
ATI WHITE PAPER

AWP-HS-01 · Revision 2.16 · June 2026

How to Design and Choose a Heat Sink for a TEC Cooling System

A Practical Engineering Guide for Heat Sink Design in TEC Cooling Systems

Thermal fundamentals, heat-sink sizing, airflow trade-offs, mounting, and validation for TEC cooling systems.

Analog Technologies, Inc.

San Jose, California · analogtechnologies.com · sales@analogtechnologies.com

Revision 2.16 · June 2026.

Table of contents

Executive summary.....	3
Quick answer — how to design and choose a heat sink for a TEC cooling system.....	4
Notation convention used in this guide.....	4
§1. Why heat-sink design matters in TEC applications.....	4
§2. What makes TEC heat-sink design different from CPU / MOSFET heat-sinking.....	5
§3. Fundamental thermal terminologies.....	6
§4. The electrical-thermal analogy — your most powerful mental model.....	11
§5. Key equations for TEC heat-sink design.....	14
§6. Six-step heat-sink design methodology.....	16
§7. Heat-sink construction and physics.....	17
§8. Heat-sink parameter studies.....	19
§9. Heat-sink types and when to use each.....	24
§10. Cooling-mode selection — passive, forced air, liquid.....	25
§11. Fan selection and the operating point.....	25
§12. Mounting interface and thermal interface material (TIM).....	27
§13. Validation envelope and pass criteria.....	28
§14. Common failure modes after correct component selection.....	29
§15. Application examples (illustrative).....	29
§16. Frequently asked questions.....	30
§17. ATI product bridge and engineering support.....	32
References.....	33

Executive summary

Heat-sink design for a thermoelectric (TEC / Peltier) cooling application is different from heat-sinking a CPU, MOSFET, or LED. A TEC does not destroy heat — it moves heat from the cold side to the hot side. The heat sink at the hot side must reject the sum $Q_{hot} = Q_c + P_{electrical}$, which for typical OEM operating points is two to five times the cold-side load (the multiple depends on COP, ΔT , I/I_{max} , and the selected TEC operating point).

In plain language: the heat the heat sink must reject at the TEC hot side equals the cold-side heat the TEC pumps plus the electrical power into the TEC.

An undersized heat sink in a TEC system does not gradually degrade the way a CPU heat sink does. As T_{hot} climbs, the TEC's Q_c capacity falls below the required load, the controller demands more current to compensate, $P_{electrical}$ rises, T_{hot} rises further. Failure tends to appear as controller saturation at worst-case conditions — long pull-down time, missed setpoint, current-limit saturation — not gentle warming.

This guide is written as a practical engineering reference. It begins with fundamental thermal terminologies and intuitive analogies; develops the formal electrical-thermal analogy that lets engineers reuse circuit-theory intuition for heat-flow problems; derives the key sizing equations; and presents detailed parameter studies of how heat-sink resistance depends on airflow, envelope volume, fin count, fin thickness, and fin spacing. Six-step methodology, fan selection, mounting interface, validation envelope, common failure modes, worked application examples, FAQ, and ATI product routing follow.

Read this paper if...

You need to design, specify, or choose the heat sink for a TEC cooling system: deriving Q_{hot} , sizing $R_{\theta SA}$, selecting cooling mode, sizing fans, mounting interface, and validation. For TEC controller selection, see companion paper AWP-TECC-02. For the full system-design view that integrates the heat sink with TEC module sizing, sensor placement, mounting, condensation control, and validation as a connected stack, see AWP-TECC-03. For TEC module fundamentals, see AWP-TECM-01. Note: ATI does not supply heat sinks, fans, or TIM; those are sourced from established thermal-management vendors.

Critical caveats

- 1) Operating-point $T_{hot,max}$, not absolute maximum.** Use $T_{hot,max}$ from the TEC datasheet performance curves at your actual operating point — not the TEC absolute maximum rating. The two values can differ by 20–40 °C; choosing the wrong one is a common and serious heat-sink-sizing error in TEC applications.
- 2) Apply margin to $R_{\theta SA}$, not to power.** A 2× margin on $R_{\theta SA}$ means half the thermal resistance at the same ΔT_{HS} (twice the cooling capacity), not twice the power.
- 3) Fan operating airflow, not free-flow CFM.** A heat sink presents static pressure to the fan; the operating point is where the fan curve meets the system-resistance curve, typically 30–60 % of free-flow CFM.
- 4) Component-level engineering guidance.** Numerical examples are illustrative patterns, not ATI specifications or product guarantees. Full enclosure, EMC, safety, service-life, and regulatory qualification for regulated end products remains the OEM's responsibility.
- 5) Sub-ambient operation requires the worst-case operating point.** When held below ambient, the TEC pumps a parasitic heat leak in addition to the design load. At the actual operating point, COP drops, Q_c rises, $P_{electrical}$ rises, and Q_{hot} is much higher than a room-temperature calculation suggests. Size for the worst-case operating point and insulate aggressively. See §2.4 and §5.4.

Quick answer — how to design and choose a heat sink for a TEC cooling system

1. Calculate the hot-side heat load: $Q_{hot} = Q_c + P_{electrical}$, where Q_c is the cold-side heat the TEC must pump and $P_{electrical} = V_{TEC} \cdot I_{TEC}$ is the electrical power into the TEC.
2. Define the worst-case ambient temperature at the heat-sink inlet, $T_{ambient,max}$.
3. Determine the allowable hot-side temperature $T_{hot,max}$ from the TEC operating-point performance curves — not the absolute maximum rating.
4. Calculate the required heat-sink-to-ambient resistance: $R_{\theta SA} \leq (T_{hot,max} - T_{ambient,max}) / Q_{hot} - R_{TIM}$.
5. Select cooling mode (passive / forced-air / liquid) based on Q_{hot} , envelope, noise, reliability, and environment.
6. Apply an environment-appropriate margin (typically 1.5–3× for OEM), then validate the assembled stack under worst-case combined conditions. The rest of this guide expands each step with terminology, equations, parameter studies, and worked examples.

Notation convention used in this guide

Symbol	Meaning	Units · notes
Q_c	Cold-side heat pumped by TEC at steady state	W · equals Q_{load} at steady state; during ramps, evaluate at worst-case operating point
$P_{electrical}$	Electrical power into TEC = $V_{TEC} \cdot I_{TEC}$	W · also written P_{elec} ; same quantity
Q_{hot}	Total heat the heat sink must reject at the TEC hot side	W · $Q_{hot} = Q_c + P_{electrical}$
T_{hot}, T_{cold}	TEC ceramic hot-side / cold-side temperatures	°C · operating-point values from datasheet curves
$T_{ambient,max}$	Worst-case air temperature at heat-sink inlet	°C · also written $T_{A,max}$
$R_{\theta SA}$	Heat-sink-to-ambient thermal resistance	K/W
$R_{TIM} (= R_{\theta CS})$	TEC-ceramic-to-heat-sink interface resistance (TIM joint)	K/W · R_{TIM} is the practical name; $R_{\theta CS}$ is the formal case-to-sink resistance

§1. Why heat-sink design matters in TEC applications

In a conventional heat-sinking problem (CPU, MOSFET, LED, power resistor), heat enters the heat sink from one face and exits to the ambient air. The component generates a roughly fixed power dissipation, the heat sink rejects it, and the component temperature rises by $Q \times R\theta JA$. An undersized heat sink makes the component run hotter, but the system still operates.

A TEC cooling system has two coupled differences:

- **The heat sink rejects more power than the load.** A TEC moves Q_c (cold-side) by adding electrical power $P_{electrical}$ to the hot side. The heat sink sees $Q_{hot} = Q_c + P_{electrical}$, which for $COP \approx 0.3-0.5$ is 2–3× the load, and at deep ΔT can reach 5× or more.
- **T_{hot} is not a free parameter — it is a TEC datasheet constraint.** The TEC's ability to pump Q_c at the required ΔT degrades as T_{hot} rises. Past a datasheet-defined limit, Q_c falls below the required load,

and the loop cannot reach setpoint. The COP also drops with rising T_{hot} , so $P_{electrical}$ rises, which raises T_{hot} further — positive feedback ending in loop saturation rather than gentle degradation.

The practical consequence: a heat sink sized for a CPU-style design margin (10–30 % over rated dissipation at typical ambient) can easily be 1.5–3× too small (worse at deep ΔT), depending on COP and operating point for the same load on a TEC. Heat-sink design has to start from Q_{hot} at the actual worst-case operating point, not from the cold-side load and not from steady-state lab conditions.

How TEC heat-sink undersizing fails

Undersized at typical ambient → product works on the lab bench → product reaches worst-case ambient (hot summer, enclosure restriction, fan aging, dust accumulation) → T_{hot} climbs → Q_c capacity derates sharply → controller increases current to compensate → $P_{electrical}$ rises → T_{hot} rises further → loop saturates and product cannot reach setpoint. Failure tends to be abrupt at worst-case conditions, not gradual at typical conditions.

§2. What makes TEC heat-sink design different from CPU / MOSFET heat-sinking

2.1 Coefficient of performance (COP) and the Q_{hot} multiplier

A TEC's coefficient of performance, defined as

$$COP = \frac{Q_c}{P_{electrical}}$$

In plain language: COP is the cold-side heat the TEC pumps divided by the electrical power into the TEC.

is rarely above 1 in OEM applications and is typically 0.3–0.6 at moderate ΔT . $P_{electrical} = V_{TEC} \cdot I_{TEC}$ is therefore 1.5–3× the load being pumped, and Q_{hot} is 2–4× the load. At deep ΔT (40–50 °C and above), COP can fall below 0.2, and Q_{hot} can exceed 5× the load.

2.2 Hot-side temperature is a TEC constraint, not a heat-sink constraint

Most heat sinks can operate at hot-side temperatures well above the TEC operating-point limit without damage. Standard Bi_2Te_3 TEC modules often cannot maintain the required pumping performance or reliability at such temperatures — they derate substantially as hot-side temperature rises, and the allowable operating-point T_{hot} is often far below the absolute maximum rating. Even more important: the TEC datasheet defines a hot-side limit beyond which Q_c may fall below the required load. Many TEC modules derate substantially as T_{hot} rises, and the heat sink must hold T_{hot} below the operating-point limit — not the absolute maximum rating — at the worst-case ambient the product will see in the field.

2.3 Worst case is hot ambient + deep ΔT + peak load + low V_{ps} — all at once

A heat sink sized for typical ambient (lab 20–25 °C) is undersized for the field worst case: hot summer ambient (35–45 °C in enclosed equipment), deepest ΔT required, peak load (pull-down transient, ramp), low supply voltage (battery, brownout) that forces the controller into longer transient operation. All four conditions can occur together in a real product.

2.4 Sub-ambient operation amplifies the problem

A TEC that holds the load 20 °C below ambient is doing two things: pumping the load AND fighting passive heat leak from the ambient back into the cold side. The effective Q_c that must be pumped grows, $P_{\text{electrical}}$ grows with it, and the heat-sink Q_{hot} grows faster than a naive calculation suggests. For sub-ambient applications, increase the heat-sink margin and insulate the cold-side enclosure.

§3. Fundamental thermal terminologies

Before diving into TEC heat-sink design, this section establishes the basic thermal vocabulary. Every TEC designer needs an intuitive working sense of seven quantities — heat flow rate, temperature difference, thermal resistance, thermal conductance, thermal conductivity, convection coefficient, heat capacity — plus the thermal time constant that ties them together. Each is defined below with units, meaning, and an analogy that makes it feel concrete.

3.1 Heat flow rate, Q

Definition: the rate at which thermal energy crosses a surface or moves along a path.

Symbol: Q **Units:** watts [W] (= joules per second, J/s)

Watts of heat are the same physical quantity as watts of electrical power. A 10 W resistor in steady state dissipates 10 W of electrical power and produces 10 W of heat that must be conducted, convected, or radiated away. In a TEC system, the heat flow rate at the hot side is Q_{hot} ; at the cold side it is Q_c ; through any segment of the heat-sink path it has its own value.

Electrical analogy: heat flow rate Q is the thermal counterpart of electrical current I — amperes of charge per second through a wire. Both are rate quantities, not static stores.

3.2 Temperature difference, ΔT

Definition: the difference in temperature between two points or surfaces — the driving force for heat flow.

Symbol: ΔT **Units:** kelvin [K] or degrees Celsius [°C], which are interchangeable for differences.

Heat flows naturally from hot to cold. The larger the temperature difference, the larger the heat flow through the same path. In a TEC heat-sink stack, the driving ΔT is $T_{\text{hot}} - T_{\text{ambient}}$, and the heat sink converts this ΔT into a rejected Q_{hot} .

Electrical analogy: ΔT is the thermal counterpart of voltage difference ΔV across a resistor — the larger the potential difference, the larger the current (or heat) that flows through the same resistance.

3.3 Thermal resistance, R_{θ}

Definition: the ΔT that develops across a thermal path per unit of heat flow rate through it.

Symbol: R_{θ} (the Greek letter theta marks it as *thermal* resistance, to distinguish it from electrical R) **Units:** K/W or °C/W (same number)

$$R_{\theta} = \frac{\Delta T}{Q} \quad [\text{K/W or } ^{\circ}\text{C/W}]$$

$R\theta$ is the single most important number for heat-sink design. A heat sink with $R\theta_{SA} = 1.0 \text{ K/W}$ rejects 1 W with a 1 °C rise of the heat-sink above ambient; the same heat sink rejecting 30 W will see a 30 °C rise. Smaller $R\theta_{SA}$ is better. $R\theta_{SA}$ depends strongly on airflow, geometry, and material — see the parameter studies in §8.

Thermal resistance — electrical analogy




Figure 1 — Electrical analogy for thermal resistance. The two circuits obey the same algebra: a voltage source V drives current I through a resistor R (Ohm's law $V = I \cdot R$); a temperature difference ΔT drives heat flow Q through a thermal resistance $R\theta$ (Fourier's law in lumped form, $\Delta T = Q \cdot R\theta$). Variables: V = voltage [V]; I = electrical current [A]; R = electrical resistance [Ω]; T_{hot} , T_{amb} = hot and ambient temperatures [°C or K]; Q = heat flow rate [W]; $R\theta$ = thermal resistance [K/W].

3.4 Thermal conductance, $G\theta$

Definition: the reciprocal of thermal resistance — how much heat the path passes per unit ΔT .

Symbol: $G\theta$ (sometimes G , U , or $1/R\theta$) **Units:** W/K

$$G_{\theta} = \frac{1}{R_{\theta}} = \frac{Q}{\Delta T} \quad [\text{W/K}]$$

Conductance is resistance flipped around. Heat-sink datasheets publish $R\theta_{SA}$, but when several paths run in parallel (e.g. the heat sink rejects some heat by convection and a smaller amount by radiation), it is cleaner to add conductances of parallel paths than to invert each resistance separately.

Thermal conductance – parallel-paths analogy

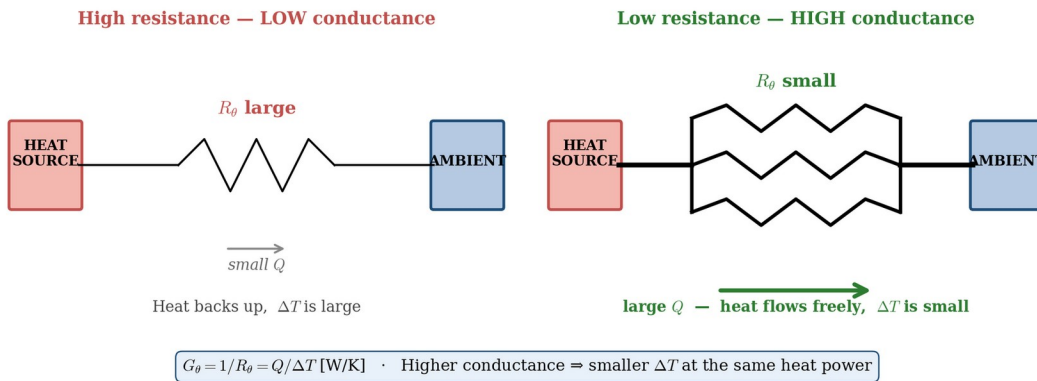


Figure 2 — Parallel-paths electrical analogy for thermal conductance. Left: a single high-resistance path produces a small heat flow Q and a large ΔT (heat “backs up”). Right: multiple resistors in parallel give an equivalent low resistance — high conductance $G\theta$ — letting large Q through with small ΔT . Same algebra as paralleled electrical resistors. Variables: $R\theta$ = thermal resistance [K/W]; $G\theta$ = thermal conductance = $1/R\theta$ [W/K]; Q = heat flow rate [W]; ΔT = temperature difference [K or °C].

3.5 Thermal conductivity, k

Definition: a property of the material itself — how much heat a 1-meter-thick slab conducts per unit area per kelvin of ΔT .

Symbol: k **Units:** W/(m·K)

Thermal conductivity is a material property; thermal resistance is a geometry-dependent quantity. For a slab of material with thickness L and cross-sectional area A :

$$R_\theta = \frac{L}{k A} \quad [\text{K/W}]$$

The underlying physics is Fourier's law of heat conduction:

$$Q = - k A \frac{dT}{dx} \quad [\text{W}]$$

Approximate thermal conductivities of common heat-sink materials at room temperature:

Material	k [W/(m·K)]	Notes
Pure copper	~400	Highest practical k; expensive, heavy, hard to extrude
Pure aluminum (1100)	~220	Standard extrusion alloy; lightweight; common HS material

Material	k [W/(m·K)]	Notes
Aluminum 6063 (alloy)	~200	Common extrusion alloy — good extrudability
Brass	~110	Rarely used for HS; common for thermal busbars
Steel (carbon)	~50	Poor HS material; structural use only
Stainless steel 304	~16	Very poor HS material; avoid in thermal path
Thermal grease (good)	5–8	Typical high-performance TIM
Air (still)	~0.026	Why heat sinks need air movement, not just air contact

Intuitive analogy: thermal conductivity is like the electrical conductivity of a wire material. Copper wire conducts electricity better than steel; copper bus also conducts heat better than steel bus.

3.6 Convection heat transfer coefficient, h

Definition: rate of heat transfer per unit surface area per kelvin of ΔT between a solid surface and a moving fluid (usually air).

Symbol: h **Units:** W/(m²·K)

$$Q = h A \Delta T \quad [W]$$

In plain language: convective heat flow equals the convection coefficient times the heat-transfer area times the surface-to-fluid temperature difference.

where Q = convective heat flow [W]; h = convection coefficient [W/(m²·K)]; A = effective heat-transfer surface area in contact with the moving fluid [m²]; ΔT = temperature difference between the solid surface and the fluid [K].

Approximate convection coefficients:

Mode	h [W/(m ² ·K)]	Notes
Natural convection in air	5–25	No fan; buoyancy-driven flow only
Forced air, modest velocity	50–150	Typical fan-cooled OEM heat sink
Forced air, high velocity	150–500	Densely finned + high CFM (server-CPU regime)
Forced water (cold plate)	1000–10000+	Liquid cooling regime
Boiling / two-phase	>10000	Heat-pipe interior, vapor chamber, phase change

The convection coefficient is the bottleneck in most OEM heat-sink designs. Aluminum is a very good conductor ($k \approx 200$ W/(m·K)); air over a finned surface is a poor conductor [1, 8]. Fins exist to multiply the surface area in contact with the air — trading material conduction (cheap, fast) for air convection (expensive in area, slow per area).

Why h depends on airflow — the boundary layer. The air molecules immediately touching the fin surface do not move; they stick because of viscosity (the “no-slip condition” from fluid mechanics). This creates a thin stationary film of air right against the metal — the thermal boundary layer. Heat has to cross this stagnant

layer by pure conduction ($k_{\text{air}} \approx 0.024 \text{ W}/(\text{m}\cdot\text{K})$) before the moving bulk flow can carry it away. The boundary layer behaves as a microscopic insulating blanket on every fin face: the thicker it is, the smaller h gets. Faster airflow, turbulent mixing, and louvered/pin/dimpled fin geometries all work the same way — they strip or thin this stagnant film, raising h and lowering $R\theta SA$. Natural convection removes the boundary layer much less aggressively than forced airflow, which is why its h is in the 5–25 $\text{W}/(\text{m}^2\cdot\text{K})$ range; forced air at modest velocity reaches 50–150, and dense-fin server-class designs get to 150–500.

3.7 Heat capacity, thermal mass, and time constant

Definition: the energy required to raise an object's temperature by one kelvin. Determines how fast or slowly the system warms up or cools down.

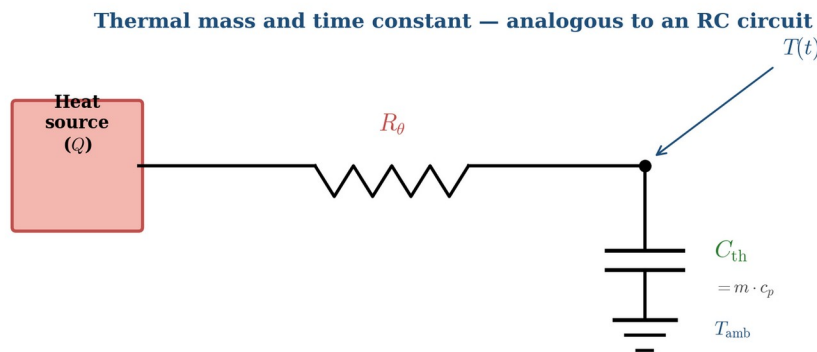
Symbols: Specific heat capacity c_p (material property, $\text{J}/(\text{kg}\cdot\text{K})$); thermal mass $C_{\text{th}} = m \cdot c_p$ (object property, J/K); thermal time constant $\tau = R\theta \cdot C_{\text{th}}$ (s).

$$C_{\text{th}} = m c_p \quad [\text{J}/\text{K}]$$

$$\tau = R_{\theta} C_{\text{th}} \quad [\text{s}]$$

The time constant τ is the time it takes the object to traverse 63.2% of the way from its initial temperature to its final temperature, under a step change in either heat input or ambient. A heavy aluminum heat sink has a large C_{th} , so it takes minutes for its temperature to settle; a small thermistor bead has a tiny C_{th} and settles in seconds. The same R-C charging curve from electrical circuits applies:

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



$$\tau = R_{\theta} \cdot C_{\text{th}} \quad \Rightarrow \quad T(t) = T_{\infty} - (T_{\infty} - T_0) e^{-t/\tau}$$

Figure 3 — Thermal RC network. A heat source Q feeds a thermal resistance $R\theta$ in series with a thermal capacitance C_{th} to ambient T_{amb} . The thermal time constant $\tau = R\theta \cdot C_{\text{th}}$ is the analog of electrical RC, governing how fast $T(t)$ reaches steady state after a load change. Variables: Q = heat flow into the network [W]; $R\theta$ = thermal resistance from heat source to ambient [K/W]; C_{th} = thermal capacitance = $m \cdot c_p$ [J/K]; $T(t)$ = transient node temperature [°C]; T_{amb} = ambient temperature [°C]; τ = time constant [s]; T_{∞} = final (steady-state) temperature; T_0 = initial temperature.

3.8 Radiation — emissivity and the Stefan-Boltzmann term

All surfaces radiate heat to their surroundings. For most OEM TEC applications at moderate temperatures (T_{hot} below ~ 80 °C), the radiation contribution from a black-anodized heat sink to ambient air is, as a typical illustrative range, around 5–15 % of total Q_{hot} dissipation, and convection dominates. Painting or anodizing the heat sink black is a cheap improvement for natural-convection designs.

$$Q_{\text{rad}} = \varepsilon \sigma A (T_s^4 - T_a^4) \quad [\text{W}]$$

Here ε is the surface emissivity (bare polished aluminum ≈ 0.05 ; anodized black ≈ 0.85 [1, 5]), $\sigma = 5.67 \times 10^{-8}$ W/($\text{m}^2 \cdot \text{K}^4$) is the Stefan-Boltzmann constant, A is the radiating surface area, and T_s and T_a are surface and ambient temperatures in absolute kelvin. The strong T^4 dependence is why radiation contributes more at higher hot-side temperatures.

§4. The electrical-thermal analogy — your most powerful mental model

Every electrical engineer already knows the rules for resistors, voltages, currents, capacitors, time constants, and Kirchhoff's laws. The good news is that heat-flow problems obey identical mathematics. Once the correspondence is set up, the standard results from circuit theory — series sums, parallel reciprocals, time constants, transient response — carry over directly to passive, lumped, linearized thermal paths. Radiation, nonlinear convection, active TEC control, and phase-change behavior require richer models beyond a simple RC network. This section makes the correspondence explicit.

4.1 Variable correspondence

Electrical quantity	Symbol · unit	Thermal quantity	Symbol · unit
Voltage (potential)	V · volts	Temperature (difference)	T (ΔT) · K or °C
Current	I · amperes	Heat flow rate	Q · watts
Resistance	R · ohms	Thermal resistance	$R\theta$ · K/W
Conductance	G · siemens (S = $1/\Omega$)	Thermal conductance	$G\theta$ · W/K
Capacitance	C · farads (C/V)	Heat capacity	C_{th} · J/K
Charge	Q_e · coulombs	Thermal energy	E_{th} · joules
Ohm's law	$V = I \cdot R$	Fourier's law (lumped)	$\Delta T = Q \cdot R\theta$
RC time constant	$\tau = R \cdot C$	Thermal time constant	$\tau = R\theta \cdot C_{\text{th}}$

One asymmetry: thermal systems have no native analog of an electrical inductor (no thermal mass-flow inertia at OEM scale) and no equivalent of negative resistance. For passive, lumped, approximately linear heat-sink problems the correspondence is very useful: series resistance, parallel paths, and RC time constants carry over directly. More advanced active-control, radiation, heat-pipe, or fluid-loop systems require richer models beyond a simple RC network.

4.2 Fourier's law — the thermal analog of Ohm's law

As introduced in §3.5, Fourier's law of heat conduction states that the rate of heat transfer through a solid is proportional to the cross-sectional area and the negative temperature gradient: $Q = -kA \cdot (dT/dx)$. The minus sign captures the physical fact that heat flows from hot to cold — down the gradient. For steady-state, one-dimensional conduction through a slab of thickness L and uniform conductivity k , integrating Fourier's law over the length gives the lumped form $\Delta T = Q \cdot R_\theta$, where $R_\theta = L/(kA)$ is the thermal resistance of the slab. This has the same algebraic structure as Ohm's law for passive, lumped, approximately-linear heat-flow paths:

$$V = I R$$

$$\Delta T = Q \cdot R_\theta$$

Read this as: the temperature drop across a thermal resistance equals the heat flow times the resistance — the single most useful equation in heat-sink design. Under the passive, lumped, approximately-linear assumptions above the correspondence is direct: current $I \leftrightarrow$ heat flow Q ; voltage drop $\Delta V \leftrightarrow$ temperature difference ΔT ; electrical resistance $R \leftrightarrow$ thermal resistance R_θ . Because the two laws share the same mathematical structure, complex thermal pathways — TEC \rightarrow TIM \rightarrow heat-sink base \rightarrow fin \rightarrow air — can be analyzed using resistor-network methods under the passive, lumped, approximately-linear assumptions above. Series resistances add directly; parallel paths combine using reciprocal-sum rules (see §4.3).

Electrical \leftrightarrow Thermal analogy

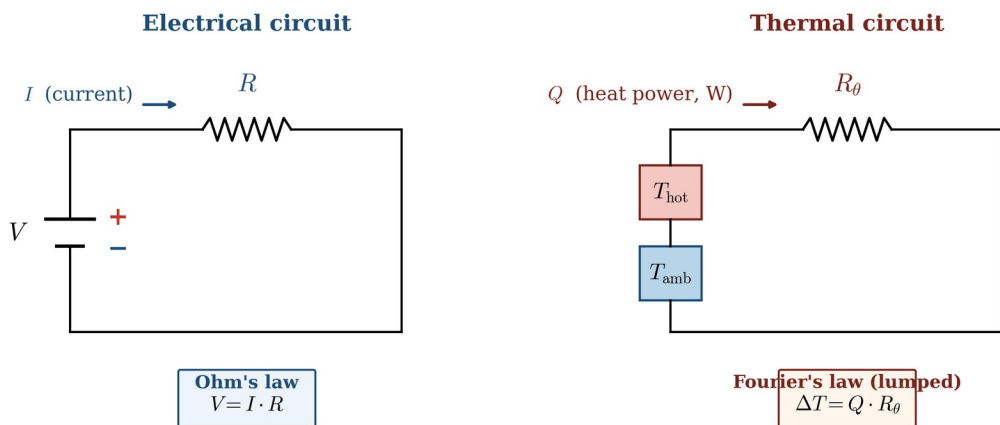


Figure 4 — Side-by-side electrical and thermal circuits. The two networks obey identical equations. Anything you know about Ohm's law applies to heat flow with one substitution: $I \rightarrow Q$, $V \rightarrow \Delta T$, $R \rightarrow R_\theta$. Variables: V = voltage [V]; I = current [A]; R = electrical resistance [Ω]; T_{hot} , T_{amb} = hot-side and ambient temperatures [$^\circ\text{C}$ or K]; Q = heat flow rate [W]; R_θ = thermal resistance [K/W].

4.3 Series and parallel thermal resistances

In a series chain of thermal paths, the same heat flow Q passes through each resistance, and the temperature drops add. In a parallel set of thermal paths, the temperature drop is the same across each branch, and conductances add (heat flows divide).

Series sum:

$$R_{\theta, \text{tot}} = R_{\theta 1} + R_{\theta 2} + R_{\theta 3} + \dots$$

Parallel sum:

$$\frac{1}{R_{\theta, \text{tot}}} = \frac{1}{R_{\theta 1}} + \frac{1}{R_{\theta 2}} + \dots$$




Figure 5 — Series and parallel thermal-resistance networks. The series-sum and parallel-reciprocal rules are identical to their electrical counterparts. Variables: $R_{\theta 1}$, $R_{\theta 2}$, $R_{\theta 3}$ = individual thermal resistances [K/W]; $R_{\theta, \text{tot}}$ = total equivalent thermal resistance [K/W]; T_1 , T_2 = endpoint temperatures.

4.4 The TEC stack as a series thermal circuit

From the TEC hot side to ambient air, heat flows through a small number of well-defined thermal resistances in series. Once you draw the network, the heat-sink design problem reduces to picking each R so the total temperature drop is acceptable for the heat flow you need to reject.

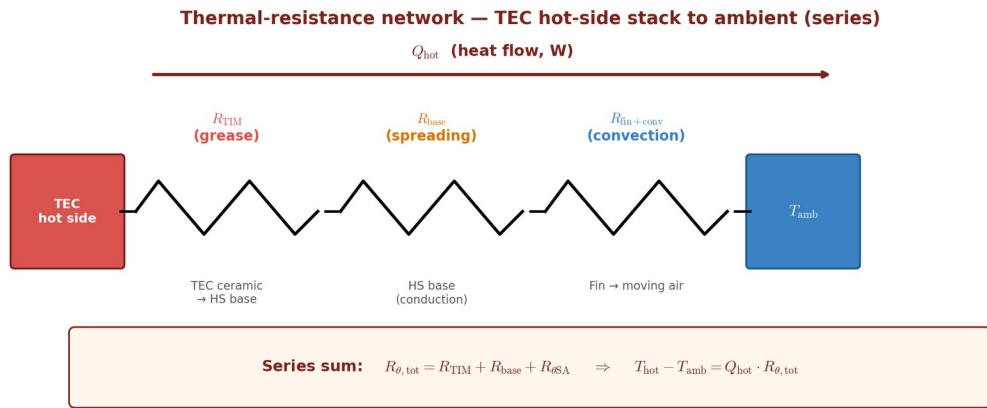


Figure 6 — Thermal-resistance network of the TEC hot-side stack: R_{TIM} (the grease layer between TEC ceramic and heat-sink base), R_{base} (conduction through the heat-sink base, sometimes called spreading resistance for high-power-density loads), and $R_{fin+conv}$ (the fin-to-air convection resistance, which is the largest term for forced-air OEM designs). All three add in series. Variables: Q_{hot} = heat flow at TEC hot side [W]; R_{TIM} = TIM resistance [K/W]; R_{base} = heat-sink base/spreading resistance [K/W]; $R_{fin+conv}$ = fin-to-air convection resistance [K/W]; $R_{\theta,tot}$ = series sum [K/W]; T_{hot} , T_{amb} = TEC hot-side and ambient temperatures [°C].

Typical magnitudes for a mid-power OEM TEC heat sink at moderate airflow: $R_{TIM} \approx 0.05\text{--}0.10$ K/W, $R_{base} \approx 0.05$ K/W (small for compact loads), and $R_{fin+conv}$ typically dominating at 0.5–1.5 K/W. The sum is the total R from TEC hot ceramic to ambient. With Q_{hot} known:

$$T_{hot} = T_{amb} + Q_{hot} (R_{TIM} + R_{\theta SA})$$

4.5 The RC time constant in the thermal world

When the load on a thermal system changes — a TEC turning on, a setpoint changing, an ambient transient — the system does not respond instantly. It responds with a single-pole exponential governed by $\tau = R_{\theta} \cdot C_{th}$, with the same first-order mathematical form as an electrical RC charging or discharging. This is why heat-sink validation needs at least one steady-state period ($\approx 5\tau$ is the rule of thumb) at each operating condition; reading T_{hot} 30 seconds into a step change is meaningless if τ is 5 minutes.

§5. Key equations for TEC heat-sink design

5.1 Variable definitions

- Q_{load} = cold-side thermal load (W) that must be removed from the controlled object — the application input.
- Q_c = heat the TEC actually pumps at the selected operating point (W). At steady state $Q_c = Q_{load}$.
- I_{TEC} , V_{TEC} = TEC operating current and voltage, read from the TEC datasheet performance curves at the selected operating point.
- $P_{electrical}$ = electrical power into the TEC = $V_{TEC} \cdot I_{TEC}$ (W).
- Q_{hot} = heat rejected at the TEC hot side, which the heat sink must dissipate = $Q_c + P_{electrical}$.
- T_{hot} , T_{cold} = TEC hot-side and cold-side ceramic temperatures (°C).
- $T_{hot,max}$ = maximum permissible T_{hot} at the operating point, from the TEC datasheet performance curves — NOT the absolute maximum rating.
- $\Delta T_{TEC} = T_{hot} - T_{cold}$.
- $T_{ambient,max}$ = maximum air temperature at the heat-sink fins under worst-case service conditions.
- $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance (K/W).
- R_{TIM} = thermal-interface-material resistance between TEC hot ceramic and heat-sink base (K/W).
- $\Delta T_{HS} = Q_{hot} \cdot R_{\theta SA}$: heat-sink temperature rise above ambient — a consequence of $R_{\theta SA}$ selection, not a design target.

5.2 Primary equations

$$P_{\text{electrical}} = V_{\text{TEC}} \cdot I_{\text{TEC}}$$

$$Q_{\text{hot}} = Q_c + P_{\text{electrical}}$$

$$\Delta T_{\text{TEC}} = T_{\text{hot}} - T_{\text{cold}}$$

The heat-sink-sizing inequality (the most-cited equation in this guide):

$$R_{\theta\text{SA}} \leq \frac{T_{\text{hot,max}} - T_{\text{amb,max}}}{Q_{\text{hot}}} - R_{\text{TIM}}$$

In plain language: electrical power into the TEC is the product of its operating voltage and current. The hot-side heat the heat sink must reject equals the cold-side heat the TEC pumps plus that electrical power. The TEC temperature difference is the hot-side temperature minus the cold-side temperature. The required heat-sink-to-ambient thermal resistance is at most the allowable hot-to-ambient temperature drop divided by the hot-side heat load, minus the interface-material thermal resistance.

Read this as: the heat sink, plus the TIM joint between it and the TEC, must hold T_{hot} below $T_{\text{hot,max}}$ at the worst-case ambient while rejecting the full Q_{hot} . R_{TIM} is typically 0.05–0.2 K/W for a well-applied 25–100 μm grease layer on a flat lapped surface; it should not be neglected when Q_{hot} is large or $R_{\theta\text{SA}}$ is small.

And the steady-state heat balance through the TEC → TIM → heat-sink → ambient stack:

$$T_{\text{hot}} = T_{\text{amb}} + Q_{\text{hot}} (R_{\text{TIM}} + R_{\theta\text{SA}})$$

5.3 Margin calculation — explicit three-step sequence

The $R_{\theta\text{SA}}$ inequality above is the marginal case. Production designs should oversize the heat sink by a margin appropriate to the operating environment. The explicit sequence is:

1. **Compute total allowed hot-side thermal resistance:** $R_{\text{total}} \leq (T_{\text{hot,max}} - T_{\text{ambient,max}}) / Q_{\text{hot}}$.
2. **Subtract the interface resistance:** $R_{\theta\text{SA,calculated}} \leq R_{\text{total}} - R_{\text{TIM}}$.
3. **Apply environment margin to the heat-sink resistance:** $R_{\theta\text{SA,selected}} = R_{\theta\text{SA,calculated}} / \text{margin}_{\text{factor}}$.

$$R_{\theta\text{SA, selected}} = \frac{R_{\theta\text{SA, calculated}}}{\text{margin factor}}$$

Operating environment	Margin factor	Rationale
Controlled lab or production floor	1.5× – 2×	Stable ambient, clean air, controlled airflow

Operating environment	Margin factor	Rationale
General OEM (office, instrument enclosure)	2× – 3×	Enclosure restriction, normal dust, occasional ambient excursions
Outdoor / kiosk / harsh environment	3× – 5×	Solar gain, dust loading, fan aging, wide ambient swing

Margin convention: margin is applied to thermal resistance, not to power. A 2× margin on R_{θSA} means the heat sink has half the calculated thermal resistance (twice the cooling capacity), not that it can handle twice the power. Heat-sink capacity scales as 1/R_{θSA} at a given ΔT_{HS}, not linearly in power.

5.4 Transient loads — the dT/dt term

For thermal-cycling loads (PCR / qPCR blocks, laser-diode pull-down, semiconductor reliability test ovens), the heat sink must reject the peak Q_{hot} during the ramp, not the steady-state average. The transient heat addition is:

$$Q_{\text{transient}} = m c_p \frac{dT}{dt}$$

In plain language: the transient heat addition during a ramp equals the mass of the controlled load times its specific heat capacity times the ramp rate.

where m is the mass of the controlled load and c_p is its specific heat capacity. During a ramp the cold-side heat demand is Q_{c,ramp} = Q_{load,steady} + Q_{transient}, and the heat sink must reject Q_{hot,ramp} = Q_{c,ramp} + P_{electrical}, with P_{electrical} read at the actual ramp operating point. The heat-sink thermal mass buffers the ramp — a large heat sink stores some of the energy and releases it later — but for accurate ramp design, size to peak Q_{hot} with margin.

§6. Six-step heat-sink design methodology

Step 1 — Derive Q_{hot} from the TEC operating point

Read I_{TEC}, V_{TEC}, and Q_c from the TEC datasheet performance curves at the worst-case operating point: worst-case ambient, peak Q_{load}, deepest required ΔT. Compute P_{electrical} = V_{TEC} · I_{TEC}, then Q_{hot} = Q_c + P_{electrical}. For ramping or thermal-cycling loads such as PCR blocks, size for the peak Q_{load} plus the transient ramp energy Q_{transient} = m · c_p · dT/dt; the heat sink must reject the peak Q_{hot} during the ramp.

Step 2 — Define the worst-case ambient envelope at the heat-sink fins

Use the actual temperature at the heat-sink fins under service conditions, not lab ambient. Account for: enclosure restriction (instrument interior can run 10–15 °C above room ambient), solar gain on outdoor or window-adjacent equipment, altitude derating, dust accumulation, and fan aging. Convert these into an effective T_{ambient,max}. For OEM products, this is typically 10–20 °C above the rated operating ambient.

Step 3 — Establish $T_{\text{hot,max}}$ from the TEC datasheet at the operating point

Locate the TEC performance-curve family in the datasheet. At your I_{TEC} and Q_c , find the T_{hot} value that still supports the required ΔT . That T_{hot} is your $T_{\text{hot,max}}$ for heat-sink sizing. Do not use the TEC's absolute maximum T_{hot} rating (often 80–85 °C in the datasheet header). At or near that limit, the TEC may no longer support the required Q_c at the required ΔT . Use the operating-point performance curves, not the absolute maximum rating, to set $T_{\text{hot,max}}$. The operating-point $T_{\text{hot,max}}$ for a typical OEM module at $I = 60\%$ of I_{max} and $\Delta T = 30\text{ °C}$ is usually in the range 50–70 °C — substantially below the absolute maximum.

Common error — using TEC absolute maximum T_{hot} instead of operating-point $T_{\text{hot,max}}$

Datasheet header lists $T_{\text{hot,max}} = 80\text{ °C}$ for the TEC. Engineer sizes the heat sink for $T_{\text{ambient,max}} = 45\text{ °C}$ and $Q_{\text{hot}} = 30\text{ W}$, computing $R\theta_{\text{SA}} \leq (80 - 45) / 30 = 1.17\text{ K/W}$ (ignoring R_{TIM} and margin for transparency). Heat sink built, product cooled to setpoint in the lab. In service, T_{hot} reaches 65 °C — below the absolute maximum, but above the operating-point limit at the required ΔT and Q_c . The TEC pumps less than the load, the controller saturates current limit, and the product fails to reach setpoint. The correct R_{total} was $\leq (60 - 45) / 30 \approx 0.5\text{ K/W}$, with R_{TIM} subtracted and margin then applied — less than half the naive calculated value.

Step 4 — Calculate required $R\theta_{\text{SA}}$

Apply the explicit three-step margin sequence from §5.3 to get $R\theta_{\text{SA,selected}}$. Verify R_{TIM} against the TIM datasheet and the assembled compression; use 0.05–0.2 K/W as a typical starting range. If $R\theta_{\text{SA,selected}}$ is impractically small for available heat-sink envelopes, revisit Step 1 (consider a different TEC operating point), Step 2 (better airflow), or Step 3 (lower $T_{\text{hot,max}}$ may be available with a different TEC SKU).

Step 5 — Select heat-sink type and cooling mode

Once $R\theta_{\text{SA,selected}}$ is known, the heat-sink type and cooling mode follow from Q_{hot} , form-factor constraints, acoustic limits, and environment. §9 covers heat-sink types; §10 covers cooling-mode selection (passive / forced-air / liquid); §8 provides the parameter studies ($R\theta_{\text{SA}}$ vs airflow, vs envelope volume, vs fin count, thickness, spacing) that guide vendor selection.

Step 6 — Validate the assembled stack against worst-case conditions

Bench validation should reproduce the worst case the product will see: hot ambient, deepest setpoint, peak load, low V_{ps} , dust-loaded fins (if the product will operate in dusty environments). Record T_{hot} at steady state and during transients, controller current and voltage, and time to reach setpoint. Pass criterion: T_{hot} stays below $T_{\text{hot,max}}$ with the controller out of saturation through the entire envelope. Detailed validation procedure is in §13.

§7. Heat-sink construction and physics

A heat sink primarily does two things at the macroscopic level: it spreads heat from a concentrated source over a large surface, and it presents that surface to a moving fluid. Almost every design decision (material, base thickness, fin count, height, thickness, spacing, surface treatment) is a trade-off between conduction (spreading the heat) and convection (rejecting it to the air).

7.1 The base — conduction and spreading resistance

The base is a flat plate of aluminum or copper, thick enough to conduct heat from a concentrated TEC footprint (typically 15 × 15 mm to 40 × 40 mm) out to the full fin-array area without large in-plane temperature gradients. Too thin a base creates spreading resistance — heat is forced to crowd through a small effective area near the load, and only the closest fins do useful work. Too thick a base adds weight and material cost without proportional benefit. Typical base thickness: 3–8 mm for compact OEM heat sinks; 8–20 mm for high-power-density loads using copper inserts or vapor chambers.

7.2 Fins — the convection surface

Fins multiply the wetted surface area exposed to moving air. For a given envelope, fin design is a balancing act: more fins means more area but tighter spacing (which restricts flow and overlaps boundary layers); thinner fins fit more in but conduct less heat up their height (lower fin efficiency η_{fin}); thicker fins conduct well but leave less air gap and less area. The fin-efficiency expression for a straight rectangular fin is:

$$\eta_{fin} = \frac{\tanh(mL)}{mL}, \quad m = \sqrt{\frac{hP}{kA_c}}$$

In plain language: fin efficiency penalizes added fin area for being less effective than the bare base area, because the fin temperature drops along its height as heat is rejected. Tall, thin, or low-conductivity fins are less efficient than short, thick, high-conductivity fins.

where L is fin height, P is fin perimeter, A_c is fin cross-section, k is the fin material's conductivity, and h is the local convection coefficient. $\eta_{fin} = 1$ means the entire fin surface is at base temperature (perfect conductor); $\eta_{fin} < 1$ means the fin tip is cooler than the base, so area near the tip contributes less. Practical fins for OEM heat sinks have $\eta_{fin} = 0.7$ –0.95.

The overall heat-sink-to-air thermal resistance for a finned heat sink is approximately:

$$R_{\theta SA} = \frac{1}{h(A_{base} + \eta_{fin} A_{fins})}$$

In plain language: the total heat-sink-to-air thermal resistance equals one divided by the convection coefficient times the effective area, where the effective area is the bare base area plus the fin area weighted by fin efficiency.

where A_{base} is the un-finned base area exposed to air and A_{fins} is the total fin surface area. The factor η_{fin} penalizes the fin area for being less effective than the base area.

7.3 Material — aluminum, copper, composites

Aluminum extrusions are the default for OEM heat sinks: good conductivity at low cost with mature manufacturing. Copper has roughly 2× the conductivity of aluminum but is 3× heavier, much more expensive, and cannot be extruded into deep fin profiles. Copper is used as a base insert under high-power-density loads (TEC > 50 W footprint) to spread heat into a finned aluminum body. Vapor chambers and heat pipes use two-phase fluid (water, methanol) inside a sealed envelope to achieve effective conductivities 10–100× that of copper at modest sizes, but only over short distances.

7.4 Surface treatments

Bare polished aluminum has emissivity $\epsilon \approx 0.05$. Anodized aluminum (black or any color) reaches $\epsilon \approx 0.85$ [1, 5]. For natural-convection designs at modest ΔT , anodizing adds 5–15 % effective cooling by enabling radiation [5, 8]; for forced-air designs at OEM ΔT it adds 1–3 % at best because convection already dominates. Anodizing is cheap and standard regardless, mainly for cosmetic and corrosion-resistance reasons. Black-painted heat sinks behave similarly to black-anodized for radiation; matte black is slightly better than gloss black.

7.5 Advanced micro/nano-textured surfaces (research-stage)

Beyond conventional anodizing, an active research area is the deliberate engineering of micro- and nano-scale roughness on the heat-sink surface — typically by chemical etching, laser texturing, or combined etching plus a high-emissivity coating. The intent is to attack both transport modes at once. At the convective end, hierarchical micro/nano features disrupt the laminar sublayer right at the wall: instead of slipping smoothly along the fin (developing a thick stagnant film, see §3.6), the airflow trips earlier into turbulence and mixes more effectively against the hot surface. The mechanism is a deliberate disruption of the thermal boundary layer at the fin wall — the near-wall layer is broken up by surface microfeatures instead of sliding past as a smooth laminar film. At the radiative end, the cavities increase effective surface area and act as small radiative traps, raising the effective emissivity above that of a flat anodized surface.

Published results across several recent academic and applied research studies (chemical etching, laser texturing, hierarchical micro/nano structures, 2020–2025) typically report 5–10 °C lower steady-state surface temperature than standard black anodizing under the same heat flux and airflow, with the largest reported gains seen in forced-convection bench tests. Reported R θ SA improvements vary substantially with surface geometry, airflow regime, contamination, and durability. These figures are published research findings, not production-validated rules; the two underlying mechanisms (higher h from boundary-layer disruption, higher effective ϵ from cavity geometry) act simultaneously. Do not include advanced-surface improvement values in design margin unless validated on the exact heat sink, airflow regime, contamination environment, and service-life profile of the target program.

Practical status for OEM TEC designs. These advanced surfaces remain at the research and pilot-production stage and are not yet recommended for production designs. Cost is materially higher than standard anodizing; process consistency, batch-to-batch repeatability, and long-term durability under dust loading and thermal cycling have not been characterized at the level required for high-volume OEM qualification. For many current production TEC programs, standard black anodizing is a practical, mature, and commonly used choice. Advanced micro/nano-textured surfaces are worth evaluating only for selected high-power or sub-ambient prototype programs where additional thermal margin is worth experimental evaluation where a 5–10 °C reduction in T_{hot} translates directly into recovered Q_c capacity, longer service life, or a smaller heat-sink envelope. For qualified OEM programs where the heat sink is part of an ATI TEC-module or TEC-controller thermal stack, ATI applications engineering can review requirements and discuss possible custom design directions; contact sales@analogtechnologies.com.

§8. Heat-sink parameter studies

This section shows how R θ SA depends on the five most important design knobs: airflow at the heat-sink face, total heat-sink envelope volume, fin count, fin thickness, and fin-to-fin spacing. Trends are universal; absolute numbers shown are illustrative and depend on geometry, alloy, base thickness, and the specific

vendor's design. Always read $R_{\theta SA}$ from the candidate heat-sink's published datasheet curves at your actual operating airflow before locking the design.

8.1 $R_{\theta SA}$ versus airflow at the heat-sink face

Forced-air heat sinks rely on a fan to push or pull air through the fin array. As airflow rises, the convection coefficient h rises (sub-linearly, roughly as $\text{airflow}^{0.55-0.6}$), and $R_{\theta SA}$ falls. The curve is steepest at low airflow and flattens at higher airflow as additional CFM yields diminishing returns. Below the natural-convection threshold (zero airflow), $R_{\theta SA}$ reverts to its still-air value, typically 2–4× the forced-air value for the same envelope.

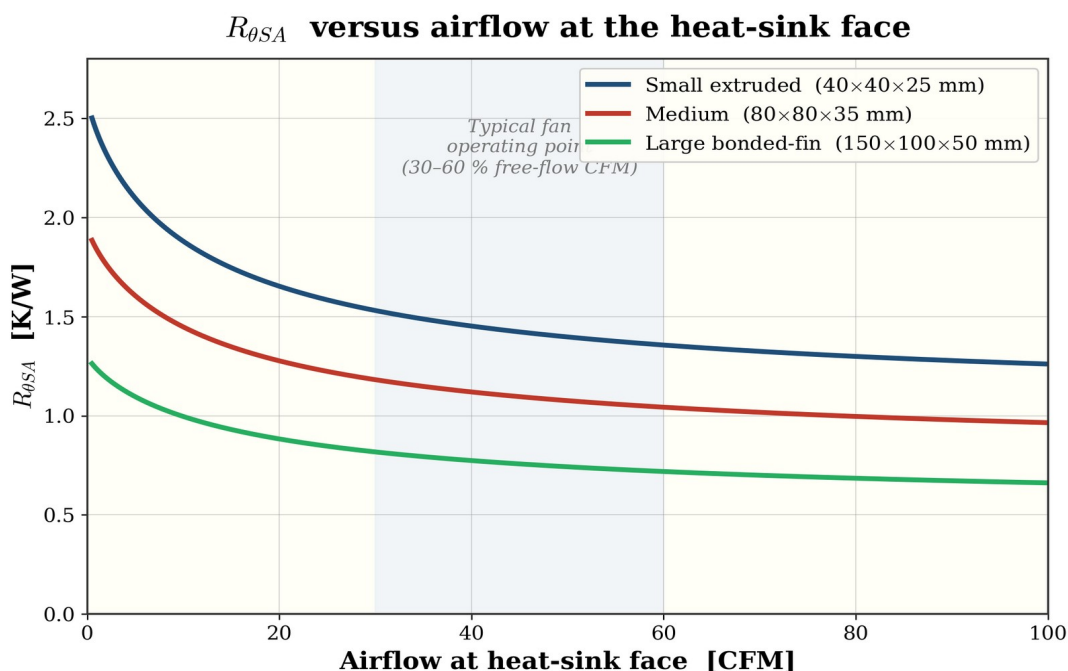


Figure 7 — $R_{\theta SA}$ versus airflow at the heat-sink face for three illustrative extruded heat sinks. The typical fan operating point lands at 30–60 % of fan free-flow CFM (the intersection of the fan curve and the heat-sink system-resistance curve, see Figure 12). Curves are illustrative; read your candidate heat sink's published $R_{\theta SA}$ versus airflow. Variables: $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance [K/W]; CFM = volumetric airflow at the heat-sink face [ft^3/min].

Practical rule — airflow

Doubling airflow typically reduces $R_{\theta SA}$ by 30–40 % [1, 3]. Going from natural convection to 25 CFM at the heat-sink face can reduce $R_{\theta SA}$ by 3–5×. Beyond 60–100 CFM on a typical OEM heat sink, gains become marginal — you reach the convection-coefficient asymptote.

8.2 $R_{\theta SA}$ versus heat-sink envelope volume

For geometrically similar designs, $R_{\theta SA}$ scales roughly as $V^{-0.4}$ in forced air and $V^{-0.45}$ in natural convection. Doubling the envelope volume (in similar form factor) reduces $R_{\theta SA}$ by about 25–30 %. The log-log slope tells you how much volume you need to add to cut R in half: about a factor of 4.7–5.7× in volume.

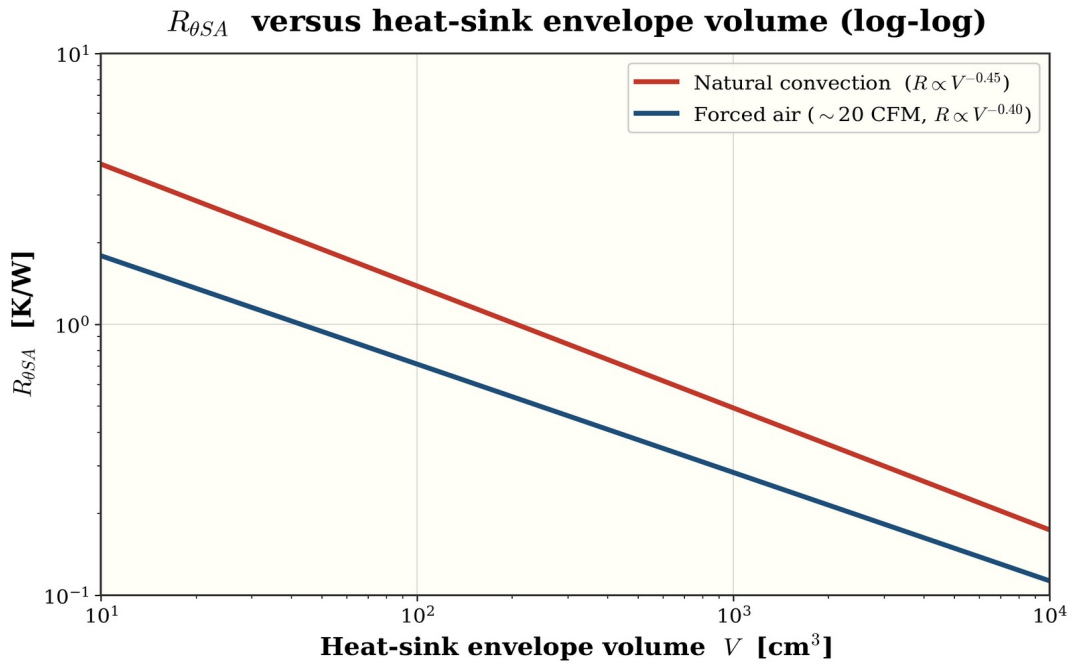


Figure 8 — $R_{\theta SA}$ versus heat-sink envelope volume (log-log) for natural convection and forced air at typical OEM airflow. Curves are illustrative trends; specific vendor designs vary by $\pm 30\%$ for the same envelope depending on fin geometry. Variables: $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance [K/W]; V = heat-sink envelope volume = length \times width \times height [cm^3].

Practical rule — volume

Cutting $R_{\theta SA}$ in half requires roughly $5\times$ the envelope volume in similar form factor [1, 3]. Adding a fan (going from natural to forced air) typically reduces $R_{\theta SA}$ by $3\text{--}5\times$ at the same volume — often a cost-effective cooling gain when the form factor and acoustic envelope allow it.

8.3 $R_{\theta SA}$ versus number of fins (fixed envelope and airflow)

More fins means more surface area, but also tighter spacing, more flow restriction, more boundary-layer overlap, and lower fin efficiency η_{fin} (because for fixed envelope, more fins forces each fin to be thinner). The net result has a clear minimum: too few fins leaves heat-rejection capacity on the table; too many fins chokes the airflow and pushes η_{fin} downward.




Figure 9 — $R_{\theta SA}$ versus fin count for a fixed envelope and fixed airflow. The optimum is typically 15–30 fins for OEM-size aluminum extrusions in forced air; the exact optimum depends on airflow, fin height, and the fan static-pressure capability. Variables: $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance [K/W]; N_f = number of fins; η_{fin} = fin efficiency (dimensionless, 0–1).

Practical rule — fin count

For typical OEM-scale aluminum extrusions under forced air, the fin-count optimum is in the 15–30 range for envelope widths of 50–150 mm — a starting range, not a universal rule; verify against the selected heat-sink $R_{\theta SA}$ -vs-airflow curves at the installed operating point. Below ~10 fins, surface area is the limit; above ~35 fins, flow choking and η_{fin} drop usually outweigh the area gain.

8.4 $R_{\theta SA}$ versus fin thickness

Fin thickness sets a trade-off between fin efficiency and surface area. Very thin fins (< 0.5 mm in aluminum) have low η_{fin} — the fin tip is much colder than the base because the thin cross-section cannot conduct heat up the fin fast enough. Very thick fins (> 3 mm) have η_{fin} near 1 but fewer fins fit in the envelope, so total area drops. The optimum is in the middle, typically 0.8–1.5 mm for OEM extruded heat sinks.

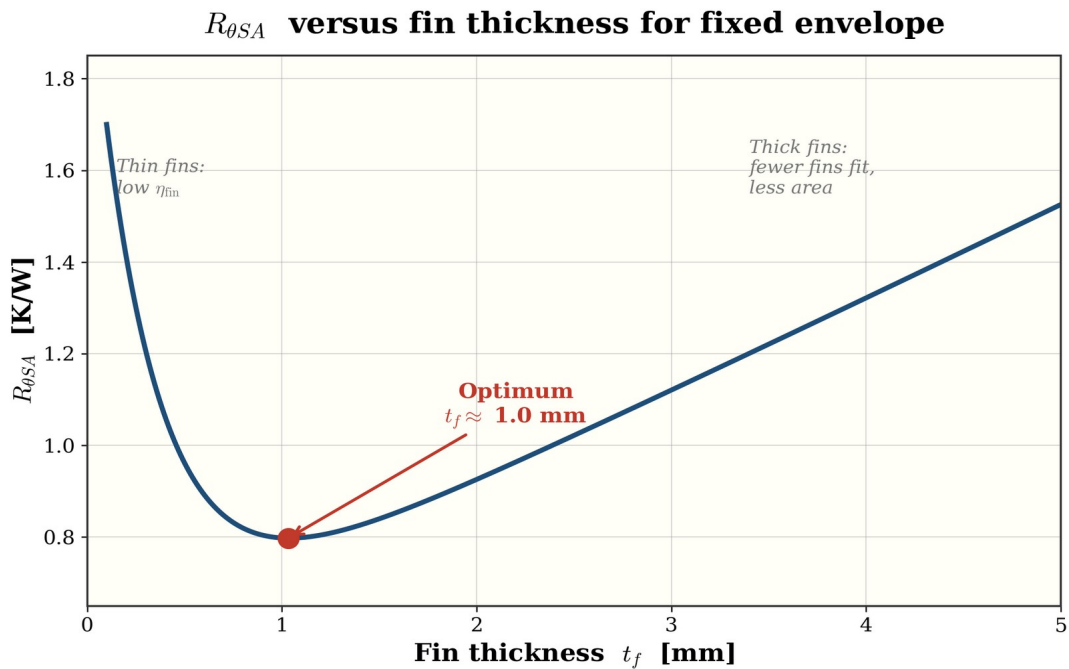


Figure 10 — $R_{\theta SA}$ versus fin thickness t_f for a fixed envelope. The optimum thickness balances fin efficiency η_{fin} (improves with thicker fins, because heat conducts more easily up a thicker fin) against total surface area (drops with thicker fins because fewer fins fit in the same envelope). For aluminum extrusions in forced air the optimum lies near 0.8–1.5 mm. Variables: $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance [K/W]; t_f = fin thickness [mm]; η_{fin} = fin efficiency (dimensionless).

Practical rule — fin thickness

Aluminum extruded fins for typical OEM forced-air designs land in the 0.8–1.5 mm thickness range; copper fins can go thinner (0.5–1.0 mm) due to higher conductivity. Skived-fin and bonded-fin processes routinely produce thinner, denser fins than extrusion can. These are illustrative starting ranges — the optimum must be verified against the selected fin material, manufacturing process, and operating airflow.

8.5 $R_{\theta SA}$ versus fin-to-fin spacing

Fin spacing sets a different trade-off. Too tight a spacing means boundary layers on adjacent fins overlap, choking flow — the inner fins see almost stagnant air. Too wide a spacing means fewer fins fit, reducing total surface area. The optimum depends strongly on airflow: forced air can use 2–4 mm because the boundary layer is thin; natural convection needs 6–10 mm because the buoyancy-driven boundary layer is thick and easily choked.

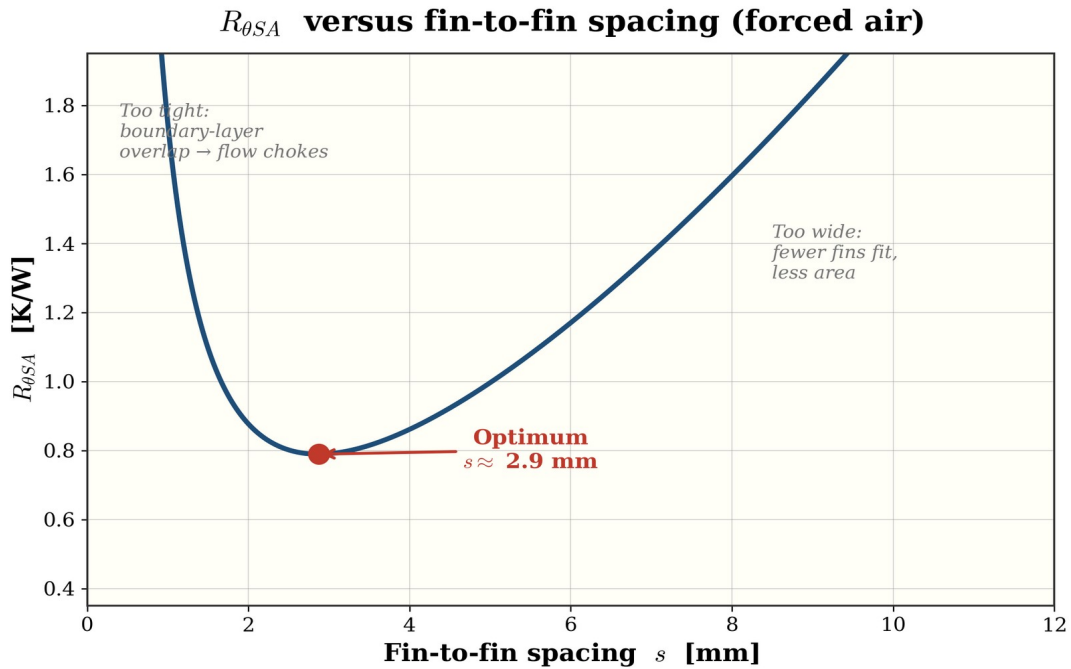


Figure 11 — $R_{\theta SA}$ versus fin-to-fin spacing s for a forced-air heat sink. The optimum for forced air is in the 2–4 mm range; for natural convection the optimum shifts to 6–10 mm because the buoyancy-driven boundary layer is thicker. Variables: $R_{\theta SA}$ = heat-sink-to-ambient thermal resistance [K/W]; s = gap between adjacent fins [mm].

Practical rule — fin spacing

For typical OEM-scale aluminum heat sinks, forced-air designs optimize around 2–4 mm fin spacing and natural-convection designs around 6–10 mm — typical starting ranges, not universal rules; the optimum depends on installed airflow, fan/system curve, and envelope. Stuffing more fins into a natural-convection design without widening the spacing makes thermal performance worse, not better.

§9. Heat-sink types and when to use each

Heat-sink type	Typical Q_{hot} range	Form factor and cost	Best fit
Extruded aluminum	5–50 W	Compact to large; low cost	General OEM, lab and instrument cooling, photonics enclosures
Bonded-fin / skived Al or Cu	20–200 W	High fin density; medium–high cost	Telecom, high-power photonics, dense rack systems
Vapor chamber + fins	30–300 W	Flat profile; high cost	High heat-flux loads, hotspot spreading, tight stack-up
Heat pipes (remote rejection)	50–500 W	Decouples source and rejection; medium cost	Tight space at TEC, larger heat sink elsewhere

Heat-sink type	Typical Q_{hot} range	Form factor and cost	Best fit
Liquid cold plate + remote exchanger	50 W – multi-kW	Compact at TEC; system complexity high	High-power lasers, sealed enclosures, sub-ambient cold plates

Ranges are typical guidelines, not hard limits. A well-designed extruded heat sink with a large fan can exceed 50 W; a vapor chamber inside a sealed industrial enclosure may be the only option at 30 W if the case constrains the form factor. Always validate the candidate against actual $R\theta_{SA}$ at your operating airflow, not against the product category.

§10. Cooling-mode selection — passive, forced air, liquid

10.1 Passive (natural-convection) cooling

Passive cooling relies on buoyancy-driven airflow over the heat-sink fins. $R\theta_{SA}$ scales roughly as $1/A^{0.5}$ for given fin geometry; doubling cooling capacity requires roughly 4× the heat-sink area. Practical upper limit is around 10 W of Q_{hot} for compact designs and 20–30 W for large heat sinks with optimized vertical fin orientation. Right choice for: acoustic-sensitive diagnostic or laboratory instrumentation, vibration-sensitive precision metrology (interferometry, optical reference cavities), high-MTBF unattended designs where a fan would be a single point of failure, and very small Q_{hot} budgets.

10.2 Forced air (fan + heat sink)

Forced air covers the 10–100 W Q_{hot} range for compact form factors. The fan establishes a convective coefficient 5–20× higher than natural convection. The trade-offs are fan acoustic noise, fan reliability (MTBF, dust loading, bearing wear), and the complexity of fan control electronics. Sizing requires the heat-sink airflow curve ($R\theta_{SA}$ vs CFM) and the fan curve (CFM vs static pressure); the operating point is the intersection, typically 30–60 % of the fan's free-flow CFM.

10.3 Liquid (cold plate + chiller or radiator)

Liquid cooling becomes attractive above 50 W when compactness at the TEC matters, acoustic limits exclude fans, ambient constraints exclude forced air (sealed enclosures, sub-ambient cold-plate operation), or thermal mass is needed to smooth a peaky duty cycle. The trade-off is system complexity: pump, hoses, fill/drain, fluid compatibility, condensation control, and pump acoustic / vibration concerns of their own. Liquid is justified by the environmental or form-factor constraint that excludes forced air, not by raw thermal performance alone.

§11. Fan selection and the operating point

Free-flow CFM (the airflow a fan moves with zero back-pressure) is not the airflow your fan delivers through a heat sink. The heat sink presents static pressure to the fan; the actual operating point is where the fan curve and the heat-sink system-resistance curve meet. For dense-fin OEM heat sinks, this typically lands at 30–60 % of the fan's free-flow CFM rating.

Fan operating point — fan curve meets system-resistance curve

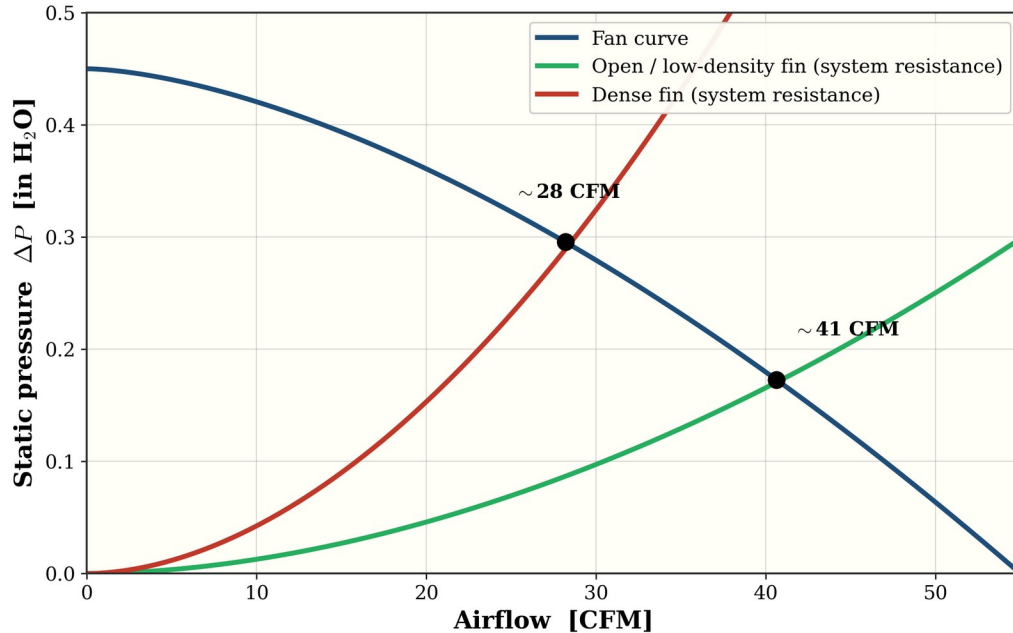


Figure 12 — Fan operating point. The fan curve (static pressure vs CFM) intersects the heat-sink system-resistance curve at the actual operating airflow. A more restrictive heat sink (dense-fin) shifts the operating point toward lower CFM and higher static pressure. Variables: CFM = volumetric airflow [ft³/min]; ΔP = static pressure across the heat sink [inches H₂O or Pa]; free-flow CFM = fan’s rated CFM at zero static pressure.

11.1 Sizing margin and aging

Specify the fan with margin against bearing wear (sleeve fans lose 20–30 % CFM over their service life; ball-bearing fans 10–20 % [5, 8] (typical planning values only — verify against fan vendor data and filter-maintenance interval)), dust loading (filters and dirty fins increase system resistance, shifting the operating point), and altitude (air density falls roughly 10 % per km, reducing both fan-curve CFM and convective coefficient h). For final sizing, determine the actual operating point from the intersection of the fan curve and the heat-sink system-resistance curve — do not use a fixed percentage of free-flow CFM as the final sizing basis. The fan operating point is the intersection of the fan curve and the heat-sink system-resistance curve, and this curve-intersection method is the authoritative sizing rule for forced-air designs. The 30–60 % of free-flow CFM observation is a planning heuristic, not a sizing criterion.

11.2 Acoustic limits

Fan acoustic noise scales steeply with rotation speed (typically as the fifth power of RPM). For acoustic-sensitive diagnostic or laboratory instrumentation, the acoustic budget may dictate a larger, slower fan rather than a smaller, faster one. Larger, slower fans often reduce broadband noise for a given total CFM compared with smaller, faster fans, and dual-fan arrangements add redundancy; final noise must be measured in the installed assembly.

11.3 Push vs pull configuration

Push (fan blowing into the heat sink) delivers slightly more airflow per fan watt; pull (fan downstream of the heat sink) gives more uniform airflow through the fin channels. For OEM designs, push is more common; pull is preferred when fins are unusually long (> 75 mm) or when the fan needs to be remote from the hot zone for acoustic or thermal reasons.

§12. Mounting interface and thermal interface material (TIM)

12.1 Surface flatness and contact pressure

Why this matters at the microscopic scale. No matter how polished an aluminum surface looks to the naked eye, at the micron scale it is a landscape of jagged peaks and valleys. When the TEC ceramic and the heat-sink base are pressed together, the two surfaces touch only at their highest peaks. The valleys between those peaks remain filled with air, and still air is a thermal insulator at roughly $0.024 \text{ W}/(\text{m}\cdot\text{K})$ — about four orders of magnitude worse than aluminum at $\sim 200 \text{ W}/(\text{m}\cdot\text{K})$. Those microscopic air pockets dominate the contact thermal resistance unless something displaces them. A thin native aluminum-oxide layer (Al_2O_3) forms on every exposed aluminum surface within seconds; it is far less conductive than pure aluminum and adds a small additional series resistance, but it is the trapped air, not the oxide, that hurts most.

The TEC ceramic and the heat-sink base must be flat and parallel to within $\pm 0.025 \text{ mm}$ (1 mil) across the contact area to avoid stress concentrations and large air gaps. Lap or machine the heat-sink base if the as-purchased flatness is worse. Apply uniform compression of $1.0\text{--}2.0 \text{ MPa}$ ($\approx 145\text{--}290 \text{ psi}$) across the TEC face — verify the manufacturer's recommended pressure window. Insufficient pressure leaves micro-gaps that act as thermal insulators; excessive pressure cracks the TEC ceramic. Even with ideal flatness and compression, peak-to-peak roughness still leaves microscopic air pockets — that is what TIM (next subsection) exists to displace.

12.2 TIM selection and thickness

How TIM works. A thermal interface material is a soft, pliable medium that flows into the microscopic valleys on both contact surfaces under clamping pressure and physically displaces the trapped insulating air. The TIM itself is not as conductive as aluminum (k typically $1\text{--}8 \text{ W}/(\text{m}\cdot\text{K})$ for grease, vs ~ 200 for Al), but it is roughly $50\text{--}300\times$ more conductive than the air it replaces — so even a $25\text{--}75 \mu\text{m}$ bond line lowers the contact thermal resistance dramatically. The thinner and more uniform the bond line, the better, because the TIM itself is a series resistance in the stack. Overspreading grease is a common error: excess thickness raises resistance, not lowers it.

TIM type	Typical k	Typical thickness	R per cm^2 (illustrative)	Best fit
Thermal grease (silicone)	$1\text{--}3 \text{ W}/(\text{m}\cdot\text{K})$	$25\text{--}75 \mu\text{m}$	$0.1\text{--}0.3 \text{ K}/\text{W}$	General OEM; most TEC datasheets assume this
High-perf grease (metal-loaded)	$5\text{--}8 \text{ W}/(\text{m}\cdot\text{K})$	$25\text{--}75 \mu\text{m}$	$0.05\text{--}0.15 \text{ K}/\text{W}$	High- Q_{hot} , dense-fin heat sinks
Phase-change pad	$2\text{--}4 \text{ W}/(\text{m}\cdot\text{K})$	$50\text{--}150 \mu\text{m}$	$0.1\text{--}0.4 \text{ K}/\text{W}$	Production-friendly, no pumping out
Indium foil	$80 \text{ W}/(\text{m}\cdot\text{K})$	$50\text{--}250 \mu\text{m}$	$0.05\text{--}0.10 \text{ K}/\text{W}$	Very high Q_{hot} , lab/photronics; reflowable
Graphite sheet	$10\text{--}20 \text{ W}/(\text{m}\cdot\text{K})$ in-plane	$70\text{--}300 \mu\text{m}$	$0.2\text{--}0.5 \text{ K}/\text{W}$	Anisotropic; in-plane spreader use

TIM thickness should be the minimum needed to fill surface micro-roughness [8]. A thinner TIM layer has lower resistance, but only if surfaces are flat enough to contact across the whole face. Overspreading grease is a common error — the excess does nothing useful and may migrate out during thermal cycling.

12.3 Mounting hardware and torque

Compression is typically applied with screws around the perimeter of the TEC or through a clamping bar. For M3 screws on a 30 × 30 mm TEC: torque of 0.2–0.4 N·m (1.8–3.5 lbf·in) per screw is a typical starting point, distributed over four perimeter screws to give the 1.0–2.0 MPa target compression. Verify the actual compression against the TEC manufacturer's mounting guidance — published torque-to-compression conversions vary with screw thread, washer presence, and the springiness of any leveling stack.

§13. Validation envelope and pass criteria

13.1 Worst-case validation envelope

Bench validation should reproduce, simultaneously, all of the worst-case conditions the product will see: maximum ambient temperature (chamber test), deepest required ΔT , peak Q_{load} (or peak ramp dT/dt), minimum supply voltage, dust-loaded heat-sink fins (if the product will encounter dust), and end-of-service-life fan CFM (use a programmable variable-CFM fan, or restrict the heat-sink inlet with a calibrated baffle). The validation envelope should also include the longest expected operating duration to confirm steady-state thermal capacity.

13.2 Instrumentation

Instrument: T_{ambient} near the heat-sink inlet, T_{HS} at the heat-sink base nearest the TEC, T_{hot} on the TEC ceramic (small thermocouple or RTD bonded with thermally-conductive adhesive), T_{cold} on the cold-side load surface, controller current and voltage, and a time stamp at controller entry to current-limit saturation (if any). Thermal imaging is useful as a survey tool but should be cross-checked against contact sensors at the points that matter.

13.3 Pass criteria

Three criteria must hold simultaneously through the entire envelope: (1) T_{hot} stays below $T_{\text{hot,max}}$ with at least 5 °C of headroom; (2) the controller has at least 10–20 % of its current capacity available (not pinned at I_{max} for sustained periods); (3) the loop reaches and holds setpoint within the application's accuracy and settling time. If any of the three fails, return to the methodology — likely Steps 1–3 — to find the wrong-sized component.

13.4 Common validation mistakes

- **Reading T_{hot} during the transient** before the heat sink has reached steady state (one full time constant $\tau = R\theta_{SA} \cdot C_{\text{th}}$ of the heat sink — often several minutes for large heat sinks). Wait at least 5τ before recording steady-state numbers.
- **Validating only at room ambient.** Chamber testing at worst-case ambient is required for OEM products targeting field service.
- **Not restricting fan airflow to end-of-life CFM.** A new fan at 100 % CFM gives a too-optimistic margin. Use a baffle, dust loading, or a programmable fan to simulate aging.
- **Not validating the deepest required ΔT .** A heat sink that passes at $\Delta T = 20$ °C may fail at $\Delta T = 40$ °C because Q_{hot} grows faster than ΔT at high ΔT .

§14. Common failure modes after correct component selection

Even with a correctly-sized heat sink on paper, the field can surface problems that did not show on the bench. Common failure modes after component selection include the following.

1. **Wrong $T_{hot,max}$ baseline** — designer used the TEC datasheet absolute maximum rating instead of the operating-point $T_{hot,max}$ at the required Q_c and ΔT . A common and serious error in TEC heat-sink design.
2. **R_{TIM} neglected** — R_{TIM} of 0.1–0.2 K/W is on the order of half of a well-designed $R\theta_{SA}$. Leaving it out of the budget makes T_{hot} exceed target by $Q_{hot} \cdot R_{TIM}$ — often 5–10 °C.
3. **Free-flow CFM used instead of operating-point CFM**. Datasheet $R\theta_{SA}$ at 50 CFM \neq actual $R\theta_{SA}$ when the fan delivers 25 CFM through the dense-fin system.
4. **Worst-case ambient is the lab ambient** — the heat sink is sized at 25 °C lab conditions and deployed in a closed instrument enclosure at 40 °C internal temperature.
5. **Fan aging or dust loading not simulated**. End-of-life fan CFM is 20–30 % below new; dust on fins adds 10–40 % to $R\theta_{SA}$ depending on environment and time.
6. **Inadequate mounting compression** — not enough screw torque or non-flat surfaces create thermal interface gaps that pump out over thermal cycles.
7. **Transient Q_{hot} ignored** — ramping loads (PCR, laser pull-down) generate peak Q_{hot} during the ramp that is significantly higher than steady-state Q_{hot} .
8. **Sub-ambient operation amplifies Q_{hot}** — holding the load below ambient adds passive heat leak from ambient into the cold side, which raises Q_c and therefore Q_{hot} beyond the naive cold-side load calculation.

§15. Application examples (illustrative)

Numbers shown are for illustration of the method, not as recommended values. Real designs require validation at worst-case conditions with the specific TEC datasheet curves, candidate heat-sink $R\theta_{SA}$ vs CFM data, R_{TIM} from the selected TIM, environment-specific margin factor, and OEM-defined $T_{hot,max}$ headroom — and OEM-level qualification for medical / IVD, automotive / LIDAR, telecom outdoor / ingress / EMC, and other regulated end products.

Example 1 — 1 W photonic detector, sub-ambient hold

- $Q_{load} = 1$ W; setpoint 15 °C; $T_{amb,max} = 30$ °C; $\Delta T = 15$ °C below ambient
- Choose small TEC at modest I/I_{max} ; from datasheet, $P_{electrical} \approx 2.5$ W; $Q_{hot} \approx 3.5$ W
- $T_{hot,max}$ from TEC datasheet at this operating point ≈ 50 °C; ΔT_{HS} budget = 50 – 30 = 20 °C
- $R_{total} \leq 20 / 3.5 = 5.7$ K/W; subtract $R_{TIM} = 0.2$ K/W $\rightarrow R\theta_{SA,calc} = 5.5$ K/W
- Lab environment, margin 2 \times $\rightarrow R\theta_{SA,selected} \approx 2.8$ K/W as a candidate starting point — a small passive aluminum heat sink with vertical fins is a plausible direction (verify in installed conditions)

Example 2 — 3 W laser diode, photonics enclosure

- $Q_{load} = 3$ W; setpoint 25 °C; enclosure $T_{amb,max} = 45$ °C; $\Delta T_{TEC} = 25$ °C
- Mid-size TEC at moderate I/I_{max} ; $P_{electrical} \approx 6$ W; $Q_{hot} \approx 9$ W
- Operating-point $T_{hot,max}$ from datasheet at 25 °C $\Delta T \approx 65$ °C; ΔT_{HS} budget = 65 – 45 = 20 °C

- $R_{\text{total}} \leq 20 / 9 = 2.2 \text{ K/W}$; subtract $R_{\text{TIM}} 0.15 \text{ K/W} \rightarrow R_{\theta\text{SA,calc}} \approx 2.05 \text{ K/W}$
- OEM environment, margin $2.5\times \rightarrow R_{\theta\text{SA,selected}} \approx 0.82 \text{ K/W}$ — forced-air extruded HS, $60 \times 60 \times 30 \text{ mm}$ with 40 mm fan

Example 3 — 10 W mid-power load, lab instrument, favorable operating point

- $Q_{\text{load}} = 10 \text{ W}$; setpoint $5 \text{ }^\circ\text{C}$; lab $T_{\text{amb,max}} = 30 \text{ }^\circ\text{C}$; $\Delta T_{\text{TEC}} = 25 \text{ }^\circ\text{C}$
- Larger TEC chosen at $I/I_{\text{max}} < 0.5$ for favorable $\text{COP} \approx 1.1$; $P_{\text{electrical}} \approx 9 \text{ W}$; $Q_{\text{hot}} \approx 19 \text{ W}$
- Operating-point $T_{\text{hot,max}}$ from datasheet $\approx 60 \text{ }^\circ\text{C}$; ΔT_{HS} budget = $30 \text{ }^\circ\text{C}$
- $R_{\text{total}} \leq 30 / 19 = 1.58 \text{ K/W}$; subtract $R_{\text{TIM}} 0.10 \text{ K/W} \rightarrow R_{\theta\text{SA,calc}} \approx 1.48 \text{ K/W}$
- Lab environment, margin $2\times \rightarrow R_{\theta\text{SA,selected}} \approx 0.74 \text{ K/W}$ — medium forced-air HS, $100 \times 100 \times 35 \text{ mm}$ with 60 mm fan

This example uses an unusually favorable $\text{COP} \approx 1.1$ (achievable only at low I/I_{max}). Most OEM operating points run with $\text{COP} < 1$; Q_{hot} is then closer to $2\text{--}3\times Q_{\text{load}}$.

Example 4 — 30 W PCR / qPCR block, thermal-cycling load

- $Q_{\text{load,steady}} = 5 \text{ W}$; ramp peak $Q_{\text{load}} = 30 \text{ W}$ during the $20^\circ/\text{s}$ ramp; setpoint range $4 \text{ }^\circ\text{C}$ to $95 \text{ }^\circ\text{C}$; ambient up to $35 \text{ }^\circ\text{C}$
- Large TEC at $I/I_{\text{max}} \approx 0.7$; ramp $P_{\text{electrical}} \approx 35 \text{ W}$ at deepest ramp; ramp peak $Q_{\text{hot}} \approx 65 \text{ W}$
- Operating-point $T_{\text{hot,max}}$ during ramp $\approx 55 \text{ }^\circ\text{C}$; ΔT_{HS} budget = $20 \text{ }^\circ\text{C}$
- $R_{\text{total}} \leq 20 / 65 = 0.31 \text{ K/W}$; subtract $R_{\text{TIM}} 0.05 \text{ K/W} \rightarrow R_{\theta\text{SA,calc}} \approx 0.26 \text{ K/W}$
- OEM environment, margin $2.5\times \rightarrow R_{\theta\text{SA,selected}} \approx 0.10 \text{ K/W}$ — large bonded-fin HS, $150 \times 100 \times 50 \text{ mm}$ with two 80 mm fans or a liquid cold plate

Example 5 — 40 W outdoor kiosk, harsh environment

- $Q_{\text{load}} = 40 \text{ W}$; $T_{\text{amb,max}} = 55 \text{ }^\circ\text{C}$ (sun-exposed kiosk in summer); $\Delta T_{\text{TEC}} = 15 \text{ }^\circ\text{C}$
- Large TEC at $I/I_{\text{max}} \approx 0.5$; $P_{\text{electrical}} \approx 50 \text{ W}$; $Q_{\text{hot}} \approx 90 \text{ W}$
- Operating-point $T_{\text{hot,max}}$ from datasheet $\approx 75 \text{ }^\circ\text{C}$; ΔT_{HS} budget = $20 \text{ }^\circ\text{C}$
- $R_{\text{total}} \leq 20 / 90 = 0.22 \text{ K/W}$; subtract $R_{\text{TIM}} 0.05 \text{ K/W} \rightarrow R_{\theta\text{SA,calc}} \approx 0.17 \text{ K/W}$
- Outdoor harsh, margin $4\times \rightarrow R_{\theta\text{SA,selected}} \approx 0.04 \text{ K/W}$ — liquid cold plate with remote heat exchanger / fan stack; consider IP-rated enclosure for the fans

§16. Frequently asked questions

Q1. Why can't I just use the TEC datasheet maximum T_{hot} rating for heat-sink sizing?

Because the absolute maximum T_{hot} is the temperature at which Q_c can fall well below the level needed to support the application — at or near the absolute-maximum hot-side rating, the TEC may no longer support the required Q_c at the required ΔT ; use the operating-point curves. The operating-point $T_{\text{hot,max}}$, read from the TEC performance curves at your actual I_{TEC} and required Q_c , is the highest T_{hot} at which the TEC still pumps the load you need. The two values can differ by $20\text{--}40 \text{ }^\circ\text{C}$.

Q2. How much margin should I add to $R_{\theta\text{SA}}$?

$1.5\text{--}2\times$ for controlled lab environments; $2\text{--}3\times$ for general OEM (office, instrument enclosure); $3\text{--}5\times$ for outdoor or harsh environments. Margin is applied to $R_{\theta\text{SA}}$, not to power.

Q3. Can I share one heat sink between multiple TECs?

Yes, when the TECs are co-located and operate together, but you must sum the individual Q_{hot} values for sizing. For low-power systems (combined Q_{hot} below ~10 W) and modest ΔT , shared heat sinks simplify the design. For higher- Q_{hot} or precision setpoints, separate heat sinks reduce cross-coupling between control loops.

Q4. How do I know what airflow my fan actually delivers through the heat sink?

Plot the heat-sink system-resistance curve (static pressure vs CFM, from the heat-sink datasheet) and the fan curve on the same axes. The intersection of the two curves is the actual operating-point airflow. This curve-intersection method is the authoritative sizing rule. The 30–60 % of free-flow CFM range is an illustrative observation, not a sizing criterion: a fan with strong static-pressure capability may deliver the required airflow even when the operating point sits well below free-flow. Verify that the delivered airflow meets the required $R\theta_{\text{SA}}$ with margin; if not, choose a fan with higher static-pressure capability or reduce system resistance.

Q5. Does anodizing improve heat-sink performance?

For natural-convection designs, yes — a modest improvement from added radiation, since the emissivity of black anodized aluminum ($\epsilon \approx 0.85$) is much higher than bare polished aluminum ($\epsilon \approx 0.05$). For forced-air designs at typical OEM ΔT the improvement is usually small because convection dominates. Quantify with vendor data on the actual heat sink before counting it in design margin. Anodizing is worth doing for corrosion and cosmetic reasons; the thermal benefit is a bonus.

Q6. Should I orient my heat sink with fins vertical or horizontal?

For natural-convection designs, vertical fin orientation is strongly preferred — horizontal fins block the buoyancy-driven air column and can significantly raise $R\theta_{\text{SA}}$. The actual penalty depends on fin spacing, height, base orientation, chimney path, and enclosure; test in the installed orientation. For forced-air designs, orientation matters less as long as the fan provides clear airflow paths through the fin channels.

Q7. Can I use thermal pads instead of grease?

Yes, but with care. The relevant interface resistance is $R_{\theta_{\text{CS}}}$ (the TIM/case-to-sink term, sometimes written R_{TIM}). Phase-change pads (graphite-loaded silicone, melting around 50–60 °C) typically approach grease performance after the first thermal cycle. Soft silicone pads usually have higher R_{TIM} than grease and are generally not recommended for high- Q_{hot} designs. Indium foil can achieve very low interface resistance in suitable assemblies but is expensive and difficult to apply. Verify against the specific TIM datasheet and your assembled-stack pressure, surface flatness, and thermal-cycle profile.

Q8. How often should I replace thermal grease?

Service life depends on the application. In stable, low-cycling environments, a well-applied silicone-grease joint can last many years — sometimes the product service life. Pump-out and dry-out are accelerated by deep thermal cycling, vibration, high ΔT across the joint, and contamination. For ramping loads (PCR blocks) and high-vibration environments, plan for shorter service intervals and verify with field data or accelerated testing.

Q9. How do I budget for a low-power system ($Q_{\text{hot}} < 5 \text{ W}$) where margin factors seem excessive?

At low Q_{hot} , ΔT_{HS} is small in absolute terms (a 2× margin on $R\theta_{\text{SA}} = 2 \text{ K/W}$ on a 5 W load means just 5 °C of HS rise). For very-low- Q_{hot} photonic loads in stable lab environments, smaller margin factors (1.3–1.5×) may be acceptable after validation because the absolute consequences of error are small — but do not reduce

margin for sub-ambient, sealed, outdoor, high-stability, or production OEM systems without assembled-system testing. For sub-ambient operation, return to full OEM margin regardless of Q_{hot} — the failure-mode amplification can be severe under deep sub-ambient setpoints, humid environments, poor insulation, or high heat leak.

Q10. What is a common and serious mistake when sizing a heat sink for a TEC system?

Using the TEC datasheet absolute-maximum hot-side temperature instead of the operating-point $T_{hot,max}$ from the TEC performance curves at the actual I_{TEC} and required Q_c . The two values can differ by 20–40 °C; the absolute-maximum rating is the temperature at which the TEC may no longer support the required ΔT at the required Q_c , while the operating-point $T_{hot,max}$ is the highest T_{hot} at which the TEC still pumps the load you need. This is a common and serious heat-sink-sizing error in TEC applications.

§17. ATI product bridge and engineering support

ATI provides TEC modules (ATE-series), TEC controllers (TEC18V family and others), thermistors and temperature sensors, and evaluation boards. Heat sinks, fans, cold plates, and TIM are typically sourced from established thermal-management vendors (examples only — verify current data, availability, and suitability with the selected supplier); ATI does not produce a complete heat-sink catalog. For qualified OEM projects, ATI applications engineering can review provided requirements and suggest a starting TEC family, controller-selection path, cooling-mode direction, and evaluation path.

17.1 Information ATI needs to review your project

To suggest a starting thermal-stack direction efficiently, prepare:

1. **Q_{load}** — cold-side thermal load in watts, including transient peaks.
2. **Ambient envelope** — $T_{ambient,max}$ at the heat-sink fins under service conditions (not just lab ambient), along with the operating-environment category (lab / general OEM / outdoor / harsh).
3. **TEC operating point + selected SKU (if known)** — setpoint, ΔT_{TEC} , and the candidate TEC module part number if one has been picked. If the TEC has not been chosen, ATI applications engineering can review project requirements and suggest a starting TEC family for qualified OEM projects.
4. **Form-factor constraints** — maximum heat-sink envelope (height × width × depth), mounting orientation, and any restrictions on fan size, acoustic budget, or moving parts.
5. **Cooling-mode preference (if any)** — passive / forced-air / liquid, or open to ATI's suggested cooling-mode direction.
6. **Service-life target** — expected operating lifetime, duty cycle (continuous vs ramping), service environment (clean office, dusty industrial, outdoor).
7. **Controller and sensor context (optional)** — controller candidate / supply voltage / sensor type and placement / stability target / heating-cooling-bidirectional.
8. **Project context (optional)** — prototype vs production, expected annual volume, custom constraints, and target schedule.

17.2 What ATI returns

A suggested thermal-stack direction (possible starting TEC family or SKU candidate if not pre-selected; cooling-mode direction; $R\theta_{SA}$ target range; controller-selection path to AWP-TECC-02), an evaluation path

(eval-board pairing and bench-test setup), and engineering review of the OEM's calculated starting point if provided. ATI feedback is a starting-point planning input; final heat-sink selection, vendor qualification, enclosure validation, and regulatory/system qualification remain the OEM's responsibility.

ATI does not pre-qualify designs for medical / IVD, automotive / LIDAR, telecom outdoor, or aerospace; that level of qualification remains the OEM's responsibility under the relevant standards (FDA / IEC 60601 / IVDR for medical; AEC-Q / ISO 26262 for automotive; relevant outdoor / EMC standards for telecom; relevant aerospace standards). ATI components do not hold these system-level certifications on their own; system qualification is the responsibility of the OEM integrator.

17.3 Companion white papers

- [AWP-TECM-01](#) — TEC module fundamentals (what a TEC is, how it pumps heat, performance curves, mounting basics).
- [AWP-TECC-02](#) — TEC controller selection (linear / PWM / hybrid topologies, control accuracy, current/voltage matching, controller datasheet interpretation).
- [AWP-TECC-03](#) — Full TEC cooling system design integrating module, controller, sensor, heat sink, mounting, and validation as a connected stack.

References

- [1] Incropera, F. P., DeWitt, D. P., Bergman, T. L., and Lavine, A. S. "Fundamentals of Heat and Mass Transfer," 8th edition, Wiley, 2017.
- [2] Lienhard, J. H. IV and Lienhard, J. H. V. "A Heat Transfer Textbook," 5th edition (open access), Phlogiston Press, 2020. Available at ahtt.mit.edu.
- [3] Kraus, A. D., Aziz, A., and Welty, J. "Extended Surface Heat Transfer," Wiley-Interscience, 2001.
- [4] Rowe, D. M. (ed.) "Thermoelectrics Handbook: Macro to Nano," CRC Press, 2006.
- [5] ASHRAE "Handbook — Fundamentals," 2021 edition.
- [6] ATI ATE-series TEC module datasheets and performance curves — see analogtechnologies.com TEC module pages.
- [7] ATI white papers AWP-TECM-01, AWP-TECC-02, AWP-TECC-03, and AWP-HS-01 (this paper) — see analogtechnologies.com/white-papers.
- [8] Sergent, J. E. and Krum, A. "Thermal Management Handbook for Electronic Assemblies," McGraw-Hill, 1998.

— End of AWP-HS-01 Revision 2.16 —