

Thermal Transient Response of Glass-Core vs. FR-4 PCBs for High-Power AI Accelerator Modules

Joyce Tong¹

1. Electrical Engineering, University of Connecticut, Storrs, CT, United States.

Abstract

Glass-core PCBs offer improved thermal conductivity over conventional FR-4, but their substrate-level thermal behavior remains largely unexplored at the board scale and power levels used in modern AI accelerator hardware. This work uses COMSOL Multiphysics to compare heat spreading and transient response between glass and FR-4 sized and loaded to represent a high-power AI accelerator module (700W/die, SXM5-class board and die geometry, direct-to-chip liquid cold-plate cooling). Steady-state and transient simulations show that peak die temperature is nearly identical between substrates once realistic die-level cooling is applied, while lateral heat spreading and off-die temperature differ substantially, providing insight for glass-core PCB thermal design in the next-generation high-power electronic systems.

Keywords: Glass-core PCB, FR-4, COMSOL Multiphysics, Thermal Time Constant, AI Accelerator Thermal Management

I. Introduction

The rapid growth of AI computing has driven sharp increases in hardware power density, generating more heat than ever in a smaller space, making thermal management critical for performance and reliability. Glass-core PCBs have emerged as a promising alternative to FR-4, offering lower dielectric loss, closer coefficient-of-thermal-expansion (CTE) matching silicon, and higher thermal conductivity [7].

Most existing thermal studies focus on package-level cooling rather than substrate-level heat transfer, and standard design guidelines such as IPC-2152 [6] are based on FR-4, leaving glass substrate thermal behavior largely uncharacterized.

This work uses COMSOL Multiphysics to compare glass-core and FR-4 substrates sized to a representative high-power AI accelerator module, under identical dual-die loading and realistic die-level liquid cooling [9]. Steady-state and transient simulations quantify peak temperature, lateral heat spreading and thermal time constants, providing a substrate-level thermal characterization to inform glass-core PCB design for next generation high-power electronics.

II. Simulation Methodology

A. Model Geometry and Materials

The two-material comparison uses matched rectangular boards representing a high-power AI accelerator module footprint, $175 \times 120 \times 1.5$ mm

(SXM5-class) with two die-scale heater elements, $33 \times 25 \times 0.5$ mm each (GH100-class, near the photolithographic reticle limit) [5], embedded in the substrate surface.

Material properties are pulled from the COMSOL library for both substrates:

$$k_{glass} \approx 1.0 - 1.4 \frac{W}{m \cdot K}, \quad k_{FR4} \approx 0.3 \frac{W}{m \cdot K}$$

B. Physics and Boundary Conditions

Each heater is modeled as a fixed-power boundary heat source, outputting a fixed power,

$$Q_{htr} = 700 \text{ W}$$

representative of a datacenter GPU TDP (H100 SXM5-class) [4]. At full 700 W TDP, the SXM5 form factor requires direct liquid cooling rather than air cooling. Conventional rack-mount air cooling is limited to roughly 350-400 W per GPU and cannot dissipate full SXM5 TDP without severe throttling [1]-[3]. Die top surfaces are therefore modeled with a convective heat flux representative of a direct-to-chip liquid cold plate,

$$h_{conv} \approx 5000 \frac{W}{m^2 \cdot K}, \quad T_{amb} = 293.15 \text{ K}$$

While the remaining exposed substrate and die side-wall surfaces use a lower ambient coefficient,

$$h_{conv} \approx 10 \frac{W}{m^2 \cdot K}, \quad T_{amb} = 293.15 K$$

Reflecting that only the dies are in direct contact with the liquid cooling. Surface-to-ambient radiation (emissivity ≈ 0.9) is included. All non-exposed boundaries are treated as thermally insulated.

C. Mesh and Solver Configuration

Mesh is locally refined around both die footprints and coarser across the bulk substrate. Two solver studies are run per material

- **Stationary solver** – steady-state field, showing peak die temperature
- **Time-Dependent solver**, $t = 0 - 20s$ – transient rise at each thermistor location, used to extract the thermal time constant for each substrate

Results are checked against a hand estimate,

$$\Delta T \approx \frac{P}{hA}$$

which for the dual die loading condition ($P = 1400 W, A_{die} \approx 0.00165 m^2$), predicts $T_{peak} \approx 463 K$. This represents the upper bound that omits the lateral-conduction and ambient-convection paths present in the full model. The COMSOL solution ($T_{peak} \approx 339K$) lies within a physically realistic range for a liquid-cooled datacenter die, confirming the model behaves as intended once die-level cold-plate cooling is separated from board-level ambient convection.

In addition to peak steady-state temperature, the time-dependent solution at each thermistor location fits a first-order exponential response,

$$T(t) = T_{\infty} - (T_{\infty} - T_0)e^{\{-\frac{t}{\tau}\}}$$

via linear regression, using data through $t = 20s$; as the transient had not fully reached steady state within this window, these values represent approximate rather than fully converged time constants [8].

III. Results and Discussion

A. Steady-State Heat Distribution

The stationary solve shows the peak die temperature reaches $\approx 339K$ for both glass and FR-4 under identical 700 W per die loading




Figure 1. FR-4 Steady-State Field




Figure 2. Glass Steady-State Field

This equal peak temperature is expected: with the die-surface cold plate ($h_{conv} \approx 5000 \frac{W}{m^2 \cdot K}$) dominating heat rejection directly at the die, the substrate material has minimal influence on the die's own peak temperature – nearly all of the 700 W leaves through the cold plate before the substrate's conductivity becomes the limiting factor.

B. Transient Thermal Response

Time-dependent solves over $t = 0-20s$ show both materials rising quickly and approaching steady state




Figure 3. FR-4 full transient




Figure 4. Glass full transient

At $t = 20\text{s}$, monitored thermistor points reach $\approx 335.35\text{ K}$ for FR-4 (Fig 3) versus $\approx 332.06\text{ K}$ for glass (Fig 4). Glass runs measurably cooler at these off-die monitoring points despite the two substrates sharing the same die peak temperature. This is consistent with glass's higher conductivity spreading heat away from the die into the surrounding substrate, lowering the temperature at fixed nearby points, while FR-4's lower conductivity keeps heat concentrated closer to the die.

Fitting each thermistor curve to $T(t) = T_\infty - (T_\infty - T_0)e^{\{-\frac{t}{\tau}\}}$ via linear regression gives a characteristic thermal time constant for each substrate:

$$\tau_{FR4} \approx 18.5, \quad \tau_{glass} = 18.1$$

IV. Conclusion

The results reveal a more nuanced thermal picture than substrate conductivity alone would suggest. Once realistic die-level cold-plate cooling is applied, peak die temperature is nearly identical between glass and FR-4 ($\approx 339\text{ K}$ for both), since it is governed primarily by the cold plate rather than the substrate — the two materials' conductivities differ by roughly 4x, yet this has almost no effect on the die's own peak temperature.

The materials diverge, however, once heat moves away from the die. At fixed off-die monitoring points, glass consistently runs cooler than FR-4 ($\approx 332\text{ K}$ vs. $\approx 335\text{ K}$ at $t = 20\text{s}$), and the transient time constants extracted for each substrate ($\tau_{FR4} \approx 18.5$, $\tau_{glass} = 18.1$) further characterize how quickly each material distributes heat through the board. This indicates that substrate material choice has limited influence on managing peak die temperature under adequate direct-to-chip cooling but meaningfully affects how heat is distributed across the rest of the board — a consideration relevant to nearby components, sensors, or connectors that are not themselves directly cooled.

This simulation isolates substrate-level conduction and does not capture complications present in a physical board: copper trace routing would itself absorb and redistribute heat, achieving symmetric routing between the two test boards is difficult given real layout constraints, and copper pour/fill for power delivery and EMI would add another heat-spreading path not modeled here. Future work should incorporate routed traces and copper fill, extend the comparison across a swept die-separation range to characterize resulting crosstalk,

and validate against an instrumented physical test vehicle.

V. References

- [1] JetCool, "JetCool Unveils Breakthrough Energy-Saving Liquid Cooling for NVIDIA's H100 GPUs," 2024.
- [2] FormulaMod, "AI Server GPU Water Cooling: Why Liquid Cooling Matters for H100, H200," 2026.
- [3] Server-Parts.eu, "NVIDIA H100 GPU: SXM vs. PCIe, Liquid vs. Air Cooling," 2024.
- [4] NVIDIA, "NVIDIA H100 Tensor Core GPU Datasheet," NVIDIA Corporation, 2023.
- [5] Spheron, "NVIDIA H100 Specs: 80GB HBM3, 3,958 TFLOPS Datasheet," 2026.
- [6] IPC-2152, *Standard for Determining Current-Carrying Capacity in Printed Board Design*, IPC, Bannockburn, IL.
- [7] Y. Sato, "Glass Core Substrate," *IEEE EPS Advanced Substrates Technical Committee Report*, November 2025.
- [8] H. Rosten and R. Viswanath, "Thermal Modeling of the Pentium Processor Package," in *Electronic Components & Technology Conference*, pp. 421–428, 1994.
- [9] M. Pedram and S. Nazarian, "Thermal Modeling, Analysis, and Management in VLSI Circuits: Principles and Methods," *Proceedings of the IEEE*, vol. 94, no. 8, pp. 1487–1501, 2006.