




Heat Beneath the Surface

Thermal Metrology for Advanced
Semiconductor Materials and Architectures



As power density rises and architectures become vertical, heat transport across materials and interfaces is becoming increasingly difficult to model and impossible to ignore.

Contents

Introduction – The Forces Reshaping Heat in Modern Electronics.....	3
1. The Semiconductor Industry at an Inflection Point	3
2. Miniaturization, Integration, and the New Thermal Reality	4
3. The Rise of Heterogeneous and 3D Integration.....	5
4. The Thermal Crisis in AI and High-Density Computing.....	6
5. New Materials and the Expanding Scope of Thermal Design	7
6. Why Measurement and Metrology Must Evolve	7
7. Thermal Failures and Reliability	8
7.1 Temperature as a Reliability Accelerator	8
7.2 Common Thermally Driven Failure Modes.....	9
7.3 Bringing Thermal Earlier into the Design Workflow.....	9
Summary	10
The Rise of Extreme Properties.....	11
1. Extreme Scales: When Thickness Rewrites the Rules	11
2. Extreme Structures: 3D Geometries Becoming Thermal Bottlenecks	12
3. Extreme Conductivity: Engineered Materials Break Traditional Limits.....	13
4. Extreme Temperature: Devices Built to Operate at 200–1000 °C.....	13
New Materials and Growth Methods.....	14
1. Advanced Materials and Architectures for Thermal Control.....	14
2. Engineering Directional and Tunable Heat Transport	16
3. From TIM Optimization to Zero-Interface Architectures	16
4. Thermal Management as a Driver of Materials Innovation.....	18
New Architecture	20
1. Rise of 3D Stacked Dies	20
2. FinFETs to GAA: Vertical Heat Paths Make Interface Resistance First-Order	21
3. Backside Power Delivery Reshapes Internal Heat Distribution and Modeling	22
4. Die-Adjacent Thermal Interfaces in Advanced Cooling Architectures	23
5. Shortening the Thermal Path Reshapes Materials, Interconnects, and Packaging....	26
Thermal Management Implications and Evolution.....	27
1. Variability Becomes a First-Order Thermal Effect	27
2. Modeling Limits in a Multi-Scale, Multi-Physics Regime	29

3. Measuring Buried Layers and Vertical Heat Paths	30
4. From Measurement Challenge to Design Requirement	30
Conclusion: From Thermal Constraint to Design Foundation.....	32
1. Reliability Is the Forcing Function	33
2. The Thermal-First Workflow.....	33
3. Measurement as a Foundation, Not a Conclusion	34
4. The Central Shift	35
About Laser Thermal	36

Introduction – The Forces Reshaping Heat in Modern Electronics

Heat has become one of the dominant constraints in advanced semiconductor systems. For more than five decades, progress in semiconductors and electronics has been defined by a consistent trajectory toward higher performance, lower power consumption, and smaller device footprints. Today that trajectory has not slowed. It has accelerated, strengthened by the rapid expansion of artificial intelligence, electrification, high-bandwidth communications, and pervasive computing. Yet the path forward looks fundamentally different than it did during the classical era of Moore's Law. Instead of relying solely on transistor scaling, the industry is now driven by system-level complexity, advanced packaging, and heterogeneous integration that drive unprecedented thermal demands that push materials and architectures to their physical limits.

Across both compute and power electronics, the next generation of devices requires materials that can function under extreme thermal, electrical, and mechanical stresses while maintaining long-term reliability. This shift is creating an urgent need for new measurement capabilities and novel approaches to thermal design. The following introduction outlines the industrial and scientific forces driving this transition. Subsequent chapters will explore each topic in more detail, but here we establish the central argument: the evolution of semiconductor and electronic systems is creating a materials and integration challenge, and heat transport has become one of its defining constraints. In many cases, our ability to measure heat flow through thin films, interfaces, and stacked materials at the relevant scales has not kept pace with architectural ambition or reliability requirements.

1. The Semiconductor Industry at an Inflection Point

The semiconductor roadmap is undergoing one of the most significant transitions in its history. Traditional transistor scaling continues to deliver progress, but geometric shrinkage alone has not been the sole driver of performance for many years. As device dimensions approach fundamental limits, improvements increasingly depend on advances in materials, interconnects, packaging, and system architecture. At the same time, demand for compute, particularly from AI, cloud infrastructure, and advanced edge devices, continues to accelerate. Meeting this demand requires more than smaller transistors. It requires system-level innovation. Industry leaders are clear about this shift. As Julien Ryckaert, vice president of Logic Technologies at imec, put it **“We still scale transistors, but the real challenge today is system scaling. You need better materials, better interconnects, and better integration if you want the system to perform.”**¹

¹ Semiconductor Research Corporation, "Microelectronics and Advanced Packaging Technologies Roadmap Version 2.0." The MAPT Roadmap 2.0, 2025.

This shift toward system-level innovation is visible across every part of the technology stack:

- Gate-All-Around transistors (GAA) are replacing FinFETs at 2 nm nodes and below, offering stronger electrostatic control and improved performance.
- Backside power delivery, beginning with Intel 18A and advancing rapidly through foundry roadmaps, reduces current resistance and improves energy efficiency but introduces significant new thermal bottlenecks.
- High-NA EUV lithography is enabling sub-2 nm patterning with improved resolution at equivalent pitches, but as feature sizes continue to scale, stochastic variability, resist performance, and defect management become increasingly critical.
- Advanced materials such as wide-bandgap semiconductors (SiC and GaN), 2D materials like MoS₂, and wafer-scale InSe are emerging to address limits in silicon scalability.

These innovations represent a profound expansion of materials complexity. They also compound the industry's thermal challenges. As nodes shrink and architectures grow more compact, thermal design is no longer an after-thought. It is an essential component of performance scaling.

2. Miniaturization, Integration, and the New Thermal Reality

Device scaling has historically delivered predictable gains in speed and energy efficiency. Today those gains are harder to achieve. Thinner wafers, denser wiring, and higher device current translate into unprecedented heat flux that must be managed within shrinking thermal budgets.

Wafers themselves demonstrate how extreme scaling has become. Wafer thickness has fallen from 150 microns in 2005 to fewer than ten microns in advanced R&D today², with R&D exploring 5-micron platforms. Such thin geometries “strip materials of their bulk properties and make them fundamentally harder to process,” as Semiconductor Engineering recently summarized.³ At these scales, surface effects dominate, interfaces multiply, and thermal behavior diverges sharply from macroscale expectations. Materials that once behaved predictably not only change, but they also become significantly more difficult to characterize.

² *Technology News*, “Infineon halves power wafer thickness”, 2024. Nick Flaherty.

³ *Semiconductor Engineering*, “Precision Under Pressure: Managing Materials Complexity in Advanced Packaging”, 2025. Gregory Haley.

These changes are amplified by the industry's transition toward 3D integration. Stacked devices, chiplets, and high-bandwidth memory architectures produce concentrated thermal hotspots that cannot be effectively dissipated through conventional package designs. The MAPT Roadmap 2.0 describes this reality clearly:

“Thermal solutions will need to adapt to energy densities that exceed the capabilities of conventional approaches. The roadmap for AI thermal management must prioritize co-optimization of cooling, packaging, and compute density to ensure that thermal limitations do not become the primary barrier to continued system performance scaling.”⁴

In other words, scaling is no longer a lithography problem alone. It is increasingly a thermal problem.

3. The Rise of Heterogeneous and 3D Integration

Emerging semiconductor architectures increasingly rely on heterogeneous and three-dimensional integration to deliver performance gains that are no longer achievable through transistor scaling alone. This trend spans advanced packaging, chiplet-based design, and true 3D stacking, driven by the need to combine logic, memory, and specialized functions within tightly coupled systems.

As one recent SRC review notes, “integrating heterogeneous components introduces constraints stemming from differences in materials, processes, interfaces, and functionalities.”⁵ These constraints extend well beyond electrical compatibility. From a thermal perspective, heterogeneous integration fundamentally reshapes heat flow by introducing additional interfaces, thin bonding layers, and confined vertical pathways that limit heat spreading and amplify local temperature rise.

Across integration styles, several thermal consequences consistently emerge:

- Heat must traverse multiple bonded interfaces, often through low-thermal-conductivity bonding oxides, where both material resistance and thermal boundary resistance become dominant contributors.
- Active layers are increasingly buried farther from heat sinks, reducing access to efficient cooling paths.
- Lateral heat spreading is constrained by thinning dies and surrounding low-conductivity materials.
- Thermal behavior becomes spatially non-uniform, with hotspots and inter-die coupling that cannot be captured by bulk material properties alone.

⁴ *Nanoscale Adv.*, “Challenges and opportunities in engineering next-generation 3D microelectronic devices: improved performance and higher integration density”, 2024. Niharika Sing et al.

⁵ *Semiconductor Research Corporation*, “Microelectronics and Advanced Packaging Technologies Roadmap Version 2.0”. The MAPT Roadmap 2.0, 2025.

“The continued development of TIMs to improve thermal management in highly integrated package devices is a vital area of research and the characterization of the bulk and thin film properties of these systems needs to continue to advance.”⁶

The need to characterize thermal behavior in these systems therefore extends beyond measuring bulk conductivity. It includes localized thermal conductivity variations, interfacial thermal resistance across bonded layers, and spatial mapping of heat flow through multi-material stacks.

The MAPT Roadmap underscores this shift, emphasizing that continued progress in advanced packaging depends on improved characterization of both bulk and thin-film thermal properties, as well as spatially resolved measurements across “2 to 4 bonded silicon layers”.⁶

These requirements reflect a broader reality: legacy thermal measurement techniques were not designed for architectures where interfaces, confinement, and vertical heat transport define system performance.

4. The Thermal Crisis in AI and High-Density Computing

Nowhere is the thermal challenge more acute than in AI accelerators and large-scale compute systems. As power densities climb, heat flux levels well above 500 W per square centimeter are becoming common, with projections climbing beyond 1,000 W/cm² for next-generation accelerators. The thermal design power of a single GPU has already surpassed the kilowatt regime. Some roadmap projections anticipate 3,000 to 5,000 W devices as early as the next manufacturing cycle.⁷

This explosive rise in heat output is forcing a reexamination of cooling at every level of the stack:

- Next-generation TIMs with higher conductivity and lower interfacial resistance
- Diamond films and Cu–diamond composites as ultra-high-performance heat spreaders
- Si substrates offering improved thermal diffusion and mechanical robustness
- Microchannel liquid cooling embedded within lids or directly into silicon
- Backside cooling architectures, including backside microfluidics

TSMC and NVIDIA are among those actively rethinking the thermal path, integrating cooling into packaging and even into the silicon substrate itself.⁷ This marks a fundamental shift. Thermal management is no longer external to the chip. It is becoming intrinsic to device design.

⁶ Semiconductor Research Corporation, “Microelectronics and Advanced Packaging Technologies Roadmap Version 2.0”. The MAPT Roadmap 2.0, 2025.

⁷ SemiVision, “TSMC x Nvidia: Breaking the Thermal Wall: How Advanced Cooling Is Powering the Future of Computing”, 2025.

As one analysis summarized, thermal improvements depend on two parallel strategies:

1. Reducing thermal resistance along the conduction path
2. Expanding the effective area for heat exchange

Both strategies require new materials, new packaging architectures, and consequently accurate measurement of thermal properties across multiple length scales.

5. New Materials and the Expanding Scope of Thermal Design

The need for thermal innovation extends beyond cooling technologies. Emerging semiconductor materials themselves are reshaping thermal design requirements.

- Wide-bandgap semiconductors (GaN and SiC) now dominate power electronics and deliver superior efficiency but operate at much higher junction temperatures.
- 2D materials, including MoS₂, InSe, graphene, and black phosphorus, unlock ultra-thin device architectures but introduce complex thermal transport behavior.
- Advanced polymer composites aim to provide lightweight, electrically insulating materials with much higher thermal conductivity, addressing challenges in electronics, wiring, and packaging.
- Diamond-based materials offer extraordinary thermal conductivity and are emerging in AI and HPC thermal designs, though challenges in large-area growth, integration, and interfacial engineering remain significant.

These materials are promising and not slow to deploy because of synthesis, but because of characterization uncertainty at relevant scales.

6. Why Metrology Must Evolve

Modern thermal design depends on accurate characterization across multiple length scales, materials, and interfaces. Thermal properties must be measured in forms that reflect how materials are actually used, not idealized bulk analogs. This includes thin films, bonded interfaces, and spatially varying structures where local behavior can dominate overall performance. In these regimes, assumptions embedded in legacy techniques break down, and measurement uncertainty propagates directly into design risk.

As a result, thermal metrology is no longer a standalone validation step. It has become an integral part of a design feedback loop: measured thermal properties inform model calibration, calibrated models guide architectural decisions, and those decisions in turn define new measurement requirements. When any link in this loop is weak, predictive accuracy degrades, forcing designers to rely on conservative margins or late stage fixes that limit performance and yield.

Recent publications emphasize the importance of precise metrology. Nanometrology is now “a cornerstone of modern manufacturing” due to its role in linking materials at the micro and nano scale to system-level performance.⁸ Without measurement data that captures interface effects, spatial variability, and scale-dependent transport, even sophisticated simulation tools cannot reliably predict thermal performance.

The MAPT Roadmap 2.0 calls for improved tools to measure thermal conductivity, interfacial resistance, and spatial heat capacity fields within advanced packages.⁶ These needs mirror the industry’s shift toward architectures where thermal design is synonymous with system performance.

7. Thermal Failures and Reliability

Thermal behavior is not only a performance constraint. It is a primary driver of reliability issues. Elevated temperature accelerates degradation mechanisms, amplifies mechanical stress in layered stacks, and can transform small variations in materials or interfaces into early-life or field failures.⁹

Importantly, many dominant failure mechanisms in semiconductors follow Arrhenius-type behavior, meaning degradation rates increase exponentially with temperature. In practical terms, a sustained increase of even 5 to 10 °C at a hotspot can reduce device lifetime by a factor of two or more, depending on the mechanism. Small thermal errors therefore translate directly into large reliability uncertainties.

This distinction matters because many of the dominant thermal resistances in modern electronic systems are no longer external to the device. They reside inside the package: bond lines, underfills, dielectrics, vias, micro-bumps, TIM0 interfaces, and buried active layers. Once thermal resistance is embedded at this level of integration, late-stage mitigation often requires architectural or materials changes rather than incremental improvements to external cooling.

As integration increases, accurate measurement of thermal conductivity and interfacial resistance at these internal layers becomes essential not just for performance optimization, but for reliability prediction.

7.1 Temperature as a Reliability Accelerator

Most reliability models are inherently temperature dependent.¹⁰ While the underlying physics may differ across failure mechanisms, the practical implication is consistent:

⁸ *Advanced Functional Materials*, “Measuring the Future – Nanometrology for Advanced Manufacturing of Miniaturized Devices”, 2025. Madhurya Chandel et al.

⁹ *Semiconductor Engineering*, “Heat-Related Issues Impact Reliability in Advanced IC Designs”, 2024. Ann Mutschler.

¹⁰ *Electronics Cooling*, “How Many Ways Can Temperature Affect Performance and Reliability of Electronic Systems?”, 2024. Abhijit Dasgupta.

modest increases in local temperature can produce disproportionate reductions in device lifetime when degradation processes are thermally activated.

In highly integrated systems, reliability is rarely governed by average package temperature. Instead, it is controlled by peak junction temperatures, spatial temperature gradients, and the magnitude and frequency of thermal cycling at critical interfaces. When heat generation is localized and dissipation paths are constrained or buried, these conditions can exist even when external thermal metrics appear acceptable.

As packaging architectures become more compact and vertically integrated, the inability to observe or accurately model these localized thermal conditions introduces significant uncertainty into reliability predictions.

7.2 Common Thermally Driven Failure Modes

Thermal-driven failures manifest differently depending on where heat is generated and how it is transported through the system. Common examples include:

- Device-level effects: self-heating–induced performance drift; accelerated bias temperature instability and hot-carrier degradation; localized thermal runaway in power devices.
- Interconnect-level effects: temperature-accelerated electromigration in interconnects, micro-bumps, and TSV structures; increased resistance leading to localized Joule heating feedback loops.
- Package and interface-level effects: delamination at bonding interfaces; TIM pump-out or dry-out; void nucleation and growth; warpage and cracking driven by thermomechanical mismatch and thermal cycling.
- System-level effects: persistent hotspots due to uneven cooling; reliability risks associated with workload-dependent thermal excursions; failure modes introduced by advanced liquid-cooling architectures under transient conditions.

In many cases, these failure mechanisms originate at interfaces or within buried layers that are difficult to characterize using conventional thermal measurement approaches.

7.3 Thermal Continuity Across the Design Stack

A thermal-first workflow does not imply locking architectures prematurely. It means identifying thermal risk early enough that materials, interfaces, and integration choices remain adjustable. This requires defining thermal requirements alongside power and performance targets and grounding early models in measured, scale-appropriate material properties.¹¹

¹¹ *Semiconductor Engineering*, “Navigating Heat in Advanced Packaging”, 2024. Gregory Haley.

Key elements of such a workflow include:

- Defining a thermal budget early in the program, including allowable junction temperatures, hotspot limits, ambient conditions, and cooling assumptions.
- Identifying the dominant thermal resistances expected in the stack, including thin films, bonding layers, dielectrics, interfaces, and TIMs, and measuring the properties that govern them, such as thermal conductivity, volumetric heat capacity, and thermal boundary resistance.
- Calibrating early-stage models with measured data and iterating rapidly while trade space remains open across materials, die thickness, bonding strategies, and cooling architectures.
- Validating designs using spatially resolved temperature and property measurements to identify non-uniformity, buried layer effects, and interface quality issues before high-volume manufacturing.

Accurate thermal metrology is therefore not a downstream validation step. It is an enabling input across the entire design stack.

Summary

Semiconductor and electronics innovation is no longer driven by device scaling alone. It is shaped by heterogeneous integration, 3D architectures, AI-driven power density increases, and a rapidly expanding palette of novel materials. These shifts not only redefine how chips are built. They redefine what must be measured to ensure reliability, performance, and manufacturability.

Thermal management sits at the center of this transformation. The ability to dissipate heat effectively will determine how far performance can scale in the coming decade. As materials become thinner, interfaces grow more complex, and architectures approach their physical limits, advanced thermal metrology becomes essential.

This eBook begins with the premise that solving the thermal challenge or reducing the risk associated with thermal uncertainty is one of the most important bottlenecks facing modern electronics. The chapters that follow explore the technologies, materials, and measurement techniques that will define the next era of semiconductor development.

“The way semiconductor technology nodes have evolved is a testament to the industry’s ingenuity. But progress can no longer rely on transistor scaling alone. The next wave of innovations will come from rethinking chip and system architectures and pairing them with advanced thermal metrology, so we can actually measure, understand, and manage the thermal behavior that now defines system performance.”

- John Gaskins, CEO and Co-Founder, Laser Thermal

The Rise of Extreme Properties

Modern electronic devices now operate in regimes far beyond what classical thermal models and bulk-property assumptions can describe. As fabrication scales shrink, architectures become more complex, materials achieve unprecedented electrical conductivity, and operating temperatures rise, thermal behavior becomes a first-order design variable. The following four trends illustrate how extreme properties are reshaping the requirements for measurement, modeling, and materials engineering in next-generation electronics.

1. Extreme Scales: When Thickness Rewrites the Rules

Advances in deposition, growth methods, and lithography have pushed semiconductor films into regimes where thicknesses are comparable to, or smaller than, the mean free paths of their primary heat carriers. At these scales, phonons and electrons experience strong boundary scattering, driving thermal conductivity far below bulk values and invalidating classical Fourier assumptions. Studies on III–V and Si/Ge thin films show that layers only 1–10 nm thick behave fundamentally differently from their bulk counterparts, with transport dominated by surface interactions rather than intrinsic lattice behavior.¹² Similar effects emerge in metal bonding layers, Cu interconnects, and amorphous systems, where size-dependent conductivity suppression complicates thermal design as IC layers shrink toward carrier mean free paths.^{13 14} To handle this shift, engineers increasingly rely on direct metrology tied to thickness and process conditions. As Hans Olson notes, semiconductor teams now use FASTR (steady-state and frequency domain thermoreflectance technique) to build databases of thermal properties for accurate device-level simulation before fabrication.¹⁵ Once films enter the nanometer regime, direct property measurement becomes foundational for correct modeling and thermal reliability.

“At molecular scales, surface effects begin to dominate bulk behavior, and materials that are well-characterized in thick films exhibit entirely different properties when reduced to a few atomic layers.”¹⁶


¹² *Journal of Applied Physics*, “Effect of film thickness on the thermal resistance of confined semiconductor thin films”, 2009. E.S. Landry et al.

¹³ *International Journal of Heat and Mass Transfer*, “Size effect on thermal transport performance of inserted Cu/Cu₃Sn bilayer”, 2023. Xiaoyi Cai et al.

¹⁴ *Physical Review B*, “Size effects on the thermal conductivity of amorphous silicon thin films”, 2016. P. E. Hopkins et al.

¹⁵ *Zurich Instruments Interview*, “Hans Olson on Thin Films, Bulk Materials, and Heat Transfer Challenges”.

¹⁶ *Semiconductor Engineering*, “Precision Under Pressure: Managing Materials Complexity in Advanced Packaging”, 2025. Gregory Haley.



2. Extreme Structures: 3D Geometries Becoming Thermal Bottlenecks

The transition from planar MOSFETs to FinFETs, nanosheet, gate-all-around (GAA) devices, stacked FET architecture, and STI-embedded channels has created geometries where heat is confined, redirected, and trapped in fundamentally new ways. Narrow fins surrounded by low-k dielectrics restrict heat spreading, elevate local temperatures, and intensify self-heating effects as transistor dimensions shrink.¹⁷

In gate-all-around nanosheet devices, the channel is fully surrounded by gate dielectric and spacer materials with relatively low thermal conductivity compared to silicon.¹⁸ This geometry improves electrostatic control but limits conductive escape paths from the active region, increasing thermal confinement. Stacked FET architectures further compound this effect by vertically integrating active devices, reducing direct thermal access to the substrate.

Backside power delivery networks (BSPDN) represent one of the most dramatic recent architectural shifts from a thermal perspective. While they improve routing efficiency and reduce IR drop, they alter vertical heat-flow pathways and introduce additional buried layers and interfaces that must be traversed before heat reaches the heat sink.

These architectural shifts are amplified at the packaging level: 2.5D and 3D integration confine heat within stacked dies, where limited lateral spreading and low-conductivity inter-die materials create hotspots and thermal coupling between tiers.¹⁹ As Semiconductor Engineering has observed, although electrical scaling continues, thermal constraints (not lithography) are increasingly the boundary on integration density.¹³

¹⁷ *Semiconductor Engineering*, “Self-Heating Issues Spread”, 2023. Brian Bailey.

¹⁸ *Journal of the Electron Devices Society*, “Impact of Self-Heating Effect on DC and AC Performance of FD-SOI CMOS Inverter”, 2024. K. H. Lee et al.

¹⁹ *IDTechEx*, “Thermal strategies for 2.5D and 3D semiconductor packaging”, 2025. Yulin Wang.

Effective modeling, hotspot-aware placement, and new materials become critical as buried and vertically integrated structures reshape thermal pathways at every level of the device.

3. Extreme Conductivity: Engineered Materials Break Traditional Limits

New classes of engineered materials have redefined how efficiently heat can be removed from semiconductor devices. Composites built from aligned fillers, sintered microstructures, and hybrid lattice architectures routinely achieve thermal conductivities five to ten times greater than previous generations of thermal interface materials at comparable thickness. Boron nitride nanotubes (BNNTs), for example, deliver $\sim 600 \text{ W/m}\cdot\text{K}$ thermal conductivity while remaining electrically insulating, an unusual and highly desirable combination for power electronics, RF devices, and aerospace systems.²⁰

Meanwhile, diamond-based composites such as the recent report of a copper–diamond material from Element Six which achieved values of $\sim 800 \text{ W/m}\cdot\text{K}$ by pairing diamond’s exceptional thermal conductivity with copper’s integrability.²¹ These materials can directly address the escalating thermal loads of AI accelerators, GaN amplifiers, and high-density logic. In GaN-on-diamond systems, integrating diamond and h-BN has been shown to reduce lattice temperature by more than 100 K and significantly enhance device performance by suppressing self-heating.¹⁷ Through these innovations, thermal interface engineering has entered an era of extreme conductivity, enabling previously unattainable power densities.

4. Extreme Temperature: Devices Built to Operate at 200–1,000 °C

As power electronics and advanced semiconductor systems push deeper into high-voltage and high-frequency regimes, operating temperatures above 200 °C have become common, particularly in SiC and GaN technologies. GaN HEMTs are especially vulnerable to localized self-heating, but advances such as GaN-on-diamond heterostructures and h-BN passivation significantly reduce channel temperature and increase device performance. At the same time, packaging and encapsulation materials are being re-engineered to survive sustained operation beyond 200 °C to align with the thermal resilience of ultra-wide-bandgap devices.

In aerospace and extreme-environment applications, BNNT composites maintain structural integrity beyond 900 °C, and ultra-high-temperature ceramics (UHTCs) support operation at or above 1,000 °C.¹⁶ These developments reflect a fundamental industry shift: thermal resilience is no longer a boundary condition, it is a performance axis, and entire material ecosystems are being designed to thrive under extreme temperature constraints.

²⁰ McKinsey Electronics, “Boron Nitride Nanotubes: The Future of Thermal Interfaces”, 2025.

²¹ McKinsey Electronics, “Emerging Thermal Management Technologies: Diamond Cooling, BNNT Materials and UVTP in Semiconductor Innovation”, 2025.

New Materials and Growth Methods

1. Advanced Materials and Architectures for Thermal Control

Thermal management has become a defining challenge across modern technology sectors, from advanced electronics and optoelectronics to transportation, energy systems, and high-performance computing. Increasing power densities, shrinking geometries, heterogeneous integration, and new architectural constraints have driven heat fluxes to levels that far exceed the capabilities of conventional materials and interfaces at the design, package, and system levels.²² At micro and nanoscales, heat transport is further complicated by size effects, boundary scattering, and deviations from classical Fourier behavior, requiring new strategies to control phonon transport and create efficient thermal pathways.²³ These pressures have accelerated research into ultrahigh thermal conductivity materials, engineered micro/nanostructures, and thermal metamaterials that can direct, confine, or redistribute heat with unprecedented precision.

A major area of progress lies in interface engineering and composite architectures, which are increasingly essential because most real systems rely on multi-layer stacks where thermal boundary resistance often dominates heat flow. Polymer composites, TIMs, and structural composites used at die attach, lid attach, cold plate interfaces, and system-level heat exchangers now leverage aligned or networked fillers to build continuous thermal pathways, while maintaining mechanical compliance and reliability.^{24 25} A notable demonstration is the use of ice-templated self-assembly to align boron arsenide into lamellar structures within a polymer matrix, producing TIM composites with thermal conductivity above 20 W/mK while keeping modulus in the kilopascal range, a combination that overcomes long-standing limits of commercial greases and epoxies.²⁶ Hybrid filler systems, hierarchical structuring, and additive manufacturing further expand design freedom, enabling anisotropic conduction, multifunctional mechanical behavior, and optimized geometries for targeted heat flow in both chip-scale and board- or module-level implementations.

In practice, composite and interface-driven package thermal performance is usually governed by a small set of controllable levers:

²² *Accounts of Materials Research (ACS)*, “Advancing Thermal Management Technology for Power Semiconductors through Materials and Interface Engineering”, 2025. Man Li et al.

²³ *Advanced Materials*, “Beyond Conventional Cooling: Advanced Micro/Nanostructures for Managing Extreme Heat Flux”, 2025. Yuankun Zhang et al.


²⁴ *Springer Nature*, “Thermal interface polymer-based composites materials: a critical review”, 2025. Mesum Abbas et al.

²⁵ *Electronics*, “A Review of Advanced Thermal Interface Materials with Oriented Structures for Electronic Devices”, 2024. Yuqian Tu et al.

²⁶ *Nature Communications*, “Flexible thermal interface based on self-assembled boron arsenide for high-performance thermal management”, 2021. Ying Cui et al.

- Filler volume fraction and percolation - does a continuous heat path form?
- Filler geometry and orientation - aspect ratio, alignment, and anisotropy
- Dispersion quality - agglomeration, voids, and local non-uniformity
- Interface thermal resistance - surface preparation, functionalization, and bonding quality
- Processing conditions - pressure, cure profile, and resulting porosity

Micro- and nano-structuring techniques are also unlocking new approaches to thermal transport control. Textured surfaces, engineered porosity, and bio-inspired topographies can reduce thermal resistance, enhance convection, or tune wettability for improved liquid cooling in cold plates, vapor chambers, and microfluidic cooling channels. In integrated systems, patterned thin films, nanoscale heat spreaders, and thermal metamaterials allow designers to shape heat flow around hotspots or confining geometries, improving stability and performance across scales.²⁷ Although manufacturing uniformity and scalability remain challenges, these innovations point toward a future where materials are no longer passive thermal conductors, but actively engineered elements that manage heat as deliberately as electrical or mechanical properties. Together, these advances form the foundation for next-generation thermal technologies across a wide range of applications.



²⁷ AZO NANO, “Keeping Cool: Nanostructures for Enhanced Electronic Performance”, 2025. Muhammad Osama.

2. Engineering Directional and Tunable Heat Transport

Modern materials growth methods now provide unprecedented control over grain size, crystal orientation, defect density, and impurity concentrations, enabling deliberate tuning of both electrical and thermal transport properties primarily at the device and thin-film level. Techniques such as molecular beam epitaxy, chemical vapor deposition, pulsed laser deposition, and advanced sintering or recrystallization processes allow microstructure to be engineered at nanometer-to-micron scales in active layers, interconnects, substrates, and engineered heat spreaders.²⁸ By controlling grain boundaries, stacking order, interfacial roughness, and dopant distribution, thermal conductivity can be enhanced through phonon channeling mechanisms or selectively suppressed through increased scattering.

These same variables also explain why tabulated thermal conductivity values increasingly fail to describe real devices. Thermal conductivity is not determined by chemical composition alone, but by the specific microstructural state produced by a given growth process. Materials with identical nominal composition can exhibit markedly different thermal behavior depending on orientation, defect density, interface quality, and thickness. As growth techniques advance, the spread in achievable thermal properties widens, while the relevance of a single reference value diminishes.

This effect is particularly pronounced in fin-based and vertically integrated architectures, where heat must traverse narrow, high-aspect-ratio features. In these confined geometries, thermal transport becomes strongly anisotropic and highly sensitive to atomic-scale disorder. Small variations in growth or post-processing can shift heat flow pathways, alter thermal boundary resistance, and change effective conductivity along different directions.²⁹ In such complex nanoscale device architecture, bulk tabulated values obscure the variations that ultimately control device temperature and reliability.

As a result, thermal conductivity must be treated as a process- and structure-dependent property rather than a fixed material constant. Accurate thermal design increasingly relies on measurements that reflect the actual growth method, geometry, and interface conditions present in the device or package structure under consideration. This perspective directly supports the later discussion of variability as a first-order thermal effect.

3. From TIM Optimization to Zero-Interface Architectures

At advanced power densities and shrinking pitches, interface resistance at die-adjacent and package-level interfaces has emerged as a first-order limiter rather than a secondary

²⁸ *Nano Letters*, “Direct Visualization of Thermal Conductivity Suppression Due to Enhanced Phonon Scattering Near Individual Grain Boundaries”, 2018. Aditya Sood & al.

²⁹ *Science Direct*, “Size-dependent phononic thermal transport in low-dimensional materials”, 2020. Moore & Shi.

loss term at both on-chip interconnect interfaces and package-level bonding layers. Both electrically and thermally, nominally “mated” surfaces interact only through a sparse population of microscopic asperities, with real contact areas often on the order of few percents even under substantial pressure.^{33b} As Semiconductor Engineering highlights for electrical contacts, small variations in surface cleanliness, roughness, compression, and metallurgy can remain hidden at low load but become catastrophic under high current density.³⁰ The same physics governs heat flow: voids filled with air dominate thermal impedance unless the interface is deliberately engineered to eliminate them.

Traditional TIMs such as greases, pads, and solders were historically sufficient because package power, pitch, and thermal gradients were modest. Today, their limitations are increasingly exposed. Greases achieve thin bond lines but suffer from pump-out, phase separation, and long-term dry-out under thermal cycling.³¹ Pads trade handling robustness for thicker bond lines and pressure-dependent conformity, while solders face reliability and scaling limits as spacing shrinks and electromigration, intermetallic growth, and thermal fatigue accelerate.³² In all cases, the interface itself, not the bulk conductivity of the materials, dominates system-level thermal resistance and long-term drift.

Phase-change TIMs represent a deliberate shift toward minimizing interface resistance rather than simply optimizing filler loading. By transitioning from a solid or semi-solid state into a viscous or molten phase at operating or reflow temperatures, PCMs can actively wet surfaces, collapse microvoids, and form bond lines approaching grease-like thickness without grease-like handling penalties. This self-adjusting behavior under heat and pressure directly targets the dominant failure modes of conventional TIMs, particularly pump-out and contact degradation during cycling. As a result, phase-change compounds are best viewed as an intermediate step toward interfaces that dynamically approach an ideal, high-contact-area state rather than static fillers of geometric gaps.

In parallel, research on advanced TIMs with oriented fillers underscores a deeper trend: conventional polymer-filler composites are nearing a performance ceiling. Even with aggressive loading, isotropic TIMs struggle to exceed $\sim 10\text{--}12\text{ W/m}\cdot\text{K}$ without compromising mechanical compliance.³³ Oriented graphene, carbon fiber, and boron nitride architectures can deliver metal-grade through-plane conductivity in laboratory demonstrations, but they remain sensitive to pressure-induced structural collapse, electrical conductivity constraints, and manufacturability at scale. The common theme is clear: as long as a heterogeneous interlayer remains, interface resistance and reliability remain difficult to fully control.


³⁰ *Semiconductor Engineering*, “The Hidden Cost of Contact Resistance”, 2025. Gregory Haley.

³¹ *ES Materials & Manufacturing*, “Recent Advances in Thermal Interface Materials”, 2020. Yongcun Zhou et al.

³² *IOP Science*, “Recent Progress in Cu–Cu Direct Bonding Technology”, 2025. Ze-Hao Zhao et al.

³³ *Nanomaterials*, “Recent Advances in Thermal Interface Materials for High-Power Electronics”, 2022.

^{33b} *Journal of Thermal Science*, “Thermal Contact Resistance and Heat Transfer Enhancement Mechanisms at Non-Smooth Contact Interfaces”, 2025. Y. Wang & al.



The logical endpoint of this trajectory is nanoscale direct bonding in advanced packaging and heterogeneous integration, where the goal is to eliminate the TIM layer altogether and approach true “zero-interface” behavior. In practice, these approaches are currently limited to specific architectures, materials systems, and cost regimes, and remain tightly coupled to manufacturing complexity and yield considerations.

In advanced packaging, this is already visible in the migration from Sn-based micro-bumps toward Cu–Cu direct and hybrid bonding, which enables finer pitch, lower electrical and thermal resistance, and improved reliability through atomic-scale diffusion across the bond interface.²⁸ Techniques such as nanotwinned or nanocrystalline copper, plasma activation, and engineered surface passivation are designed to make the bonded interface disappear as a distinct thermal and electrical barrier at BEOL-compatible temperatures.³⁴ From a thermal perspective, this represents the ultimate extension of the TIM roadmap: moving from materials that manage interface imperfections, to materials that collapse them in operation, and finally to architectures that remove the interface as a dominant resistance term altogether.

4. Thermal Management as a Driver of Materials Innovation

Recent excitement around wide-bandgap semiconductors and ultra-high-thermal-conductivity materials underscores a broader shift in how materials are evaluated for advanced electronics: thermal management is no longer a secondary constraint, but a defining parameter of material innovation at both the device substrate and heat-spreader level. In GaN and SiC power devices, escalating power densities have pushed self-heating to the forefront, motivating the integration of materials such as diamond, not for novelty, but because conventional substrates and heat-spreading approaches are no longer sufficient. Demonstrations of GaN HEMTs on diamond substrates show more than 2× improvement in heat dissipation compared to GaN-on-SiC, achieved by explicitly engineering the thermal interface to reduce boundary resistance and exploit diamond’s extreme thermal conductivity.³⁵

³⁴ *Nature Communications*, “Nanocrystalline copper for direct copper-to-copper bonding with improved cross-interface formation at low thermal budget”, 2024. Chuan He et al.

³⁵ *Semiconductor Today*, “GaN HEMTs on diamond demonstrates twice the heat dissipation of GaN-on-SiC”, 2024.

Industry-facing case studies now frame diamond heat spreaders as enabling higher RF power density, longer duty cycles, and improved reliability, rather than incremental thermal optimization.³⁶

At the same time, emerging materials like boron arsenide (BAs) illustrate how thermal performance is now driving fundamental semiconductor research itself for potential use as substrates, heat spreaders, or integrated functional layers. Recent experimental work shows high-quality BAs crystals achieving thermal conductivity exceeding 2,000 W/m·K, rivaling or surpassing diamond while retaining favorable semiconductor properties such as a mid/wide bandgap, high carrier mobility, and compatible thermal expansion.³⁷ Crucially, the excitement around BAs is not limited to heat spreading alone; it reflects a recognition that future semiconductor materials must intrinsically combine electronic functionality and exceptional heat transport. As devices approach limits set by power density and 3D integration, these developments signal a reframing of materials innovation: success is increasingly defined by how effectively a material manages heat at the device and interface level, not merely how well it switches or conducts charge.

Examples of conductivity improvements of materials between 2015 and 2025

Material / class	~2015 thermal conductivity (W/m·K)	~2025 thermal conductivity (W/m·K)
Cu (bulk)	400	400 (no change)
Cu thin film (10 nm)	~50-100	~150-200
Sintered Ag TIM	~20	60-100
Graphite sheet (APG/TPG)	~1,000	1,500-2,000
Diamond	~2,200	2,500-3,200
Polymer composite (TCP)	< 2	5-20

“Many applications need high thermal conductivity to effectively take heat out of the devices so that they work properly.”
— Zhifeng Ren, University of Houston

³⁶ *Element Six*, “Diamond heat spreaders for GaN-on-SiC RF power amplifiers (Industry case study)”, 2023–2024.

³⁷ *Laser Focus World*, “Boron arsenide surpasses diamond, silicon in thermal conductivity”, 2025. Sally Cole Johnson.


New Architecture

1. Rise of 3D Stacked Dies

3D integration represents a structural shift in semiconductor design because it moves performance gains from transistor scaling toward system-level integration density. Stacking logic, memory (especially HBM), and specialized chiplets increases bandwidth and reduces interconnect length, enabling architectures that would be difficult to realize in 2D layouts. Industry coverage increasingly frames advanced packaging as a “front-end design imperative,” particularly for AI and HPC systems where packaging decisions directly shape performance and manufacturability.³⁸

However, the same vertical integration that improves electrical behavior introduces a thermal penalty. Thin dies reduce lateral heat spreading, buried tiers sit farther from the heat sink, and inter-die materials such as bonding layers and dielectrics form low-conductivity barriers that trap heat. Analyses of 3D stacks consistently highlight heat accumulation in middle dies and limited thermal escape paths as a core bottleneck, motivating interest in thermal TSVs, local heat spreaders, and embedded or liquid cooling despite their added complexity.³⁵

As adoption accelerates, large-scale logic stacking continues to face fundamental constraints that blend physics with manufacturing realities. Mechanical requirements often force thicker or more compliant bonding layers to protect interconnect reliability, but those same layers increase thermal resistance.³⁹ As a result, thermal behavior becomes inseparable from yield, reliability, and process integration, elevating the importance of metrology and design-for-manufacturing as integration density rises.⁴⁰



³⁸ *IDTechEx*, “Thermal Strategies for 2.5D and 3D Semiconductor Packaging”, 2025. Yulin Wang.

³⁹ *Semiconductor Engineering*, “Getting Rid of Heat in Chips”, 2023. Brian Bailey.


⁴⁰ *Royal Society of Chemistry*, “Challenges and Opportunities in Engineering Next-Generation 3D Microelectronic Devices”, 2024. Niharika Singh et al.

2. FinFETs to GAA: Vertical Heat Paths Make Interface Resistance First-Order

FinFETs were the industry's first non-planar transistor, and they made self-heating impossible to ignore. Compared to planar devices, fins are partially isolated by surrounding dielectrics, which limits heat dissipation into the substrate and raises local device temperatures. While chip-level temperatures may still be governed by package limits, local hot regions increasingly drive variability and long-term reliability concerns.⁴¹

Gate-All-Around (GAA) architectures extend this trend further. As devices transition from FinFETs to GAA, heat escape becomes even more difficult because the active region is smaller and more completely surrounded by low-thermal-conductivity materials. Industry coverage emphasizes that while GAA enables continued scaling beyond FinFETs, it also intensifies local self-heating at a point in the roadmap where both device density and power density are at their peak.

The implication is that interface thermal resistance and nanoscale heat flow paths become first-order design constraints. Small changes in interfacial conductance, dielectric thermal resistance, or local heat spreading can now translate directly into performance loss, accelerated aging, and tighter operating margins. Thermal design and thermal metrology must therefore move closer to the device and interface scale, rather than relying solely on bulk package-level measurements.³⁸




⁴¹ Semiconductor Engineering, "Self-Heating Issues Spread", 2023. Brian Bailey.

3. Backside Power Delivery Reshapes Internal Heat Distribution and Modeling

Backside Power Delivery Network (BSPDN) re-architects the power distribution network by routing power rails on the backside of the wafer rather than through frontside interconnect layers. This decouples power delivery from signal routing and enables thicker, lower-resistance power lines, improving voltage integrity and freeing frontside routing resources for logic density and performance.^{42 43}

While the primary motivation is electrical, the architectural consequence is thermal. Relocating power rails and vias alters where heat is generated and how it propagates vertically through the stack. In conventional planar and FinFET architectures, the silicon substrate beneath the transistor acts as a primary heat spreading medium. With BSPDN and associated wafer bonding and thinning steps, that bulk silicon path is significantly reduced or reconfigured, and the transistor becomes increasingly embedded between interconnect stacks. In effect, active devices are sandwiched between BEOL layers on both sides, where low-k dielectrics and metallization, rather than bulk silicon, dominate the immediate thermal environment.



⁴² imec, "How to Power Chips from the Backside: Benefits and Building Blocks of a Backside Power Delivery Network", 2022. Naoto Horiguchi et al.

⁴³ McKinsey Electronics, "The Best Kept Secret in Semiconductor Innovation: Backside Power Delivery", 2023.

As BSPDN approaches production readiness at sub-2nm nodes, manufacturers must manage new process steps such as wafer bonding thinning, nano-TSV formation, and backside metallization, each of which introduces new thermal paths and integration sensitivities.^{44 45}

Importantly, BSPDN does not eliminate thermal challenges; it redistributes them. Reduced resistive losses can lower some heat generation, but altered material stacks and vertical pathways can intensify localized hotspots if not carefully co-designed. As a result, BSPDN reinforces the need for coupled electrical-thermal modeling and advanced metrology that can resolve vertical heat flow and localized dissipation in complex architectures.⁴⁶

4. Die-Adjacent Thermal Interfaces in Advanced Cooling Architectures

Advanced integration technologies such as chiplets, Chip-on-Wafer-on-Substrate (CoWoS), and large multi-die modules are reorganizing the thermal stack to reduce resistance and manage hotspots. This has driven increased emphasis on TIM0 interfaces at the die level and configurations in lidless packages, where the cooling solution couples more directly to the die. Industry definitions consistently describe TIM0 as the interface between a bare die and a heatsink or cold plate, distinct from traditional TIM1 and TIM2 layers.⁴⁷

In parallel, cooling architectures themselves are evolving. Air-cooled heatsinks are increasingly supplemented or replaced by liquid-cooled cold plates, direct-to-chip liquid cooling, and in some cases backside-cooled or embedded microfluidic approaches. These architectures shorten the thermal path and increase local heat transfer coefficients, but they also intensify mechanical, materials, and interface constraints at the die-adjacent layer.

This evolution increases both the number and criticality of interfaces while shrinking available spreading area. Removing lids and collapsing interfaces can shorten the thermal path, but it also places stricter demands on coplanarity, bondline thickness control, void mitigation, and long-term reliability, especially in large, high-value assemblies.⁴⁸ Under liquid cooling, these constraints become even tighter due to higher clamping pressures, dynamic thermal cycling, and in some cases fluid-induced mechanical stresses. These challenges elevate TIM selection and process control from secondary considerations to architectural enablers.

⁴⁴ *Lam Research Newsroom*, “The Other Side of the Wafer: The Latest Developments in Backside Power Delivery”, 2022. Sandy Wen.

⁴⁵ *Semiconductor Engineering*, “Backside Power Delivery Nears Production”, 2025. Laura Peters.

⁴⁶ *IEEE Spectrum*, “Intel Is All-In on Backside Power Delivery”, 2023. Samuel Moore.

⁴⁷ *I-Connect007*, “Beyond Thermal Conductivity: Exploring Polymer-Based TIM Strategies for High-Power-Density Electronics”, 2025. Padmanabha Shakthivelu and Nico Bruijnjs.

⁴⁸ *Indium Corporation Blog*, “Solder TIMs: Decoding TIM1 vs. TIM1.5 for Advanced Packaging”, 2024. Jim McCoy.

At the same time, material strategies are evolving rapidly. Industry and research sources document advances in solder-based TIMs, liquid metals, polymer composites with aligned fillers⁴⁹, and interfacial engineering approaches designed to reduce thermal boundary resistance while preserving mechanical resilience.⁵⁰ In direct liquid-cooled configurations, TIM materials must maintain stability under higher clamping pressures, withstand aggressive thermal cycling, and support increasingly thin bondlines required for high heat-flux operation. As cooling performance improves, the relative contribution of the TIM layer to total thermal resistance increases, placing greater emphasis on conductivity, coplanarity control, and long-term mechanical reliability.

At the most aggressive end of the roadmap, direct backside liquid cooling and in-chip microfluidic cooling architectures aim to reduce or eliminate conventional polymer TIM layers by bringing the cooling medium closer to the active silicon. Demonstrations from leading research programs and advanced packaging platforms show meaningful reductions in thermal resistance when lids are removed or when cooling is integrated directly at the silicon interface. However, these approaches are not yet broadly deployed in high-volume manufacturing and remain tightly coupled to yield, reliability, cost, and integration complexity constraints.^{62 63}

In practice, TIM layers do not simply disappear. Rather, they are replaced by alternative bonded interfaces, metallic joints, or engineered silicon structures that still introduce thermal boundary resistance and mechanical design challenges. As cooling moves closer to the transistor, the interface becomes thinner, more structurally integrated, and more sensitive to coplanarity, stress, and defect control. Interface engineering therefore does not diminish in importance; it becomes more critical, shifting from bulk polymer layers to precision bonding and silicon-level integration strategies.⁶⁴

⁴⁹ *Carbon*, “Carbon-Based 3D Array-Reinforced Thermal Interface Materials with Highly Oriented Structure and Superior Thermal Conductivity”, 2025. Sijin Yan et al.





⁵⁰ *CS MANTECH Conference*, “Characterization of a Novel Thermal Interface Material Based on Nanoparticles for High Power Device Package Assembly”, 2023. Zeina Abdallah et al.

⁶² *Engineering*, “A Review of Recent Developments in “On-Chip” Embedded Cooling Technologies for Heterogeneous Integrated Applications”, 2023. Srikanth Rangarajan et al.

⁶³ *Techovedas*, “TSMC Breaks the Thermal Wall: 5 Game-Changing Facts About Its Direct-to-Silicon Liquid Cooling”, 2025. Kumar Priuadarshi.

⁶⁴ *IEEE 74th Electronic Components and Technology Conference (ECTC)*, “An Energy-efficient Si-integrated Micro-cooler for High Power and Power-density Computing Applications,” 2024. Y. -J. Lien et al.

Emerging Liquid Cooling Methods



5. Shortening the Thermal Path Reshapes Materials, Interconnects, and Packaging

Across device scaling, packaging, and system integration, a common architectural objective has emerged: minimize the distance and resistance between heat generation and heat removal. As geometries shrink, current density and localized heat generation increase. At the same time, newer materials, thinner layers, and more interfaces can raise thermal resistance, making heat extraction more difficult as device and packaging architectures become dimensional.³⁸

At the package level, this principle drives designs toward lidless approaches and more direct coupling between die and cooling hardware. Eliminating intermediate layers reduces junction-to-case resistance, but it also forces tighter control over interface planarity, material compliance, and metrology. Packaging decisions increasingly reflect a trade-off between electrical optimization and thermal viability rather than purely mechanical convenience.⁵¹

In 2.5D and 3D systems, the tension is even sharper. Shorter electrical interconnects improve bandwidth and efficiency, but vertical stacking lengthens thermal escape paths and increases hotspot risk, particularly for buried dies. This conflict is why the industry is simultaneously investing in new cooling concepts, advanced interface materials, and high-resolution thermal measurement techniques. The push to shorten electrical paths is now inseparable from the need to deliberately engineer and measure thermal paths across every scale of the system.⁵²

⁵¹ *MicrochipUSA*, “The Future of Semiconductor Miniaturization”, 2024.

⁵² *IDTechEx*, “Thermal Strategies for 2.5D and 3D Semiconductor Packaging”, 2025. Yulin Wang.

Thermal Management Implications and Evolution

1. Variability Becomes a First-Order Thermal Effect

Thermal measurement sensitivity reflects how strongly measured thermal properties depend on factors that are often assumed to be negligible. Two samples that are nominally identical can yield meaningfully different results due to differences in surface roughness, boundary and interface conditions, growth conditions, and sample history. In thermal transport, these details matter. Assuming that a thin film produced by one supplier will exhibit the same thermal behavior as an ostensibly identical film from another supplier is rarely a safe assumption.

Silicon provides a clear example. At room temperature, its thermal conductivity can vary by roughly 15–25% as doping concentration increases from lightly doped to heavily doped regimes, even when crystal structure and geometry remain unchanged. This variability implies that thermal conductivity values used in models may frequently be inaccurate. In some cases, the thermal conductivity of electronically relevant materials listed in property databases within simulation software can be incorrect by an order of magnitude. These discrepancies highlight a broader point: thermal conductivity is not an absolute material constant. It is a property that depends on processing, structure, interfaces, and measurement conditions, and must be measured in the context in which the material is actually used.


The Impact of Doping on Silicon Thermal Conductivity

Thermal measurement sensitivity in silicon depends on several interconnected variables, including intrinsic thermal conductivity, dopant concentration, sample preparation, and interface conditions. Doping introduces impurity atoms into the silicon lattice, typically using group III or V elements such as boron or phosphorus, which adjust the number of available charge carriers and shift the Fermi level to create desired electrical behavior. These impurities also alter the phonon transport landscape by disturbing the periodicity of the lattice. As a result, phonon scattering increases, the phonon mean free path decreases, and thermal conductivity is reduced. Studies using time-domain thermoreflectance measurements show that doped silicon at concentrations near $2 \times 10^{19} \text{ cm}^{-3}$ exhibits thermal conductivity reductions of approximately 22% relative to intrinsic silicon, with similar trends observed across different dopant species.⁵³

⁵³ *International Journal of Thermal Sciences*, “Systematic investigations on doping dependent thermal transport properties of single crystal silicon by time-domain thermoreflectance measurements”, 2022. Xuanhui Fan et al.

The magnitude of this reduction depends on dopant type, concentration, and mass variance. Higher impurity densities amplify impurity scattering, especially in both n-type and p-type doping regimes, and similar mechanisms have been observed across other semiconductor systems, such as silicon-doped GaN. Structural imperfections introduced during wafer growth or device processing further influence thermal transport. Dislocations, stacking faults, oxidation-induced defects, and vacancy clusters all reduce conductivity by creating additional scattering centers. In thin silicon layers below roughly one hundred nanometers, boundary scattering dominates, suppressing thermal conductivity even when doping levels remain moderate.

Thermal conductivity in Si as a function of the doping level for various dopants.



Hahn, Konstanze & Melis, Claudio & Bernardini, Fabio & Colombo, Luciano. (2021). Engineering the Thermal Conductivity of Doped SiGe by Mass Variance: A First-Principles Proof of Concept. *Frontiers in Mechanical Engineering*. 7. 712989. 10.3389/fmech.2021.712989.

Thermal boundary conditions also influence measurement sensitivity. The presence of interface layers, surface roughness, or oxide films can modify thermal boundary resistance, which must be considered when interpreting measurements. Temperature is another important variable, since bulk silicon's thermal conductivity decreases with temperature above 30 K due to intensified phonon-phonon scattering. Together, these factors mean that accurate thermal characterization of doped silicon requires control of dopant species, concentration, and distribution, as well as careful attention to sample preparation and interface engineering. Understanding these interactions is essential for assessing device-level heat flow in microelectronics, power devices, and nanoscale structures where thermal transport plays a critical role.^{54 55}

⁵⁴ *University Wafer, Inc.*, "Doped Vs Undoped Silicon Wafers: A Comprehensive Comparison".

⁵⁵ *Frontiers in Mechanical Engineering*, "Engineering the Thermal Conductivity of Doped SiGe by Mass Variance: A First-Principles Proof of Concept", 2021. Konstanze Hahn & al.

2. Modeling Limits in a Multi-Scale, Multi-Physics Regime

Across advanced semiconductor systems, the relationship between material structure and measured performance is increasingly difficult to model accurately. Simulation tools remain essential, but persistent gaps exist between modeled behavior and real-world thermal and electrical response, particularly as designs span nanometer-scale devices through package- and system-level assemblies.

Simulation accuracy degrades as systems move into regimes dominated by multi-scale physics and quantum-scale effects. As Silicon Semiconductor notes, modern heterogeneous systems span “from nanometer-scale transistors to centimeter-scale packages,” requiring simulation tools to transition across scales “while maintaining accuracy and computational efficiency.”⁵⁶ This challenge is amplified as electrical and thermal effects converge, introducing behaviors that are not well captured by classical or single-domain modeling approaches.

Even when simulations are capturing the general thermal behavior, nominal assumptions can obscure real variability. Semiconductor Engineering emphasizes that early simulations are typically run under idealized conditions, and that discrepancies emerge once real-world variations are introduced.⁵⁷ As Teradyne’s Nitza Basoco notes, “Initial predictions will always have gaps, and test results must continuously feed back into refining the model.”⁵⁵ Temperature gradients, process variation, and material inconsistencies routinely expose limits in predictive accuracy once devices move from design intent toward fabricated reality.⁵⁷

Thermal modeling is particularly sensitive to interfaces, surface conditions, and stacking-induced stress. In 3D integrated structures, stacked and thinned dies increase thermal resistance and create non-uniform heat flow paths, complicating both simulation and interpretation of results.⁵⁸ As noted in Semiconductor Engineering, thermal expansion mismatches and stress accumulation in stacked devices often cannot be anticipated in simulation alone, requiring multiphysics analysis and empirical validation.⁵⁹ These issues are further compounded by variations in surface smoothness, wafer thinning, underfill application, and interface quality—seemingly minor deviations that can accumulate and significantly alter performance.

⁵⁶ *Silicon Semiconductor*, “How simulation is helping to usher in tomorrow’s chips”, 2025.

⁵⁷ *Semiconductor Engineering*, “Simulation Closes Gap Between Chip Design Optimization and Manufacturability”, 2025. Gregory Haley.

⁵⁸ *Design News*, “Managing Thermal Challenges in 3DIC Designs”, 2024. Lee Wang.

⁵⁹ *Semiconductor Engineering*, “Testing for Thermal Issues Becomes More Difficult”, 2024. Gregory Haley.

3. Measuring Buried Layers and Vertical Heat Paths

These modeling challenges are compounded by the fundamental difficulty of measuring buried layers, interfaces, and vertical heat paths. Thermal transport is inherently non-local. Any thermal measurement integrates the response of multiple layers, interfaces, and boundaries, not just the region of interest. As a result, two material stacks that appear nominally identical can produce meaningfully different measured thermal properties due to differences in surface condition, interfacial quality, growth method, or processing history.

Depth-sensitive thermal measurements do not directly image subsurface features. Instead, they probe a thermal length scale over which heat diffuses during the experiment, meaning the measured signal reflects an averaged response over a finite depth. Sensitivity to buried layers or interfaces is inferred through modeling rather than observed directly. As probing depth increases, the thermal response becomes less uniquely attributable to any single layer or interface, and uncertainty grows as multiple material properties influence the measurement simultaneously.

As lateral features shrink to only a few nanometers, thermal transport occurs over similarly small length scales, making direct measurement inherently difficult. Thermal diffusion lengths, optical spot sizes, and detector sensitivities are typically much larger than the native dimensions of individual fins, channels, or interfaces. Consequently, optically measured signals inevitably represent spatial and volumetric averages rather than localized nanoscale behavior. However, emerging metrology approaches are beginning to bridge this gap by integrating high-resolution probe architectures with quantitative thermoreflectance methods. These developments point toward a new class of nanoscale thermal measurements capable of resolving buried features and vertical heat paths with sub-micron precision while maintaining compatibility with realistic device structures.

4. From Measurement Challenge to Design Requirement

As device dimensions shrink and integration density increases, thermal characterization requirements have shifted accordingly. Modern systems demand measurements that are sensitive to nanoscale structure, repeatable across samples, non-contact, and robust against variability in surface and interface conditions. Because thermal signals inherently combine local and distributed heat transport, effective characterization must capture spatial variation while still providing averaged, physically meaningful properties suitable for use in design and modeling workflows.

In this context, spatial thermal conductivity mapping and quantitative evaluation of interface thermal resistance have become essential capabilities rather than specialized techniques. For thin films and multilayer structures with characteristic dimensions below 10 nm, even highly localized variations in interface quality or near-surface structure can

dominate overall thermal behavior. Measurement approaches must therefore resolve lateral non-uniformity while maintaining sensitivity to buried layers and interfaces that control vertical heat flow.

These requirements are most acute in advanced packaging and 3D stacked architectures, where cross-plane heat transport through multiple bonded interfaces, dense integration, and heterogeneous materials place vertical heat transport at the center of performance, reliability, and lifetime considerations. Accurate characterization of buried layers, interfaces, and cross-plane thermal transport is no longer a purely research-driven objective. It has become a design requirement, necessary for validating models, constraining simulation assumptions, and ensuring that thermal decisions made early in development remain valid as devices move toward fabrication and deployment.

Conclusion: From Thermal Constraint to Design Foundation

Semiconductor scaling is no longer limited primarily by lithography. It is constrained by heat. As dies stack, interfaces multiply, and power density rises, the dominant thermal resistances increasingly sit inside the package. Bond lines, dielectric stacks, micro-bumps, TSVs, TIM0 and TIM1 interfaces, and buried active layers now determine peak junction temperature and long-term reliability. External cooling still matters, but it is no longer where the most consequential thermal decisions are made.

The MAPT Roadmap 2.0 describes the situation directly:

“Advanced 3D packaging creates unique thermal challenges. The stacking of dies reduces the area for heat extraction which creates an additive-effective power density requiring careful inter-stack floorplan optimizations to mitigate power densities to the extent that the architecture can support. The stacking also increases the thermal resistance between dies in the stack and the coldplate or heatsink.”⁶⁰

Once heat is trapped within vertical stacks and buried interfaces, late-stage mitigation becomes expensive or impossible. Thermal is no longer something that can be tuned at the end. It must be defined at the beginning.

Concrete Examples: When the Bottleneck Is Buried

Consider a 3D stacked AI accelerator using HBM and a lidless TIM0 configuration. Electrically, vertical integration improves bandwidth and shortens interconnects. Thermally, however, the middle logic die sits farther from the cold plate while thinning limits lateral heat spreading. Bonding layers and underfills introduce additional low-conductivity barriers. As a result, the dominant thermal resistance may lie within the stack such as a dielectric layer below active devices, a Cu-Cu bonded interface with imperfect planarity, variations in silicon conductivity, or a TIM0 bond line with voiding or thickness variation. In this case, improving the heat sink may deliver limited gains, while reducing resistance at a buried interface can produce larger temperature reductions.

A similar challenge is emerging with backside power delivery networks (BSPDN). To route power from the backside, the silicon wafer must be significantly thinned, removing much of the original heat-spreading substrate. At the same time, power delivery structures introduce additional heat sources below the device layer. The remaining silicon, historically optimized for electrical performance, must now also serve as a primary thermal pathway, and that is a balance that current materials and processes do not yet fully resolve.

This illustrates the architectural shift described throughout this eBook: thermal bottlenecks are moving inward.

⁶⁰ *Semiconductor Research Corporation*. "Microelectronics and Advanced Packaging Technologies Roadmap Version 2.0". The MAPT Roadmap 2.0, 2025.

1. Reliability Is the Forcing Function

Thermal behavior is not only a performance constraint. It is a reliability accelerator. Peak junction temperatures, steep vertical gradients, and cyclic heating at interfaces drive:

- Electromigration in micro-bumps and interconnects
- Delamination and void growth at bonded interfaces
- TIM pump-out and dry-out
- Self-heating-induced performance drift
- Thermomechanical stress accumulation in stacked dies

Small local temperature differences can produce disproportionate reductions in lifetime. In advanced nodes, thermal and electrical effects are coupled tightly enough that temperature drift feeds back into electrical degradation and vice versa. This is why industry leaders increasingly emphasize thermal co-design.

“What is clear is that solving the industry’s heat problem will be an interdisciplinary effort. It’s unlikely that any one technology alone, whether that’s thermal-interface materials, transistors, system-control schemes, packaging, or coolers—will fix future chips’ thermal issues. We will need all of them.”⁶¹

- James Myers, Program Director for Thermal Cross Technology Optimization at imec

When reliability margins shrink and architectures become more complex, assumptions become risks. Measurement reduces these risks.

2. The Thermal-First Workflow

A thermal-first design workflow does not lock architecture prematurely. It makes thermal constraints visible early enough that design choices remain adjustable. A practical framework looks like this:

1. Define the Thermal Budget at Architecture Stage

- Maximum junction temperature
- Allowable hotspot delta
- Ambient and workload assumptions
- Cooling boundary conditions

Thermal requirements should sit alongside power and performance targets, not downstream from them.

⁶¹ *Commun. Association for Computing Machinery*, “Semiconductor Manufacturers Feel the Heat”, 2026. Samuel Greengard.

2. Map the Thermal Path and Identify Dominant Resistances

- Thin films and dielectric stacks
- Bonded interfaces and micro-bumps
- TIM0, TIM1 and TIM2 layers
- Vertical heat paths in 3D stacks
- Substrate and spreader layers

The goal is not to measure everything. It is to measure what dominates.

3. Measure Properties at Relevant Length Scales

- Thermal conductivity of process-specific thin films
- Interfacial thermal resistance across bonded layers
- Volumetric heat capacity where transient effects matter
- Spatial thermal conductivity or temperature non-uniformity

Property variability due to thickness, grain structure, porosity, bonding quality, and processing history must be treated as first-order.

4. Calibrate Models with Measured Inputs

Models remain indispensable. But they are only as reliable as their inputs. Measured properties reduce uncertainty bands and improve predictive confidence while the design space is still open.

5. Close the Loop

- Simulation identifies sensitivity
- Sensitivity defines what to measure next
- Measurement refines simulation
- Design iterates with reduced risk

The fastest path to reliable performance is not late-stage troubleshooting. It is an early, measurement-informed loop.

3. Measurement as a Foundation, Not a Conclusion

Historically, evaluating thermal properties of new materials has been slow, complex, and sample-intensive. That friction has delayed validation cycles and slowed material innovation.

That constraint is changing. Advances in automation, instrument integration, and workflow simplification now allow:

- Faster sample-to-data cycles
- Higher repeatability
- Reduced operator dependency
- Scalable property databases tied to process conditions

As design cycles accelerate and material diversity increases, streamlined and data-driven thermal metrology is becoming foundational to next-generation electronics development. The ecosystem is moving toward:

- Higher automation
- Near-line measurement capability
- Eventually, in-line or process-integrated thermal metrology tools
- Nanoscale sensitivity where geometry demands it

As architectures become more vertically integrated and interface-driven, this evolution is not optional. It is logical. Thermal metrology will increasingly sit closer to fabrication, not just in research labs. The goal is not only to characterize materials after they are chosen, but to inform which materials and architectures should be chosen in the first place.

4. The Central Shift

The FinFET era marked the point where traditional scaling encountered fundamental thermal limits. Today, heterogeneous integration, 3D stacking, backside power delivery, and AI-driven power density have made thermal constraints inseparable from system performance. Thermal limits are shaping architecture. In this environment:

- Variability is not noise.
- Interfaces are not secondary.
- Models cannot float on assumed properties.
- Reliability is inseparable from heat flow.

"The next decade of semiconductor progress will depend not only on new materials and new architectures, but on the ability to measure thermal behavior accurately, at the scales where it matters, fast enough to guide design decisions. Measurement is no longer the final validation step. It is the foundation. "

- John Gaskins, CEO and Co-Founder, Laser Thermal

ABOUT LASER THERMAL

Laser Thermal is a precision instrumentation company born from academic research and driven by the mission to innovate thermal property measurements for researchers, engineers, and product developers.

Our suite of thermal analysis tools NTM, TOPS, and FASTR addresses longstanding gaps in the market: accurate, high-throughput, and user-friendly measurements of thermal conductivity, thermal resistance, and volumetric heat capacity. These capabilities span an unprecedented range of material properties and length scales, from nanoscale thin films to bulk composites.

Laser Thermal was founded on decades of research and innovation originating in the ExSiTE Lab at the University of Virginia, grounded in a strong foundation of scientific excellence and engineering precision. Our team, led by experts in thermophysics, brings over a century of combined R&D experience, with focused expertise in microelectronics, optoelectronics, and thermal metrology.

With over ten issued or pending patents, our work has drawn recognition from DARPA, the U.S. Air Force, and the Department of Defense. From our facility in Charlottesville, Virginia, we carry out all research, development, and manufacturing, ensuring full control over quality and innovation. Today, Laser Thermal is trusted by leading organizations pushing the boundaries of material performance, where precise thermal insight is critical to success.

Laser Thermal empowers innovators and researchers across semiconductors, power electronics, and thermal management industries with accurate, efficient, and scalable thermal metrology. With NTM, TOPS, and FASTR, we've eliminated the technical and operational barriers of legacy techniques, providing precise measurements over a wide range of length scales in the lab and production settings.

**For those redefining thermal performance,
we provide the tools to measure what matters.**