellians* (Cornell University Press, 1991).** Focused history of how Maxwell's theory was refined and accepted by FitzGerald, Heaviside, Hertz and Lodge after Maxwell's death.
- **Pietro Greco — *Quanti. La straordinaria storia della meccanica quantistica* (Carocci editore, collana "Le sfere" 153, Settembre 2020; ISBN 978-88-430-9956-6, 320 pp.).** Italian-language narrative history of quantum mechanics from Planck and Einstein's light-quantum through Bohr, the "old quantum theory", the Heisenberg/Schrödinger mechanics of 1925–26, and the long Einstein–Bohr debate on realism and completeness. Beyond the scope of this pamphlet (which stays in classical electromagnetism), but the natural next step for any reader who wants to follow what happens to the field picture once quantisation enters the story.

12.6 Standards

- **IEC 60364** — Low-voltage electrical installations. The international parent standard behind most national wiring codes.
- **IEC 60757** — Letter / numerical / colour codes for marking and identifying conductors.
- **IEC 61000 series** — Electromagnetic compatibility (EMC). Particularly IEC 61000-4-x for immunity testing and 61000-6-x for generic emission / immunity limits.
- **IEEE Std 142 ("Green Book") — Grounding of Industrial and Commercial Power Systems.** Extremely practical ref-

erence on industrial grounding architecture.

12.7 A suggested “watch next” sequence for Maxwell

If the Maxwell chapter (chapter 8) has left you wanting to *really* get it:

1. 3Blue1Brown — *Divergence and curl* (10 minutes).
2. Veritasium — *The Big Misconception About Electricity* (15 minutes).
3. Feynman Lectures, Vol. II, chapter 15 (*The vector potential*) and chapter 27 (*Field energy and field momentum*).
4. Open Falstad’s simulator, pick “Plane wave” and “Dipole radiation”, and watch E and B regenerate each other in real time.

That sequence, in that order, takes about two evenings and leaves you with a permanent mental model of why Maxwell’s equations are what they are.

License

Copyright (c) 2026 rf aabrewery@proton.me.

This project is released under two licenses, each matched to what it actually covers:

Pamphlet text and figures — CC BY 4.0

The pamphlet text, the rendered PDF, and every figure authored for this pamphlet (cover and in-text diagrams under `images/`) are licensed under the **Creative Commons Attribution 4.0 International License (CC BY 4.0)**.

You are free to **share** (copy, redistribute) and **adapt** (remix, transform, build upon) the material for any purpose, including commercially, provided you give **appropriate attribution** (name: *rf*, a link to this pamphlet, and an indication of any changes). No additional restrictions.

Full legal text: https://creativecommons.org/licenses/by/4.0/legalcode. Human-readable summary: https://creativecommons.org/licenses/by/4.0/.

THE WORK IS PROVIDED "AS-IS" AND WITHOUT WARRANTIES OF ANY KIND.

Build helpers — MIT License

The Python build scripts `build_em_fields_pamphlet.py`, `make_cover_figure.py`, `make_back_cover_figure.py` and the other in-text `make_*_figure.py` helpers are released under the MIT License:

> Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:
>
> The above copyright notice and this permission notice shall be included in all copies or substantial portions of

the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Scope

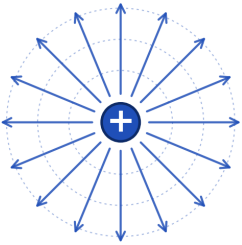
Neither license extends to the third-party references cited in the Bibliography chapter (books, videos, articles, standards); those works retain their own copyright and licenses and are cited for reference only.

Why split?

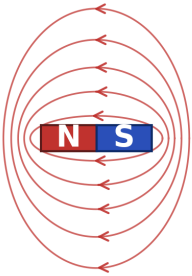
CC BY 4.0 is the standard license for prose, diagrams, and educational content — it is written in the vocabulary of creative works (“share”, “adapt”, “attribution”). The MIT License is written in the vocabulary of source code (“the Software”, “merchantability”, “non-infringement”) and fits the two small Python helpers. Keeping each under its natural license is cleaner than stretching either one beyond its intended scope.

Fields, everywhere

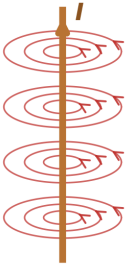
From a single charge to the geomagnetic field and the light between



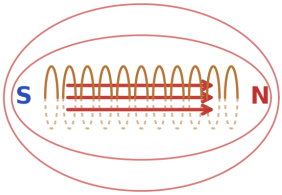
Point charge — radial E-field



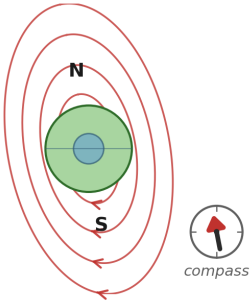
Bar magnet — magnetic dipole



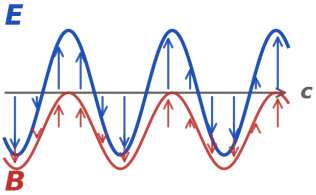
Current-carrying wire — B rings around I



Solenoid — B concentrated inside the coil



Earth — planetary magnetic dipole



EM wave — E and B, propagating at c

The electromagnetic field is the physical medium: always there, at every scale.