A CRYPTOSOWHAT TECHNICAL BRIEF — FRONTIER MATERIALS SERIES

The Thermal Wall

Copper–Carbon Composites and the Hidden Bottleneck of the AI Compute Buildout

Dr. Gregory S. Carmichael, DBA · Lt. Col., USAF (Ret.)
Founder & CEO, Advanced Nano-Materials Manufacturing LLC · Editor, CryptoSoWhat

May 2026

Abstract. The economics of artificial intelligence are increasingly written in watts, and watts become heat. As accelerator power densities climb past the point where bulk copper can move heat away fast enough, the binding constraint on compute is migrating from the transistor to the heat path. This brief examines copper–carbon composite structures—the deliberate marriage of copper's manufacturability with the extreme thermal conductivity of carbon allotropes—as a route past that thermal wall. The central finding is that the performance of every such system is governed not by the carbon, whose intrinsic conductivity is spectacular, but by the copper–carbon *interface*, which is intrinsically hostile: molten copper does not wet carbon, and an untreated junction throttles heat flow. We review four integration architectures—diamond–copper particulate composites, metal-encapsulated pyrolytic graphite, vertically aligned carbon-nanotube interface layers, and electrodeposited graphene–copper—and map each to the thermally critical use case it actually serves. We then isolate the engineering lever that converts a hostile interface into a conductive one (a thin, continuous bonding interlayer—a reactive carbide where one forms, and wetting, percolation, or contact-area control where it does not) and the two constraints that decide field reliability (anisotropy and thermal-expansion matching). We close with the strategic implication: the defensible position is not the composite but the interface chemistry and the deposition process, and that position is most naturally occupied by a domestic, electrochemistry-led manufacturer.

Keywords: thermal management; copper–carbon composites; interfacial thermal conductance; data center cooling; diamond–copper; pyrolytic graphite; graphene; carbon nanotubes; electrodeposition.

---

1 Introduction: The Thermal Wall

For two decades the limiting reagent of computing was the transistor. That era is closing. The leading AI accelerators now dissipate on the order of a kilowatt per package, and the roadmap points upward. The constraint that decides whether a chip can be clocked, a rack can be filled, and a data center can be financed is no longer how many transistors fit on a die—it is how fast the heat they generate can be removed. Cooling is no longer a packaging afterthought; it is the rate-limiting step of the AI buildout.

Copper has been the default heat-conduction metal because it is cheap, abundant, ductile, electroplatable, and conducts heat at roughly $400\,\mathrm{W\,m^{-1}K^{-1}}$. But copper is heavy, and it has a hard physical ceiling. Several allotropes of carbon blow past that ceiling: chemical-vapor-deposited diamond reaches on the order of $2000\,\mathrm{W\,m^{-1}K^{-1}}$ while being electrically insulating, and oriented graphitic carbon conducts heat in-plane above $1500\,\mathrm{W\,m^{-1}K^{-1}}$ at roughly a quarter of copper's density. The temptation is obvious—replace copper with carbon and break through the wall.

The temptation is also a trap. Carbon's headline conductivities are intrinsic, single-crystal, often anisotropic numbers that do not survive the transition to a manufacturable three-dimensional part. There is, today, no electrochemical or additive process that deposits pristine bulk carbon into arbitrary cooling geometries at production scale and cost. The realistic engineering question is therefore not "carbon *instead of* copper" but

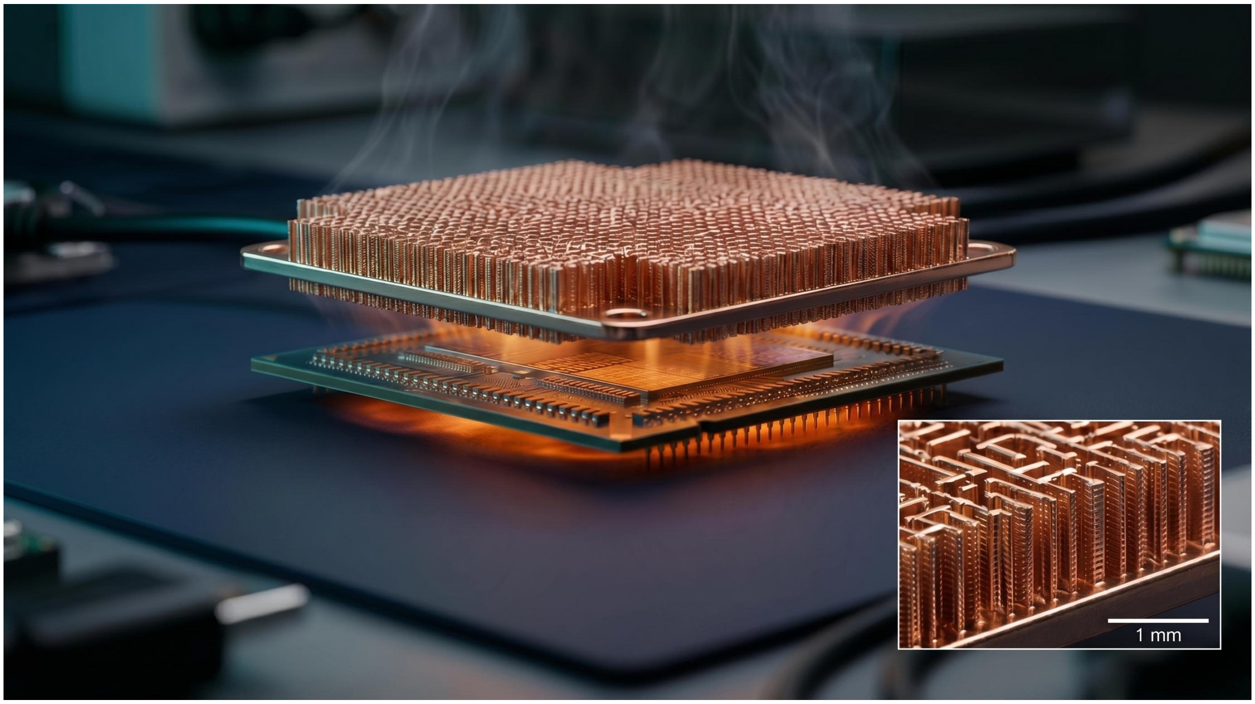

Figure 1: **The object of the exercise.** A copper micro-pin-fin heat exchanger lifted clear of a high-power AI accelerator; heat glows from the exposed silicon die below. The inset shows the additively manufactured pin-fin microarchitecture that does the work (scale bar 1 mm). Fins this fine are most natural for liquid cold plates; related geometries adapt to forced air, but the hydraulic and pressure-drop constraints differ sharply. The estimate in Section 6 treats an idealized forced-air case only; a liquid cold plate would require its own pressure-drop and coolant-side heat-transfer model.

“carbon *integrated with* copper”—and that integration is where the physics gets unforgiving.

At the Dinner Table

Think of copper as a wide, well-paved highway for heat: not the fastest possible road, but one you can build anywhere, cheaply. Some forms of carbon are more like a maglev train—far faster, but they only run on special track and you cannot easily lay that track wherever you want. The smart move is not to rip out the highways. It is to build maglev lines where they pay off and connect them cleanly to the existing roads. The whole problem, it turns out, is the on-ramp where the two systems meet.

2 First Principles: Why the Interface Decides Everything

2.1 The hostile junction

Copper and carbon do not like each other. Copper does not form a stable carbide, and molten copper does not wet a carbon surface: contact angles remain near 130° even at temperatures around 1100 °C, so a copper–carbon junction tends not to bond chemically at all [1]. A purely mechanical contact leaves nanoscale gaps and a population of defects exactly where heat must cross. The result is a large *thermal boundary resistance* (Kapitza resistance) localized at the interface, which can dominate the heat path even when both constituents are excellent conductors.

Figure 2: **The carbide bridge, in concept.** Copper (left) and graphitic carbon (right, rendered as stacked graphene platelets) cannot bond directly—the bare junction is non-wetting, and heat stalls there. The luminous interlayer spanning the gap is the few-nanometer carbide bridge: the thin, continuous phase that couples the two materials and lets heat cross. It is the small structure on which the entire composite's performance turns, and the locus of the defensible engineering described in Section 4.

2.2 Quantifying the penalty

The cleanest way to see why the interface governs the system is the Hasselman–Johnson effective-medium model for a particulate composite whose inclusions of radius a and conductivity k_p sit in a matrix of conductivity k_m, separated by an interfacial conductance h_c [2]. Defining the dimensionless interface parameter

$$\alpha = \frac{k_m}{a h_c}, \tag{1}$$

the effective conductivity of the composite is

$$\frac{k_{\text{eff}}}{k_m} = \frac{\left[k_p(1+2\alpha)+2k_m\right]+2V_p\left[k_p(1-\alpha)-k_m\right]}{\left[k_p(1+2\alpha)+2k_m\right]-V_p\left[k_p(1-\alpha)-k_m\right]}, \tag{2}$$

where V_p is the inclusion volume fraction. The instructive limit is the *Kapitza radius*

$$a_k = \frac{k_m}{h_c}. \tag{3}$$

When the inclusion size a is comparable to a_k, the interface term α approaches unity and the high conductivity of the carbon (k_p) is effectively switched off: the composite conducts no better than—and can conduct worse than—plain copper. Only when h_c is driven high enough that $a_k \ll a$ does the carbon's conductivity actually appear in k_{eff}. In plain terms, the equations say the carbon is wasted until the interface is fixed.

So What

Every dollar of value in a copper–carbon thermal product is unlocked at the interface, not in the filler. A diamond particle conducting at $2000\,\mathrm{W\,m^{-1}K^{-1}}$ trapped behind a bad boundary performs like an expensive rock. This is good news strategically: it means the differentiator is a thin engineered layer and the process that deposits it—knowledge and IP—rather than access to exotic raw material. The moat is chemistry, not mining.

3 Four Architectures for Copper–Carbon Integration

No single copper–carbon structure serves all thermally critical applications. The correct architecture is dictated by whether the job is bulk conduction, lateral spreading, crossing a die-level gap, or upgrading an existing copper surface.

These four architectures are usually framed as competitors. They are better understood as complementary layers of a single *graded* stack, each material placed at the depth where its advantage is decisive and its weakness irrelevant (Figure 3). Carbon goes where heat flux is high and mass is expensive—the die interface and the spreader; copper or graphene–copper carries the convective duty in the fin field, where surface area and manufacturability dominate; and a thin graphene skin couples heat into the coolant at almost no mass penalty. The result is the rare design that *raises* dissipation while *lowering* weight, because carbon's true advantage is gravimetric: per kilogram, diamond and oriented graphite conduct heat roughly an order of magnitude better than copper. Notably, in a micro-pin-fin array the fins are already near unit thermal efficiency, so the leverage lives not in the fins but in the spreader and die interface beneath them—exactly where carbon is placed. The subsections below treat each layer in turn.

3.1 Diamond–copper particulate composites

Here micron-scale diamond grit is dispersed in a copper matrix and consolidated by spark-plasma sintering, hot pressing, or pressure infiltration. The diamond is first metallized—or the copper is alloyed—to grow a carbide bridge that defeats the wetting problem. Performance scales with how well that bridge is controlled. Titanium-coated diamond in copper, forming an interfacial titanium-carbide phase, has reached roughly $457\,\mathrm{W\,m^{-1}K^{-1}}$ at 40 volume percent diamond [3]. Tungsten metallization, annealed into a tungsten-carbide transition layer, has pushed a 50-volume-percent composite to about $640\,\mathrm{W\,m^{-1}K^{-1}}$—the highest reported for hot-press sintering in that range [1]. Crucially, these composites also let the engineer dial in a low coefficient of thermal expansion, which matters enormously for direct attachment to silicon. This is the architecture for a bulk heat sink or heat-spreader baseplate that must beat copper on conductivity *and* expansion match.

3.2 Metal-encapsulated pyrolytic graphite

The most commercially mature copper–carbon architecture does not disperse carbon at all; it encapsulates a solid core of thermally annealed pyrolytic graphite (TPG/APG) inside a metal shell. The graphite core carries heat in-plane above $1500\,\mathrm{W\,m^{-1}K^{-1}}$ at a density near $2.25\,\mathrm{g\,cm^{-3}}$, while the copper or aluminum shell supplies mechanical integrity, solderability, and a tunable expansion coefficient [4]. The payoff is decisive where mass is costly: an encapsulated-graphite spreader has demonstrated more than a 50% reduction in thermal resistance together with roughly a 48% reduction in weight versus a solid copper spreader of the same function [5]. Commercial implementations are flight-qualified on military aircraft and spacecraft [6]. The central design caveat is anisotropy: the $> 1500\,\mathrm{W\,m^{-1}K^{-1}}$ figure is in-plane only, and through-plane conductivity is more than an order of magnitude lower, so copper thermal vias must be inserted to force

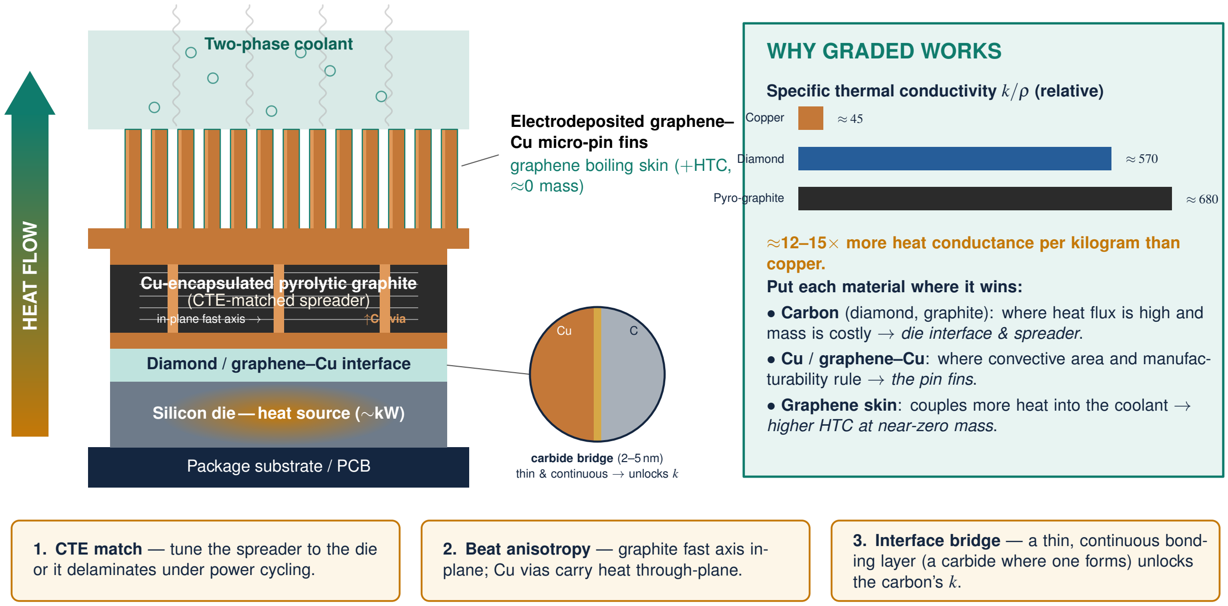

≈12–15× more heat conductance per kilogram than copper.

Put each material where it wins:

- **Carbon** (diamond, graphite): where heat flux is high and mass is costly → *die interface & spreader.*
- **Cu / graphene–Cu**: where convective area and manufacturability rule → *the pin fins.*
- **Graphene skin**: couples more heat into the coolant → *higher HTC at near-zero mass.*

1. **CTE match** — tune the spreader to the die or it delaminates under power cycling.
2. **Beat anisotropy** — graphite fast axis in-plane; Cu vias carry heat through-plane.
3. **Interface bridge** — a thin, continuous bonding layer (a carbide where one forms) unlocks the carbon's k.

Figure 3: **A graded copper–carbon thermal stack.** Heat flows upward from the silicon die through a graded sequence of materials, each chosen for the duty at its depth: a diamond or graphene–copper die interface (highest heat flux); a CTE-matched, copper-encapsulated pyrolytic-graphite spreader that moves heat laterally at low mass, with copper vias defeating the graphite's through-plane anisotropy; an electrodeposited graphene–copper micro-pin-fin field that supplies convective area and manufacturability; and a graphene boiling-enhancement skin. The three governing constraints—expansion matching, anisotropy orientation, and a thin, continuous carbide interface bridge—are summarized below the stack. Bars at right show specific thermal conductivity k/ρ, the metric on which carbon outperforms copper by roughly an order of magnitude.

heat into the plane of the core [7]. This is the architecture for lateral heat spreading and for weight-critical aerospace and defense systems.

3.3 Vertically aligned carbon-nanotube interface layers

At the die-to-heat-sink junction, the enemy is the thermal interface material—the grease or pad whose resistance often dwarfs everything around it. Vertically aligned carbon-nanotube (CNT) forests, with copper backfilled into the spaces between tubes, attack this gap directly: the tubes provide through-plane conduction paths while the electroplated copper fills voids, lowers porosity, and acts as a lateral spreader [8, 9]. The hard limit is bundling. An individual multi-walled nanotube conducts near $600\,\mathrm{W\,m^{-1}K^{-1}}$, but a bundle falls to roughly 150 and a sheet to about $50\,\mathrm{W\,m^{-1}K^{-1}}$ as tube-to-tube contact resistance and defects accumulate [10]. CNT–copper is therefore a thermal-interface and thermal-via technology, not a bulk-material one—a precision tool for the specific, stubborn gap at the chip.

3.4 Electrodeposited graphene–copper

The final architecture is the one most aligned with an electrochemistry-led manufacturer. Graphene is co-deposited with, or layered into, copper by electroplating—the same fundamental process used to grow copper in advanced additive manufacturing. A patented direct-current electrodeposition route produces a copper-based graphene composite coating, 30 to 300 micrometers thick, with thermal conductivity reported between 390 and $1112\,\mathrm{W\,m^{-1}K^{-1}}$ [11]. Related strategies build a graded interface by alternating copper and

Architecture	Role & form	Reported thermal performance	Best-fit use case
Diamond–Cu	Bulk particulate composite; carbide-bridged	$\sim$457–640 $\mathrm{W\,m^{-1}K^{-1}}$ bulk; tunable low CTE	Chip-level heat sinks / baseplates
Encapsulated TPG	Solid graphite core in metal shell	$>$ 1500 $\mathrm{W\,m^{-1}K^{-1}}$ in-plane core; -50% resistance, -48% mass vs. Cu	Lateral spreaders; aerospace / defense
CNT–Cu	Aligned forest, Cu backfill; interface layer	Tube $\sim$600 $\mathrm{W\,m^{-1}K^{-1}}$; bundle $\sim$150; system gains via Cu spreading	Die-level TIMs and thermal vias
Graphene–Cu	Electrodeposited coating / laminate	$\sim$390–1112 $\mathrm{W\,m^{-1}K^{-1}}$ coating; $+40\%$ diffusivity	Surface upgrades; electroplating-adjacent lines

Table 1: Copper–carbon integration architectures mapped to thermally critical use cases. Values are representative published results, not specification limits; see text and references.

graphene layers, and a plasma-CVD graphene network grown on copper powder has raised thermal diffusivity by about 40% over pure copper while cutting porosity to under 1% [12]. Aligned copper–reduced-graphene-oxide composites add a structural dividend, with flexural strength up roughly 22% over sintered copper at under one volume percent loading, alongside strongly anisotropic conduction [13]. This is the architecture for upgrading existing copper surfaces and for processes that are cousins to copper electroplating.

At the Dinner Table

There is no one "copper-carbon material." There are four different tools for four different jobs. Diamond-in-copper is a solid, high-performance block you bolt under a chip. Encapsulated graphite is a lightweight sheet that whisks heat sideways—it flies on fighter jets. Nanotube layers are a thin gasket that fixes the worst spot, the squishy interface right at the chip. And graphene-in-copper is a coating you can plate onto parts you already make. Picking the wrong tool for the job is the most common and most expensive mistake.

4 The Interface Toolbox: Where the Defensible IP Lives

Because the interface governs the system (Section 2), the engineering action concentrates in a single design variable: the thin carbide layer that bridges copper and carbon. The literature now spans a family of these bridges—chromium carbide, silicon carbide, titanium carbide, tungsten and its carbides, boron carbide, and most recently two-dimensional MXene phases such as Ti_3C_2. The quantitative leverage is large. A boron-carbide interlayer only about two nanometers thick has been shown to raise the interfacial thermal conductance by a factor of roughly 14, to a peak near 286 $\mathrm{MW\,m^{-2}K^{-1}}$ [14]. A Ti_3C_2 MXene interlayer improved the conductivity of a diamond–copper composite about 1.6-fold over the unmodified case [15]. And across systems the governing rule is consistent: the highest interface conductance accompanies a layer that is *thin and continuous*; too thick a carbide adds its own resistance, while a discontinuous one leaves the wetting problem unsolved [16]. One caution against over-generalizing: the carbide bridge is the right mechanism wherever the carbon will react to form one—diamond– and graphite–copper above all—but it is not universal. Where carbides do not form readily, the same interface-conductance role is played by other measures: wetting agents and graded layers in graphene–copper, contact-area and backfill control in CNT–copper, and percolation rather than chemical bonding in the phase-change composites of Section 7. The unifying variable is interfacial thermal conductance; the carbide bridge is its most common embodiment, not its only one.

So What

This is the patent territory. The composite recipes are broadly known and increasingly crowded; the durable, ownable invention is a specific interlayer chemistry plus a scalable, preferably electrochemical, method to deposit it thin and continuous at production rates. A claim that reads on “a controlled sub-five-nanometer carbide bridge produced by [process X] yielding interfacial conductance above [threshold Y]” has teeth. Whoever owns the on-ramp owns the highway interchange—and licenses it to everyone building the road.

5 The Two Constraints That Decide Field Reliability

5.1 Anisotropy is a layout problem, not a materials problem

Graphitic carbon and graphene conduct superbly in-plane and poorly through-plane; the ratio can exceed an order of magnitude. A design that assumes isotropic conduction will fail. The discipline is to treat anisotropy as a routing problem: orient the carbon so its fast axis aligns with the dominant heat-flow direction, and use copper vias or copper backfill to carry heat across the slow axis. The encapsulated-graphite-with-copper-vias architecture (Section 3.2) is precisely this principle made into a product.

5.2 Thermal-expansion mismatch is the reliability killer

The failure that ends careers is not low conductivity—it is delamination. Copper, carbon, and the carbide interlayer expand at different rates; every power cycle strains the interface, and a mismatched stack debonds over time, after which thermal performance collapses. This is why the ability to tune the composite’s coefficient of thermal expansion to match silicon or a power substrate is as important as raw conductivity, and why the encapsulated and carbide-bridged architectures, which expose that knob, dominate demanding applications.

At the Dinner Table

Two things quietly decide whether one of these parts survives in the real world. First, carbon moves heat brilliantly in one direction and badly in the others, so you have to point it the right way—like laying floorboards along the hallway, not across it. Second, when things heat up and cool down, copper and carbon swell by different amounts, and if you do not match them they eventually peel apart, like a sidewalk cracking through summers and winters. Get those two right and the part lasts; get them wrong and it fails no matter how good the headline number was.

6 How Much Better? A First-Order Estimate for Air Cooling

To size the prize, model the path from junction to ambient as three thermal resistances in series,

$$R_{ja} = R_{\text{int}} + R_{\text{spr}} + R_{\text{conv}}, \tag{4}$$

the die-to-spreader interface, the spreader’s conduction and lateral spreading, and the fin-to-air convection. The decisive fact about *air* cooling is where the resistance lives: air carries little heat per degree, so the convective term R_{conv} dominates and the solid-conduction terms are comparatively small. That sets a hard ceiling on this architecture, because the carbon improves the conduction terms and leaves convection untouched.

6.1 The carbon does almost nothing for the fins

At this scale a micro-pin fin is already a near-perfect conductor. Fin efficiency is $\eta = \tanh(mL)/(mL)$ with $m = \sqrt{4h/(kd)}$. For a representative copper pin ($d = 0.2\,\mathrm{mm}$, $L = 1.5\,\mathrm{mm}$, forced-air $h \approx 200\,\mathrm{W\,m^{-2}K^{-1}}$), $m \approx 100\,\mathrm{m^{-1}}$, $mL \approx 0.15$, and $\eta \approx 99.3\%$. Replacing the copper with diamond ($k = 2000$) lifts η only to $\approx 99.85\%$—half a percentage point. Copper, or graphene–copper for the weight saving, is the correct fin material; the leverage lives entirely beneath the fins.

6.2 Where the carbon earns its place

Carbon attacks the two conduction terms. A diamond or graphene–copper die interface with a thin carbide bridge roughly halves R_{int}. A copper-encapsulated pyrolytic-graphite spreader—conducting in-plane near $1500\,\mathrm{W\,m^{-1}K^{-1}}$ against copper's 400—spreads the concentrated die heat into the fin field far faster and cuts R_{spr} by more than half. Table 2 gives a representative breakdown for a $\sim$20 mm die, taking the conservative single-phase view that the graphene boiling-skin term in Figure 3 contributes nothing in dry air.

Resistance element	Solid Cu (K/W)	Graded (K/W)	What changes
Die + interface, R_{int}	0.08	0.04	diamond / graphene–Cu + carbide bridge
Spreader / base, R_{spr}	0.07	0.03	encapsulated pyrolytic graphite ($\sim$1500 in-plane)
Fin → air, R_{conv}	0.15	0.15	unchanged: air-side limited; fins already $\sim$99% efficient
Total, R_{ja}	**0.30**	**0.22**	$\approx -27\%$

Table 2: Representative first-order junction-to-ambient resistance budget for forced-air cooling of a $\sim$20 mm die. Values are illustrative engineering estimates for comparison, not measured specifications.

6.3 The bottom line

Summed, the graded stack lowers junction-to-ambient resistance from about 0.30 to 0.22 K/W—a **25–30% reduction**, and one that improves as the air side strengthens (a more aggressive $R_{\mathrm{conv}} = 0.10$ widens the gap to $\sim$32%; a weaker 0.20 narrows it to $\sim$23%). Translated into what an engineer feels: at a fixed $\sim$200 W load the junction runs roughly **16 K cooler**; held instead to the same junction temperature, the stack carries about **36% more power** (from $\approx$200 to $\approx$273 W against a 60 K rise budget, consistent with the $0.30 \to 0.22$ K/W drop in Table 2). All of this while shedding close to half the spreader mass and on the order of a fifth to a third of the total heat-sink assembly.

At the Dinner Table

Picture cooling a hot skillet. In air, the slow step is the air lazily lifting heat off the surface—polishing the metal so heat reaches the surface faster barely helps, because the air is the bottleneck. That is why trading copper for exotic carbon in an air-cooled fin stack buys only about a quarter less temperature: you have sped up everything except the step that was actually slow. What you *do* get, almost for free, is a much lighter part. And if you ever switch to liquid—a fast-moving river instead of lazy air—those same carbon upgrades suddenly matter enormously.

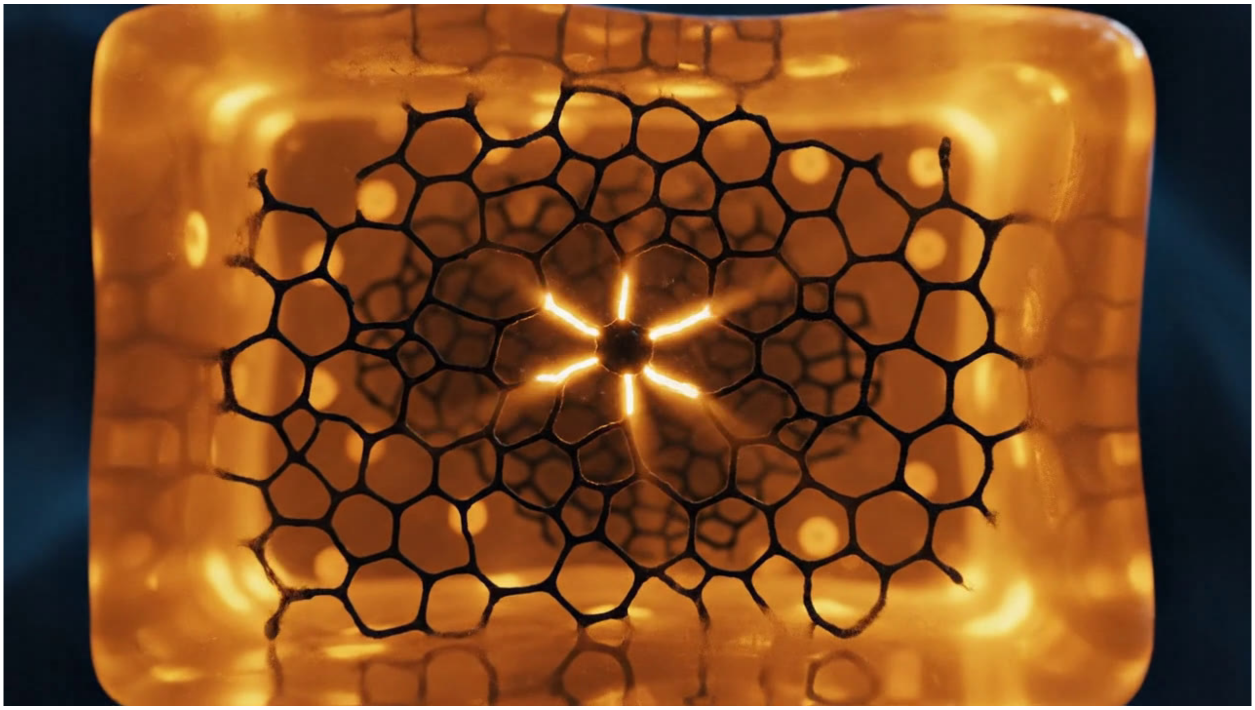

Figure 4: **The idea, made visible.** A carbon network threaded through a block of beeswax: heat entering at the center races outward along the conductive lattice instead of stalling in the surrounding wax. The hexagonal motif is fitting—graphene is itself a honeycomb lattice of carbon atoms, the same geometry the bees build—and it captures the mechanism of Figure 5 in a single frame: a few percent of carbon turns an inert heat sponge into a fast, distributed conductor.

So What

In air, that 16 K and 36% are real and worth capturing—but the headline is weight, not temperature. The thermal gain is genuine yet capped near 25–30% because the air—the one term carbon cannot touch—now dominates the budget; the structural wins (roughly half the spreader mass, and a CTE-matched interface that survives power cycling) are the stronger case for an air-cooled deployment. The thermal upside compounds the moment the system goes wet: in a cold-plate or two-phase loop R_{conv} falls by an order of magnitude, the conduction terms become essentially the whole budget, and the same graded stack—boiling skin now active—delivers a far larger total improvement. The architecture is *air-ready and liquid-optimal*: it can earn its place in forced-air systems on weight and spreading alone, while offering a substantially larger thermal gain in liquid or two-phase loops.

7 Beyond the Cold Plate: Carbon-Enhanced Phase-Change Materials

Not every thermal problem is a steady-state cold plate. A large class of systems—missile seekers, avionics modules, directed-energy pulses, anything that must be lightweight, sealed, and free of pumps or external heat rejection—instead faces a *transient*: a short, intense burst of heat that has to be absorbed in place. The classical answer is a phase-change material (PCM): a wax that banks heat in its latent heat of fusion, melting silently with no moving parts. It is an elegant, flight-proven trick—and it has one crippling weakness, which happens to be the same one this entire brief is about.

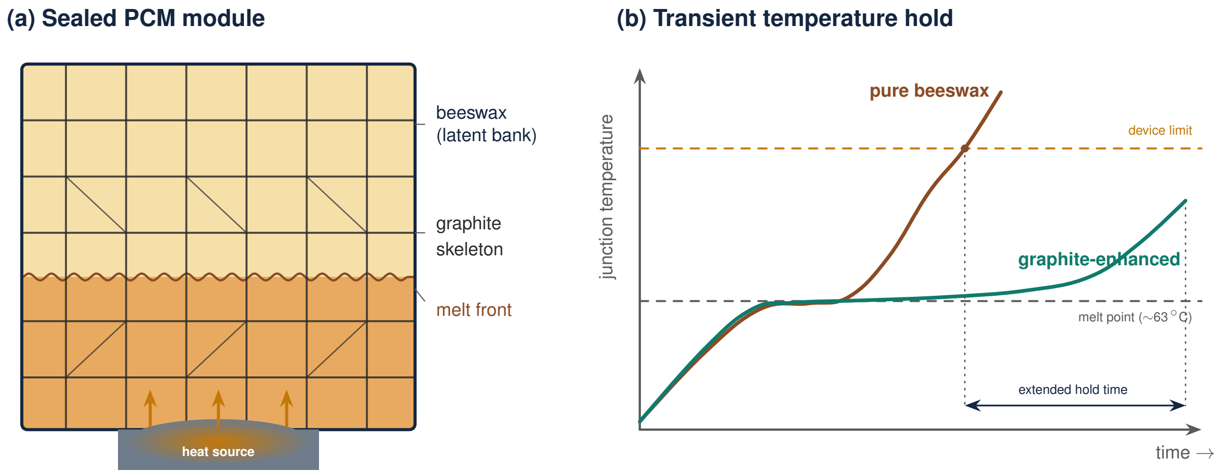

Figure 5: **Carbon-enhanced phase-change cooling for a sealed, pumpless load.** (a) A graphite skeleton threaded through beeswax inside a sealed enclosure; the heat source melts the wax along an advancing front. (b) Transient junction temperature: unenhanced beeswax melts only near the source and the temperature soon spikes past the device limit, whereas the graphite-enhanced block engages its whole latent-heat bank, holding a long flat plateau well below the limit and extending the usable hold time.

7.1 Why the wax is conduction-starved

A wax such as beeswax (melt point $\sim$62–64 °C, latent heat $\sim$150–190 J/g) stores a great deal of energy in its melt transition but conducts heat abysmally—around 0.2–$0.3\,\mathrm{W\,m^{-1}K^{-1}}$, three orders of magnitude below copper. The result is a cruel irony: in a fast transient only the thin shell of wax against the heat source actually melts and absorbs energy, while the bulk of the latent-heat bank sits cold and untouched and the source overheats anyway. The capacity is there; the wax simply cannot move heat *into* itself fast enough to use it. Low conductivity and leakage of the molten phase are exactly the limits that have constrained wax PCMs in spacecraft thermal control [17].

7.2 The carbon fix—and the percolation rule

The cure is to thread a continuous, high-conductivity carbon network through the wax (Figure 4), and the governing design rule echoes Section 4: a *percolating network* beats dispersed high-conductivity particles, because isolated flakes hit the same interfacial wall and, worse, never connect into a path. **Expanded graphite** is the practical champion—its loading builds a three-dimensional interconnected skeleton of conduction pathways while simultaneously confining the molten wax [18]. The numbers are striking—though a sourcing note is owed first: the quantitative enhancement results below are measured for *paraffin*-wax/expanded-graphite composites, not beeswax. Beeswax is the illustrative material here (its own melt point and latent heat differ, and would need direct measurement); the mechanism transfers, but the figures are the paraffin authors'. In the cited study, paraffin with $\sim$10 wt.% expanded graphite reached a radial thermal conductivity near $10\,\mathrm{W\,m^{-1}K^{-1}}$—reported as about 20× the authors' own paraffin baseline—with under 13% loss of latent heat and over 99% mass retention across 50 thermal cycles [17]. Graphene nanoplatelets give a similar but smaller effect at low loading [19]. The trade is favorable—surrender roughly an eighth of the energy bank to make the *entire* bank reachable within the time the transient allows. Figure 5 shows the consequence: a long, flat temperature plateau where the unenhanced wax would have spiked past the device limit.

7.3 Three dividends that fit a sealed, weight-critical load

First, leak control comes free: the same graphite that fixes conduction also acts as an encapsulating scaffold, holding the molten wax by capillary action so it cannot pool or migrate [20]—decisive for a sealed module that may operate in any orientation. Second, the filler is light (graphite $\sim 2.2\,\mathrm{g\,cm^{-3}}$), so the conductivity gain costs almost no mass, unlike a metal foam or fin insert. Third, the network accelerates re-solidification between bursts, shortening the recharge interval and allowing more duty cycles per mission. The approach is flight-relevant: an expanded-graphite/wax block has held electronics at 68–71 °C and cut peak temperature by at least 23 °C under a 10 W transient [17]. Two cautions: there is an optimum loading—5–10 wt.% is the electronics sweet spot, beyond which latent-heat loss begins to bite [18]—and the graphite is anisotropic, so its fast axis should point toward the heat source. Diamond, by contrast, is the wrong material here: costly, particulate, and unable to form the percolating path the physics demands.

At the Dinner Table

A block of wax is a superb heat sponge with a terrible straw. It can soak up an enormous amount of heat as it melts—but heat crawls into wax so slowly that, during a quick pulse, only the outer skin melts while the cold core does nothing and the electronics cook anyway. Lacing the wax with a fine graphite web is like running a network of fast little straws through the sponge: now the whole block drinks at once. You give up a sliver of capacity to the graphite that does not melt, and in return the part holds its temperature far longer—and, since graphite weighs almost nothing, it stays light enough to fly.

So What

This is the cleanest fit for carbon of any case in this brief, because a PCM's only real flaw is low conductivity—exactly what a graphite network repairs—and the use case rewards low mass, where carbon dominates. For a defense portfolio it lands squarely in the sealed, pumpless, weight-critical niche: seekers, avionics, and pulsed-power loads. And it carries an unusual supply-chain story: the wax can be bio-based and domestically sourced, the graphite is abundant, and the product is a passive, sealed, made-in-region thermal block with no fluorinated fluids and no pumps to fail—a dual-use part whose entire bill of materials can be controlled onshore.

8 Manufacturing, Scale, and the Economics of Heat

A laboratory record conductivity is not a product. The architectures separate sharply on manufacturability. Encapsulated pyrolytic graphite is a mature, qualified production technology, but it is dominated by entrenched incumbents and the graphite core is expensive. Diamond–copper is producible by sintering and infiltration but is capital-intensive and constrained by diamond cost. Electrodeposited graphene–copper is, by contrast, an extension of a process already operated at industrial scale—copper electroplating—which is what makes it attractive as an entry point: the unit operation exists, the capital is comprehensible, and the differentiation lives in the electrolyte chemistry and the interlayer control rather than in exotic furnaces.

Several engineering questions must be closed before any of this ships, and the resistance model above is silent on all of them. Fine-pitch fins exact a pressure-drop and fan- or pump-power penalty; the forced-air case in particular trades pumping power for surface area, and a fair comparison must carry that cost. The copper–carbon and carbide interfaces must be shown to resist oxidation and galvanic corrosion over service life, and to survive thermal cycling without delamination—the failure mode, driven by expansion mismatch, that ends programs. Where electrical isolation is required the choice of allotrope is not free: diamond is dielectric and safe against live circuitry, whereas graphite and graphene are conductive and must be kept clear of it. And the electrodeposited graphene–copper route, attractive on cost, still has to demonstrate yield and run-to-run repeatability at production volume. None of these is disqualifying; all of them are gating, and a credible program treats them as the real work.

The macroeconomic frame is what makes this worth a thought-leadership audience's attention. The AI buildout is, at the margin, a contest for power and for the ability to reject heat. Cooling overhead is a direct tax on the useful compute extractable from a fixed power envelope, and it compounds: better heat rejection lets a chip run faster, lets a rack pack denser, and lets a given substation support more revenue-generating silicon. Materials that move the heat-rejection frontier are therefore not a niche components story—they are leverage on the single largest capital-formation event in the technology economy. And because the heat path is also a supply-chain and security surface, who can manufacture these structures, and where, is a question of industrial sovereignty, not merely cost.

So What

The investable thesis is not "sell a better heat sink." It is: own the interface IP, enter through electroplating-adjacent manufacturing where capital is rational, and position the output as strategic infrastructure for domestic AI and defense compute. The thermal path is a chokepoint on the AI economy. Chokepoints are where pricing power and policy support concentrate—and a domestic, dual-use, electrochemistry-led producer is unusually well placed to hold one.

9 Conclusion

Copper–carbon composites are a real answer to a real wall, but the answer is counterintuitive. The carbon is not the hard part: its spectacular conductivity is freely available, and the raw material is not where the defensible value lies. The hard part—and therefore the valuable part—is the interface: the few nanometers of interface chemistry that decide whether the carbon's conductivity ever reaches the chip, and the manufacturing process that lays that layer down thin, continuous, and at scale. Choose the architecture by the job (bulk conduction, lateral spreading, the die-level gap, or surface upgrade), respect the two reliability constraints (anisotropy and expansion matching), and concentrate invention on the interlayer and its deposition. Do that, and a heat-management material becomes what it should be in the age of compute-as-capital: a position on a chokepoint, not a commodity part.

References

[1] Shah, A.W., et al. "High Thermal Conductivity Diamond–Copper Composites Prepared via Hot Pressing with Tungsten-Coated Interfacial Layer Optimization." *Materials* (2025). PMC12387508.

[2] Hasselman, D.P.H., and Johnson, L.F. "Effective Thermal Conductivity of Composites with Interfacial Thermal Barrier Resistance." *Journal of Composite Materials* 21 (1987): 508–515.

[3] "Interfacial Characterization and Thermal Conductivity of Diamond/Cu Composites Prepared by Liquid–Solid Separation Technique." *Materials* (2023). PMC10005683.

[4] "Analysis and Evaluation of Thermally Annealed Pyrolytic Graphite Heat Spreader for Power Modules." Conference paper (2020).

[5] "Analysis and Evaluation of Thermally Annealed Pyrolytic Graphite Heat Spreader for Power Modules." OSTI 1649438 (2020).

[6] Boyd Corporation; Eaton. *k-Core Encapsulated Graphite Heat Spreaders — Technical Datasheets.* Retrieved 2026.

[7] "Annealed Pyrolytic Graphite." Technical reference summary, including copper/aluminum thermal-via practice.

[8] "Carbon Nanotube/Copper Composites for Via Filling and Thermal Management." IEEE Conference Publication.

[9] "Nano-engineered Multiwall Carbon Nanotube–Copper Composite Thermal Interface Material for Efficient Heat Conduction." NASA Technical Reports Server, NTRS 20050081904.

[10] "Characterization of Transferred Vertically Aligned Carbon Nanotube Arrays as Thermal Interface Materials." *International Journal of Heat and Mass Transfer*.

[11] "Preparation Method of Copper-Based Graphene Composite with High Thermal Conductivity." U.S. Patent 11,834,751.

[12] "Research Progress on the Preparation and Properties of Graphene–Copper Composites." *Metals* 15(10) (2025): 1117.

[13] "A Novel Fabrication Method of Copper–Reduced Graphene Oxide Composites with Highly Aligned Reduced Graphene Oxide and Highly Anisotropic Thermal Conductivity." (2022). PMC9064611.

[14] Cui, S., et al. "Enhancing Interfacial Heat Conduction in Diamond-Reinforced Copper Composites with Boron Carbide Interlayers for Thermal Management." *Composites Part B: Engineering* 287 (2024): 111871.

[15] "Preparation of Diamond/Copper Composites Modified by $Ti_3C_2T_x$ as Interlayer with Enhanced Thermal Conductivity." (2021).

[16] "Influence of Interfacial Carbide Layer Characteristics on Thermal Properties of Copper–Diamond Composites." *Journal of Materials Science* (2012).

[17] "Thermally Conductive Shape-Stabilized Phase Change Materials Enabled by Paraffin Wax and Nanoporous Structural Expanded Graphite." *Nanomaterials* 15(2) (2025): 110.

[18] "Preparation and Characterisation of Expanded Graphite/Paraffin Composite Phase Change Materials with Enhanced Thermal Properties." *Energy Sources, Part A* (2025).

[19] "Studies on the Thermal Characteristics of Nano-Enhanced Paraffin Wax Phase Change Material for Thermal Storage Applications." *Journal of Energy Storage* (2023).

[20] "Preparation and Characterization of Paraffin/Expanded Graphite Composite Phase Change Materials with High Thermal Conductivity." *International Journal of Heat and Mass Transfer* (2022).

Disclaimer. This brief is technical and strategic commentary for the CryptoSoWhat readership. Reported performance figures are representative published laboratory results and are not engineering specifications. Nothing herein is investment, legal, or engineering advice. Quantitative claims are attributed to the cited primary sources; the strategic interpretation is the author's.