

Part 2: Design

Server Cold Plates

BY AUSTIN M. SHELNUTT, P.E.; DAVID QUIRK, P.E., MEMBER ASHRAE

In the September issue of *ASHRAE Journal*, this column started a three-part series on liquid-cooled server processor cold plates. Part 1 concentrated on the history of cold plates used in cooling computers. Part 2 focuses on the technical design of cold plates—to provide HVAC designers and others with a more intuitive understanding of the devices they are (or soon will be) charged with cooling.

Cold Plate Design Fundamentals

The value and impact of cold plate use for cooling high power data center processors has been discussed in the first column of this series.¹ Part 2 presents an overview of cold plate design fundamentals and provides the reader with the new metrics for consideration on how to best “optimize” a cold plate design.

Cold plates invoke two primary modes of heat transfer: conduction and convection. Conduction, of course, is the mode where heat is transferred away from the information technology equipment (ITE) processor through solid-to-solid interfaces including the thermal interface material, the cold plate base and up into the cold plate fin structure. Convection is responsible for transporting heat from the cold plate surface into the liquid coolant.

Thermal resistance (typically measured in K/W, °C/W or °F/Btu/h) is the most widely practiced metric for evaluating cold plate performance. Both conduction and convection contribute to the total thermal resistance of a cold plate and can be generalized with the following relationships:

$$R = \frac{\Delta T}{\dot{q}} \quad (1)$$

$$R_{\text{cold\_plate}} = R_{\text{conduction}} + R_{\text{convection}} \quad (2)$$

$$R_{\text{conduction}} = \frac{t}{kA_{\text{CS}}} \quad (3)$$

$$R_{\text{convection}} = \frac{1}{hA_{\text{surface}}} \quad (4)$$

where (Equations 1–4):

R is the variable used for thermal resistance (°C/W)

ΔT is the temperature drop from heat source to heat sink for the segment of heat transfer being analyzed (°C)

$\dot{q}$ is the rate of heat being transferred (W)

t represents the thickness of a solid material for heat conduction (m)

k is the material thermal conductivity (W/mK)

A_{CS} is the cross-sectional area of the conduction solid (m²)

h is the heat transfer coefficient (W/m²K)

A_{surface} is the total area of the heat transfer surface exposed to the liquid coolant (m²).

Most modern artificial intelligence (AI) processor cold plate designers focus heavily on reducing the

Austin M. Shelnuttt is CEO of Strategic Thermal Labs in Georgetown, Texas. David Quirk, P.E., is CEO of DLB Associates in Neptune, N.J.

convection heat transfer resistance ($R_{\text{convection}}$), which means looking for ways to increase either the heat transfer coefficient (h), the surface area (A_{surface}) or both when possible. For forced convection heat transfer using a given fluid, flow velocity is the primary lever for increasing heat transfer coefficient (HTC). Increasing flow velocity results in an increase in HTC. Unfortunately, increasing flow velocity has three primary drawbacks:

1. Increasing the cold plate pressure drop at a quadratic rate:

$$\Delta P \sim \Delta v^2 \quad (5)$$

2. Typically (though not always) increasing the volumetric flow rate of the fluid. A larger volumetric flow rate of the fluid means a lower temperature rise of the fluid for a given heat dissipation and consequently less efficient operation for facility infrastructure:

$$\Delta T = \frac{\dot{q}}{\rho Q C_p} \quad (6)$$

3. The HTC increases at a rate that is proportional to the square root of velocity:

$$\Delta h \sim \sqrt{\Delta v} \quad (7)$$

where (Equations 5–7):

ΔP is the pressure drop across the cold plate (Pa)

v is the velocity of fluid traveling through the cold plate (m/s)

ΔT is the temperature drop from heat source to heat sink for the segment of heat transfer being analyzed (°C)

h is the heat transfer coefficient (W/m²K)

$\dot{q}$ is the rate of heat being transferred (W)

Q is the total volumetric flow rate (m³/s)


ρ is the density of the fluid (kg/m³)

C_p is the specific heat capacity of the liquid at constant pressure (J/kg-K)

On the flip side, increasing surface area results in a linear decrease in convection thermal resistance without requiring additional flow, up until the point where the conductive resistance or caloric (advective) resistance of the cold plate becomes the limit (aka caloric limit).

While the possible combinations of geometric characteristics for processor cold plates are limitless, the overwhelming majority of high-performance cold

FIGURE 1 Depiction of the physical parameters of a liquid-cooled cold plate.



plates rely upon rectangular, parallel fin arrays with hydraulic diameters less than 1 mm (0.04 in.), also known as microchannels. The reason this configuration is so popular is that tightly-packed rectangular fins provide extremely high surface area to volume ratios, low pressure drop and are easy to manufacture by comparison to many other possible geometries.

As shown in *Figure 1*, the primary design features for a microchannel cold plate are:

- Fin height;
- Fin spacing; and
- Fin thickness.

Figures 2 and 3 compare the impacts of changing the primary microchannel attributes on total heat dissipation capacity (fixed temperature rise divided by thermal resistance) and pressure drop.

Decreasing the fin spacing on a microchannel cold plate results in an increase in the total heat transfer surface area, corresponding to a reduction in the convection thermal resistance, but with the unfortunate penalty of increased pressure drop due to the loss in overall flow area (increased channel velocity). *Figure 4* illustrates this relationship.

In summary, increasing the heat transfer coefficient by way of flow velocity increase results in parabolic increases in pressure drop for a square root improvement in HTC. Increasing the surface area of a cold plate by decreasing fin spacing has a similar effect on pressure drop increase but provides a linear improvement in convection heat transfer. This leads to a very important rule for cold plate design: within practical design limits, surface area always wins. In

FIGURE 2 Cold plate heat dissipation capacity as a function of fin spacing and height (50 mm × 50 mm base with 3 mm height).





FIGURE 3 Cold plate pressure drop as a function of fin spacing and height (50 mm × 50 mm base with 3 mm height).



most cases, if the cold plate designer trades off surface area to try and increase flow velocity, the decision will negatively impact performance for a given flow rate.

Beyond cold plate thermal resistance, two additional performance metrics should be considered in cold plate design:

- Caloric (or advective) efficiency (CE); and
- Coefficient of performance (COP).

Caloric Efficiency


Caloric efficiency (CE) is a rarely considered but incredibly powerful term that evaluates how effective a convection heat transfer device is at using the sensible cooling potential of a given mass flow rate. This ratio considers the temperature rise of the coolant vs. the temperature rise of the convective surface by dividing the caloric resistance by the convection heat transfer resistance:

$$R_{\text{caloric}} = \frac{1}{\dot{m}C_p} \quad (8)$$

$$\eta_{\text{caloric}} = \frac{R_{\text{caloric}}}{R_{\text{convection}}} \times 100 \quad (9)$$

Using this relationship, a cold plate with a caloric efficiency of 100% would indicate that the exiting fluid is the same temperature as the surface being cooled and is capable of providing no further cooling work at that flow rate. It is operating at the caloric limit of the

FIGURE 4 Microchannel cold plate surface area and pressure drop vs. fin spacing (50 mm × 50 mm base with 3 mm height).



coolant flow rate. Said different differently, at 100% CE:

$$T_{\text{return}} - T_{\text{supply}} = T_{\text{cold plate surface}} - T_{\text{supply}} \quad (10)$$

$$T_{\text{cold plate surface}} = T_{\text{return}} \quad (11)$$

Increasing the surface area of a cold plate by decreasing fin spacing causes the caloric efficiency to increase (Figure 5).

Cold plates that have high CE require less flow rate and create larger fluid-side temperature rises (better facility cooling efficiency) than cold plates that have low CE. A cold plate that has low CE would require much higher coolant flow rates to achieve a comparable thermal resistance. Higher coolant flow rates require more facility pumping power, produce lower return water temperature and reduce the efficiency of facility cooling infrastructure.

FIGURE 5 CE optimization as a function of fin spacing and height (50 mm × 50 mm base with 3 mm height).





FIGURE 6 COP as a function of fin spacing and height (50 mm × 50 mm base with 3 mm height).



Coefficient of Performance

The final term for evaluating cold plate designs is coefficient of performance (COP), which is a semi-dated and underrated term from the early days of industrial thermodynamics. The goal of COP is to determine the amount of cooling power provided divided by the amount of energy consumed (power input):

$$\text{COP} = \frac{\text{Heat Rejected (Watts)}}{\text{Energy Consumed (Watts)}} = \frac{\dot{q}_{out}}{P_{in}} \quad (12)$$

Heat rejected for a cold plate can be calculated as:

$$\text{Heat Rejected} = \dot{q}_{out} = \frac{\Delta T_{CF}}{R_{CF}} \quad (13)$$

where ΔT_{CF} is the temperature gradient from case to fluid and R_{CF} is the thermal resistance from case to fluid of the cold plate.

Energy consumed for a cold plate can be calculated as:

$$\text{Energy Consumed} = P_{in} = Q \times \Delta P \quad (14)$$

where Q is volumetric flow rate (typically m^3/s) and ΔP is differential pressure from inlet to outlet of the cold plate (typically Pa).

COP can then be expressed more simplistically as:

$$\text{COP}_{\text{cold\_plate}} = \frac{\Delta T_{CF}}{Q \Delta P R_{CF}} \quad (15)$$

Maximizing COP of a cold plate for any application should be the objective. A larger COP means that more heat can be rejected using less energy input (typically from the facility).

Using COP allows for a cold plate designer to extract meaning from the parametric trade-offs of flow rate, pressure drop and thermal resistance. Decreasing fin spacing of a microchannel cold plate results in significant improvements in heat dissipation capacity, but the increases in pressure drop result in a net decrease in COP (Figure 6).

Finally, COP is furthermore useful in data center liquid cooling conversations because it can easily be converted into the more commonly expressed term partial power utilization effectiveness (pPUE).

$$\text{pPUE}_{\text{cold\_plate}} = 1 + \frac{1}{\text{COP}_{\text{cold\_plate}}} \quad (16)$$

Or

$$\text{pPUE}_{\text{cold\_plate}} = 1 + \frac{Q \Delta P R_{CF}}{\Delta T_{CF}} \quad (17)$$

Cold Plate Design Metric Summary

The intent of this column is to provide valuable insight into how to use three unique heat transfer metrics to design and optimize processor cold plates. Table 1 summarizes these metrics, which are thermal resistance, caloric efficiency and coefficient of performance.

Cold plate design choices not only affect processor

TABLE 1 Cold plate design metrics.

METRIC	PERFORMATIVE DEFINITION	DESIGN GOAL
Thermal Resistance	How much heat the cold plate can remove at a given temperature rise	Reduce resistance (both conduction and convection)
Caloric Efficiency (CE)	How efficient the cold plate is at using the coolant flow rate	Increase CE
Coefficient of Performance (COP)	The amount of heat rejected per unit of energy consumed	Increase COP

thermal resistance but also influence facility energy performance. Low CE designs require higher mass flow rates, driving up pump energy and lowering return water temperatures. This reduces the efficiency of chillers, heat exchangers and free cooling. By contrast, designs with higher CE and optimized COP improve overall data center partial PUE (pPUE), linking cold plate design directly to operator energy costs.

The design performance calculated using the parameters and equations in this column can only be obtained if the technology cooling system (TCS) loop operates as designed. The most significant

components of this desired operation are associated with maintaining high fluid quality and avoiding issues in the TCS loop that could result in corrosion, erosion or biofouling. Part 3 of this series will start with a discussion of materials compatibility and continue with the commissioning (Cx) and maintenance of the liquid cooling systems serving cold plates, all to help ensure long-term cooling performance.

References

1. Demetriou, D., D. Quirk, 2025. "Liquid Cooling Cold Plates, Part 1: History." *ASHRAE Journal* (9). ■