ENERGY EFFICIENCY WITH SUSTAINABLE PERFORMANCE: TECHNIQUES, TOOLS, AND BEST PRACTICES (EESP)

Towards Direct-to-Chip Liquid-Cooled Megawatt Racks: Characterization of a Multi-Kilowatt Chip-Scale Cooler & the Implications on Facility Energy Use

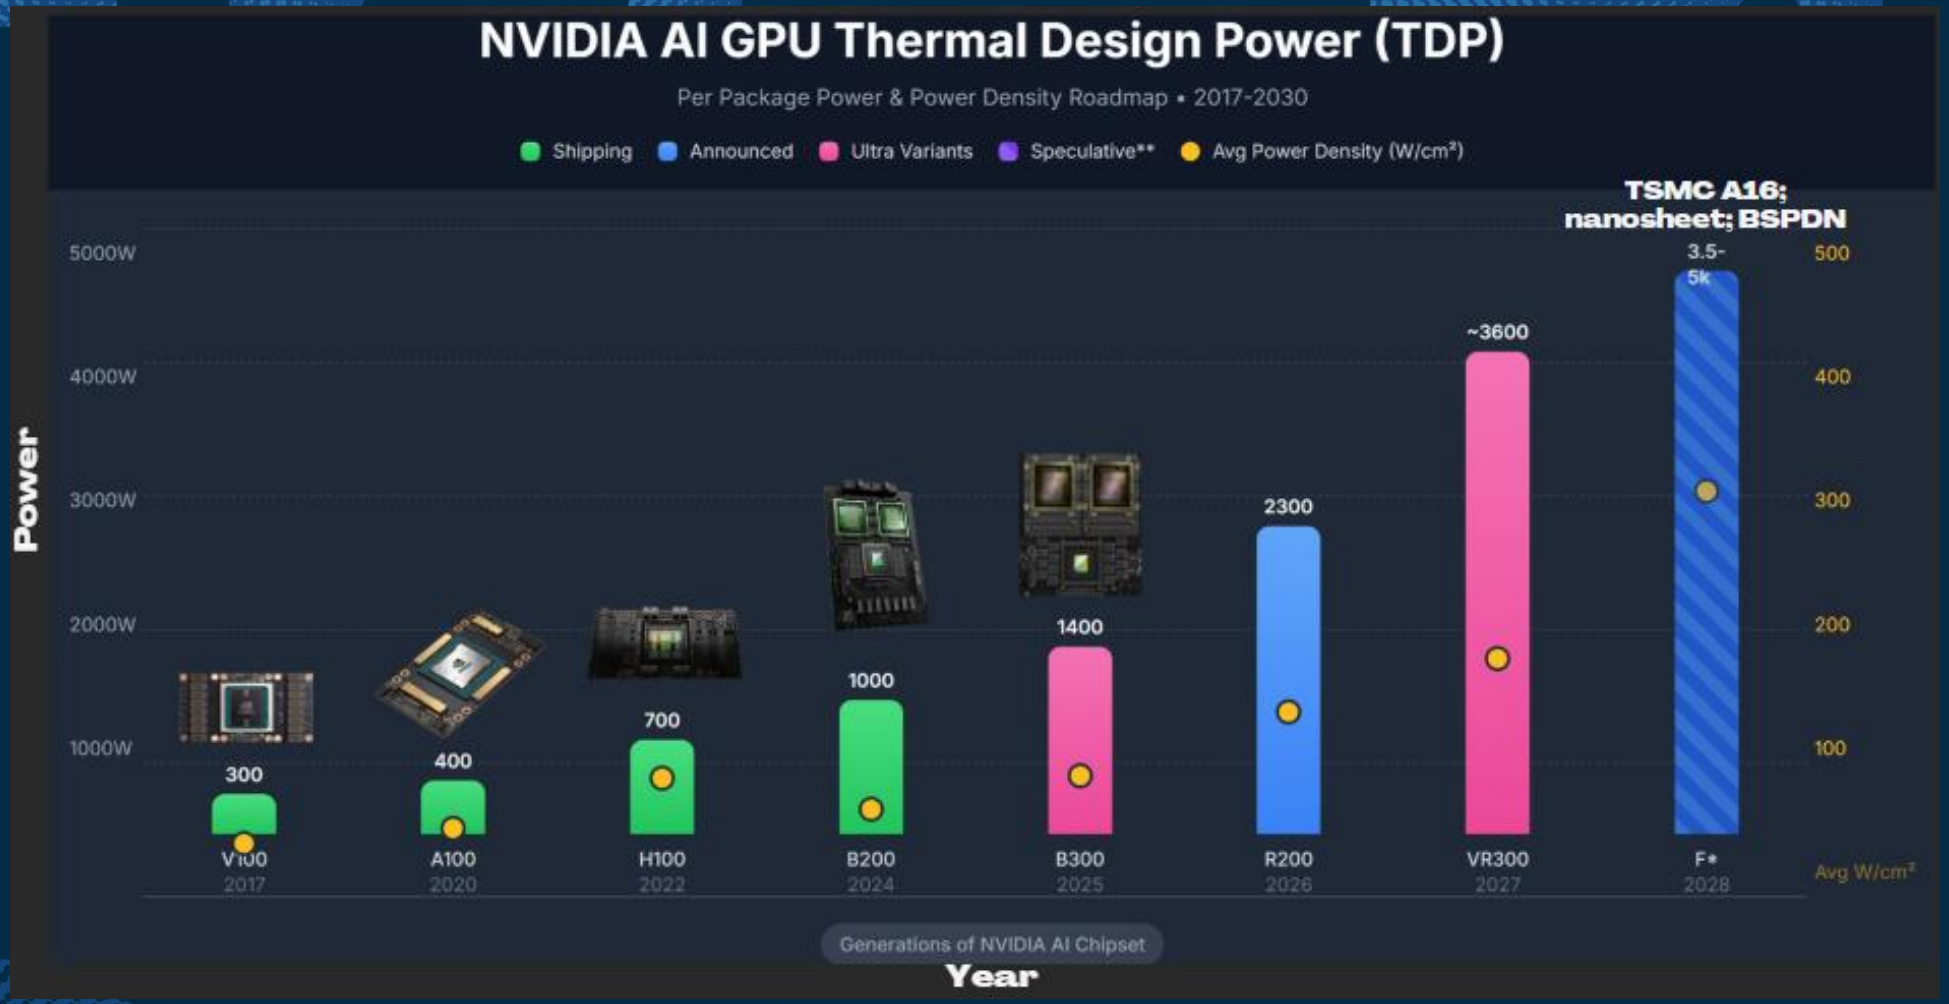
Daniel Lorenzini [a], Yogendra Joshi [a], Thomas Steen [b], David Sickinger [c], and Aaron Andersen [c]

[a] EMCOOL Inc, Atlanta, GA, USA
[b] Nortek Data Center Cooling, Oklahoma City, OK, USA
[c] National Laboratory of the Rockies, Golden, CO, USA



SC26 hpc unites | Chicago

The rising thermal wall for ML/AI GPU superchips



NVIDIA AI GPU chipset power and average power density roadmap, illustrating the growing thermal challenge for next-generation AI accelerators and the need for advanced liquid cooling solutions.

1 MW rack-scale ecosystem demonstration

Nortek
Dry Cooler



EMCOOL chip scale coolers can remove up to 5 kW per chipset



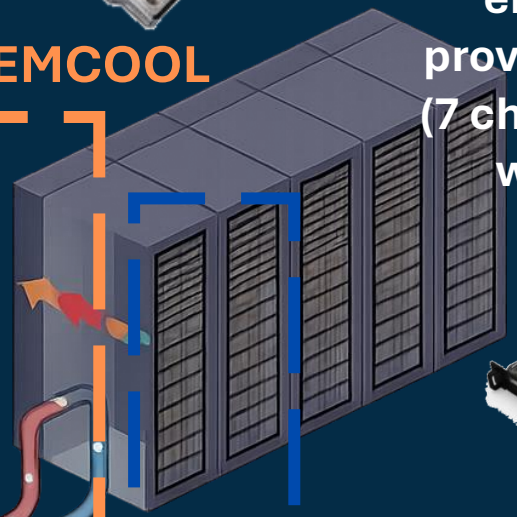
Thermal Pixels server emulator TTVs can provide 28 kW per blade (7 chipsets at 4 KW each within 1 U blade).



Nortek
CDU



EMCOOL

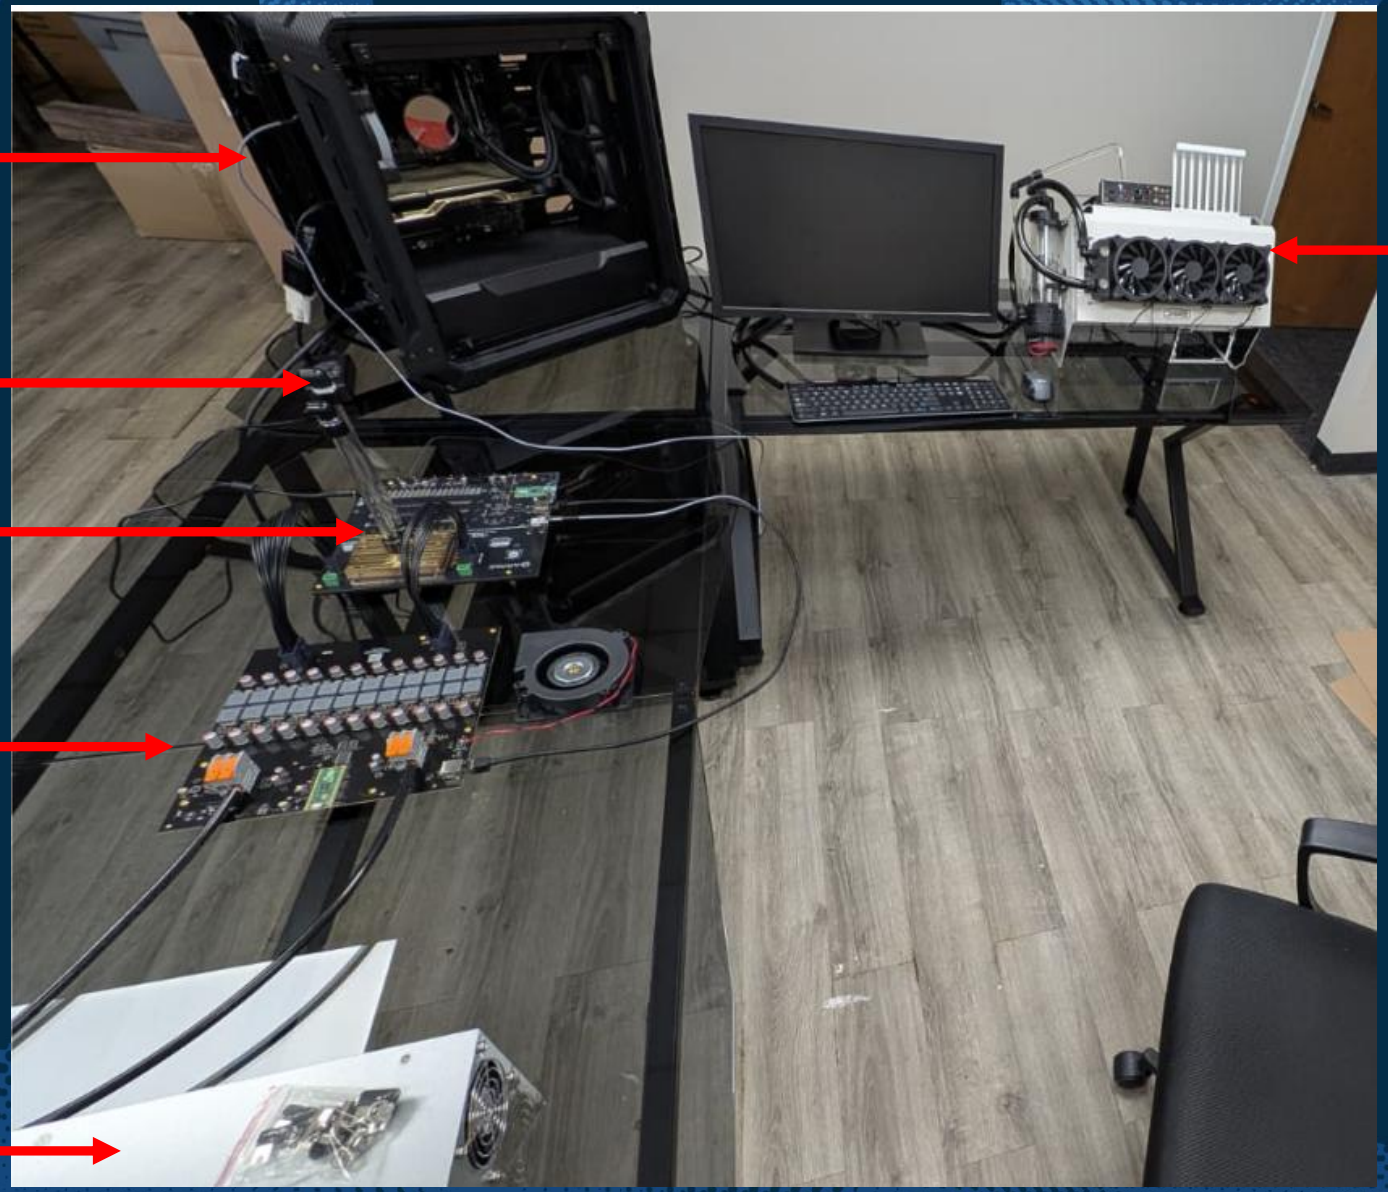


Thermal Pixels TTVs

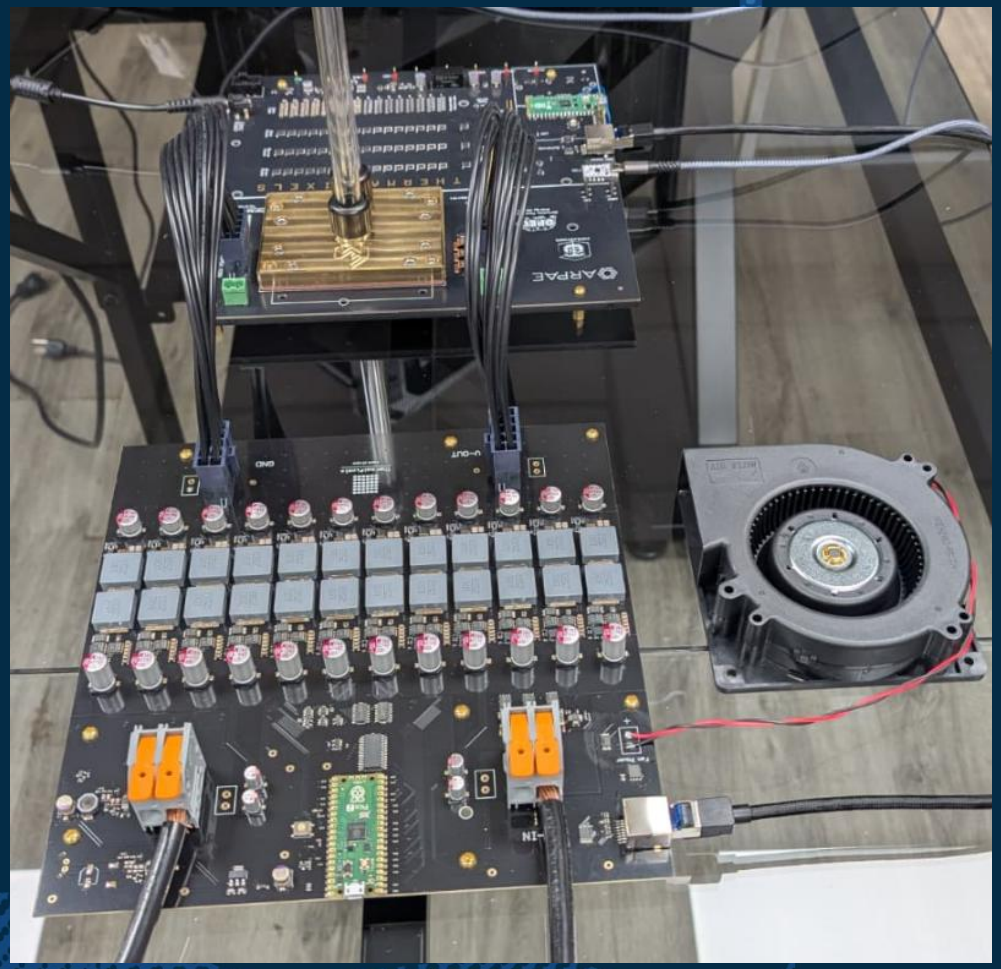
Proposed 1 MW rack-scale ecosystem demonstration integrating Thermal Pixel TTVs, EMCOOL chip-level thermal management, Nortek CDU, and rack-to-ambient heat rejection.

SC26

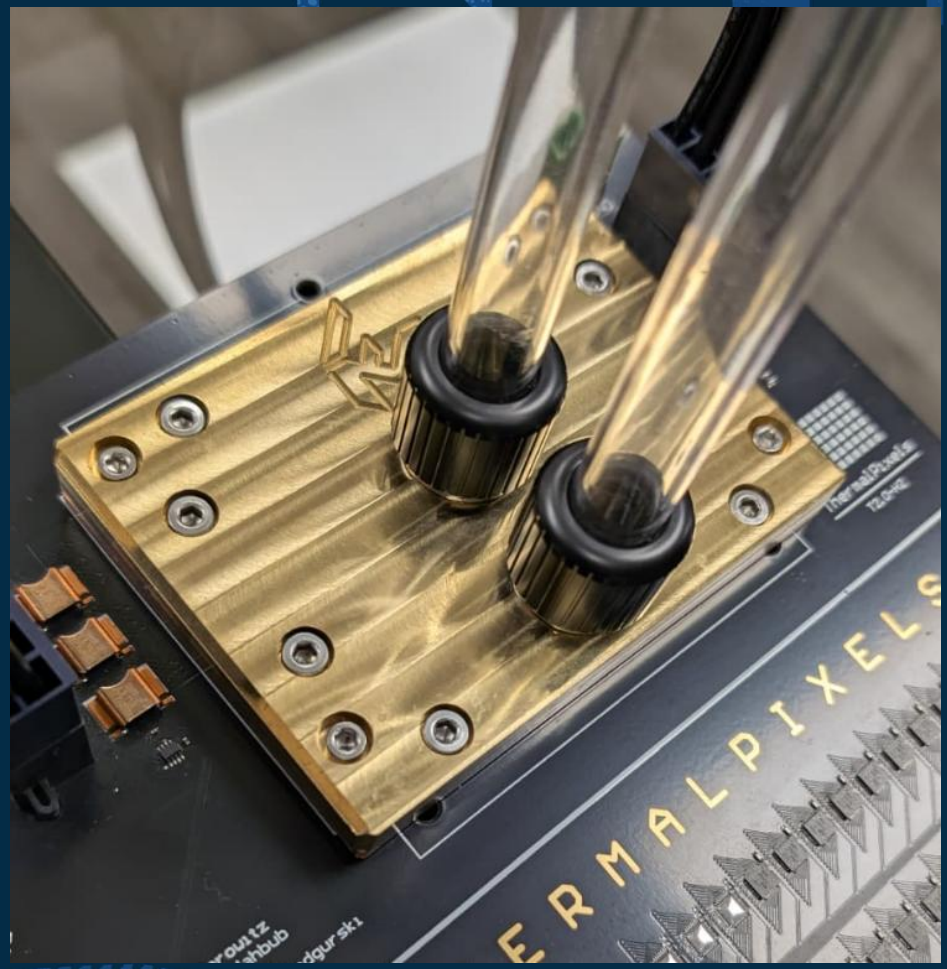
Experimental setup



Experimental setup

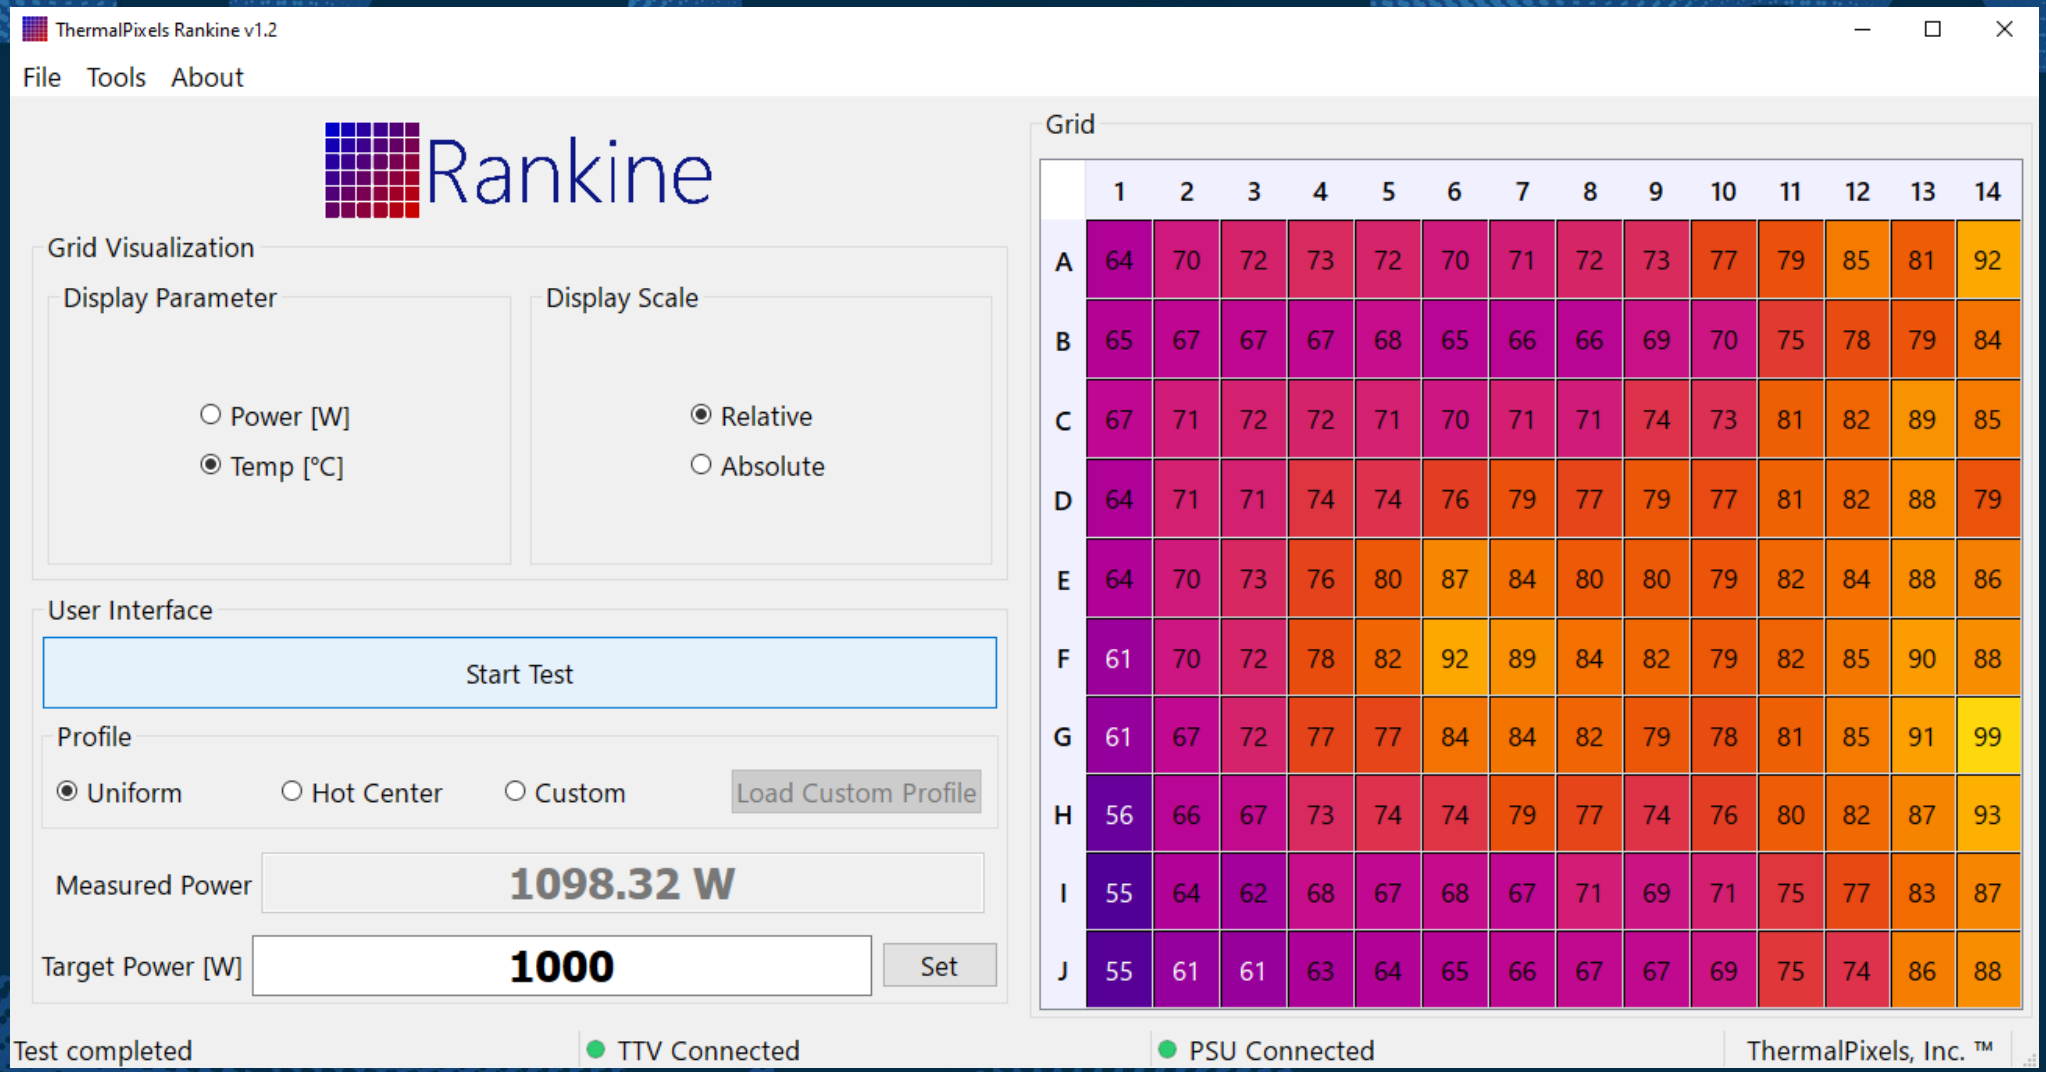


TTV and power management board



Cooling block

User interface



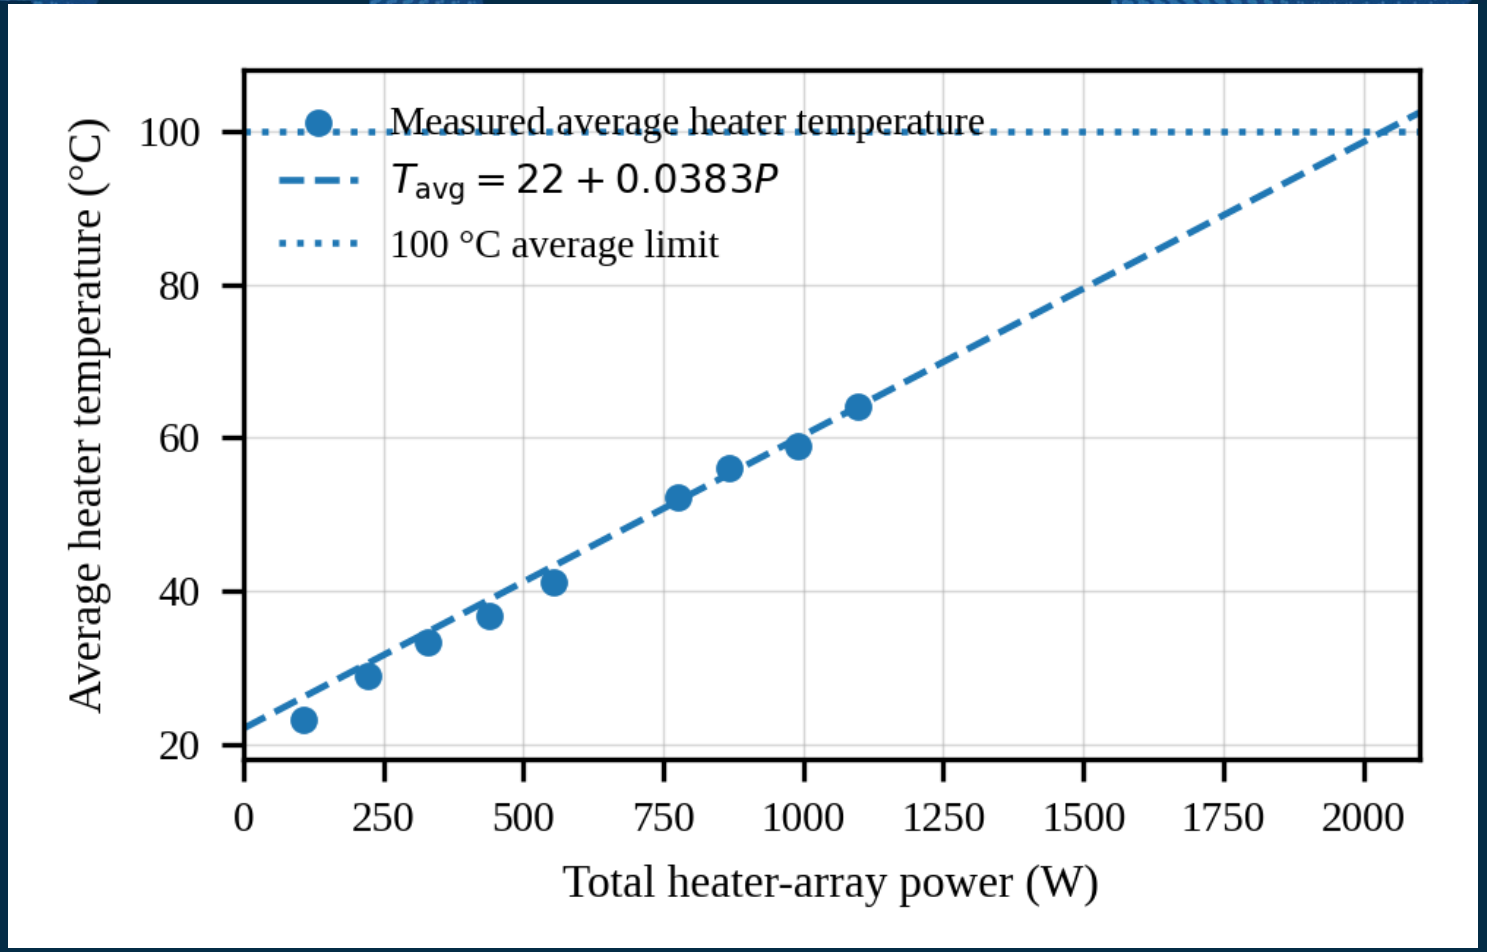
The TTV has 140 discrete heating zones that can be independently controlled.

Operating conditions

Experimental configuration

Parameter	Value
Active heater area	34 mm x 55 mm (18.70 cm^2)
Heater zones	140 (14 x 10)
Coolant	Water
Flow-rate range	1.0-5.0 L/min
Inlet temperature	22 °C
Reported temperature	Average over all heaters
TIM1	PTM7950, 8.5 W/m-K
Intermediate layer	700 micrometer silicon
TIM2	TC-3080, 7 W/m-K

Results



Measured average heater temperature at 1 L/min and 22 °C inlet. The high-power slope corresponds to 0.0383 K/W.

Results

Central thermal resistance decomposition

Model term	Value	Interpretation
R_{fixed}	0.0200 K/W	TIMs, silicon, heater, contact, spreading
R_{cooler}	0.0111 K/W	Cooler convection
Flow exponent	0.60	Microchannel flow scaling
R_{water}	0.00718 K/W	Mean coolant warming
Total at 1 L/min	0.03828 K/W	Matches measured baseline

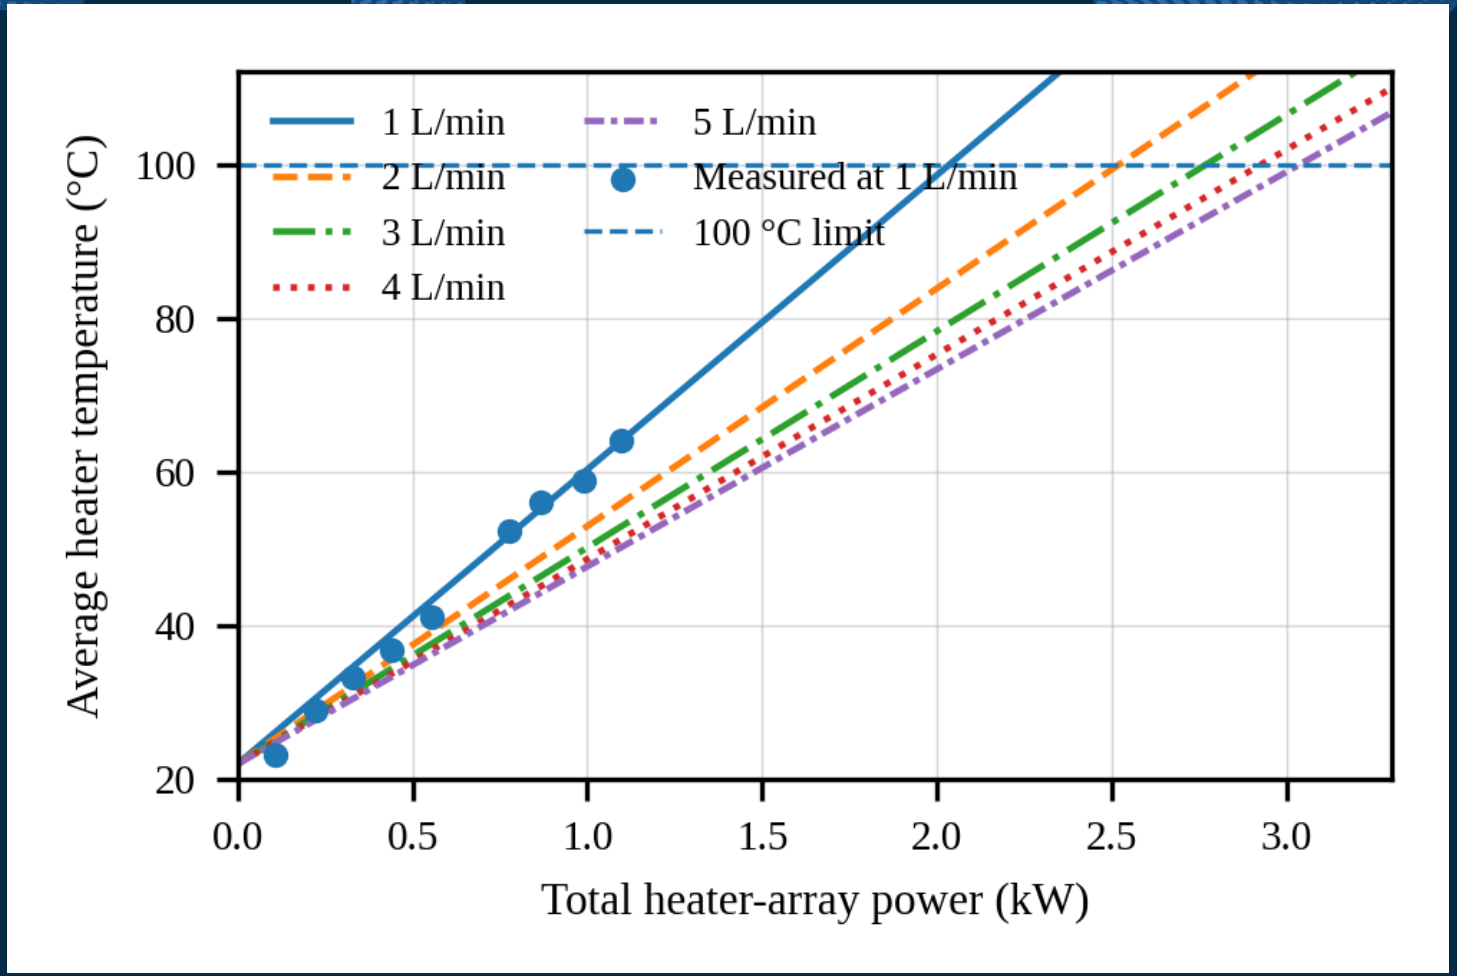
$$R_\theta(Q) = R_{fixed} + R_{cooler} Q^{-0.60} + 0.00718/Q$$

$$T_{avg}(P,Q) = T_{in} + P\,R_\theta$$

$$P_{max}(T_{limit},Q) = (T_{limit} - T_{in}) / R_\theta$$

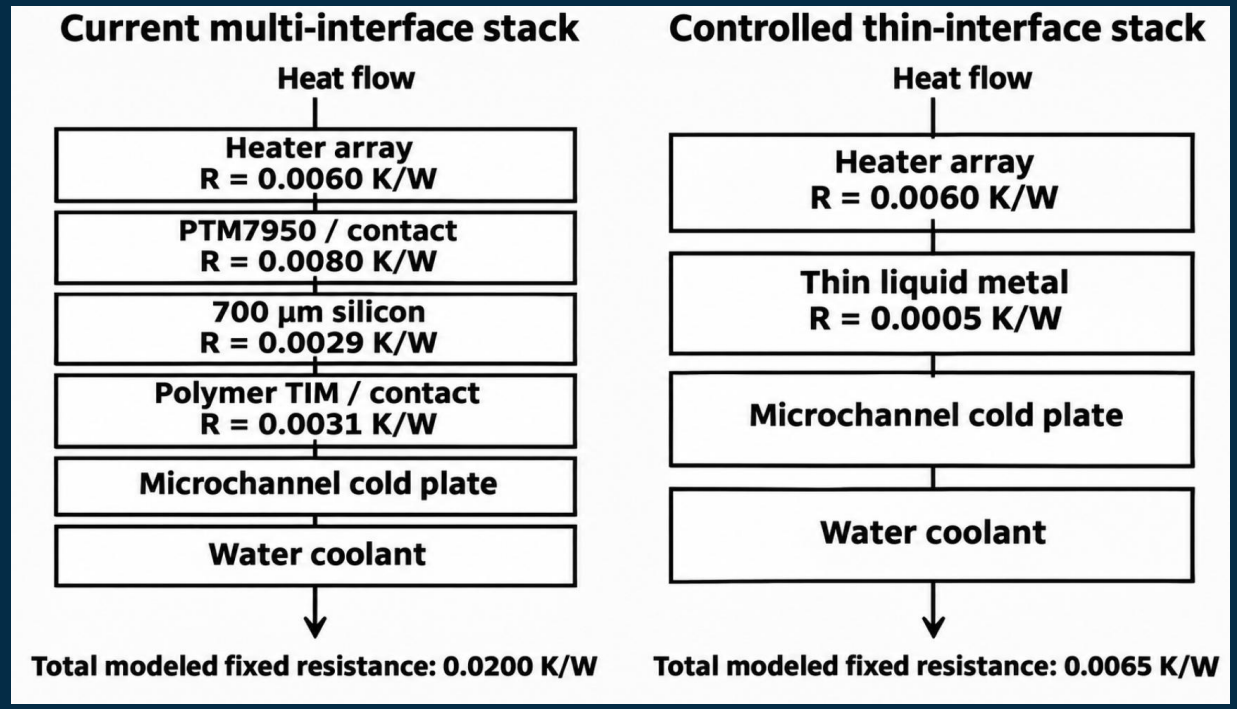
The split is a calibrated, physically plausible decomposition rather than a layer-by-layer measurement.

Results



Current-stack average heater temperature versus total power. Points at 1 and 2 L/min are measured; 3-5 L/min curves are model projections.

Results

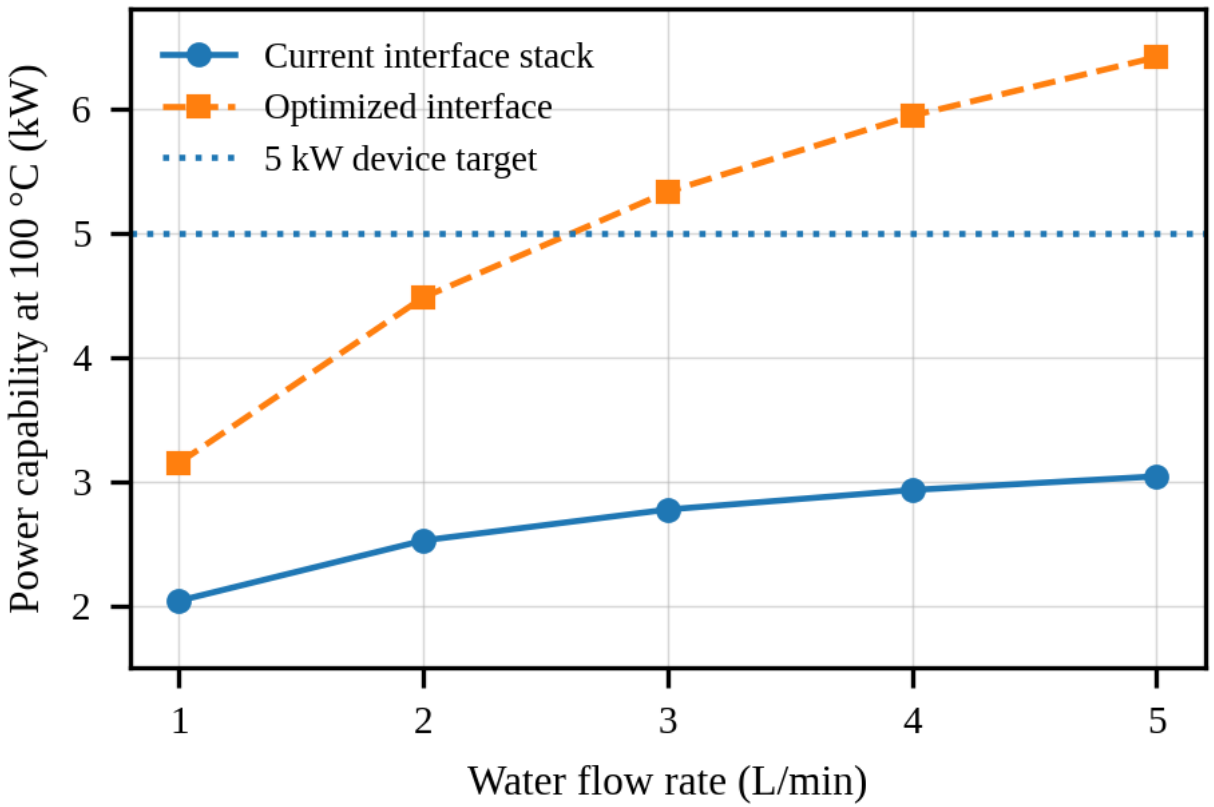


Modeled heat-flow path for the measured multi-interface stack and the controlled thin-interface scenario. Cold-plate geometry and coolant are held unchanged.

Flow-independent resistance budget

Contribution	Current (K/W)	Optimized (K/W)
TIM1/contact	0.0080	--
700 micrometer silicon	0.0029	--
TIM2/non-flat contact	0.0031	--
Controlled liquid metal	--	0.0005
Residual heater/spreading/contact	0.0060	0.0060
Total R_{fixed}	0.0200	0.0065

Results

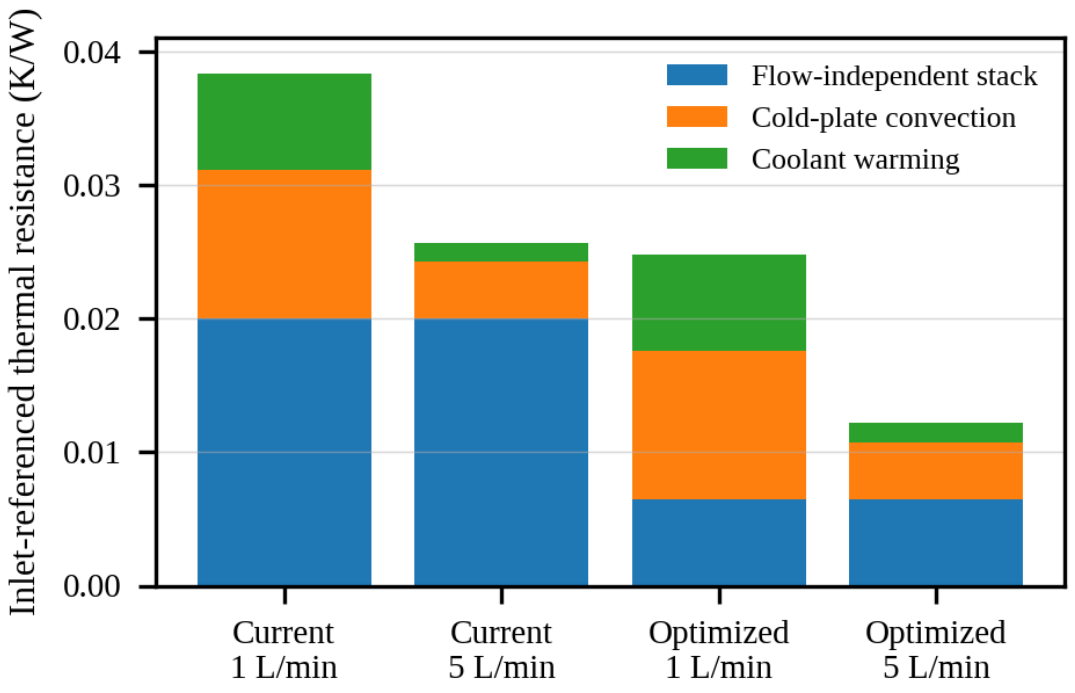


Optimized interface power capability

	Power capability (kW) at different TTV max temperatures			
Flow (L/min)	70 °C	80 °C	90 °C	100 °C
1	1.94	2.34	2.74	3.15
2	2.76	3.33	3.91	4.48
3	3.28	3.96	4.65	5.33
4	3.66	4.42	5.18	5.94
5	3.95	4.77	5.59	6.41

Modeled 100 °C power capability of the current and optimized interface stacks.

Results

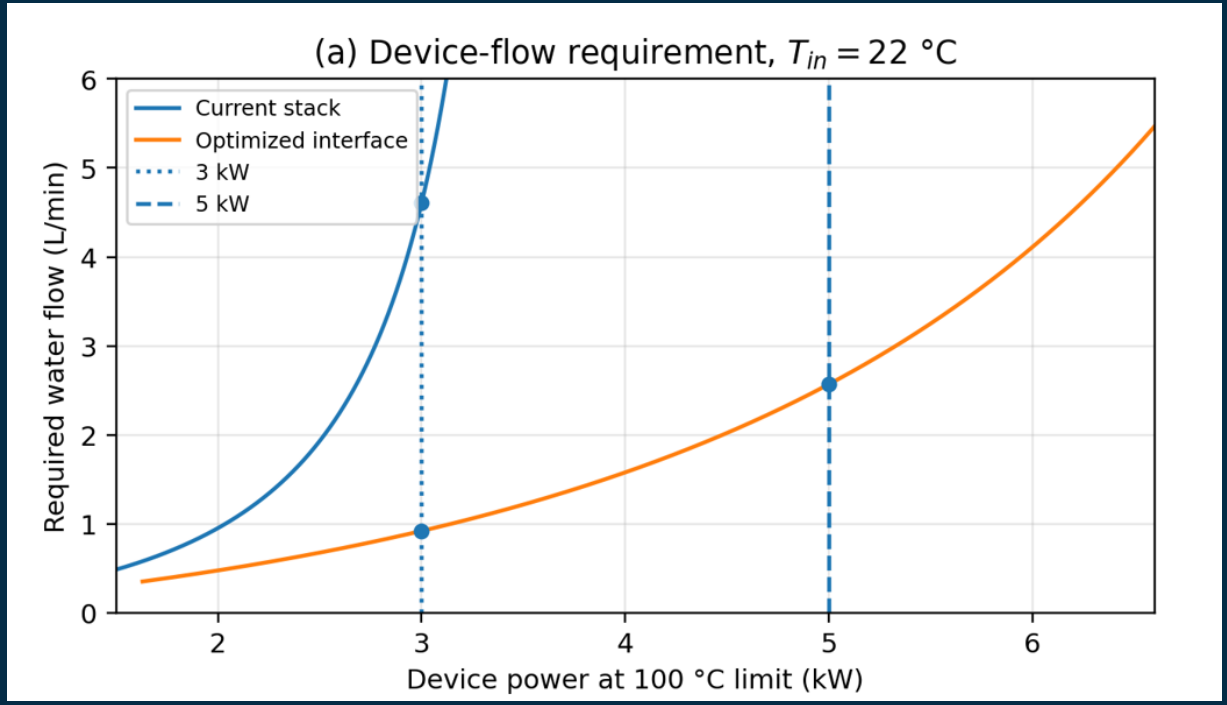


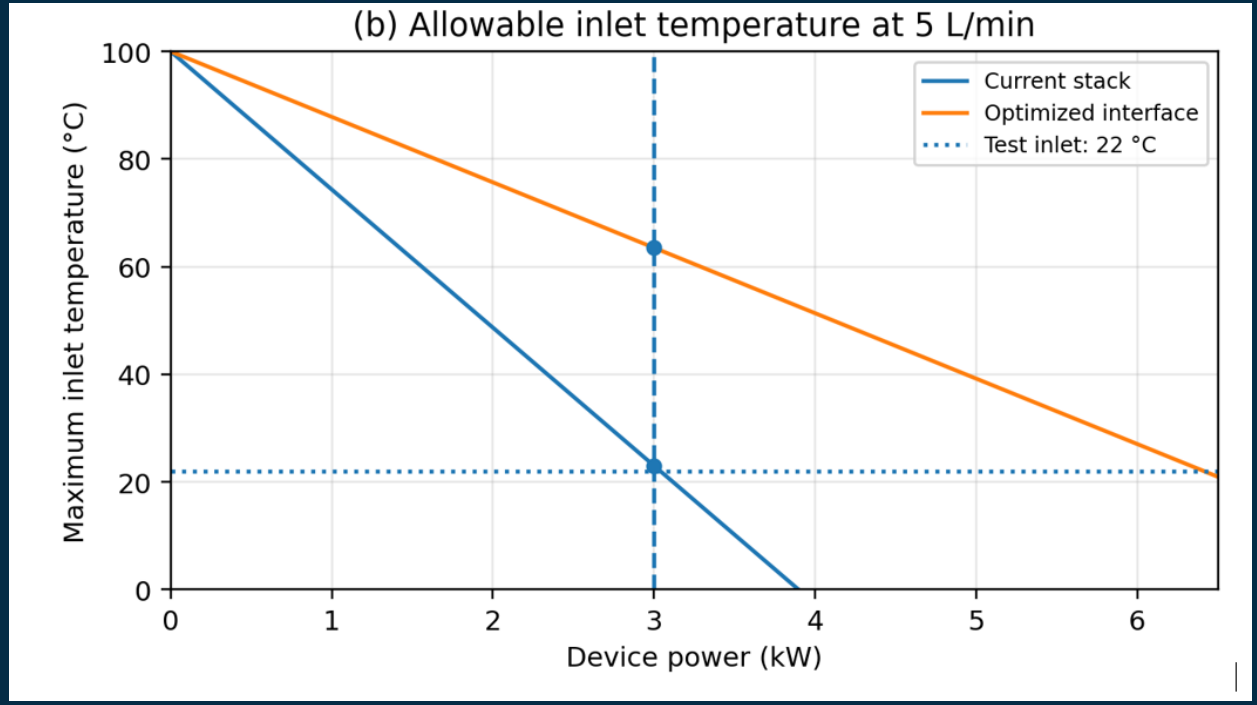
Modeled inlet-referenced resistance budget at 1 and 5 L/min. Interface control changes both the total resistance and the value of higher flow.

Modeled temperature-rise allocation at 5 L/min

Case	Fixed (°C)	Convection (°C)	Coolant (°C)	Total rise (°C)
Current, 3 kW	60.0	12.7	4.3	77.0
Optimized, 3 kW	19.5	12.7	4.3	36.5
Current, 5 kW	100.0	21.1	7.2	128.3
Optimized, 5 kW	32.5	21.1	7.2	60.8

Results





Device-level sustainability maps derived from the reduced-order model: (a) flow required to maintain a 100 °C average-heater limit at a 22 °C inlet; (b) maximum allowable inlet temperature at 5 L/min.

Final remarks

~ 2x
Improved cooling capacity
(3.04 kW → 6.41 kW)

67% lower
Interface thermal resistance
(0.02 → 0.0065 K/W)

1 MW Rack
Enables 28 kW blades with
next-gen 5 kW AI GPUs

- **Validated hardware baseline:** Testing on a 140-zone AI chip emulator confirms heat removal of up to 3.04 kW with standard thermal layers.
- **Interfaces are the real bottleneck:** Switching from a conventional multi-layer approach to a thin, highly-conductive interface more than doubles the heat removal to 6.41 kW.
- **Lower facility power and water usage:** Superior chip-level heat transfer is one of the key enablers for warmer coolant temperatures and reduced pumping power, cutting facility energy use.
- **Next step – 1 MW rack integration:** Ongoing collaboration between EMCOOL, Nortek Data Center Cooling, NLR, Siemens, and Georg Fischer to integrate 28 kW liquid-cooled blades into a 42U rack, targeting a ~ 1 MW demonstration.