The Problem with Graphene

Why the Most Discussed Material of the Century Is Still the Worst Specified

Dr. Gregory S. Carmichael, DBA

Quantum Reserve Press

A buyer’s field guide to engineered carbons: nomenclature, geometry, functionalization, carbon nanotubes, fullerenes, aerogels and twisted bilayers, where strength, conductivity and heat actually come from, dispersion and testing protocols, and a six-section purchase specification.

Contents

ABSTRACT

Graphene is not overhyped. Graphene is under-specified. Those are different failures with different remedies, and confusing them has cost the field roughly fifteen years.

This paper makes four arguments. First, "graphene" is a noun where a specification should be, and the international standards that would fix this now exist and are almost unused. Second, the celebrated intrinsic records — stiffness, strength, thermal conductivity, optical absorption — belong to near-ideal monolayer specimens that virtually nobody buys, and each step toward a commercial powder subtracts a quantifiable amount from those particular numbers. Not every useful property works that way: barrier, sorption, catalysis, and shielding are often *created* by defects, porosity, functionalization, and interfaces rather than degraded by them. Third, the same material serves three engineering regimes that require loadings differing by two orders of magnitude and formulations that actively conflict. Fourth, and most consequentially, outcome is set by process rather than by purchase: the worked calculation in Part VI shows the same nominal material at the same nominal loading producing a **+182%** modulus improvement, a **+48%** improvement, or an **+10%** improvement. Strength is a separate argument with a separate model, and this paper keeps the two apart.

A fifth argument runs underneath the other four. The properties that make these materials famous belong to architectures rather than to substances, and Part VIII takes that to its limit: the same carbon spans roughly **75,000-fold** in thermal conductivity between an aligned film and an aerogel, and two graphene sheets rotated by about one degree become a superconductor. If arrangement can do that, no chemical noun can ever be an adequate purchase specification.

Every load-bearing quantitative claim in this paper is derived from a single parameter model and checked dimensionally at build time. Illustrative ranges and descriptive figures that are not model outputs are registered as explicit exemptions rather than passed off as derived. Where a value rests on an assumption, or on a model applied outside its derivation conditions, that assumption is named in the text.

PART I — WORDS

1.1 The layer ladder

Everything downstream depends on this, and it is genuinely simple.

Graphene is one layer of carbon atoms bonded in a hexagonal lattice, each atom bound to three neighbors. The standards permit the synonyms monolayer graphene and single-layer graphene, abbreviated 1LG, precisely because the bare word is misused so consistently.

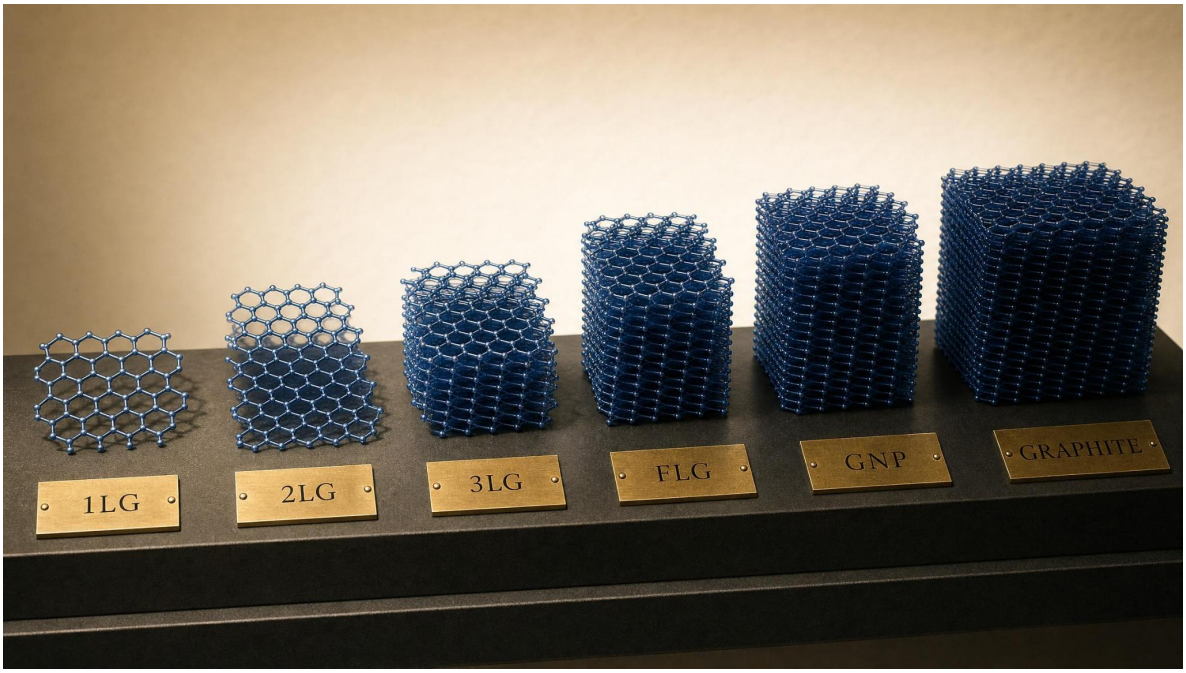

Figure 1. The layer ladder. Left to right: monolayer, bilayer and trilayer graphene, few-layer graphene, graphene nanoplatelet, graphite. The ladder is ordinal and not to scale. The GNP rung follows common usage rather than the standard: ISO fixes the graphene nanoplate by *thickness*, 1 to 3 nm, which overlaps the few-layer regime rather than sitting above it, and material thicker than about 3 nm is a graphite nanoplatelet outside the GR2M family altogether. See §1.1 and §1.3.

Term	Abbrev.	Layers	Note
Graphene	1LG	1	The reference material for every famous number
Bilayer graphene	2LG	2	Stacking registry specified separately
Trilayer graphene	3LG	3	
Few-layer graphene	FLG	3-10	Well-defined stacked layers
Graphene nanoplate / nanoplatelet	GNP	1–3 nm thick	Defined by thickness and lateral size, not layer count; overlaps FLG
Graphite nanoplatelet	—	>3 nm thick	Not a GR2M; what most products sold as GNP actually are
Graphite	—	bulk	Three-dimensional; distinct electronic structure

The 2024 vocabulary standard adds the collective term **graphene-related 2D material (GR2M)**, covering one to ten layers and including graphene oxide, reduced graphene oxide, and functionalized variants. GR2M is the honest umbrella noun, and it is the word most product literature should use and almost never does.

1.2 Why ten is the number

The boundary at ten layers is not administrative convenience. It is approximately where the electronic behavior underlying the standard's electrical definition converges toward bulk graphite.

Graphene's electronic behavior comes from a band structure in which conduction and valence bands meet at a point, with charge carriers behaving as though massless. Add a second layer in registry and that structure changes. Add a third and it changes again. By roughly ten layers the electronic structure has converged to within measurement resolution of bulk graphite. The standard encodes exactly this reasoning: a two-dimensional material of up to **10 layers** for electrical measurements, beyond which the properties are not distinct from the bulk.

Convergence is not uniform across properties, and this matters commercially. Electronic behavior converges quickly. Specific surface area falls as one over layer count and keeps falling well past ten. In-plane stiffness is nearly unchanged with layer count, which is why platelet products still reinforce.

A fifteen-layer platelet is a poor electronic material and a perfectly reasonable mechanical one. The word covers both; the specification must not.

1.3 The GNP problem, stated fairly

First, the term itself carries two meanings, and the field commonly uses the looser one. ISO defines a graphene nanoplate as a nanoplate consisting of graphene layers with a thickness between 1 and 3 nm and lateral dimensions from roughly 100 nm to 100 µm. That is a **morphological** definition keyed to thickness and lateral size, not a layer-count bin, and at 1–3 nm it overlaps the few-layer regime rather than sitting above it. Material thicker than about 3 nm is a graphite nanoplatelet, which is outside the GR2M family entirely. Most powder sold commercially as GNP is 5 to 50 nm thick and is therefore, by the standard's own terms, graphite nanoplatelet. The acronym is the same; the material is not.

With that settled, graphene nanoplatelets carry most commercial volume, and they deserve a defense before a critique. GNPs are produced at scale, handle in existing equipment, and genuinely improve barrier, percolation, tribology, and stiffness in the right formulations. There is nothing wrong with GNP. There is a great deal wrong with selling GNP as graphene.

Three harms follow from the substitution. **Property expectations transfer illegitimately**: a buyer who read the 1.0 TPa figure and received a thirty-layer platelet has no framework to predict its behavior. **Loading calculations break**: reinforcement efficiency and percolation both scale with aspect ratio, so a 1 nm flake and a 15 nm platelet of equal lateral size differ roughly fifteenfold in the loading required. **Failed programs poison the well**: most enterprises that ran a graphene trial, saw two percent improvement, and concluded the material does not work had in fact run a GNP trial at a monolayer-appropriate loading with a dispersion method that never exfoliated anything.

The vocabulary standard addresses this head-on. The term "graphene-enhanced" is **deprecated** except where the filler is actual single-layer graphene; the correct construction is GR2M-enhanced or, for platelets, GNP-enhanced. A standards body does not deprecate a marketing term casually.

1.4 Graphene oxide and its reduction

Graphene oxide is a different material from graphene, not a grade of it. The dominant route oxidizes graphite in strong acid, forcing oxygen-containing groups onto and between the layers: epoxide and hydroxyl on the basal plane, carboxyl and carbonyl at edges. The oxygen pushes layers apart and makes the sheets hydrophilic, which is why GO disperses in water when graphene does not.

The cost is total. Every oxygen attached to the basal plane converts an sp^2 carbon to sp^3 and removes it from the conjugated network. Graphene oxide is an electrical insulator by many orders of magnitude.

Reduced graphene oxide is GO after chemical, thermal, or electrochemical oxygen removal. Reduction restores conductivity but never to pristine values, and the reason is pathway-dependent in a way that is usually glossed over. Some routes are relatively benign: chemical reduction can open an epoxide ring and restore the sp^2 carbon without removing it. Others are not: thermal reduction drives off oxygen as carbon monoxide and carbon dioxide, taking lattice carbon with it and leaving vacancies behind. Either way, residual functional groups, disordered domain boundaries, and — for the carbon-losing pathways — permanent vacancies prevent full recovery. rGO conductivity typically sits two to four orders of magnitude below CVD graphene.

The practical consequence for a buyer is that **"rGO" without a stated reduction route is an incomplete specification**, because the route determines how much of the lattice survived.

Practical rule. A carbon-to-oxygen ratio below about 4:1 means graphene oxide regardless of the label. Above roughly 10:1 you are in rGO or lightly functionalized territory. Pristine graphene sits above 20:1 and is usually reported at 50:1 or better.

1.5 Dimensionality determines function

The most useful organizing idea in engineered carbons is that the interesting materials differ by dimensionality, and dimensionality determines what a filler can physically do.

Class	Dimensionality	Geometry	What it can do in a matrix
Fullerene (C_{60})	0D	~0.7 nm sphere	Molecular; electron acceptor; interphase modification only, no crack bridging
Carbon nanotube	1D	1–50 nm dia., 1–100 µm long	Percolating networks at very low loading; load along one axis; barrier by network rather than tortuosity
Graphene / GNP	2D	1–100 nm thick, 0.1–100 µm wide	Most efficient tortuosity filler; reinforces in two axes; conducts in-plane
Graphite / foam / fiber	3D	bulk	Bulk property contribution by volume

A sub-nanometer sphere has an aspect ratio of one and cannot bridge a crack, so a fullerene cannot reinforce by the load-transfer mechanism that governs fibers and platelets, though it can still alter a polymer through nanoscale interphase effects. A cylinder does not obstruct a diffusion path the way a plate does, so per unit volume a nanotube is a far less efficient tortuosity filler than a platelet, though nanotube composites do achieve large barrier improvements by other routes including network formation and interphase densification.

The distinction worth holding is between **mechanism** and **possibility**. Geometry determines which mechanism is available, and mechanism determines how much material you need and how the benefit scales. It rarely determines that a result is unreachable. Claims of the form "X cannot do Y" are usually claims about efficiency wearing the costume of physics, and this paper is arguing against exactly that habit.

1.6 Who governs the vocabulary

Four distinct kinds of body are routinely conflated, and the conflation is contractually material.

Standards bodies. ISO Technical Committee 229 and IEC Technical Committee 113 produce the technical specifications defining the field's vocabulary and measurement methods. Their output is the only source of definitions with contractual force.

The trade association. The Graphene Council, grown over roughly twelve years to a community exceeding forty thousand materials professionals, relaunched in 2025 as the **Advanced Carbons Council**, headquartered in Zug, with scope widened to nanotubes, nanofibers, carbon fiber, synthetic graphite, nanodiamonds, fullerenes, MXenes, and reclaimed carbons. It runs the Verified Graphene Producer program, the only third-party verification combining material characterization with an in-person facility audit, and it drove the Graphene Classification Framework into **ISO/TS 9651:2025**.

Research institutions. The University of Manchester holds the origin claim and hosts the National Graphene Institute and the Graphene Engineering Innovation Centre. These are research and scale-up facilities with world-class measurement capability. They are not a certifying body.

Metrology institutes. The National Physical Laboratory and its international counterparts publish good practice guides that are, in practice, more immediately usable than the standards themselves, because they explain what goes wrong.

Referring to a single "graphene council in Manchester" merges all three. In procurement the distinction decides what can be written as an acceptance criterion: a Verified Producer designation is a supply-chain assurance, a Manchester report is a measurement, and an ISO number is a definition.

1.7 What the audit actually found

The 2018 study by Kauling and colleagues remains the field's most important commercial document. The team built a systematic protocol combining electron microscopy, atomic force microscopy, Raman spectroscopy, elemental analysis, and X-ray photoelectron spectroscopy, then applied it uniformly to material from **sixty producers** across three continents. They found that worldwide graphene quality was poor and not optimal for most applications, that most companies were in fact producing graphite microplatelets, and that the market's classification scheme was erroneous. The majority of samples contained under ten percent of material meeting the definition.

The proposed remedy is the part usually omitted from the retelling. The authors did not call for enforcement. They called for classification by **distribution function**, reporting layer number and

flake size as statistical distributions rather than single values. That recommendation is now embedded in ISO/TS 21356-1 and ISO/TS 9651.

This is the deepest point in the section. A powder is not a material with a layer count. It is a population with a layer distribution. A drum averaging five layers might be uniformly five-layer material, or half monolayer and half nine-layer, and those two drums behave completely differently in a formulation. **A single number is the wrong shape of answer.**

> **Dinner Table.** Imagine buying lumber where "oak" could mean a solid oak plank, an oak-veneered board, or a bag of oak sawdust, and all three are legally labeled oak at the same price. That is the graphene market. When someone audited 60 producers, most of what was being sold as the miracle material was ordinary graphite ground very fine. The industry then spent years wondering why the miracle material kept disappointing people.

> **So What.** The word on the drum tells you nothing. Buy against ISO/TS 9651:2025, which specifies what has to be measured and reported, and treat a supplier's inability to fill in that data sheet as the most informative thing you will learn about them.

PART II — GEOMETRY

Geometry does more explanatory work in this field than chemistry.

2.1 Aspect ratio is the master variable

Aspect ratio is lateral dimension over thickness. A 10 µm platelet 2 nm thick has an aspect ratio of 5,000; the same lateral size at 30 nm thick gives 333.

Nearly every property of interest scales with it. Electrical percolation threshold falls roughly as one over aspect ratio. Barrier improvement rises roughly linearly with it. Reinforcement efficiency rises until the flake exceeds the critical load-transfer length, then saturates. And melt viscosity rises steeply, which is the constraint that ends most formulation programs.

The last item is the trap. The same geometry that makes a filler effective makes it unprocessable, and the viscosity wall arrives before the property target. Part IX treats this quantitatively.

For barrier performance the **Nielsen tortuosity model** gives a usable estimate: composite permeability relative to the neat matrix is approximately $(1 - \varphi)$ divided by $(1 + \alpha\varphi/2)$, with φ the volume fraction of aligned impermeable platelets and α the aspect ratio.

Aspect ratio	Loading	Permeability reduction
1,000	1 vol%	**6.1-fold**
100	1 vol%	**1.5-fold**

Same weight of filler, same chemistry, four times the benefit, from geometry alone.

Assumption flagged. The Nielsen model assumes perfectly aligned, perfectly impermeable, non-overlapping platelets. Real systems achieve a fraction of the predicted benefit. Use it as an upper bound and a scaling law, never as a design value.

2.2 Specific surface area is a layer-count lie detector

This is the single most useful diagnostic in the field, it costs roughly a hundred dollars per sample, and it is routinely omitted from datasheets.

The graphene unit cell has an area of **0.0524 nm^2** and contains **two carbon atoms**. Dividing the mass of two carbon atoms by that area gives an areal mass density, and the reciprocal, doubled to count both faces, gives the theoretical specific surface area:

> **SSA(monolayer) = 2 ÷ (2 × M_C ÷ (N_A × A_cell)) = 2,627 m^2/g**

Derived from the lattice constants above, the value is **2,627** m^2/g, universally quoted in the literature as 2,630 after rounding. That number is fixed by the lattice, cannot be improved on, and every additional layer in a stack divides it. So:

> **Average layers ≈ 2,627 ÷ measured BET surface area (m^2/g)**

Quoted SSA (m^2/g)	Implied layers	Honest description
750	3.5	Genuine few-layer graphene
500	5.3	Few-layer graphene
300	8.8	Upper edge of FLG
150	17.5	Graphene nanoplatelet
80	32.8	Graphite nanoplatelet
30	87.6	Fine graphite powder

Two honest caveats, because a test this powerful invites misuse. **Restacking**: nitrogen adsorption measures accessible surface, so genuinely exfoliated monolayers that have collapsed in the drum will report the restacked area. High SSA is strong evidence of thin material; low SSA is not proof of thick material, though it is proof that the material as delivered will not present thin-material surface to your matrix, which is the thing you were buying. **Micropore artifacts**: BET can overstate area for microporous carbons, so cross-checking against methylene blue adsorption on a qualification sample is worth the marginal cost.

If a supplier will not report BET surface area, that refusal is itself the datum.

2.3 Lateral size and the critical length

Load does not enter a reinforcing particle at its edge. It enters through interfacial shear and accumulates along the particle's length until it reaches what the particle can carry. The distance required is the **critical length**. A flake shorter than that can never be loaded to capacity; it slips before it stresses.

The Manchester group established the numbers using Raman spectroscopy of the 2D band, whose frequency shifts under strain, allowing stress inside the flake to be mapped against position. Interfacial shear stress for van der Waals adhesion to a polymer was reported across a range of **0.3** to **2.3 MPa** in that work, with the interface breaking down and transferring stress by friction at the upper end, giving a critical length near **3 µm**. Later measurements across graphene/polymer systems span roughly 0.15 to 2.3 MPa depending on layer count, substrate, and strain level. Efficient reinforcement, meaning the flake is loaded to a substantial fraction of capacity along most of its length, requires roughly ten times that.

> **Design rule: below about 3 µm lateral size a flake contributes little; efficient reinforcement wants 30 µm and above.**

Now consider what liquid-phase exfoliation actually produces. Sonication-assisted routes typically yield flakes between 0.3 and 2 µm, because the cavitation that separates layers also fractures the basal plane. The most common laboratory route to thin graphene produces flakes an order of magnitude below the reinforcement threshold.

This is not a contradiction in the literature. Small thin flakes are excellent where surface area dominates: electrochemistry, catalysis, sensing. Large thick platelets are the reverse. There is no grade good at both, and a supplier claiming one is describing an unresolved trade-off rather than a breakthrough.

2.4 Stacking order: the variable nobody asks about

Two flakes can have identical layer counts, lateral sizes, and chemistry, and disperse completely differently, because of how their layers are rotationally aligned.

In **AB (Bernal) stacking**, layers sit in crystallographic registry, as in natural graphite, at an interlayer spacing of **0.335 nm**. Registry maximizes van der Waals binding and makes the stack genuinely hard to separate.

In **turbostratic** stacking, layers are rotationally misaligned. There is no registry, spacing expands to roughly 0.342 to 0.344 nm, and interlayer attraction drops substantially. Turbostratic material exfoliates far more readily in solvents and in compounding equipment. Flash-graphene literature values it explicitly for this: rapid exfoliation on mixing during composite formation, which is not achievable from AB-stacked material.

Stacking order is measurable and belongs on the datasheet. Raman gives the diagnostic — the M peak arising from interlayer coupling is present in AB-stacked material and absent in turbostratic material, which instead shows TS1 and TS2 features. XRD confirms through the expanded spacing and the collapse of three-dimensional reflections into two-dimensional bands.

For bulk composites, **turbostratic six-layer material may outperform AB-stacked three-layer material**, because what reaches the polymer matters more than what left the reactor.

2.5 Provenance determines capability

Route	Layers	Lateral size	Defects	Scale	Best use	Watch for
Mechanical exfoliation	1	10–100 µm	Near zero	Micrograms	Research reference only	Not a commercial route
CVD on metal foil	1–3	Continuous film	Low, grain boundaries	m^2/yr	Transparent electrodes, membranes	Transfer introduces tears, residue, wrinkles
Liquid-phase exfoliation	1–10	0.3–2 µm	Low basal, high edge	Tonnes	Inks, coatings, electrochemistry	Surfactant residue; below reinforcement threshold
Electrochemical exfoliation	1–8	1–10 µm	Moderate	Tonnes	Conductive additives	Intercalant residue; batch variability
Oxidation–reduction	1–3	0.5–10 µm	High, permanent	Tonnes	Barrier, adsorbents	Residual oxidant metals; conductivity never recovers
Thermal expansion of graphite	5–50	5–50 µm	Low	Tonnes	GNP for bulk composites, tribology	Aspect ratio below claim; heavily restacked
Flash Joule heating	3–10, turbostratic	0.1–5 µm	Low	Kg to tonnes	Bulk composites, cement, plastics	Feedstock variability; halogen carryover
Plasma / arc / pyrolysis	2–15	Sub-µm	Variable	Tonnes	Conductive additives, energy storage	Amorphous fraction; wide distributions

Two entries deserve comment.

CVD produces the highest-quality material in the world and stands behind most record property measurements. It also produces film on a metal substrate, so every application requires a transfer step, and transfer is where quality goes to die: tearing, polymer residue from the support layer, wrinkling. When a device paper reports performance well below intrinsic values, transfer damage is the first suspect.

Flash Joule heating deserves attention from anyone holding a waste stream. Short high-current pulses convert carbon-rich feedstock — coal, coke, carbon black, tires, food waste, mixed plastics — into turbostratic graphene in milliseconds without solvent, at a reported electrical energy cost near **7.2 kJ** per gram. The output is not display-grade and is not trying to be. It is bulk composite filler with unusually favorable dispersion behavior, made from something otherwise going to landfill.

The strategic dimension is worth stating plainly. Natural flake graphite and its downstream processing are concentrated in a small number of jurisdictions, and graphite sits on the United States critical minerals list. A route converting domestic waste carbon into composite-grade material is a supply-chain position as much as a manufacturing process. It does not produce electronic-grade graphene and should not be marketed as though it does.

> **Dinner Table.** Two things decide whether a carbon filler works, and neither is what it is made of. The first is shape: how wide the flake is compared to how thick. The second is how easily the stack comes apart when you stir it. A wide thin flake blocks moisture, stiffens plastic, and carries current at a fraction of the loading a fat one needs. And there is a hundred-dollar test — surface area — that tells you within minutes whether you bought thin flakes or ground-up pencil lead, which is why so few sellers volunteer it.

> **So What.** Put three geometric requirements in the purchase order before you discuss chemistry: BET surface area with an implied layer count, lateral size as a distribution rather than a single value, and stacking order by Raman. Those three numbers predict formulation behavior better than anything else on a datasheet.

PART III — CHEMISTRY

3.1 Pristine graphene does not want to be in your resin

A perfect graphene sheet is about as inert as carbon gets. Every carbon is satisfied by three sigma bonds and a delocalized pi system: no dangling bonds, no polar sites, nothing for a polymer to grip beyond weak dispersion forces.

Two consequences follow, and both are fatal to naive formulation. **It does not wet**, because its surface energy is poorly matched to most polar polymers and to water, and a material that does not wet does not disperse. **It restacks**, because two sheets in contact are bound across their entire face, and for large flakes that integrates to a substantial energy. Any dispersion of pristine graphene is driving thermodynamically uphill, and the moment shear stops the system starts walking back down.

This is why solvent selection dominates liquid-phase exfoliation, with N-methylpyrrolidone, dimethylformamide, and ortho-dichlorobenzene recurring throughout the literature: their surface tension approximately matches graphene's surface energy, minimizing the enthalpy of mixing. It is also why they are unusable industrially, given boiling points, toxicity, and disposal costs that dwarf the value of the material dispersed.

3.2 The trade-off that governs everything

> **Every *covalent basal-plane* modification that improves dispersion destroys some fraction of the conjugated network that gave graphene its properties.**

Attaching a functional group to a basal-plane carbon converts it from sp^2 to sp^3, removing it from the delocalized electron system. Electrical conductivity depends entirely on that system. Thermal conductivity depends on a coherent phonon lattice the same modification disrupts. Mechanical strength depends on an unbroken bond network.

Every strategy is therefore a position on a spectrum:

Strategy	Mechanism	Dispersion	Conductivity retained	Interface	Where it belongs
None (pristine)	—	Very poor	~100%	Very weak	Nowhere, without help
Surfactant	Physisorption	Good	Intrinsic kept, junctions blocked	Weak	Aqueous inks where residue can be removed
Pyrene / π–π stacking	Non-covalent aromatic	Good	~95%	Moderate	Conductive composites needing dispersion
Polymer wrapping	Non-covalent	Good	~90%	Moderate	Nanotube sorting, stable dispersions
Silane grafting	Covalent, edge-preferential	Good	50–80%	Strong	Epoxy, cementitious systems
Amine functionalization	Covalent into cure network	Good	40–70%	Very strong	Structural epoxy
Maleic anhydride grafting	Covalent via compatibilizer	Good	Variable	Strong	Polyolefins
Oxidation (GO)	Covalent, basal plane	Excellent	<0.1%	Strong	Barrier, adsorbents, never conductors

The table answers the question of why certain things seem to work and others do not, and it does so as a spectrum rather than a wall. A formulator who chose an oxidized grade because it dispersed beautifully and then measured conductivity has not discovered that graphene fails to conduct. He has discovered that he bought an insulator, from the far end of a spectrum whose middle he never examined.

The qualifier matters, because there is an exception and it is the useful one. **Non-covalent routes are the underused middle**. Pyrene-based dispersants adhere to the basal plane through pi-pi stacking without converting a single sp^2 carbon to sp^3, so the conjugated network survives intact while dispersion improves substantially. The cost reappears elsewhere rather than vanishing: an adsorbed layer sits between flakes and adds junction resistance, which §6.4 shows is worth orders of magnitude at the contact even when the sheet itself is untouched. The trade is real, but it is a trade between *lattice damage* and *junction penalty*, not a law that dispersion must cost conjugation. The trade is that attachment is reversible and can be defeated by temperature, solvent, or time, so it must be qualified against service conditions rather than mixing conditions.

3.3 Matching chemistry to matrix

Functionalization is not a quality grade. It is a matching problem whose correct answer changes entirely with the matrix.

- **Epoxy:** amine functionalization, because the amine participates in the cure and produces a covalent bridge from filler into network. Silanes as the workhorse alternative.
- **Polyolefins:** nothing polar will work. Use maleic-anhydride-grafted polyolefin as compatibilizer, or long-alkyl-chain functionalization.
- **Polyamides and polyesters:** carboxyl or amine groups can react during polycondensation; in-situ polymerization often beats melt compounding.
- **Cementitious systems:** carboxyl and hydroxyl are favorable because the mechanism is nucleation, not load transfer. Watch for sulfonate dispersants interfering with superplasticizers already present.
- **Elastomers:** silane coupling as in silica-filled rubber; the mechanism is hydrodynamic reinforcement plus filler network.
- **Water-based coatings:** surfactant or GO, accepting conductivity loss, or a pyrene dispersant if conductivity is required.

The failure mode is buying “functionalized graphene” as though functionalization were a single thing. Ask what group, at what surface density, attached where, measured how. A supplier unable to answer at that resolution designed the material for a datasheet rather than for your matrix.

Dinner Table. Graphene is chemically close to a nonstick pan: very little adheres to it, which is exactly why it resists mixing into anything. The common fix is to roughen the surface chemically, and roughening it is precisely what destroys the electrical and thermal magic. That is the central tension. There is one partial way around it — coat the sheet rather than scratch it, using molecules that cling flat to the surface without breaking any bonds — and it works, at the price of a thin insulating film at every point where two sheets touch. So the trade-off can be moved and reshaped, but not deleted, and a salesperson who claims to have deleted it has stopped describing where it went.

So What. Decide which single property you are actually buying before you choose a grade, because the chemistry that delivers dispersion for a barrier film disqualifies the same material for a conductor. If you need both, look at non-covalent pyrene dispersants and qualify them against service temperature, not mixing temperature.

PART IV — TUBES

4.1 Chirality: you are buying a statistical mixture

A single-wall nanotube is a graphene sheet rolled into a cylinder, and the direction of the roll, described by chiral indices (n,m), determines the electronic structure entirely. The rule is arithmetic: when (n − m) is divisible by three the tube is metallic or nearly so, otherwise it is semiconducting with a bandgap varying inversely with diameter.

In unsorted synthesis, chirality is essentially random. **Roughly one third of the tubes are metallic and two thirds semiconducting.** This is statistics, not a defect.

The implication rarely appears in commercial literature. Buying a kilogram of unsorted single-wall tubes for a conductive application means two thirds of the purchase sits in the percolating network as a resistive element. Chirality-sorted material exists, produced by density-gradient ultracentrifugation, polymer wrapping, or aqueous two-phase extraction, and costs one to three orders of magnitude more. For bulk composites the unsorted mixture is the correct economic choice; for anything where electronic character actually matters, unsorted material cannot work and buying it wastes the program.

4.2 Walls: why double-wall tubes exist

Type	Walls	Diameter	Cost	Percolation loading	Character
SWCNT	1	0.7–2 nm	Highest	0.01–0.05 wt%	Best properties, hardest to disperse, mixed chirality
DWCNT	2	2–4 nm	High	0.05–0.2 wt%	Functionalize outside, preserve inside
MWCNT	5–50	10–100 nm	Lowest	0.1–1 wt%	Workhorse; sword-in-sheath failure

The middle row exists for a specific and elegant reason. Functionalizing a single-wall tube to make it disperse necessarily damages the only wall it has. A double-wall tube separates the jobs: **the outer wall takes the chemistry and the interface while the inner wall retains an intact electronic and mechanical structure.** For applications needing both dispersion and retained conductivity, DWCNTs are frequently the right answer and are almost never considered, because most buyers know only single-wall and multi-wall.

4.3 Sword-in-sheath

A multi-wall tube is concentric cylinders held by van der Waals forces with no covalent bonding between walls. Loaded in tension inside a composite, the matrix grips only the outer wall. That wall carries the load, reaches its limit, and fails; the inner walls then slide out of the broken sheath carrying essentially nothing.

The arithmetic is unforgiving. A twenty-wall tube has roughly five percent of its cross-section available to carry tensile load. Every reported MWCNT composite strength falling far below prediction is partly explained here before dispersion is even discussed.

The consequence is that MWCNTs are excellent for electrical percolation, where only the outer wall needs to conduct and conduction runs along the network rather than through walls, and structurally inefficient in tension.

4.4 The sonication paradox

Nanotubes ship as entangled agglomerates, because high aspect ratio means tubes wrap around each other during synthesis. The standard laboratory response is probe sonication, where collapsing cavitation bubbles tear the tangle apart.

Cavitation also cuts tubes. Extended sonication reduces a 10 µm tube to a 1 µm tube, dropping aspect ratio from 1,000 to 100. Since the **percolation threshold** scales inversely with aspect ratio, the loading required rises by roughly an order of magnitude. **The dispersion step destroyed the property that justified the material.**

Assumption flagged. The inverse-aspect-ratio scaling of percolation threshold is reliable, but excluded-volume theory does not pin the prefactor, which varies by more than an order of magnitude with orientation and contact assumptions. Published thresholds are an order-of-magnitude guide; the threshold must be measured for the specific system.

The signature is diagnostic and common: conductivity rises with sonication time, peaks, then falls. The rise is dispersion winning; the fall is tube-cutting winning. Almost nobody maps this curve, and it is a two-day experiment that determines the entire formulation.

The same logic applies to graphene by a different mechanism, with cavitation fracturing the basal plane below the reinforcement threshold. Notably, ISO/TS 21356-1 caps bath sonication at ten minutes for sample preparation, warning that longer times cause basal-plane scission and further exfoliation. The preparation step changes the material being measured.

Alternatives that shear without cutting: three-roll milling, which applies high shear in a controlled gap and is standard for epoxy; high-shear rotor-stator mixing, itself a demonstrated scalable exfoliation route; and industrially, purchasing a pre-dispersed masterbatch from someone who has already solved the problem.

4.5 Catalyst residue

Nanotubes grow on metal catalyst particles, typically iron, cobalt, or nickel on an alumina or magnesia support. Cheap material has not been thoroughly purified, and residual metal in commercial multi-wall material commonly runs from under one percent to over fifteen percent by weight.

This matters four separate ways. It **corrupts electrical measurements**, since metal conducts and shifts percolation curves. It **catalyzes polymer degradation**, particularly oxidative degradation of polyolefins during melt processing. It **disqualifies the material for electronics**, where mobile metal ions are a reliability problem. And it **dominates the toxicology**, since much of the biological reactivity reported in early studies tracked with iron content rather than with carbon.

Thermogravimetric analysis gives residual ash for a few hundred dollars and ICP-MS speciates it. Neither is optional on a qualification sample, and ash content should be a contractual specification rather than a curiosity.

> **Dinner Table.** Nanotubes arrive as a knot, like a bag of tangled Christmas lights, and the standard way to untangle them is to blast them with ultrasound. The problem is that the blasting also chops them into shorter pieces, and length was the entire reason they worked. So people untangle their material into uselessness and conclude the material failed. Meanwhile a third of the expensive kind are semiconductors rather than metals no matter what you paid, and the cheap kind fail like a sword pulling out of its sheath, with only the outermost layer ever doing any work.

> **So What.** Map conductivity against sonication time before fixing a process, because the curve peaks and then falls and the peak is your answer. Better, buy a masterbatch and let someone else own the

> dispersion problem. And require residual ash content by thermogravimetric analysis on every lot.

PART V — BALLS

5.1 What a fullerene is

C_{60} was identified in 1985 and earned the 1996 Nobel Prize in Chemistry. **Sixty carbon atoms** arranged as **twenty hexagons** and **twelve pentagons** form a closed cage roughly **0.7 nm** across.

The pentagon count is not a design choice. Euler's polyhedron formula requires exactly twelve pentagons to close a hexagonal sheet into a sphere regardless of how many hexagons are present. Every closed fullerene, from C_{60} through C_{84} to a nanotube endcap, has exactly twelve.

The pentagons also create the chemistry. Curvature forces carbon atoms out of planarity, and this pyramidalization strains the pi system and makes C_{60} electron-hungry in a way flat graphene is not. C_{60} readily accepts up to six electrons.

5.2 What fullerenes are good for, and what they are not

That electron affinity produced the one genuinely successful industrial fullerene application. PCBM, a soluble C_{60} derivative, served for two decades as the electron acceptor in bulk-heterojunction organic photovoltaics. It was displaced from the research frontier around 2018 by non-fullerene acceptors with better absorption and tunability, but the chemistry was real and it worked.

What fullerenes cannot do is reinforce by the mechanism that governs fibers and platelets. A 0.7 nm sphere has an aspect ratio of one: it cannot bridge a crack, cannot form a percolating network at any sensible loading, and cannot carry load transferred along its length.

What it can do is change the polymer around it. Experimental C_{60}/epoxy work reports modulus increases in the range of roughly 6 to 20 percent and tensile strength gains in the low teens, achieved through interphase effects — restricted chain mobility near the particle, altered local crosslink density, and nucleation — rather than through load transfer. Those are real gains obeying a different scaling law. They do not benefit from a high-aspect-ratio load-transfer mechanism, because a sphere's aspect ratio is fixed at one, and they saturate quickly with loading rather than scaling with volume fraction the way a percolating or load-bearing filler does.

The useful framing is therefore mechanistic rather than categorical: a fullerene is an interphase modifier, not a reinforcement, and it should be specified, dosed, and measured as one.

Related structures worth naming so they are not confused: **endohedral fullerenes** with an atom trapped inside, of genuine interest for quantum sensing and as candidate qubits; **carbon nano-onions**, concentric shells used in lubrication and electrochemistry; **carbon quantum dots**, which are mostly not fullerenes despite the shared size regime; and **graphene quantum dots**, small graphene fragments where edge chemistry rather than the basal plane dominates.

Two cautions. Fullerenes remain a specialty chemical priced by the gram, not a bulk material. And the marketing of C_{60} in oil as a human supplement traces to a single 2012 rodent study that was a toxicology control producing an unexpected result; it does not constitute a human health finding.

> **Dinner Table.** Buckyballs are soccer balls made of sixty carbon atoms, and they are genuinely beautiful chemistry. They are the wrong shape for crack bridging: you cannot reinforce concrete with marbles, however strong the marbles are, because that job needs something long enough to span the crack. They can still stiffen a plastic somewhat, by changing how the polymer behaves in the thin shell of material around each ball — a different trick with a different ceiling, and one that stops paying as you add more. Buckyballs found one real industrial job, as the electron catcher in early plastic solar cells, and have mostly been replaced there.

> **So What.** Classify any proposed carbon additive by shape first — sphere, tube, or plate — then ask which mechanism the claimed benefit requires. Do not reject a spherical additive merely because it cannot bridge a crack. Reject it when the claimed benefit depends on fiber- or platelet-style load transfer, which its geometry cannot supply, and price it as an interphase modifier when that is what it actually is.

PART VI — WHERE THE PROPERTIES ACTUALLY COME FROM

6.1 The reference numbers and their fine print

Property	Value	Specimen	What breaks it
Young's modulus	**1.0 TPa** (1,000 GPa)	Suspended monolayer, AFM indentation	Little; stiffness survives stacking
Intrinsic strength	**130 GPa**	Defect-free suspended monolayer	Any defect; any grain boundary
Breaking strain	**25%**	Same	Same
Thermal conductivity	**4,840–5,300** W/m · K	Suspended monolayer	Contact with anything
Electrical conductivity	**~10^6 S/m**	Clean supported monolayer	Junctions, adsorbates, defects
Optical absorption	**2.293%** per layer	Suspended monolayer	Nothing; set by a fundamental constant

The mechanical figures come from pressing an AFM tip into a graphene membrane suspended over a hole about one micron across, on a specimen selected to be defect-free over that area. The number is the strength of a perfect crystal at micron scale. It is not the strength of anything in a drum.

The optical figure is the most remarkable of the set. A single layer absorbs π times the fine structure constant of incident visible light, **2.293%**, essentially independent of wavelength. Optical absorption of graphene is fixed by a fundamental constant of nature. It also means a monolayer is 97.7 percent transparent, which contradicts everyone's intuition. Graphene powder looks black because a rough powder scatters light many times before it escapes, not because the material absorbs strongly. Hold that for Part VII.

6.2 Strength and stiffness: the worked calculation

Take the modified rule of mixtures:

> **E_c = η_o · η_l · E_f · V_f + E_m · (1 − V_f)**

with η_o the orientation efficiency factor, η_l the length efficiency factor, E_f filler modulus, E_m matrix modulus, V_f filler volume fraction. Orientation factors for platelet reinforcement are 1 for perfectly aligned platelets and **8/15** for randomly oriented platelets in three dimensions. The corresponding value for randomly oriented **fibers** in three dimensions is **1/5**, and the two are frequently confused in the composite literature because the same symbol carries different derivations for different filler geometries. Length efficiency approaches 1 when flakes greatly exceed critical length and falls toward 0.5 as they approach it.

Assumption flagged. The 8/15 factor is derived for randomly oriented platelets under a specific orientation distribution and stress-transfer model. It is used here as the realistic-case value. A genuinely in-plane-random state, such as a calendered film or an aligned laminate, sits between 8/15 and 1 and would need its own derivation rather than reuse of this factor.

Set the case. Epoxy at **3.0 GPa**, filler at 1,000 GPa, densities **2.2** and **1.2** g/cm^3, so one percent by weight converts to **0.55** percent by volume.

Case	η_o	η_l	Effective V_f	E_c	Change
A — Perfect aligned, large flakes, full exfoliation	1.00	1.00	0.55 vol%	**8.46 GPa**	**+182%**
B — Good practice 3D random platelets, moderate flakes	8/15	0.50	0.55 vol%	**4.44 GPa**	**+48%**
C — Illustrative poor dispersion assumed **20%** effective exfoliation, remainder at matrix modulus	8/15	0.50	0.11 vol%	**3.29 GPa**	**+10%**

Assumption flagged. Case C's 20% effective exfoliation fraction is a chosen illustrative value, not an empirical population statistic. It is the only *unsupported* parameter in that row, and deliberately so: the unexfoliated remainder is assigned the **matrix** modulus. That is not the absence of an assumption — it is a conservative null assumption, namely that unexfoliated material contributes no stiffness beyond the matrix volume it displaces. The sensitivity table below is what that choice costs if it is wrong. An earlier version of this model gave the remainder an assumed effective modulus of 10 GPa, which meant two unsupported parameters jointly produced the published number while the text credited only the exfoliation fraction. That is the failure mode this paper is about, committed by this paper, and the correction is to remove the parameter rather than to footnote it.

Because the assumption is worth seeing rather than asserting, here is what happens if the agglomerated remainder is treated as a stiff inclusion after all:

Effective modulus assigned to the unexfoliated remainder	Case C result
Matrix modulus, **3.0 GPa** — conservative null assumption (headline case)	**+10%**
10 GPa	**+11%**
30 GPa	**+14%**
60 GPa	**+18%**

The conclusion survives the whole range, which is the point of showing it. Substitute your own measured exfoliation fraction and the row recomputes.

Note also what this table does **not** show. It is a stiffness model only. The tensile strength consequence of poor dispersion comes from a separate fracture-mechanics argument in §6.3 and does not follow from the rule of mixtures, which contains no flaw population and cannot predict strength. The two arguments are independent and should be evaluated separately.

Read the table carefully, because it is the paper's central argument. The **nominal feedstock and the loading are identical** in all three rows; what differs is the post-process state. Row A is aligned with large flakes, row B is randomly oriented with moderate ones, row C is mostly unexfoliated. Those are outcomes of the process, not properties of the drum, which is exactly the distinction the paper is arguing for. The difference between a **+182%** improvement and an **+10%** improvement is process, not purchase — orientation, effective flake length, and how much of the material was actually exfoliated. Those are variables the buyer controls and rarely measures.

The gap to the intrinsic 1.0 TPa has a second component: the interface. Load enters the flake by shear, and interfacial shear strength for van der Waals adhesion runs on the order of a megapascal while the flake is capable of 130 GPa. **The composite is not limited by the graphene. It is limited by the glue.** This is why covalent functionalization buys real strength even at the cost of conductivity, and why the honest answer to "should I functionalize" is "which property are you selling."

6.3 Why strength falls while modulus rises

This pattern appears constantly and confuses people badly. It is not a contradiction but two measurements probing different things.

Modulus is a volume-averaged, small-strain property. Every particle contributes stiffness in proportion to its volume and its own stiffness, and an agglomerate is still stiffer than epoxy, so modulus rises anyway. Modulus is tolerant of bad dispersion.

Strength is controlled by the single worst flaw. Linear elastic fracture mechanics gives failure stress as K_{IC} divided by Y times the square root of pi times flaw size, where Y is a geometry

factor depending on crack shape and position. Y = 1 is the central through-crack in a wide plate; a simple edge crack is closer to Y = 1.12; $\mathbf{Y = 2/\pi \approx 0.64}$ is the standard factor for a small internal penny-shaped crack. For epoxy at $\mathbf{0.6\ MPa \cdot m^{1/2}}$, taking the central through-crack as the reference case:

Agglomerate size	Predicted failure stress
20 µm	**75.7 MPa**
50 µm	**47.9 MPa**
100 µm	**33.9 MPa**
200 µm	**23.9 MPa**

Neat epoxy tensile strength is typically **70 to 80 MPa**. Setting failure stress equal to the neat strength gives a critical flaw size that depends on which crack idealization you adopt: roughly **16 µm** for an edge crack, **20 µm** for a central through-crack, and **50 µm** for an internal penny-shaped crack. Note the direction: the edge-crack factor is the least forgiving of the three, so Y = 1 is a reference case rather than a conservative one. The defensible statement is a band, not a line:

> **The agglomerate size at which strength begins to fall sits in the tens of micrometers for this epoxy baseline. Above it, strength falls as the inverse square root of flaw size.**

Assumption flagged. Three assumptions carry this result and each should be checked before it is used as a specification. The geometry factor Y changes the answer by more than a factor of three across the three standard idealizations above, and nothing in the physics selects one of them for you.

The figures are also not directly interchangeable with a measured particle diameter, because a denotes a different dimension in each idealization: half the crack length for a central through-crack, the full crack length for an edge crack, and the radius for a penny-shaped crack. Read **16, 20 and 50 µm** as characteristic flaw dimensions a under their respective idealizations, and convert deliberately before comparing them to anything counted on a micrograph.

Finally, an agglomerate is not a sharp crack. Treating it as one is a sharp-crack surrogate whose bias should be calibrated rather than assumed conservative, since agglomerate morphology, internal voiding and interfacial debonding all determine how crack-like the real defect behaves. And the calculation uses a single baseline epoxy; a tougher matrix moves the band upward as the square of its fracture toughness.

What survives the caveats is the scaling and the method. Strength is flaw-controlled and falls as the inverse square root of the worst agglomerate, so a dispersion specification should be written as a maximum agglomerate size derived from your own matrix toughness and strength, not as a visual judgment. Run the calculation for your system, state Y, and put the number in the specification. That is the transferable part. It also explains why a composite can show a 40 percent modulus gain and a 30 percent strength loss in the same specimen, with both results correct.

At low loadings, a poorly dispersed filler is not a reinforcement. It is a population of engineered defects.

6.4 Electrical conductivity: junctions, not material

Graphene conducts at roughly 10^6 **S/m**; epoxy at roughly 10^{-13} S/m. Nineteen orders of magnitude separate them, and a composite typically lands near 10^{-2} to 10^0 S/m.

Current does not flow through graphene in a composite. It flows through a network of graphene, and at every junction it crosses a thin polymer layer. Electrons cross by tunneling, and tunneling resistance rises exponentially with gap width. For a barrier height of **one electron volt** the decay constant computes to **10.2** per nanometer, so:

> **Each additional nanometer of polymer between two flakes multiplies junction resistance by roughly 28,000; each additional 0.1 nm multiplies it by 2.8.**

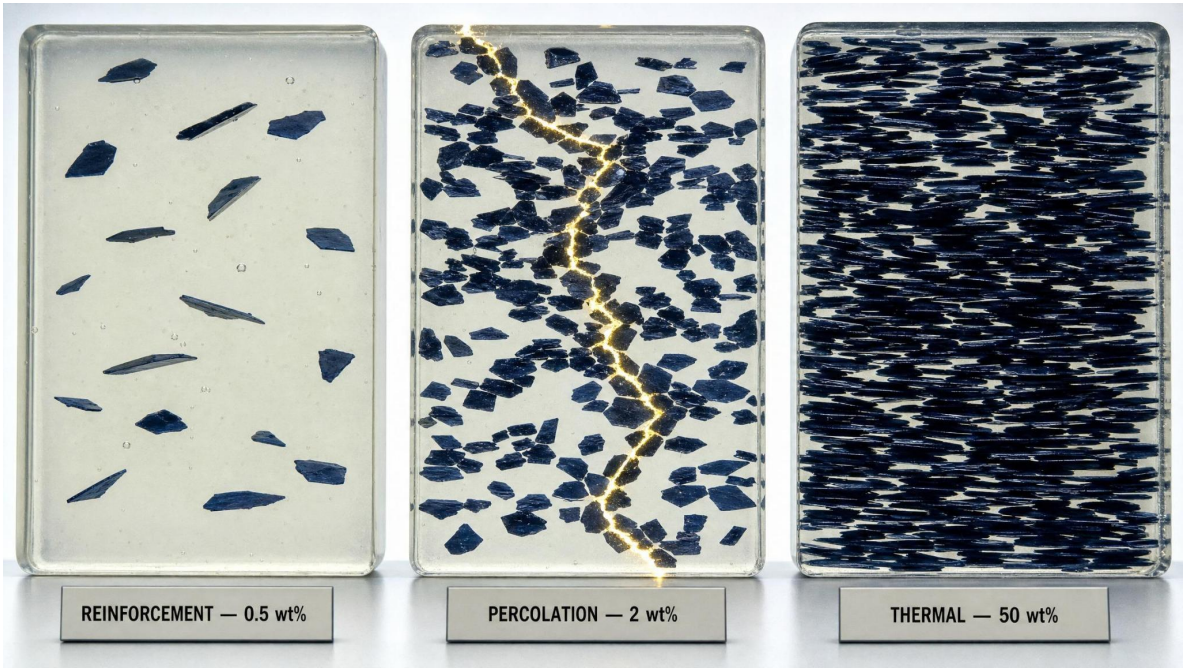

Figure 2. Three regimes, three loadings, one material. Mechanical reinforcement, electrical percolation and thermal conduction call for loadings two orders of magnitude apart, and a formulation optimized for one is usually disqualified for the others. The illuminated path in the centre panel is a single connected route across the specimen: below the percolation threshold no such route exists, and above it the junctions rather than the filler set conductivity. See §6.4.

Which is why a surfactant that disperses beautifully can destroy conductivity. A one-nanometer surfactant shell is invisible in a micrograph, thermodynamically excellent, and worth four orders of magnitude of resistance. This single mechanism explains most "great dispersion, no conductivity" results in the field.

Above threshold, conductivity follows a power law in excess volume fraction, with exponents near 1.3 in two dimensions and 2.0 in three, though tunneling-dominated systems routinely measure much higher.

Filler	Aspect ratio	Typical threshold
SWCNT	>3,000	0.01–0.05 wt%
MWCNT	~1,000	0.1–1 wt%
High-AR GNP	~1,000	0.1–0.5 wt%
Standard GNP	50–200	1–3 wt%
Carbon black	~1	10–20 wt%

And then the target question, which almost nobody asks first:

Application	Requirement	Realistic loading
Static dissipation	10^6–10^9 Ω/sq	0.5–2 wt%
Antistatic	10^4–10^6 Ω/sq	1–3 wt%
EMI shielding, 30 dB	>1 S/m bulk, adequate thickness	5–15 wt%
Resistive heating	10–1,000 S/m	10–25 wt%
Replacing a metal conductor	$>10^5$ S/m	Not achievable in a polymer composite

Most real commercial requirements are electrostatic dissipation, comfortably achievable at low loading with modest material. Programs fail when a project scoped for dissipation is sold a solution formulated for shielding, or when someone quietly hopes to displace copper.

One measurement trap. Aligned platelets produce anisotropic conductivity, commonly ten to a hundred times higher in-plane than through-thickness. An in-plane four-point measurement and a through-

thickness two-point measurement on the same sample can differ by three orders of magnitude, and both can be reported honestly. Always specify direction, method, and contact scheme.

6.5 Thermal conductivity: the property that behaves differently

This is the most misunderstood area in engineered carbons, with a single root cause: people assume heat behaves like electricity.

> **Electrons tunnel. Phonons do not.**

That sentence generates every difference that follows.

Thermal percolation is weaker, less abrupt, and less reliable than electrical percolation. Electrical conductivity jumps ten orders of magnitude at a sharp threshold, because once flakes are within a nanometer or two electrons cross the gap by tunneling. Heat carried by lattice vibration has no equivalent mechanism, and the matrix already conducts heat reasonably well, so there is no insulator-to-conductor transition for the filler network to trigger. In most filled polymers thermal conductivity therefore rises smoothly and roughly linearly with loading.

Assumption flagged. The absence of a sharp thermal threshold is the common case, not a universal law. If your architecture creates continuous covalent paths rather than point contacts, expect nonlinear behavior and measure for it.

The literature is genuinely divided on whether a thermal threshold exists at all. Many systems show none. Others, generally those with well-connected or covalently welded filler networks, show threshold-like nonlinear increases. The safe planning assumption is that **you will not get an electrical-style jump**, and that any nonlinearity you do get will be modest and will depend on achieving genuine network connectivity rather than mere proximity.

The contrast ratio is far smaller. Graphene-to-epoxy electrical contrast is around 10^{19}. Thermal contrast against a **0.2 W/m · K** matrix is around 10^4 at absolute best and closer to 10^3 for a real platelet. Physics with a contrast ratio of a thousand behaves entirely differently from physics with a contrast ratio of ten quintillion.

The dilute limit is bounded. For randomly dispersed inclusions of infinite conductivity, Maxwell's classical result caps enhancement at $1 + 3\varphi$. At one volume percent that is a **+3%** improvement, with a perfect filler, perfectly dispersed, and zero interface resistance. High-aspect-ratio platelets beat the sphere bound through shape factor, but the ceiling is low and real interfaces claw back much of the gain.

The famous number degrades in stages. Suspended monolayer graphene measures near 5,000 W/m · K. Place the same monolayer on silicon dioxide and it falls to roughly **600 W/m · K**, because substrate coupling scatters the flexural phonons carrying much of the heat. Encase it fully in a matrix and it falls toward **160 W/m · K**. The material has not changed; the boundary conditions have. By the time your graphene is embedded in epoxy, which is the only condition in which a composite exists, the relevant conductivity is one to two orders of magnitude below the headline.

System	Thermal conductivity
Neat epoxy	0.2 W/m · K
Epoxy + 1 wt% GNP	0.25–0.35 W/m · K
Epoxy + 10 wt% GNP	0.6–1.2 W/m · K
Epoxy + 30 wt% GNP	2–4 W/m · K
Epoxy + 50–60 wt% GNP, aligned	5–15 W/m · K (through-plane far lower)
Aligned GNP film, no matrix	500–1,500 W/m · K in-plane
Copper, for reference	400 W/m · K

The last rows show where graphene genuinely wins thermally, and it is not as a filler. Compressed aligned graphene and graphite films used as heat spreaders in consumer electronics beat copper in-plane at a fraction of the weight and have shipped in enormous volume. They work because they are

nearly all graphene with almost no matrix, oriented along the direction heat must travel. That is a material, not a composite, and the distinction is the whole lesson.

> **Between roughly 2 and 15 W/m · K a filled polymer can reach the target only at 30 to 60 wt% loading, paying for it in viscosity, processability, cost, and mechanical properties. Above roughly 15 W/m · K it cannot get there at all. In both regimes the question to ask is whether a continuous graphitic path — aligned film, foam, freeze-cast network, or graphite sheet — is the better architecture.**

6.6 Modulus, compression, and anisotropy

Graphite's elastic constants tell the story. In-plane stiffness is roughly **1,060 GPa**, governed by covalent bonds. Out-of-plane stiffness is roughly **36 GPa**, governed by van der Waals forces. A factor of **29** separates the two directions in the same crystal.

The consequence is direct. Aligned platelets stiffen a composite substantially in the plane of alignment and barely at all perpendicular to it. Compressive properties through thickness improve far less than tensile properties in-plane, and often not at all. A datasheet reporting a large modulus improvement without stating direction, alignment state, and test method has reported an unusable number.

For elastomers the **Guth-Gold** relation estimates hydrodynamic reinforcement as $1 + 0.67 \cdot f \cdot \varphi + 1.62 \cdot f^2 \cdot \varphi^2$, with f the aspect ratio. At aspect ratio 100 and one volume percent that predicts a **3.3-fold** modulus increase, with the quadratic term contributing more than the linear one. That squared dependence is why elastomer reinforcement is exquisitely sensitive to how much filler was broken down during mixing.

Assumption flagged. Guth-Gold was derived for rod-like fillers and is applied to platelets by analogy. It captures the scaling and is unreliable as an absolute prediction.

Cementitious systems work by a different mechanism entirely. Graphene in concrete at 0.01 to 0.1 percent by weight of cement does not reinforce by load transfer. It nucleates calcium silicate hydrate, refines pore structure, and accelerates hydration. The mechanism is chemical seeding, which is why dosages sit two orders of magnitude below polymer loadings and why the optimum is sharp: excess filler interferes with hydration and entrains air. Anyone applying polymer intuition to concrete will overdose by a factor of a hundred and conclude the material does not work.

> **Dinner Table.** Same bag of material, same amount, three different outcomes: nearly triple the stiffness, half again the stiffness, or almost no stiffness benefit at all. Nothing about the purchase changed. Only the mixing changed. And if that badly mixed batch also left clumps above a certain size, a **separate** calculation — about cracks, not stiffness — explains why the part can end up snapping sooner as well. Two different questions, two different models, and it is worth keeping them apart. And stiffness and strength are not the same test: stiffness is an average over the whole part, so clumps still count, while strength is decided by the single worst clump in it, the way a chain is decided by one link. That is why a composite can get stiffer and weaker at the same time, and why both numbers are honest.

> **So What.** Audit the compounding step, not the certificate of analysis. Derive your own maximum agglomerate size from your matrix's fracture toughness and strength as in §6.3, state the crack geometry factor you assumed, and put that number in the specification rather than judging dispersion by appearance. And know which of the three regimes — reinforcement, percolation, or thermal — you are buying before choosing a loading.

PART VII — LIGHT

Vantablack is the most instructive case study in this paper, because the property everyone attributes to carbon nanotubes does not belong to carbon nanotubes.

7.1 What it is

Vantablack, an acronym for Vertically Aligned NanoTube Array, was developed by Surrey NanoSystems and announced in 2014. It absorbs up to **99.965** percent of incident radiation from ultraviolet through near-infrared. In 2019 an MIT team reported a coating reaching **99.995** percent. The material is a forest of aligned, equally spaced, high-aspect-ratio nanotubes roughly **30 micrometers** tall, grown by chemical vapor deposition directly on a substrate at around **400** °C in the Surrey process.

7.2 Why it is black, which is not why you think

Three effects combine, and none is exceptional material absorptivity.

Index matching at the entrance. The forest is mostly empty space, with a fill factor of a few percent, so its effective refractive index sits very close to that of air. Light arriving at the surface meets almost no index discontinuity and almost nothing reflects. A polished black surface, by contrast, reflects four percent or more before absorption is even in play.

Multiple internal scattering. Light that enters is deflected among the tubes many times, each encounter absorbing a fraction, until essentially nothing survives to exit.

Geometric trapping. The spacing and height of the forest make the probability of a photon finding its way back out vanishingly small.

Blackness is therefore a structural property of the architecture, not a chemical property of the carbon.

7.3 The lesson

Test the claim. If blackness came from the material, dispersing the same nanotubes in a binder and painting it should reproduce it. It does not, and the reason is documented in the patent literature: coating a nanotube in polymer makes absorption far less efficient and produces far higher reflectance. The polymer raises the effective index, eliminating the match, and creates a smooth surface that reflects specularly.

Surrey's own sprayable variant uses randomly oriented rather than vertically aligned tubes and delivers performance approaching but not equal to the grown array. Commercial ultra-black paints reach around **99.4** percent by brush. That sounds close to 99.965 percent. It is a factor of seventeen more reflected light.

> **A property demonstrated in one architecture does not transfer to another architecture of the same material.**

Suspended graphene conducts heat at 5,000 $\mathrm{W/m \cdot K}$; embedded graphene does not. Defect-free graphene reaches 130 GPa; agglomerated graphene concentrates stress. An aligned nanotube forest absorbs 99.965 percent of light; the same tubes in paint do not come close. In every case the chemistry is constant and the architecture carries the property.

Recall the fine-structure result from §6.1: a single graphene layer absorbs 2.293% of visible light and transmits the rest. Graphene is nearly transparent. Carbon nanomaterials are not intrinsically black; certain arrangements of them are nearly perfect light traps.

> **Dinner Table.** The blackest material ever made is not made of unusually black stuff. It is a lawn of microscopic tubes, and light that wanders into the lawn simply never finds its way back out, the way a shout dies in a dense forest. Grind up that same lawn and stir it into paint and you get ordinary black paint, because you destroyed the forest and kept only the trees. This is the whole field in one example: the magic is in the arrangement, not the ingredient.

> **So What.** Whenever a supplier cites a record property, ask what architecture produced it and whether your product can reproduce that architecture. If the record came from an aligned, suspended, or contin-

> uous form and you are buying powder for a matrix, the number does not transfer and should not enter your business case.

PART VIII — ARCHITECTURE: AEROGEL AND THE MAGIC ANGLE

Part VII argued that architecture carries the property. Two systems demonstrate it at the extremes, and both are worth understanding by anyone specifying carbon materials, because between them they set the outer bounds of how much of a material's behavior geometry can own.

Graphene aerogel takes the best thermal conductor ever measured and turns it into a thermal insulator. Magic-angle twisted bilayer graphene takes two ordinary graphene sheets and turns them into a superconductor when one is rotated by about one degree, the carrier density is tuned by electrostatic doping, and the device is held within a couple of degrees of absolute zero.

8.1 Graphene aerogel

What it is. A three-dimensional open network of graphene or reduced graphene oxide sheets, mostly empty space. Densities run from roughly **0.16 mg/cm^3** at the record end — light enough that the solid is lighter than the air filling it — up to a hundred milligrams per cubic centimeter for structural grades.

How it is made, and why that decides everything. Three routes dominate, and they differ in the one variable that matters:

- **Freeze-casting.** A graphene oxide hydrogel is frozen so ice crystals template the pore structure, then freeze-dried and reduced. Cheap, scalable, tunable pore architecture. The sheets meet at **physical contacts** held by van der Waals forces and residual functional groups.
- **CVD on a sacrificial template.** Graphene is grown on nickel foam and the nickel is etched away, leaving a **seamless covalent network**. Expensive, geometry limited by the template, and electrically and thermally far superior.
- **Direct ink writing.** Extruded graphene oxide inks, then reduction. Architectural control at the hundred-micron scale, junction quality between the other two.

The junction is the whole story, and it is the same story as §6.4 and §6.5 in a different costume. A freeze-cast aerogel is a network of superb conductors connected by poor contacts. A CVD foam is a network of superb conductors connected by covalent bonds. They look identical in a photograph and behave nothing alike.

The result that should be on the wall of every formulation lab. Bulk graphene aerogel has a thermal conductivity of roughly 0.02 to **0.2 W/m · K**. It is used as **insulation**. The endpoint used here is deliberately conservative: a measured **4 mg/cm^3** aerogel has been reported at **0.0047** to **0.0059** W/m · K at room temperature, roughly eighty percent below still air, with the authors attributing it to dominant interface thermal contact resistance — the same junction mechanism this section is about, arrived at independently. Meanwhile an aligned graphene film, made of the same carbon, reaches up to **1,500 W/m · K** in-plane, with reported values spanning roughly 700 to 1,500 W/m · K depending on annealing schedule, thickness and sheet size.

> **The same elemental carbon framework, radically different material state and architecture, and a span of about 75,000-fold in thermal conductivity between these two selected endpoints.**

Assumption flagged. Architecture dominates that comparison but is not the only thing that differs between the endpoints. The aligned film is a densified, largely defect-free graphitic stack; low-density aerogels are typically reduced graphene oxide carrying residual functional groups and a defect population of their own, and reported aerogel conductivities span a wide range with density, reduction chemistry, and junction quality. The ratio is an honest illustration of how far geometry can move a property. It is not a controlled experiment isolating geometry.

That single comparison makes the paper's argument more forcefully than anything else in it. If you specify "graphene" and expect thermal performance, you have specified nothing at all — the same purchase order can deliver a heat spreader or a firebreak.

The counter-case, which matters. Anisotropic freeze-cast aerogels, aligned and then densified, reach far higher conductivity along the alignment direction. So porosity alone is not the villain. The governing variables are **junction quality and orientation**, and an architecture that fixes both recovers much of the performance. This is the constructive version of the same lesson.

Mechanics. Graphene aerogels are remarkably compressible, with many formulations recovering from 90% compressive strain. Modulus sits in the kilopascal range, and it scales with density raised to a power typically between 2 and 3. Density, not chemistry, is the master mechanical variable. A supplier improving their aerogel's modulus by densifying it has not improved the material.

Where aerogels genuinely win. Electromagnetic shielding, where the porous structure promotes multiple internal reflection and produces absorption-dominated shielding at very low density. Sorption, where oil uptake of tens to hundreds of times the aerogel's own weight is routine. Electrodes and catalyst supports, where accessible surface area is the point. And piezoresistive sensing, where the very junction weakness that ruins thermal transport becomes the transduction mechanism, since compressing the network changes contact resistance dramatically.

Two measurement traps specific to this class.

Specific properties flatter. Aerogel literature reports density-normalized values — specific modulus, specific shielding effectiveness, specific strength. Dividing by a density one-thousandth of the competing material produces spectacular numbers that say nothing about whether the part will work. **Always require the absolute value, the density, and the sample thickness together.** A shielding effectiveness figure without thickness is meaningless.

Monolith properties do not transfer to reinforced forms. A free-standing aerogel monolith and the same aerogel encapsulated in a fiber blanket are different materials. Reinforcement changes thermal conductivity, opacity, and compression response, and any optical property of the monolith is generally lost entirely. Specify and test the form you will actually deploy, not the form the datasheet was measured on.

> **Dinner Table.** Take the best heat conductor ever measured, whip it into a foam that is 99.9% air, and you get one of the better insulators available. The element is the same. The bonding, the defect state, the density and the architecture are not: the foam is usually made by a route that leaves oxygen and broken bonds behind. What dominates the difference, though, is how the pieces connect, because heat only travels as fast as it can cross the joints. It is the difference between a copper bar and a box of copper filings: same metal, and only one of them will burn your hand.

> **So What.** For any aerogel or foam, specify junction formation route — freeze-cast, templated CVD, or printed — because it predicts transport performance better than the material name does. Reject density-normalized figures and require absolute value, density, and thickness together. And qualify the encapsulated or reinforced form you will actually use, never the bare monolith.

8.2 The magic angle

Naming first, because the field's term is specific. Two graphene monolayers stacked with a relative rotation are called **twisted bilayer graphene**. At certain special rotations they are called **magic-angle twisted bilayer graphene**, abbreviated MATBG. "Twisted pair" is a cabling term and will send you to the wrong literature.

The result. In 2018 Cao, Jarillo-Herrero and colleagues at MIT stacked two graphene monolayers with a relative twist of about **1.1°**. The rotation creates a moiré superlattice, and at that angle the electronic band structure develops nearly flat bands near zero Fermi energy. Flattening a band quenches the kinetic energy of the electrons, which leaves Coulomb interaction dominant, which produces strongly correlated behavior. The stack showed correlated insulating states at half filling, and on electrostatic

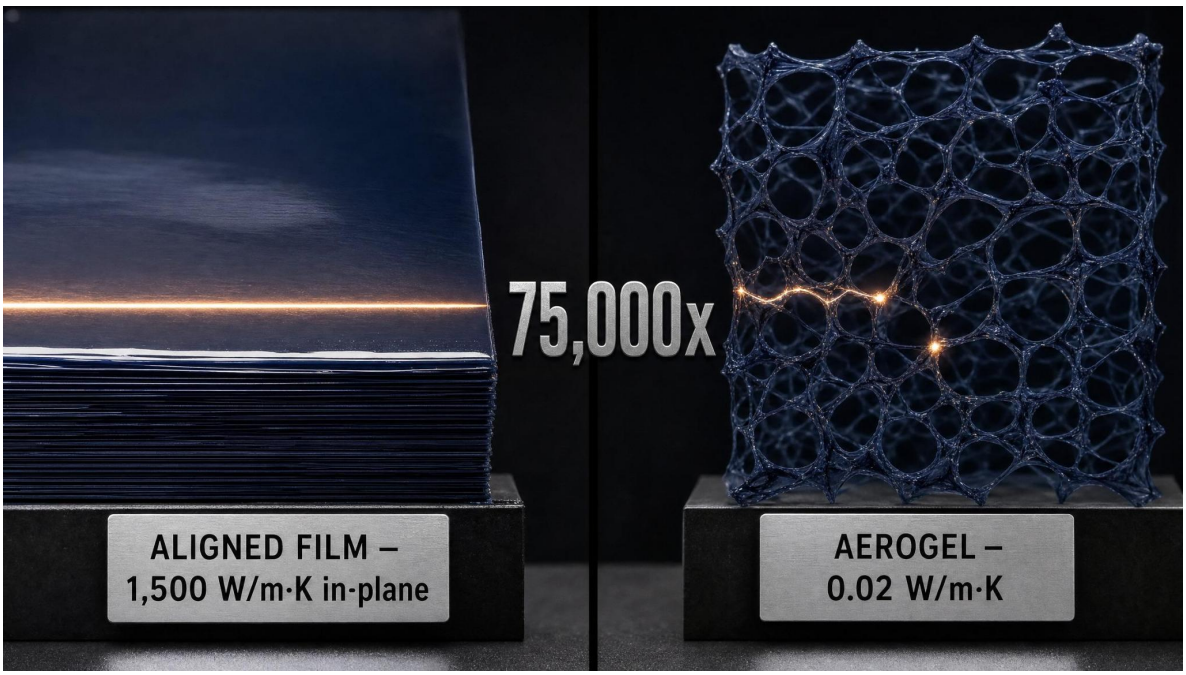

Figure 3. The architecture span. The same element and the same bonding class, arranged two ways, separated by roughly 75,000-fold in thermal conductivity. Architecture dominates the comparison but is not the only difference between the endpoints: the film is a densified, largely defect-free graphitic stack, while low-density aerogels are typically reduced graphene oxide carrying residual oxygen and a defect population of their own. It is an illustration of how far geometry can move a property, not a controlled experiment isolating it. See §8.1.

doping it became a **superconductor with a critical temperature up to 1.7 K**.

The angle was not discovered by trial and error. Bistritzer and MacDonald had predicted in 2011, from a continuum model of the moiré bands, a discrete series of magic angles with the largest and most severely flattened near 1.05°.

The sensitivity is the point. One device at **1.33°**, well off the magic angle, reported superconductivity with a critical temperature near 150 mK, roughly **11-fold** below the 1.7 K of the original magic-angle devices.

Assumption flagged. That ratio compares two specific devices and should not be read as a coefficient on angle alone. Critical temperature in these systems also depends on hydrostatic pressure, heterostrain, carrier filling, and dielectric environment, and devices off the magic angle have exceeded 3 K under pressure. What the comparison supports is that twist angle is a first-order variable, not that a quarter of a degree by itself accounts for the whole difference.

> **Two carbon sheets, identical chemistry, identical layer count. The difference between an ordinary semimetal and a superconductor is a rotational alignment held to about a tenth of a degree.**

What has followed. Magic-angle twisted trilayer graphene, reported in 2021, shows stronger and more tunable superconductivity. Superconductivity has been reported in twisted quadrilayers and pentalayers. And rhombohedral, or ABC-stacked, multilayer graphene has emerged as a separate crystalline route to correlated superconducting states that requires no twist at all, which is worth watching precisely because it removes the alignment problem that currently keeps MATBG confined to laboratory devices.

What this is not, stated plainly. A critical temperature of 1.7 K is 1.7 degrees above absolute zero, colder than liquid helium. MATBG is not a room-temperature superconductor, not a power-transmission technology, and not a materials-supply story. Fabrication of the established devices uses two mechanically exfoliated monolayers, assembled by a tear-and-stack technique that gives both halves

a common crystallographic origin, rotated to about a tenth of a degree, encapsulated in hexagonal boron nitride, at device dimensions of microns. Scalable routes using CVD-grown material are an active research direction and are not the basis of the results quoted here. Twist angle drifts during fabrication and is not uniform across a device. There is currently no scalable powder analogue that reproduces the controlled twist, encapsulation, carrier density, and cryogenic conditions responsible for the effect, and none is in prospect on the timescale of a procurement decision.

The claim this section exists to preempt. Turbostratic graphene, discussed in §2.4, is randomly twisted multilayer graphene. A commercial turbostratic powder therefore contains, somewhere in its angular distribution, regions near the magic angle. That fact is true and it is worthless. Those regions are nanoscale, isolated, undoped, unencapsulated, disconnected from one another, and at room temperature. The powder does not superconduct, and no processing step available at industrial scale supplies the angular control, encapsulation, doping, and temperature the effect requires. Any marketing that connects turbostratic powder to magic-angle physics is trading on a shared word, which is precisely the failure this paper is about.

Why a specification paper should care about a cryogenic laboratory curiosity. Because it settles the ceiling on the argument. If one degree of rotation separates a semimetal from a superconductor in a two-layer stack of identical atoms, then no chemical noun can ever be an adequate specification for a carbon material. Layer count, lateral size, defect density, stacking order, and orientation are not supplementary details attached to the word "graphene." They are the material. The word is a family name.

> **Dinner Table.** Stack two sheets of graphene and rotate one by about one degree — roughly the angle between the twelve and the one on a clock face, divided by five — and the pair becomes a superconductor, carrying current with no resistance at all. In one device rotated a quarter degree further, the effect came back about elevenfold weaker. Nothing was added and nothing was removed; the atoms are identical. Alignment is the first-order variable, though pressure, strain, and how the sheet is electrically tuned move the result as well. It also only works a couple of degrees above absolute zero, and only when the sheet is electrically tuned to the right carrier density, so nobody is wiring a city with it. As a demonstration that arrangement beats ingredient, though, nothing else comes close.

> **So What.** Treat any claim connecting bulk powder to magic-angle physics as disqualifying, since the effect requires angular control and encapsulation that no powder process can deliver. More broadly, use this case as the ceiling argument in any procurement conversation: if one degree of rotation changes everything, a chemical name can never be a sufficient specification, and the distribution-based data sheet is the only defensible alternative.

PART IX — DISPERSION AND PROCESS

9.1 Why dispersion fails, mechanism by mechanism

Thermodynamic reversal. Dispersion is uphill. Shear supplies the energy; when shear stops, van der Waals attraction and Brownian motion begin reassembling agglomerates. A dispersion perfect off the mill can be measurably worse after eight hours in a drum and worse again after a cure cycle.

Cure-phase flocculation. In thermosets, viscosity drops sharply on heating before gelation, and that low-viscosity window is exactly when particles have the mobility to reaggregate. It occurs after all mixing has ended, so dispersion measured in liquid resin does not predict dispersion in the cured part. This mechanism alone accounts for a large share of irreproducible epoxy results.

The viscosity wall. Rheological percolation, the loading at which the filler network dominates flow, occurs at lower loading than electrical percolation, typically by a factor of two to five. Viscosity becomes unmanageable before conductivity arrives. Above roughly two percent GNP in epoxy, viscosity rises steeply enough to prevent wetting, degassing, and infusion. **The process limit arrives before the property target**, and this is the most common hard stop in real programs.

Void entrainment. High-viscosity systems trap air, and one percent void content costs roughly five to ten percent of interlaminar shear strength in a laminate. A nanofiller raising viscosity enough to

add two percent voids has cost more strength than it added, invisibly.

Filtration. In liquid composite molding, fiber reinforcement acts as a filter. Agglomerates strain out at the flow front, leaving the inlet loaded and the outlet starved. The part carries a gradient nobody specified, and coupon testing on a central section will not find it.

9.2 Method selection

Method	Shear	Cuts particles	Scale	Best for
Probe sonication	Very high, local	Severely	Lab only	Sample prep, with time limits
Bath sonication	Low	Moderately	Lab	Redispersion, per ISO limits
Three-roll mill	High, controlled gap	Minimally	Kg to tonnes	Epoxy, pastes, inks
High-shear rotor-stator	High	Moderately	Tonnes	Aqueous and solvent dispersions
Ball / planetary mill	High impact	Severely	Tonnes	When defects are acceptable
Twin-screw extrusion	Moderate, brief	Moderately	Tonnes	Thermoplastics; screw design critical
Masterbatch let-down	Low	No	Any	The correct industrial answer in most cases
In-situ polymerization	N/A	No	Batch	Highest quality; least flexible

Masterbatch deserves emphasis. Buying a pre-dispersed concentrate and letting it down in existing equipment transfers the hardest problem to whoever solved it. An entire midstream sector now exists to occupy this position, and its existence is the market's verdict on where the difficulty lies. The trade is a higher price per kilogram of contained filler and reduced formulation freedom. In most programs that is the right trade, and the instinct to buy powder because powder is cheaper is usually an expensive instinct.

9.3 The scale-up discontinuity

A typical laboratory result uses 100 mg of filler, probe-sonicated in solvent for two hours, cast and cured slowly, at a specific energy input on the order of 10^6 to 10^7 joules per kilogram of filler. A typical production run uses 100 kg compounded in a twin-screw extruder with a thirty-second residence time, at 10^3 to 10^4 joules per kilogram.

Three to four orders of magnitude separate them. The published result was produced under conditions that cannot be reproduced in manufacturing at any price. This is not a criticism of researchers investigating material physics. It is a criticism of readers treating the paper as a recipe.

> **No composite result should be believed as a manufacturing prediction unless dispersion energy is reported and is achievable in the intended process.**

Requiring specific energy input in the methods section would improve the literature's industrial usefulness more than any other single change.

9.4 Batch-to-batch variability

The failure that kills industrial adoption is not that the material does not work. It is that it works in June and not in September, from the same supplier under the same part number.

Liquid-phase exfoliation is a statistical process whose output distribution drifts with feedstock graphite, sonication amplitude, temperature, and centrifugation. Nobody publishes control charts, and almost no buyer demands them.

The remedy is ordinary manufacturing discipline applied to a material the industry treats as exotic: incoming inspection on every lot against a fixed specification, statistical process control on key parameters, retained samples, and a contractual right of rejection. This is how carbon black, silica, and

titanium dioxide are bought. That engineered carbons are bought differently is a symptom of an immature supply chain, not of an unusual material.

> **Dinner Table.** Stirring nanomaterial into resin is like stirring cocoa into cold milk. It looks mixed while you are stirring and separates the moment you stop, and it separates again when you heat it. Worse, the amount of stirring energy a laboratory uses is thousands of times what a factory line can deliver in the seconds it has, so a published recipe simply cannot be followed at scale. The honest industrial answer is usually to buy the material already mixed into a concentrate and dilute it, which is why an entire industry now exists to do exactly that.

> **So What.** Demand specific energy input in any supplier trial data and reject results you cannot reproduce in your own equipment. Default to masterbatch rather than powder, and put lot-level incoming inspection and a right of rejection in the supply agreement.

PART X — TESTING

Testing splits into two questions with two different answers: what did I buy, and what happened to it in my process.

10.1 Incoming: what did I buy

The sequence follows ISO/TS 21356-1 for structure and ISO/TS 23359 for chemistry.

Test	Measures	Diagnostic value
Raman	Layer number, defect density, stacking order	Primary tool. I(2D)/I(G) above 2 with a single Lorentzian 2D peak near 30 cm^{-1} width indicates monolayer; broadening and multi-component 2D indicates stacking. I(D)/I(G) gives defect density and crystallite size via Tuinstra-Koenig. Absent M peak with TS1/TS2 present indicates turbostratic stacking.
XRD	Stack thickness, interlayer spacing	Fastest screen. A sharp strong (002) peak near 2θ = **26.5°** means graphite is present. Scherrer gives stack height; divide by **0.335 nm** for layer count. Spacing expanded toward 0.343 nm indicates turbostratic.
BET surface area	Accessible surface	Layer-count cross-check against 2,627 m^2/g. See §2.2 caveats.
XPS	C/O ratio, sp^2/sp^3 fraction, functional groups	Confirms or refutes oxidation and functionalization claims. C1s deconvolution separates sp^2 near 284.5 eV, sp^3 near 285.3, C-O near 286.5, C=O near 287.8, O-C=O near 289. The π–π^* shake-up near 291 eV is direct evidence of intact conjugation.
TGA	Oxidation onset, ash content	Residual catalyst and support. Essential for nanotubes. Oxidation onset discriminates GO from rGO from graphene.
ICP-MS	Elemental impurities	Speciates the ash. Required for electronics and for toxicology.
AFM	Flake thickness distribution	Reference method for layer count, but apparent monolayer height varies from 0.4 to 1.0 nm with substrate and tip. Statistical only; measure hundreds of flakes.
TEM + diffraction	Layer count, crystallinity	Definitive. Edge counting resolves layers; diffraction spot intensity ratios distinguish monolayer from bilayer.
SEM image analysis	Lateral size distribution	Measure 500+ particles. Laser diffraction assumes spheres and misreports platelets badly.
Zeta potential / sedimentation	Dispersion stability	Predicts dispersion shelf life and aqueous behavior.

Two things to insist on contractually. Every structural result reported as a **distribution with a stated sample size**, never a single number, per the recommendation the 2018 audit made and the standards adopted. And the **ISO/TS 9651:2025** technical data sheet template as the reporting format, which exists precisely to make supplier claims comparable.

10.2 After mixing: what happened to it

Test	Measures	Why it matters
Optical microscopy, thin section	Agglomerates above 5 µm	Cheapest dispersion check. Count and size against the matrix-specific fracture-mechanics limit derived in §6.3, not against appearance.
SEM, cryo-fracture surface	Agglomerates, voids, pull-out vs. fracture	Distinguishes interfacial from cohesive failure directly.
TEM, microtomed section	Nanoscale dispersion, exfoliation state	Ground truth on whether exfoliation actually happened in your process.
Micro-CT	3D agglomerate and void map	Non-destructive; catches gradients and filtration effects coupon testing misses.
Raman 2D-band shift under strain	Actual stress transfer	The definitive test for whether load reaches the filler. A strong shift rate means the interface is transferring load; a weak one means it is slipping. Very few organizations do this, and it answers the question every other mechanical test only approaches.
Rheology, low-frequency G'	Filler network state	A low-frequency storage modulus plateau is the fingerprint of a percolated network. Sensitive, quantitative, fast.
DMA / DSC	Glass transition shift	Indicates constrained polymer at the interface, giving an independent read on interfacial quality.
Resistivity (ASTM D257, D4496)	Conductivity	Specify direction, contact scheme, and geometry or the number is uninterpretable.
Laser flash (ASTM E1461)	Thermal diffusivity	Combine with density and specific heat. Measure both in-plane and through-plane.
Mechanical (ASTM D638, D790, D5045)	Strength, stiffness, toughness	Report coefficient of variation, not only the mean.
Weibull modulus	Flaw population consistency	The underused metric. Dispersion quality shows up as reduced scatter before it shows up as an improved mean. A filled composite with a lower Weibull modulus than the neat matrix has been given a worse flaw population regardless of what the average says.
OTR / WVTR	Barrier	The most forgiving property: needs orientation, not interfacial bonding. Often the first application to succeed.

A program reporting mean tensile strength and nothing else is discarding the most sensitive available signal about its own dispersion quality. Scatter is data.

Dinner Table. There are two questions and people usually ask neither. What did I actually buy, and what did my factory do to it? The first is answered by a few hundred dollars of standard laboratory tests that most buyers skip. The second is answered by looking at a thin slice of the finished part under a microscope and counting clumps. And here is the underused trick: when a material is well mixed, the test results stop scattering. Consistency shows up before the average improves, so the spread in your numbers is telling you about your mixing long before the mean does.

So What. Run the incoming panel once per supplier and per lot thereafter, and require the results as distributions. On the outgoing side, add Raman strain mapping to confirm load transfer and track Weibull modulus alongside mean strength, because scatter is the earliest available signal of dispersion quality.

PART XI — THE PURCHASE ORDER

Everything above reduces to a document.

Section 1 — Identity, using ISO/TS 80004-13:2024 terms only. Material class (graphene, FLG, GNP, GO, rGO, SWCNT, DWCNT, MWCNT); production route; complete ISO/TS 9651:2025 technical data sheet.

Section 2 — Structure, as distributions with sample sizes. Layer number: mean, median,

standard deviation, method, n. Lateral size: mean, D10/D50/D90, method, n. Specific surface area by BET, nitrogen at 77 K. Stacking order, AB or turbostratic, by Raman M-peak and XRD spacing. Defect density as I(D)/I(G) at a stated excitation wavelength.

Section 3 — Chemistry. Carbon-to-oxygen ratio by XPS. Fraction sp^2/sp^3 with C1s deconvolution reported. Functional groups: identity, surface density, attachment location. Residual ash by TGA. Metals by ICP-MS covering iron, cobalt, nickel, manganese, sulfur, and any process-specific species. Moisture content.

Section 4 — Form and handling. Powder, dispersion, or masterbatch. If dispersion: solvent, solids loading, dispersant identity and loading, shelf life, sedimentation data. If masterbatch: carrier resin, loading, let-down ratio, recommended process window. Bulk density.

Section 5 — Acceptance and control. Certificate of analysis per lot against all of Sections 2 and 3. Statistical process control data on request. Retained sample, minimum 100 g, held two years. Right of rejection on specification failure with a defined arbitration laboratory.

Section 6 — Application data, which must state its own conditions. Every claimed property improvement accompanied by matrix, loading, dispersion method, specific energy input, cure or process schedule, test standard, measurement direction, sample size, and standard deviation. *Claims not meeting this standard are marketing and are excluded from the agreement.*

That final clause is worth writing into contracts verbatim. It costs nothing, changes the conversation immediately, and separates suppliers who characterize their material from suppliers who do not.

Dinner Table. All of this collapses into one page you attach to a purchase order. It says: tell me what it is in the standard's own words, give me the measurements as ranges rather than single numbers, tell me what is left over from making it, and if you claim it improved something, tell me under exactly what conditions. Any supplier who can fill that page in is worth talking to. Any supplier who cannot has just told you something more useful than their price.

So What. Adopt the six-section specification as a standing attachment to every engineered-carbon purchase order, and include the exclusion clause verbatim so that undocumented performance claims carry no contractual weight.

PART XII — COST, SAFETY, AND POSITION

12.1 The cost argument is usually the wrong argument

Industrial-grade graphene nanoplatelets now trade in the range of roughly **50** to **75** dollars per kilogram, approaching parity with conventional additives such as carbon black. Research-grade few-layer material runs from hundreds to thousands. Multi-wall nanotubes run from tens to a few hundred dollars per kilogram; single-wall material remains far higher. C_{60} is priced by the gram, and CVD film is priced by area and not comparable to any of them.

Assumption flagged. The range of fifty to seventy-five dollars per kilogram reflects trade reporting from early 2026 and has not been verified against primary quotes. It varies widely with grade, volume, and region, and should be refreshed before entering any financial model.

Now the arithmetic that matters. At one percent loading and the upper end of that range, filler adds about seventy-five cents to the cost of a kilogram of compound. Against a resin system at four to eight dollars per kilogram, that is a rounding error.

The barrier to adoption was never the price of the powder. It is qualification: requalifying a part, requalifying a process, re-running the aging and environmental matrix, and accepting product liability for a material change. In regulated sectors that cost is measured in years and millions, and it is invariant to filler price.

This reframes the commercial question. Cost reduction in graphene production, while welcome, has been addressing a constraint that was not binding. What would actually accelerate adoption is **re-**

duced qualification burden, which comes from consistency, standards conformance, third-party verification, and suppliers who can demonstrate that lot 47 matches lot 12. The industry has spent fifteen years optimizing the wrong variable.

12.2 Safety, stated proportionately

Two distinct hazard profiles are frequently and unhelpfully merged.

High-aspect-ratio rigid fibers. Long, rigid, biopersistent multi-wall nanotubes have shown asbestos-like pathogenic responses in animal mesothelium studies, consistent with the fiber paradigm: a shape and durability profile macrophages cannot clear. The World Health Organization respirable fiber definition — longer than **5 µm**, thinner than **3 µm**, aspect ratio above 3:1 — is the relevant geometric screen. **NIOSH** recommends an exposure limit of **1 µg/m^3** elemental carbon as an eight-hour time-weighted average for nanotubes and nanofibers. Short, tangled, or flexible tubes do not fit the paradigm equally, and hazard is not uniform across the class.

Platelets. Graphene platelets do not fit the fiber paradigm. Inhalation data are less alarming and less complete. The prudent posture is standard nanomaterial precaution rather than either complacency or fiber-equivalent controls.

The controlling variable in practice is form, not chemistry. Dry powder at low bulk density aerosolizes readily and is the dominant exposure route. Liquid dispersions and solid masterbatches substantially reduce that route but do not eliminate it: NIOSH identifies aerosol-generating operations on liquids, sonication among them, as capable of releasing airborne engineered nanomaterial. A facility buying masterbatch instead of powder has removed its largest exposure pathway as a side effect of a decision made for process reasons, and still needs controls wherever it sonicates, sprays, machines, or thermally processes the material. That is worth knowing before designing an expensive containment system for a problem a purchasing decision solves.

12.3 Freedom to operate

Tens of thousands of patents cover graphene and related materials, with filing concentrated in a few jurisdictions and much of the volume of uncertain enforceability. A freedom-to-operate opinion in this space is expensive and inconclusive, and the defensible positions are usually not the material itself.

Defensibility lives in the application-specific formulation, in the process for achieving dispersion at scale, and in the qualification data package for a regulated use. The last is the most durable and least appreciated. A competitor can replicate a formulation. Replicating three years of aging data and a customer's internal qualification is a different problem.

> **Dinner Table.** The material is cheap. At the amounts actually used, it adds pennies to a kilogram of plastic. What is expensive is proving to a customer, a regulator, or an insurer that changing the recipe did not break anything, and that cost does not fall when the powder gets cheaper. Meanwhile the real safety question is not what the material is but what shape it is in when it reaches you: loose powder floats and is worth worrying about, the same material already stirred into a resin is not.

> **So What.** Stop negotiating on price per kilogram and start negotiating on lot-to-lot consistency and documentation, because qualification is the binding constraint. Buy masterbatch rather than powder where the process allows, which reduces both dispersion risk and inhalation exposure in one decision.

SO WHAT

The problem with graphene is that "graphene" is a noun where a specification should be.

Every downstream failure traces to that. Buyers purchase a word and receive a distribution. Formulators apply loadings derived from monolayer physics to thirty-layer platelets. Programs measure conductivity on material that was oxidized so it would disperse. Executives conclude the material does not work when what did not work was the transaction.

Three actions follow, and all three are available today.

First, stop buying the word. ISO/TS 9651:2025 provides a technical data sheet template, a naming syntax, and a defined measurement method for every characteristic that matters. It was built over several years by more than a hundred subject matter experts specifically to end this problem. Demand it in the purchase order. Reject material arriving without it. A supplier who cannot populate that template does not know what they are selling, which is a more important fact about them than their price.

Second, identify your regime before specifying anything. Mechanical reinforcement, electrical percolation, and thermal conduction require loadings two orders of magnitude apart and formulations that actively conflict. There is no universal grade, and a supplier claiming improvement across all three is describing three different experiments. Decide which property you are buying, then let that decision drive layer count, lateral size, functionalization, and loading, in that order.

Third, qualify the process, not the powder. The worked calculation in Part VI is the most important number in this paper: the same feedstock at the same loading yields **+182%**, **+48%**, or **+10%**, depending on the three variables the model actually moves: orientation efficiency, effective flake-length and load-transfer efficiency, and exfoliated fraction. Interface strength is a fourth lever, treated separately in §6.2, because the rule of mixtures contains no interface term and cannot speak to it. The separate fracture-mechanics argument in §6.3 says the poorly dispersed case also carries a strength penalty, but that conclusion needs its own flaw-population model and is not produced by the stiffness calculation. Diligence spent on the certificate of analysis while the compounding step goes unmeasured is diligence spent on the smaller variable. Report specific energy input. Measure agglomerate size against the fracture-mechanics threshold rather than against appearance. Track Weibull modulus, because scatter reveals dispersion quality before the mean does.

There is a larger point underneath the technical one.

The famous properties of engineered carbons are real. They are also properties of **architectures**, not of substances. Suspended graphene conducts heat at 5,000 W/m · K and embedded graphene does not. A defect-free lattice reaches 130 GPa and an agglomerated one concentrates stress. A vertically aligned nanotube forest absorbs 99.965 percent of incident light and the same tubes in a binder do not come close. The same carbon reaches 1,500 W/m · K as an aligned film and insulates at 0.02 W/m · K as an aerogel. Two graphene sheets rotated to about one degree superconduct when cooled and electrically tuned, while a device a quarter degree off that angle reported a critical temperature roughly elevenfold lower.

In every one of those cases the chemistry is constant and the architecture carries the property. The industry has spent two decades selling chemistry and delivering architecture by accident.

Which is why the durable contribution here is not a formulation. It is a measurement discipline.

Apply the hundred-year test. A specific epoxy formulation will be obsolete within a decade. A supply agreement will expire. A patent will lapse. What survives is the standard: the shared definition of what the material is, the shared method for measuring it, and the shared expectation that a claim arrives with the conditions under which it was produced. Those outlast every product built on them, and they are what allows the next generation to build without re-litigating the vocabulary.

The market has the instruments. ISO/TS 80004-13 for the words. ISO/TS 21356-1 and ISO/TS 23359 for the measurements. ISO/TS 9651:2025 for the specification. Third-party verification for the supply chain. Nothing further needs to be invented.

What remains is the decision to insist.

Whoever defines proof governs the industry. In engineered carbons the proof has already been defined, and is waiting for buyers willing to demand it.

STANDARDS AND REFERENCES

Standards

- ISO/TS 80004-13:2024 — Nanotechnologies — Vocabulary — Part 13: Graphene and other two-

dimensional (2D) materials
- ISO/TS 9651:2025 — Nanotechnologies — Classification framework for graphene-related 2D materials
- ISO/TS 21356-1:2021 — Structural characterization of graphene — Part 1: Graphene from powders and dispersions
- ISO/TS 21356-2 — Structural characterization of graphene — Part 2: Graphene sheets on a substrate *(stage 60.00, under publication as of August 2026 — verify status before citing as published)*
- ISO/TS 23359:2025 — Chemical characterization of graphene-related 2D materials from powders and liquid dispersions
- ISO/TS 23879 — Structural characterization of graphene oxide flakes
- ISO/TR 19733:2019 — Matrix of properties and measurement techniques for graphene and related 2D materials
- ASTM D257, D638, D790, D3039, D4496, D5045, E1461

Literature

- Kauling, A. P., Seefeldt, A. T., Pisoni, D. P., Pradeep, R. C., Bentini, R., Oliveira, R. V. B., Novoselov, K. S., Castro Neto, A. H., "The Worldwide Graphene Flake Production," *Advanced Materials* 30(44), 1803784 (2018)
- Bøggild, P., "The war on fake graphene," *Nature* (2018)
- Lee, C., Wei, X., Kysar, J. W., Hone, J., "Measurement of the Elastic Properties and Intrinsic Strength of Monolayer Graphene," *Science* 321, 385 (2008)
- Balandin, A. A. et al., "Superior Thermal Conductivity of Single-Layer Graphene," *Nano Letters* 8, 902 (2008)
- Seol, J. H. et al., "Two-Dimensional Phonon Transport in Supported Graphene," *Science* 328, 213 (2010)
- Nair, R. R. et al., "Fine Structure Constant Defines Visual Transparency of Graphene," *Science* 320, 1308 (2008)
- Gong, L. et al., "Interfacial Stress Transfer in a Graphene Monolayer Nanocomposite," *Advanced Materials* 22, 2694 (2010)
- Hernandez, Y. et al., "High-yield production of graphene by liquid-phase exfoliation of graphite," *Nature Nanotechnology* 3, 563 (2008)
- Paton, K. R. et al., "Scalable production of large quantities of defect-free few-layer graphene by shear exfoliation in liquids," *Nature Materials* 13, 624 (2014)
- Luong, D. X. et al., "Gram-scale bottom-up flash graphene synthesis," *Nature* 577, 647 (2020)
- Poland, C. A. et al., "Carbon nanotubes introduced into the abdominal cavity of mice show asbestos-like pathogenicity," *Nature Nanotechnology* 3, 423 (2008)
- Kroto, H. W., Heath, J. R., O'Brien, S. C., Curl, R. F., Smalley, R. E., "C60: Buckminsterfullerene," *Nature* 318, 162 (1985)
- Nielsen, L. E., "Models for the permeability of filled polymer systems," *Journal of Macromolecular Science* A1, 929 (1967)
- Xie, Y. et al., "Interface-mediated extremely low thermal conductivity of graphene aerogel," *Carbon* 98, 381 (2016)
- Papageorgiou, D. G., Kinloch, I. A., Young, R. J., "Mechanical properties of graphene and graphene-based nanocomposites," *Progress in Materials Science* 90, 75 (2017)
- Cao, Y. et al., "Unconventional superconductivity in magic-angle graphene superlattices," *Nature* 556, 43 (2018)
- Bistritzer, R., MacDonald, A. H., "Moiré bands in twisted double-layer graphene," *PNAS* 108, 12233 (2011)
- Park, J. M. et al., "Tunable strongly coupled superconductivity in magic-angle twisted trilayer graphene," *Nature* 590, 249 (2021)
- NIOSH Current Intelligence Bulletin 65: Occupational Exposure to Carbon Nanotubes and Nanofibers (2013)
- National Physical Laboratory, Good Practice Guides on graphene measurement

---

Load-bearing quantitative values in this paper are derived from a single parameter model and verified dimensionally at build time; every numeric token in the body text is either model-derived or carries a

registered exemption. Assumptions, and models applied outside their derivation conditions, are flagged in the text where they occur.

Verification. Every load-bearing quantitative value in this paper is derived from a single parameter model and checked dimensionally at build time. Source of record: SHA-256 `06ed7d46729082435c03993fa657d2f5…`. Build gate: 85 regression tests across value provenance, reverse numeric coverage, claim-registry enforcement, structural integrity and markup. Assumptions, and models applied outside their derivation conditions, are flagged in the text where they occur.