DF Tech Note 2604-01v10

Not Just a Heat Spreader: Diamond Enables 100x Cooling with Dramatically Less Water

Thermally in a high-power chip package, it turns out diamond integrated as a substrate at the chip level is not only effective as a heat spreader – but by redistributing the total thermal resistance stack and the temperature drops across it, it uniquely enables evaporative cooling directly at the chip level. This can deliver 10–100x better cooling performance and up to 55x less water flow required.

Cooling performance is not determined only by total thermal resistance but by how that resistance is distributed:

- In conventional designs, most of the temperature drop occurs in the semiconductor substrate. The backside interface to be cooled is already relatively cold.

- In contrast, the integration of diamond as a substrate with thinned silicon reduces upstream thermal resistance and shifts a larger fraction of the temperature drop downstream to be available for higher performance cooling. By making available an interface with higher temperature at the backside of the chip, evaporative cooling is enabled.

Evaporative cooling with (two-phase) water can enable orders-of-magnitude increases (10–100×) in heat transfer while reducing liquid flow requirements by up to three orders of magnitude under appropriate operating conditions.

It has not been used at the chip level in AI data centers because silicon wafers and other heat spreader materials offer low-temperature surfaces that require complex and often impractical engineering tradeoffs, in particular deep sub-atmospheric pressures, to enable evaporation with a state-of-the-art thermal resistance stack.

Diamond integrated at the chip level redistributes the total system resistance stack and makes available a hotter backside surface at a temperature readily suitable for two-phase cooling in ways that were not practical and/or not efficient before, whether in the form of jet impingement, spray, or passive porous surfaces.

Further just like the diamond is very effective at spreading heat from hotspots on the transistor facing side, it is very effective at spreading heat from hotspots that may occur (and in prior work not using diamond have caused issues with) localized dryouts, etc. on the cooling facing side.

By enabling near-atmospheric pressures for two-phase water, the vapor volumes become manageable too– a key issue for fitting piping, etc. into the required 1U packaging form factor.

Effectively commonplace facility level evaporative cooling can now be brought directly to chips.

It is also a solution that readily handles silicon die warpage issues especially in larger packages – an issue that requires conventional coldplate solutions to make use of lots of really bad TIM interlayers.

Temperatures in silicon chips are generally managed to not exceed 105ºC at the hottest spot and not exceed 85ºC on average across the chip. The former is linked to chip failure whereas the latter is linked to managing leakage loss. For a class W4 datacenter with 45ºC inlet liquid temperature at the facility level (and practically 5-7ºC higher once it reaches the racks), this offers a delta of 60K for single-phase liquid cooling. Diamond substrated chips shift the temperature distribution from 50-65ºC to 80-100ºC, enabling different approaches to cooling that have distinct advantages.

Standard coldplate cooled silicon: Big temperature drop yields a backside too cold for phase change

For starters, let us understand the total thermal resistance stack of a *state-of-the-art liquid cooled high-power silicon chip:*

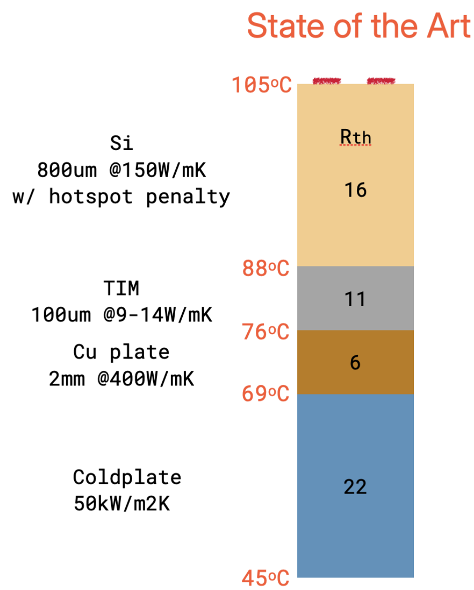

From bottom to top, we have a reference stack adding up to 55 mm^2K/W of thermal resistance corresponding to a max power possible given a max temperature delta of 60K for a reticle-area sized chip.

- A coldplate with lamella etc. to efficiently transfer heat from copper into liquid may perform at $50kW/m^2K$ and thus have a thermal resistance for a reticle-limit sized, 858mm^2 chip of $1/(50kW/m^2K)=20$ mm^2K/W thermal resistance (22 with hotspot penality);
- A copper plate part of the coldplate that may be 1-3mm thick, with Cu having a resistance of 400W/mK, thus a thermal resistance added of 2mm / 400W/mK = 5 mm^2K/W (6 with hotspot penalty);
- A Thermal Interface Material (TIM) of a poor conductivity of 8-14W/mK that on top has to be a fairly sizable 100um thick to handle the warpage of a silicon chip and the flatness of the copper plate of the coldplate – adding 100um / 14W/m·K = 7 mm^2K/W in thermal resistance (11 with hotspot penalty);
- The silicon at 800um thickness and 150W/m·K conductivity adding 800um/150W/m·K = 5.3 mm^2K/W in thermal resistance (16 with hotspot penalty determined with precise 3D thermal simulations – see Figures below – in case of an nvidia H100 type GPU tensor core layout).

The resultant temperature waterfall is annotated in red in the above chart too: The 60-degree delta spreads in proportion to the respective effective thermal resistances across the stack. So by example the coldplate with 45ºC liquid according to a W4 datacenter classification sees a 60/52*20 = 23 degrees higher temperature towards the next layer in the stack.

Notably, in this total thermal resistance stack, the temperatures at the bottomside of the silicon is 88ºC and at the backside of the Cu plate it is just 69ºC.

At those temperatures, water does not evaporate – not at atmospheric pressure and only at deep sub-atmospheric pressures. This is fundamentally why commonly just single-phase (water) liquid is used in the state of the art.

For achieving two-phase cooling with its 10-100x higher cooling capacity – see below – one would need to lower the pressure to have it evaporate, to below 0.3 bar:

Pressure (bar abs)	Pressure (kPa)	Saturation temp (°C)	Latent heat h_fg (kJ/kg)	Vapor density (kg/m³)
0.10	10.0	45.8	2393	0.068
0.15	15.0	54.0	2373	0.099
0.20	20.0	60.1	2358	0.130
0.25	25.0	65.0	2346	0.160
0.30	30.0	69.1	2336	0.191
0.40	40.0	75.9	2319	0.249
0.50	50.0	81.3	2305	0.307
0.58	58.0	85.0	2295	0.351
0.70	70.0	90.0	2283	0.422
1.013	101.3	100.0	2257	0.598

Table: Water phase transition temperatures at various pressures. Resultant latent heat and vapor density.

Compare diamond chip: Because upstream resistance drops, backside temperature is *higher*

The thermal performance of single-crystal diamond (SCD) is well known – being 16x over Si and ~6x over SiC, Cu, AlN, Ag, Au, etc. This thermal factor determines the factor by which power can be scaled in a chip with a well designed hotspot distribution. But how this plays a role as part of the total thermal resistance stack in a CoWoS or EMIB packaged chip is generally not as well understood.

The junction-to-coolant temperature difference (ΔT_jc) available to drive evaporation scales with how conductive the substrate is. Silicon has $k \approx 150$ W/m·K; diamond at 2,200 W/m·K means the same power density requires roughly 15× less ΔT to conduct the same flux through the same thickness. That leaves more temperature budget for the evaporative process itself, which translates directly to higher critical heat flow (CHF) before the surface exceeds the critical superheat.

Integrating diamond at the chip level goes along with reducing the semiconductor thickness from ~800 µm to <30 µm to further dramatically lower thermal resistance. For silicon thinned to 20um and substrated by 600um of SCD, the combined thermal resistance is just 0.4 mm^{2}K/W, with strong hotspot spreading yielding less than 1mm^{2}K/W in effective thermal resistance.

This means even for a 2500W reticle-area sized chip, the temperature drop for this Si/SCD stack is only 2500W/858*1mm^{2}K/W = 3 degrees Kelvin, meaning if the backside of the SCD substrate will have a 100ºC or higher max temperature, it will allow for water to evaporate at >0.7 bar or higher.

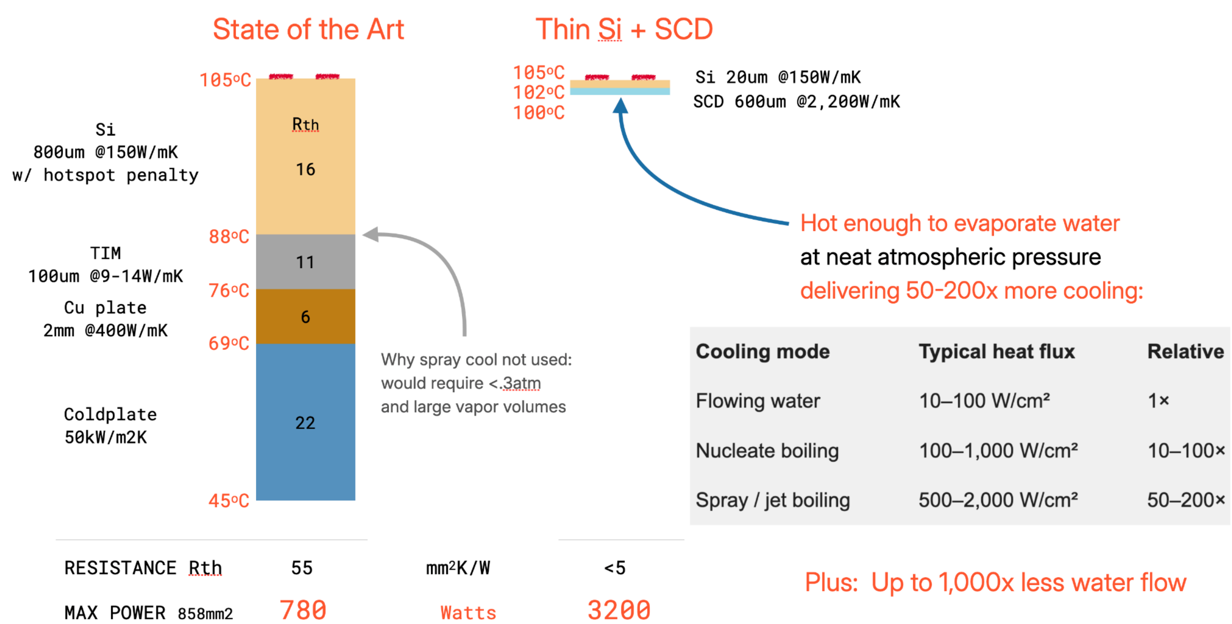

Cooling mode	Typical heat flux	Relative
Flowing water	10–100 W/cm²	1×
Nucleate boiling	100–1,000 W/cm²	10–100×
Spray / jet boiling	500–2,000 W/cm²	50–200×

Figure: The temperatures for state-of-the-art versus diamond based chips are 69ºC versus 100ºC as available to cooling liquids. More than double the power level is supported by the diamond substrated chip while maintaining the same max hotspot temperature.

A pressure of 0.7 bar is equivalent to the atmosphere at 3,000m altitude. Engineering is readily available to support such slightly lower pressures. (Or some may choose to build a data center at 3,000m altitude on Earth…)

The backside of the diamond reaches temperatures evaporating water even at standard atmospheric pressure for a good percentage of the surface area of the chip. On average it is 87ºC, which is very close to the commonly targeted average temperature for managing chip leakage losses.

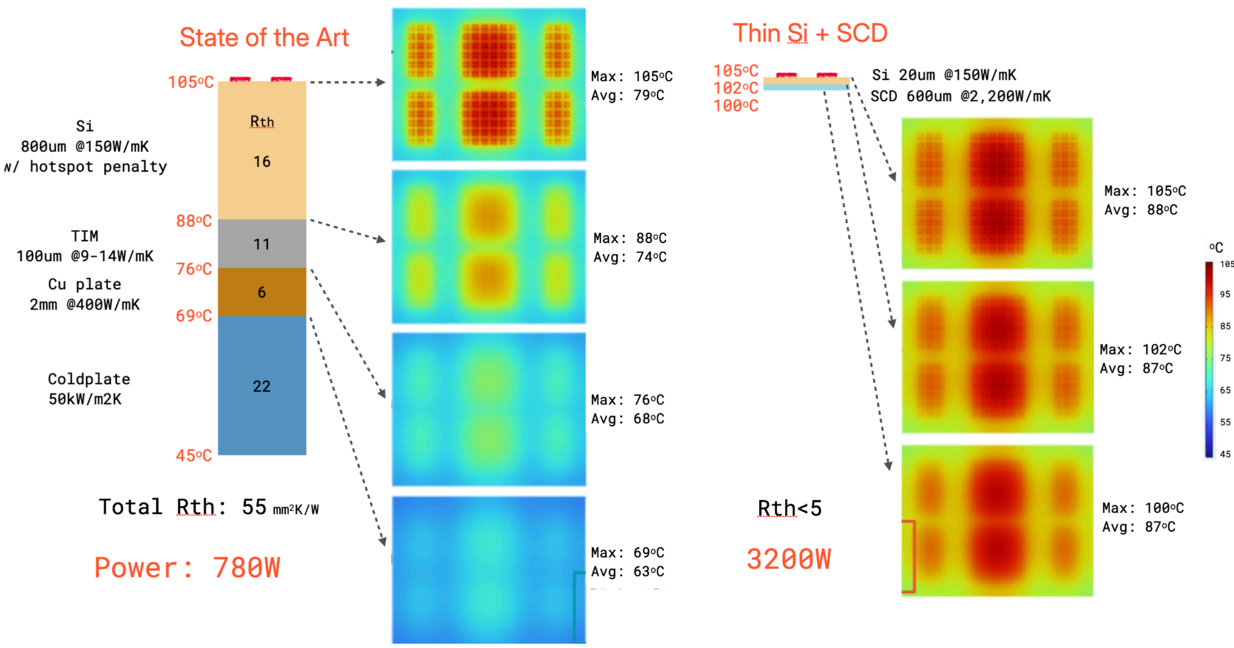

Figure: Thermal maps at each of the interfaces of the total system stack resulting with an nvidia H100 inspired tensor core thermal map at max power – for a 780W standard versus a 3,200W diamond chip each powered maximally to for hottest spot not to exceed a max temperature. (The heat transfer coefficient (HTC) for the diamond substrated chip is assumed to be 100 kW/m²K.)

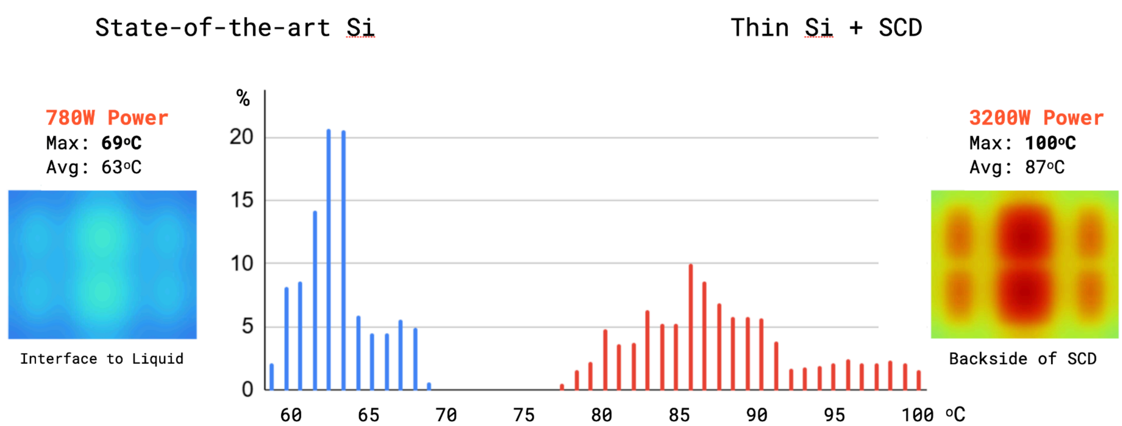

Figure: Backside surface temperature distributions for a standard silicon chip versus thinned silicon substrate with diamond.

The key insight is that, at chip packaging scale, diamond does not just spread heat — it changes the waterfall of interface temperatures to uniquely enable the use of the most powerful heat sinking mechanism. In a nutshell:

ΔT_interface = Q × R_interface.

If R_upstream ↓ → ΔT_interface ↑

Diamond does not just reduce thermal resistance — it shifts where the temperature drop occurs, enabling phase change at the cooling interface.

The Power of Two-Phase

Evaporative (boiling) cooling can remove ~10× to 100× more heat per unit area than single-phase liquid water flow over the same surface—if it's engineered well.

In flowing liquid (single-phase cooling), heat removal is limited by the heat transfer coefficient (HTC): typically 1–3.5 $W/cm^2 \cdot K$ and the temperature difference (ΔT).

In evaporative cooling (two-phase, boiling), now one taps into Latent Heat of Vaporization, which for water is ~2260 J/g absorbed during phase change (at 100°C). That's huge—it is orders of magnitude more energy than raising water temperature. It enables heat transfer coefficients of 1–10+ $W/cm^2 \cdot K$ and heat fluxes of ~0.1–1+ kW/cm^2 (and even higher in optimized systems).

Cooling mode	Typical heat flux	Relative performance
Flowing water	10–100 W/cm^2	1×
Nucleate boiling	100–1,000 W/cm^2	10–100×
Spray / jet boiling	500–2,000 W/cm^2	50–200×

Boiling is so powerful because energy goes into phase change, not just temperature rise; vapor bubbles disrupt boundary layers for much higher heat transfer; and liquid continuously re-wets the surface, maintaining efficient heat transfer.

This is why spray cooling + diamond substrates is so powerful: One can keep junction temps low while removing extreme heat densities. It is not just hotspot heat spreading; it is enabling much better heat sinking.

Furthermore, in case of localized drying out of some spray cool surface areas, as it may occur, the diamond acts reversely as a heat spreader on that as well, increasing the operating process windows.

There are a host of ways of implementing evaporative cooling directly at the chip level, each operated with slightly sub-atmospheric pressure, including jet injection, spray cooling, and porous membranes that work passively with a capillary effect.

Diamond + thin-film water is a powerful combination

The use of diamond fundamentally redistributes the thermal resistance stack, elevating the cooling interface temperature and enabling efficient two-phase spray cooling using water at near-atmospheric pressures. The result is a step-function improvement in cooling capability, enabling substantially higher heat flux removal while dramatically reducing liquid flow requirements.

Spray cooling directly against a diamond surface is very effective — and in some configurations one of the most powerful cooling approaches available for high heat-flux applications.

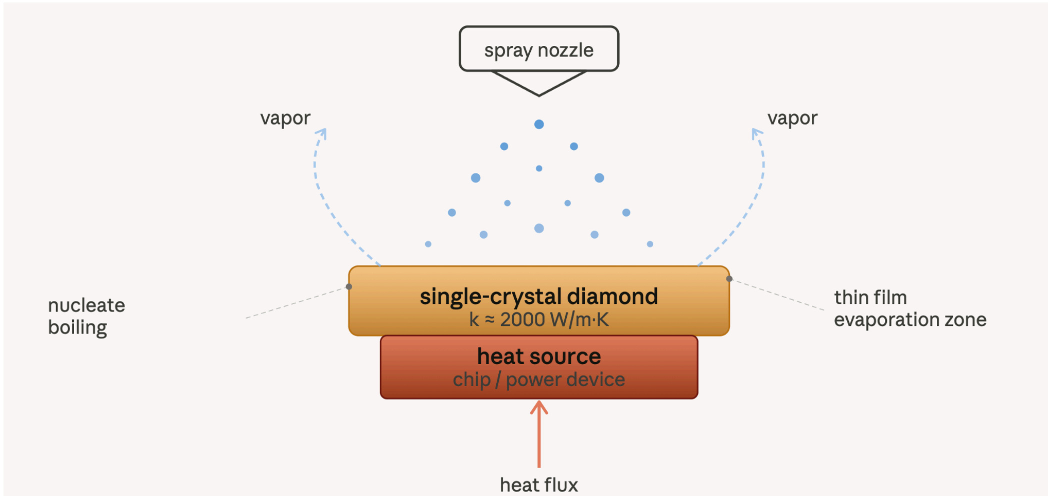

Spray versus jet impingement

Here's the comparison laid out as numbers first, then the mechanism behind why they diverge.

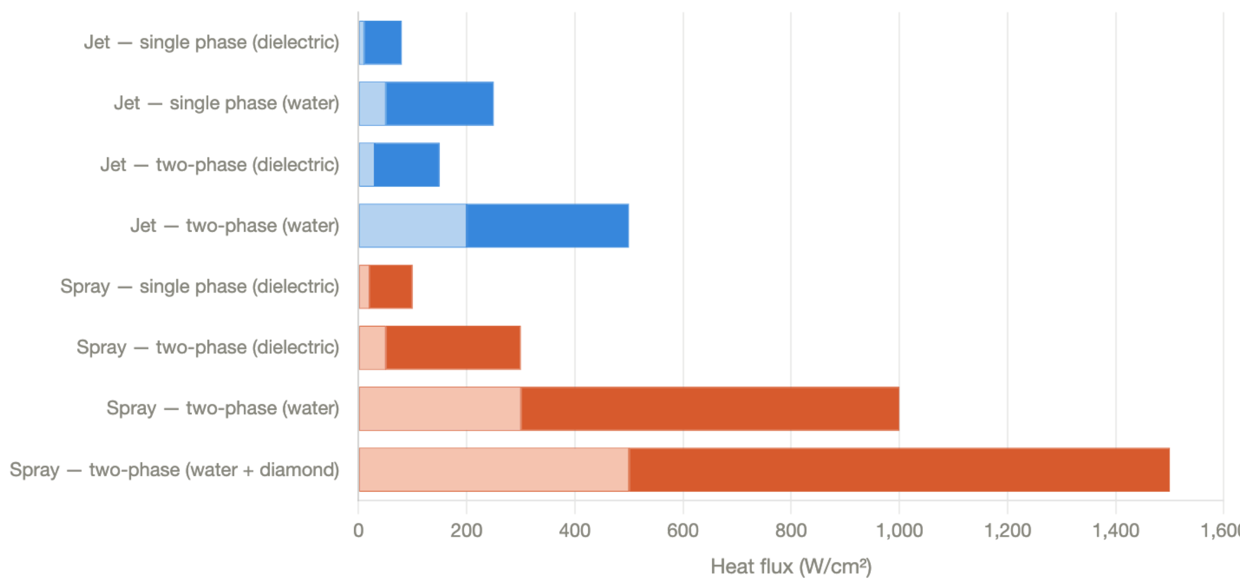

The numbers tell a clear story, but understanding why spray outperforms jet at the top end requires looking at the mechanisms.

Where jet impingement excels

A submerged or confined jet delivers a very high local HTC right at the stagnation point — typically 2–8 $W/cm^2 \cdot K$ for water — because the thin hydrodynamic boundary layer at impingement is extremely effective. The problem is that this advantage is spatially concentrated. As the liquid spreads radially outward (the "wall jet" region), the boundary layer thickens rapidly and HTC drops off sharply, often by 50–70% within 2–3 jet diameters. For a diamond tile, this creates a lateral temperature gradient that the diamond spreads remarkably well, but it does mean the jet isn't extracting heat uniformly across the footprint.

Two-phase jet impingement (operating near or in nucleate boiling) pushes CHF to 200–500 W/cm² for water, which is excellent. But the CHF limit for a submerged jet tends to arrive as a vapor crown that shields the stagnation point — and once it forms, the thermal resistance spikes suddenly. It is a less graceful failure mode than spray.

Where spray gains the advantage

Spray cooling produces a distributed thin-film evaporation regime across the entire wetted surface rather than a single impingement zone. The key advantages:

- The droplet kinetic energy continuously disrupts the vapor boundary layer, preventing the stable vapor blanket that limits jet CHF
- Thin-film evaporation at the three-phase contact line (liquid-vapor-solid) is the most efficient heat removal mechanism known, with local HTCs potentially exceeding 10 kW/cm²·K in the evaporating meniscus
- Because the mechanism is distributed, the spatial uniformity is much better — important for diamond tiles where one wants the full surface contributing
- CHF for water spray sits in the 0.5–1+ kW/cm² range depending on spray parameters, and with diamond spreading underneath it one is effectively raising the apparent CHF further because the diamond prevents the local surface temperature from hitting the threshold that triggers film boiling.

The practical operating window for spray in the nucleate/thin-film regime is also wider — one can maintain relatively stable surface temperatures (ΔT of 20–40°C above saturation) across a broad flux range, whereas jet impingement in two-phase mode can be harder to control near CHF.

The diamond multiplier

For both methods, diamond substrate changes the comparison in a specific way: it shifts the bottleneck away from the surface-to-coolant interface and back toward the coolant's intrinsic CHF. A bare silicon die might hit CHF at a local hotspot while most of the surface is still cool; a diamond spreader ensures the spray/jet sees a nearly isothermal surface, so the coolant's CHF limit is reached much more uniformly across the whole footprint. This effect is proportionally larger for jet (which has strong spatial non-uniformity) than for spray (which is already more spatially uniform), so diamond gives jet a bigger boost in relative terms — but spray on diamond still leads in absolute heat flux capacity.

Diamond does not just passively receive the water mist spray — it actively participates. Its extreme thermal conductivity (~2,200 W/m·K) spreads the heat laterally before the spray even touches the surface. This means:

- The temperature distribution across the spray footprint is much more uniform than it would be on silicon, copper, or SiC. Hotspots are blunted before they reach the spray interface.
- The spray sees a lower peak temperature, which keeps us further from the critical heat flux (CHF) limit and prevents premature transition to film boiling.
- Heat fluxes that would trigger boiling crisis on conventional substrates can be managed on diamond because the spreading is doing real work.

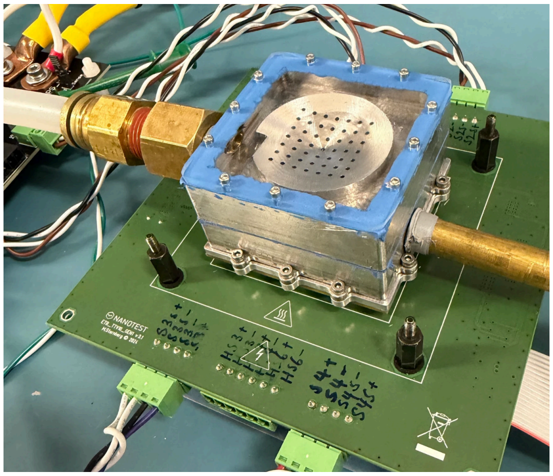

Photo: DF’s San Jose, CA chip lab is testing advanced chip cooling using thermal test vehicles.

Deep Sub-Atmospheric (~0.25 bar) versus Near-Atmospheric (0.7 bar)

Going to 0.25 bar (65°C saturation) is an interesting operating point — low enough that junction temperatures become very attractive for silicon longevity, but the sub-atmospheric penalties discussed above compound significantly. Here's how each effect scales from our 0.7 bar baseline.

Latent heat drop — still negligible. h_fg at 0.25 bar is 2346 J/g versus 2283 at 0.7 bar. A 2.7% improvement actually — one gets marginally more cooling per gram evaporated. Still essentially a non-issue in either direction.

Vapor volumetric flow — now severe. Vapor density at 0.25 bar is 0.160 kg/m³ versus 0.422 at 0.7 bar. That's 62% less dense — so for the same 500 W chip load our vapor handling system must move 2.6× the volume it handles at 0.7 bar, and 4.3× of what it would handle at atmospheric. Vapor ducts, plenum sizing, and pump inlet area all scale with this. It becomes a dominant design constraint.

Operating point	Inlet Pressure	Discharge (~1 bar)	Pressure Ratio
1.013 bar (ref)	1.013 bar	1.013 bar	1.0
0.7 bar	0.7 bar	~1.0 bar	1.4
0.25 bar	0.25 bar	~1.0 bar	4.0

Table: Pump compression ratio – jumps hard.

The compression ratio of 4.0 at 0.25 bar is outside the comfortable single-stage range for a water-ring pump. Standard water-ring pumps operate efficiently up to pressure ratios of about 2.5–3.0 in a single stage. Beyond that, one either needs a two-stage configuration or to accept sharply degraded volumetric efficiency as internal slip (vapor leaking back from high-pressure to low-pressure side) increases. Parasitic pump power for the same 500W load could reach 60–120W — a 12–24% system overhead versus 3–8% at 0.58 bar. This is where 0.25 bar starts to feel genuinely costly from a system efficiency standpoint.

Non-condensable gas — much more problematic. Water's vapor pressure at boiling is equal to the operating pressure, per definition of boiling. Trying to achieve 65C boiling requires a pressure that's only

25 kPa (or 0.25 bar), i.e. 25% of ambient pressure. If air leaks into the system at even a slow rate, its partial pressure builds relative to a very low total system pressure. A non-condensable fraction that would be tolerable at 0.7 bar becomes debilitating here. At 0.25 bar total, just 2.5 kPa of dissolved/ingressed air represents 10% of the total pressure — enough to suppress condenser HTC by 40–60% and cause progressive system degradation. Degassing becomes not just good practice but a continuous operational requirement. One would likely need either a permanent small purge flow or an inline non-condensable vent on the condenser.

Clausius-Clapeyron slope — now a real design input. dP/dT at 0.25 bar is approximately 0.90 kPa/°C, versus 1.55 at 0.58 bar and 3.62 at 1 bar and even better at 0.7 bar. This means the evaporating film needs nearly twice the local superheat at 0.25 bar to drive the same evaporative flux as at 0.7 bar. For thin-film evaporation on diamond, where one is operating with only a few degrees of superheat above saturation, this noticeably narrows our usable ΔT window before the film either goes subcooled (no evaporation) or exceeds nucleate boiling onset (transitions away from thin-film regime).

Atomizing spray nozzle behaviour. This is specific to our system design. At 0.25 bar ambient pressure the spray droplets evaporate more aggressively in flight before reaching the surface — flash evaporation begins as soon as the droplet exits the nozzle into the low-pressure environment. Depending on standoff distance and droplet size, one may lose a meaningful fraction of our liquid to pre-surface evaporation rather than surface thin-film evaporation. At 0.7 bar this effect is minor; at 0.25 bar with a 65°C saturation temperature, any droplet whose surface temperature exceeds 65°C during flight will begin flashing. Nozzle-to-surface distance would need to be minimised, or the liquid supply temperature kept several degrees below saturation (subcooled feed) to suppress in-flight evaporation.

Condenser sizing. The condenser must reject heat at 65°C rather than 85°C. If your ultimate heat sink is ambient air or facility cooling water at say 25–40°C, the available ΔT for condenser heat rejection actually increases slightly going to 0.25 bar (65 − 40 = 25°C driving force versus 85 − 40 = 45°C at 0.7 bar). So the condenser gets harder to size not because of ΔT but because of the much larger vapor volume it must condense — the vapor inlet velocity to the condenser will be high and pressure drop through the condenser becomes a significant fraction of the total system pressure, which at 0.25 bar is very small to begin with. A 1 kPa pressure drop through a condenser is trivial at 1 bar; at 0.25 bar (25 kPa absolute) it represents 4% of total pressure and meaningfully shifts the saturation temperature at the condenser inlet versus outlet.

Parameter	**0.58 bar / 85°C**	**0.25 bar / 65°C**	**Change**
Latent heat h_fg	2295 kJ/kg	2346 kJ/kg	+2% (better)
Vapor density	0.351 kg/m³	0.160 kg/m³	−54% (worse)
Pump pressure ratio	1.72	4.0	+133% (much worse)
dP/dT slope	1.55 kPa/°C	0.90 kPa/°C	−42% (worse)
NCG sensitivity	Moderate	Severe	Qualitative jump
Junction temp benefit	Baseline	−20°C	Significant gain

Water Volume

The volumetric flow rates differ by orders of magnitude across these three approaches because the heat transfer mechanisms are fundamentally different:

Cooling Mode	Heat Removal Mechanism	Typical dT Driving Flow	Flow rate for 500W, 1cm2	Liquid System
Evaporative	Latent heat of vaporisation	~5-15°C superheat	13 mL/min	Trivially small
Microchannel two-phase	Latent heat + convection	~10-30°C	150mL/min	Intricate
Coldplate single-phase	Sensible heat only	~10°C rise	710 mL/min	Major

The numbers are stark. The evaporative system moves roughly 55× less liquid mass than a coldplate for the same heat load. This is the entire thermodynamic argument for phase-change cooling in one comparison.

Why evaporative flow is so small

The latent heat of water at 85°C is 2295 kJ/kg. To remove 500 W continuously one needs to evaporate exactly 500/2295 = 0.218 g/s of water. That is 0.218 mL/s — about 13 mL/min. For a 1 cm^2 chip that is a surface flux of 13 $mL/min/cm^2$, which is a very gentle spray or wick feed. The liquid delivery system is trivially small.

The coldplate calculation is the opposite. Single-phase water with a 10°C temperature rise carries heat via sensible capacity: $Q = \dot{m} \cdot Cp \cdot \Delta T$, so $\dot{m} = 500 / (4182 \times 10) = 0.012$ kg/s = 12 mL/s = 0.71 L/min. That requires a real pump, real flow regulation, real pipe diameter, and significant pressure drop through the coldplate channels.

Vapor volumetric flow — the evaporative catch

The liquid side of evaporative cooling is tiny, but the vapor side is not. At 0.58 bar the vapor density is 0.351 kg/m^3, so 0.218 g/s of vapor occupies $0.218/0.351 \times 1000 = 621$ mL/s = 37 L/min of vapor volume. That vapor has to go somewhere — condenser, pump inlet — and that volume flow is what sizes the vapor-side infrastructure.

System Pressure	Liquid Flow @ 500W	Vapor Volume Flow	Ratio Vapor/Liquid Vol
1.013 bar (100°C)	13 mL/min	13.1 L/min	1,600x
0.7 bar (90°C)	13 mL/min	30 L/min	2,200x
0.25 bar (65°C)	13 mL/min	78 L/min	6,000x

This ratio — the expansion ratio from liquid to vapor — is the fundamental sizing driver for the vapor-side components. It gets dramatically worse as one goes sub-atmospheric because vapor density falls while liquid density barely changes.

Practical pump sizing implications

For a 0.7 bar evaporative system targeting a 10 mm × 10 mm diamond tile at 500 W: The liquid feed pump is almost irrelevant in size — 13 mL/min at modest pressure (a few hundred mbar to overcome delivery head and nozzle pressure drop). A small peristaltic or diaphragm pump rated at 50–100 mL/min gives a 4–8× safety margin for flow control. Power consumption under 1W.

The water-ring vacuum pump handling the vapor is the dominant component. At 0.7 bar inlet, 30 L/min of vapor, with a compression ratio of 2.2, a single-stage water-ring pump in the 0.25–0.5 kW shaft power range is appropriate. This is the real parasitic load — not the liquid feed.

At 0.25 bar the vapor volume jumps to 78 L/min and the compression ratio to 6.0, pushing us toward a 0.75–1.5 kW two-stage pump for the same 500 W chip load. The pump becomes larger than the heat load it is supporting, which is the practical argument against going too deep sub-atmospheric unless junction temperature reduction is the overriding constraint.

Evaporative Cooling versus Microchannels

The architecture of microchannels is well-understood mechanically. Etched into or bonded to the back of the die, they carry fluid at pressures just below saturation. As the fluid absorbs heat, it nucleates and boils in the channels.

The problem is flow instability — two-phase flow in narrow channels creates pressure oscillations that can destabilize the boiling front and cause local dryout. This is especially pronounced at high heat flux. Parallel channel designs with inlet restrictions and diverging geometries help, but it remains a fundamental challenge.

As a result, two-phase microchannel approaches remain challenging due to flow instability, pressure oscillations, and dryout risks at high heat flux – apart from all the other clogging, fouling, and particle contamination events that microchannels face generally.

Cooling	**Evaporative / 2-Phase Spray**	**Microchannel**
	• Uses phase change of water • Heat removal dominated by latent heat of vaporization (~2257 J/g) • Can achieve very high effective heat transfer coefficients (HTC)	• Typically single-phase liquid flow • Heat removal via convective heat transfer • Limited by fluid heat capacity and flow rate
Heat Transfer Coefficient (HTC)	5–10× higher	
Max Heat Flux	2x higher	
Energy Efficiency (Pumping Power)	• Lower flow rates needed (latent heat does the work) • Pumping power relatively low • But requires vapor handling + condensation	• Requires high flow velocity + pressure drop • Pumping power can be significant • Simpler fluid loop
Operational Risk	• Surface degradation	• Clogging • Fouling • Particle contamination events

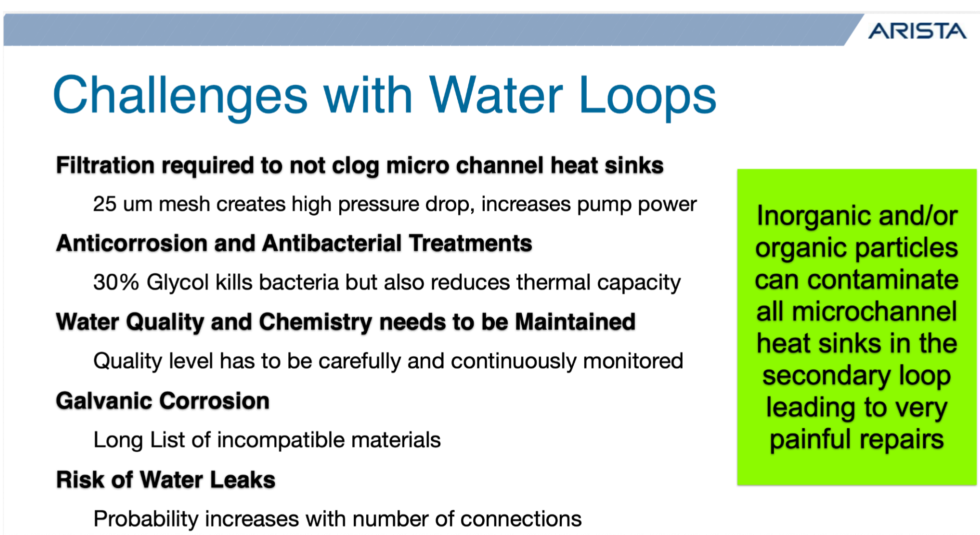

OCP 2003 Keynote Slide (Arista): Water-loop challenges for microchannels and 1-phase liquid cooling.

Summary

Diamond integration at the chip level enables a fundamentally new thermal operating regime.

Rather than only reducing thermal resistance, it reallocates the temperature drop across the thermal stack, enabling phase-change cooling at the interface.

This allows practical implementation of two-phase cooling using water at near-atmospheric pressure, unlocking:

- order-of-magnitude (10–100×) increases in heat removal capability
- significant reductions (55x) in liquid flow requirements by up to three orders of magnitude
- new scaling pathways for high-power semiconductor systems

This shift represents not an incremental improvement, but a structural change in thermal architecture, with implications for the future of high-performance computing and data center design.

Other materials with a fraction of diamond's thermal conductivity and heat spreading performance fall short of enabling this because their backside temperature is too low, requiring very low pressures to be realized and very large vapor volumes to be managed.

Evaporative cooling is commonly used in AI data centers anyways – at the facility level. We can now apply it directly at the chip level.

A broad set of approaches exist to implement thin films of water on a hot plate including but not limited to spray cooling, jet impingement, and porous membranes.