

# Efficient Chip-cooling using Embedded Biomimetic Microfluidics

**Remco van Erp**

remco.vanerp@corintis.com

Corintis SA

**Malik Fahrni**

Corintis SA

**Miguel Salazar de Troya**

Corintis SA

**Athanasios Boutsikakis**

Corintis SA

**Emile Soutter**

Corintis SA

**Andrea Francescon**

Corintis SA

**Florent-Valery Coen**

Corintis SA

**Jordan Doytchinov**

Corintis SA

**Ali siddiqui**

Corintis SA

**Saurabh Tomar**

Corintis SA

**Bharath Ramakrishnan**

Microsoft

**Cam Turner**

Microsoft Corporation

**Husam Alissa**

Microsoft <https://orcid.org/0000-0001-5018-2605>

**Ashish Raniwala**

Microsoft

**Brijesh Warriar**

Microsoft

**Ricardo Bianchini**

Microsoft

**Jim Kleewein**

Microsoft

**Selim Bilgin**

Microsoft Corporation

**Christian Belady**

Microsoft

---

## Physical Sciences - Article

### Keywords:

**Posted Date:** January 30th, 2025

**DOI:** <https://doi.org/10.21203/rs.3.rs-5814747/v1>

**License:**  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

**Additional Declarations:** **Yes** there is potential Competing Interest. A patent corresponding to this work has been filed (Application EP22386065.1A). corresponding Author, Remco van Erp, is co-founder of Corintis, a company that supports development of liquid cooling applications for data centers. The other authors declare no competing interests.

---

# 1 **Efficient Chip-cooling using Embedded Biomimetic Microfluidics**

2 *Power dissipation in the latest accelerator chips for artificial intelligence has exceeded 1 kW and is*  
3 *increasing with each successive generation, while simultaneously introducing new thermal resistances and*  
4 *spatially varying temperature requirements due to heterogeneous integration. In contrast, cooling design*  
5 *has remained largely homogeneous, relying predominantly on straight channels or arrays of fins which are*  
6 *not tailored to the underlying chip. As data center operators target to increase coolant supply*  
7 *temperatures of data centers to reduce electricity and water consumption, more power needs to be*  
8 *extracted from increasingly complex chips with reduced temperature differentials between the junction*  
9 *and coolant. Hence, cooling is becoming the main driver for both chip performance and data center*  
10 *efficiency. Microfluidic cooling, with coolant directly flowing through channels embedded inside the silicon*  
11 *chip, has been proposed as a candidate to overcome the performance limitations of standard cold plates,*  
12 *as it aims to minimize thermal resistance by eliminating interfaces between the chip and the coolant. In*  
13 *addition, the close coupling between the heat source and cooling, combined with mature silicon-based*  
14 *microfabrication methods, enables more design freedom for cooling structures beyond straight channels*  
15 *and fins. However, while prior demonstrations of microfluidic cooling on function integrated circuits have*  
16 *shown the feasibility of this cooling integration, they suffered from high pressure drops, large temperature*  
17 *gradients and localized hot-spots, in addition to packaging and integration concerns. In this work, we show*  
18 *that the chip-level spatial power distribution is the source of these discrepancies, and that microfluidic*  
19 *topology optimization can simultaneously address the temperature and pressure drop challenges with*  
20 *microfluidic cooling, by creating a hierarchical network of channels that balance heat transfer and pressure*  
21 *drop mimicking the arteries, veins and capillaries in the circulatory system. Our result demonstrates that,*  
22 *by utilizing a bio-inspired chip-aware microfluidic cooling architecture, up to 18 °C reduction in hot-spot*  
23 *temperature rise on functional CPUs can be achieved, or a reduction in pressure drop by over 67%*  
24 *compared to a canonical design of pin fins. We demonstrate a 3x reduction in core temperature spread*  
25 *versus unoptimized designs, and furthermore observe that average thermal resistance can be reduced by*  
26 *up to 55% compared to standard cold plate cooling. These results validate that including microfluidic*  
27 *cooling as an integral part of the chip design process may be a viable path to allow for a further extension*  
28 *of the roadmap of silicon CMOS chips with increasing power levels, while simultaneously addressing the*  
29 *sustainability concerns at the data center scale.*

## 30 Introduction

31 Dennard scaling in silicon complementary metal-oxide semiconductor (CMOS) chips has led to a  
32 continuous increase in transistor density while maintaining a consistent device-level heat flux until the  
33 early 2000s<sup>1</sup>. The end of Dennard scaling initiated a steady increase in heat flux generation by generation,  
34 leading to dark silicon and eventually the need for multicore architectures<sup>2,3</sup>. Between 2010 and 2024,  
35 power of integrated circuits has risen from 250 W to 1500 W, driven simultaneously by an increase in heat  
36 flux and silicon surface area per chip<sup>4</sup>. Adequate thermal management is a critical part of chip design, as  
37 high junction temperatures negatively affect device reliability and performance, and a single region of  
38 high heat flux can create local hot-spots that limit the maximum clock speed of a chip. This thermal  
39 challenge is further aggravated by advanced packaging and 3D integration: co-packaging of high  
40 bandwidth memory (HBM) and compute dies introduces variations of maximum junction temperature  
41 requirements within a package, as memory needs to operate at a lower temperature (85 °C) than the  
42 compute die (100 °C), while having a lower effective thermal conductivity than the solid silicon substrate  
43 of the compute die due to its multi-layer design<sup>5</sup>. Consequently, thermal management for these chips is  
44 increasingly challenging and energy consuming, especially in data centers where large quantities of chips  
45 are densely packed, limiting the advancement of global compute capacity. As AI data centers are on track  
46 to consume more electricity than a country the size of the Netherlands<sup>6</sup>, addressing cooling needs in an  
47 energy-efficient way is paramount.

48 Air-cooling has long been the standard in data center cooling, but increasing power levels require reduced  
49 air temperatures and increased fan speeds, both having a significant impact on data center electricity  
50 consumption. Cooling air to sub-ambient temperatures using chillers is energy intensive and should be  
51 avoided to enable more sustainable data centers. Furthermore, for thin 1 Rack Unit (RU) central  
52 processing unit (CPU) servers, fan power can already account for up to 25% of the total server power,  
53 though this percentage is lower for GPU servers. Increasing fan speeds will amount to a further increase  
54 in power consumption, as well as an increase in noise levels that are regulated by occupational health and  
55 safety regulations in many countries. In addition, increasingly taller heat sinks required to maintain  
56 acceptable temperatures with air cooling increases the physical distance between GPUs, which should be  
57 minimized for low-latency chip-to-chip communication. The usage of evaporative cooling is less energy  
58 intensive than the use of chillers, but requires large amounts of fresh water, which can create water  
59 scarcity in neighboring communities<sup>7</sup>. It is estimated that data center coolant supply temperature should  
60 be increased to 60 °C to rely on dry-coolers throughout the year independently of the location in the  
61 United States<sup>8</sup>, which would eliminate the majority of all data center water consumption. The data center  
62 industry is thus faced with three competing forces: An increase in chip power, an increase in package  
63 thermal resistance due to heterogeneous integration, and a reduction in available thermal budget. The  
64 only way to extend the roadmap for silicon CMOS chips in a sustainable way is to reduce thermal  
65 resistance.

66 To address these challenges, liquid cooling is increasingly used at scale to meet the efficiency and density  
67 targets of new data centers to support chips with thermal design points (TDPs) exceeding 750 W, such as  
68 new deployments of graphical processing units (GPUs) and AI accelerator chips used for artificial  
69 intelligence (AI) and the training of large language models. For single-phase liquid cooling using cold plates  
70 aqueous coolants are favored due to their high specific heat, containing propylene glycol and corrosion  
71 inhibitors to ensure long-term reliability with aluminum or copper channels. These channels are typically  
72 fabricated using skiving, where a blade deforms metal to create an array of straight fins. This skiving  
73 process can create parallel microchannels down to 60 micrometers in size, with aspect ratios as high as  
74 50. However, this process lacks design freedom beyond simple parallel fins, resulting in a uniform channel  
75 design that cannot be tailored to the heterogenous nature of chips. Reducing the thermal resistance of a

76 single-phase cold plate is possible by increasing the flow rate, as it reduces the coolant temperature rise  
77 due to sensible heat, but comes at the expense of a quadratic increase in pumping power, which scales  
78 with the product of flow rate and pressure drop. Due to the homogeneous design of a skived cold plate,  
79 only a fraction of the additional flow rate will be targeted to the critical areas of the chip. Given the system-  
80 level constraints on these two parameters, it's desirable to minimize junction-to-inlet thermal resistance  
81 with limited pressure drop and flow rate using a targeted cold plate design that maximized heat transfer  
82 in the critical areas of the chip, while minimizing pressure drop and flow rate in the areas of least concern.

83 Both jet impingement and manifold microchannel (MMC) coolers allow for such optimization due to the  
84 increased heat transfer in developing flow regions, but typically at the millimeter-scale, which is far  
85 beyond the length scale of local hot-spots in a chip. Furthermore, the increase in heat transfer at the  
86 impingement zones comes at the expense of reduced heat transfer near recirculation zones, which needs  
87 to be carefully balanced. Alternative technologies such as pumped two-phase cooling, or flow boiling can  
88 achieve a high local heat transfer coefficient due to phase change. However, to avoid reaching critical heat  
89 flux, sufficiently small channel dimensions are required to increasing surface area for heat transfer which  
90 introduces stability challenges, such as increased superheat before onset boiling, backflow, and a rapid  
91 increase in pressure drop as phase transition occurs. When operating a server rack with multiple parallel  
92 chips cooled down from a single supply, such stability issues can disrupt the operation of the complete  
93 system. As such, pumped two-phase systems often operate at a low outlet vapor quality, reducing the  
94 benefit of the high latent heat of two-phase cooling. In addition, the lack of simple generalized models  
95 that include the nucleation of flow boiling makes it challenging to optimize a two-phase cold plate a-priori  
96 to a specific heterogeneous cooling problem. Finally, the performance increases attainable by any type of  
97 cold plate optimization is ultimately limited by the thermal interface materials (TIM), which attaches the  
98 cooling system to the chip.

99 Embedded microfluidic cooling, where the cooling is directly integrated in the silicon substrate eliminates  
100 the thermal resistance introduced by the TIM, enabling die-level heat fluxes exceeding  $700 \text{ W/cm}^2$ .  
101 Furthermore, the usage of lithography and etching unlocks complete micro-scale design freedom down  
102 to sub-micrometer resolution with aspect ratios exceeding 100. The performance of embedded  
103 microfluidic cooling has been extensively studied on thermal test vehicles (TTVs) with uniform heat  
104 sources<sup>9-11</sup>. Several studies demonstrated the benefits of optimizing cooling structures tailored to a  
105 specific spatial power distribution following a co-design approach<sup>12</sup>. However, only a limited number of  
106 studies evaluated embedded microfluidic cooling on functional chips<sup>13-16</sup>, often with significantly lower  
107 cooling performance than anticipated by the TTV studies, and large temperature variations across the  
108 chip. The latter can be attributed to the highly non-uniform power distribution on a chip and the reduced  
109 heat spreading in embedded microfluidic cooling. As the cooling gets closely coupled with the chip, it  
110 becomes essential to optimize the precise flow of coolant to address the critical regions of the chip, to  
111 minimize junction temperatures and maximize temperature uniformity.

112 Topology optimization for embedded microfluidic cooling allows for systematic reduction of hot-spot  
113 temperature rise given a set of constraints on flow rate and pressure drop due to the predictable behavior  
114 of laminar flow, while maintaining the operational simplicity of single-phase flow. Instead of explicitly  
115 defining the cooling channels, the geometry in topology optimization is parametrized in a set of pixels that  
116 define the presence of solid or liquid material, and through an iterative process a geometry is formed to  
117 minimize an objective function such as junction temperature rise<sup>17-22</sup>. Based on the die-level heat flux (Fig.  
118 1a), and the design freedom offered by lithography-based manufacturing it has the potential to create a  
119 chip-aware cooling design that balances high heat transfer with narrow cooling fins in hot-spot regions  
120 while satisfying the low pressure drop requirements of data center cooling (Fig. 1b). This design approach

121 mimics the circulatory system of the human body, which is similarly optimized to increase local heat and  
122 mass transfer while minimizing the pumping power of the heart.

123 In this work, we investigate the usage of topology optimization for microfluidic cooling to reduce pressure  
124 drop and thermal resistance, demonstrate a robust packaging approach to guide the coolant into the chip  
125 (Fig. 1c), and experimentally benchmark their performance versus air and liquid cooling. Two chips were  
126 selected as part of this study: The Intel i9 (Fig. 1d) was selected as a comparison with a prior study that  
127 benchmarked the cooling performance of air-cooling, cold plates and immersion cooling<sup>23</sup>. The Intel i7  
128 8700K (i7) and Intel i9 9900K (i9). The Intel i7 (Fig. 1e) was selected as a comparison point with a prior  
129 study that evaluated the cooling performance of embedded micro pin fins<sup>16</sup>. Due to confidentiality of the  
130 precise power maps, we will base our study on the power maps in Fig. 1f,g that are derived from publicly  
131 available sources as described in the methods section.

## 132 **Results**

### 133 **Simulation and optimization**

134 Figure 1h,i show the cooling designs generated using the topology optimization approach outlined in the  
135 methods section for the i9 and i7, respectively. The design for the i9 (Fig. 1h) shows longer and more  
136 narrow channels, which can be explained due to the higher substrate thickness. Narrow channels are  
137 mostly concentrated on hot-spot regions on the chip, whereas larger branches are formed upstream and  
138 downstream of the hot areas. In contrast, the i7, due to its shallower channels, results in a more sparsely  
139 populated fin structure (Fig. 1i). One can see large open channels that create a direct connection between  
140 the inlet and the outlet on both the i7 and the i9. This can be explained by the fact that the optimization  
141 is only defining inlet pressure but does not constrain the flow rate. Hence, excess flow of coolant that  
142 bypasses the hot-spots is not penalized in this method. The resulting velocity fields of the i9 and i7 are  
143 shown in Fig. 1j,k, respectively, and the junction temperatures in Fig. 1l,m.

### 144 **Numerical performance comparison for i9 using different cooling methods**

145 Figure 2a,b shows a comparison of thermal performance of the i9 across the three examined designs:  
146 straight channels, pin fins, and optimized geometries. For a design pressure of 300 mbar and an inlet  
147 temperature of 20 °C, the optimized design for the i9 achieves a simulated maximum junction  
148 temperature of 67 °C, corresponding to a junction-to-inlet temperature rise of 47 °C and a hot-spot  
149 thermal resistance of 0.16 K/W. At identical pressure, we manage to achieve a 6 °C and 18 °C temperature  
150 reduction versus straight channels and pin fins, respectively (Fig. 2a). Alternatively, with the optimized  
151 design, a temperature increase of 70 °C can be maintained with 50% and 67% less pressure drop for  
152 straight channels and pin fins, respectively. In prior work on the i9, thermal resistance was measured for  
153 the i9 9900K under various cooling configurations, such as air cooling, cold plate liquid cooling, and 2-  
154 phase immersion cooling<sup>23</sup>. Their study observed a thermal resistance of 0.45 K/W to 0.57 K/W for air  
155 cooling, a thermal resistance between 0.32 K/W and 0.46 K/W for cold plate water cooling, and a thermal  
156 resistance of 0.25 K/W for immersion cooling. Results for microfluidic cooling on the i9 9900K, using  
157 optimized designs, resulted in thermal resistances as low as 0.13 K/W as shown in Extended Data Fig. 6.  
158 Hence, we can state that, by carefully designing microfluidics in conjunction with the power map, 257%  
159 more heat can be extracted from the i9 9900K than using air-cooling, 178% more heat can be extracted  
160 than using cold plate liquid cooling, and 78% more heat can be extracted from the chip than using  
161 immersion cooling, as shown in Fig. 2b.

### 162 **Numerical performance comparison for i7 using different cooling designs**

163 Figure 2c shows a comparison the thermal performance of the i7 between the three examined designs:  
164 straight channels, pin fins, and optimized microfluidic geometries. For reference, we also include a pin  
165 fin design with one inlet-outlet configuration, as evaluated in the literature<sup>16</sup>. For a design pressure of  
166 380 mbar and an inlet temperature of 20 °C, the i7 achieves a simulated maximum junction temperature  
167 rise of 68 °C, corresponding to a junction-to-inlet temperature rise of 68 C and a hot-spot thermal  
168 resistance of 0.32 K/W. This corresponds to an 8 °C and 7 °C temperature decrease with respect to  
169 straight channels and pin fins, respectively. To maintain a temperature rise of 70 °C, 40% and 45% less  
170 pressure drop is required compared to straight channels and pin fins, respectively. Prior studies on the i7  
171 evaluated the pin fin configuration extending over the full width of the chip<sup>16</sup>. In their study, average  
172 core temperatures and maximum core temperatures were reported. When the thermal resistance was  
173 calculated based on average core temperatures, an outstanding average junction-to-inlet thermal  
174 resistance of 0.20 K/W was obtained. However, when instead calculating the thermal resistance based  
175 on maximum core temperature, a significantly higher thermal resistance value of 0.45 K/W was  
176 obtained (Extended data Fig. 4). Ultimately, the hottest part in the chip will limit the performance of the  
177 device, so achieving uniform core temperatures is essential in the design of microfluidic cooling.

178 Simulation results that include die-level power distribution can give new insights into the origin of these  
179 discrepancies and highlight how optimized topologies improve temperature uniformity within the die.  
180 Figures 2c-f show the channel and junction temperature of the i7 8700K with pin fins and optimized  
181 designs side-by-side. The channel temperature of the pin fin designs (Fig. 2e,f) shows a clear temperature  
182 gradient across the chip. Two cores are aligned in parallel to the flow path. Hence, coolant heats up as it  
183 passes through the device, causing elevated core temperatures at the outlet region. A 57% higher core  
184 temperature is observed at the outlet cores compared to the inlet (Fig. 2e). In comparison, the optimized  
185 design in Fig. 2f,h shows very uniform core temperatures and reduced temperature gradients. The flow  
186 path (Fig. 2h) shows that hot coolant, as it passes close to the cores, is evacuated laterally to prevent hot  
187 coolant from heating up the outermost part of the cores. This comparison clearly shows the need for  
188 careful design of microfluidic cooling.

## 189 **Etching and fabrication**

190 The cooling design was transferred into the silicon die of the i7 using a lithography and etching process  
191 (Fig. 3a), further described in the methods section. A close-up inspection (Fig. 3b) shows vertical sidewalls  
192 and with a mirror-like smooth surface finish on all etched surfaces (Fig. 3c). This process could not be  
193 repeated on the i9, because it contains a die metallization to enable adhesion of a solder TIM, which  
194 prevents deep reactive ion etching. Therefore, a laser-machining method was utilized for the i9 (Fig. 3d)  
195 which successfully ablated both this metallization layer and the silicon underneath. However, a clear layer  
196 of redeposition was visible as a light glow around the etched structures in fig. 3e,f for the i9 and i7,  
197 respectively. Fig. 3g shows another i7, with the same design laser etched as in Fig. 3d, but with visibly  
198 rougher surface finish on the etched regions. The final package was installed on both CPUs as shown in  
199 Fig. 3i as detailed in the methods section, leading to a final assembly with barb fittings that interface with  
200 standard computer liquid cooling equipment. The package shown in Fig. 3i shows three ports: One central  
201 inlet and two separate outlets, to maintain the possibility of measuring the flow distribution between the  
202 two outlet streams, although this functionality was not utilized in this work. The final package was  
203 qualified before using press-decay testing at 5 bars, to ensure reliable operation.

## 204 **Experimental results**

205 Figure 4a shows the results of the overclocking campaign on the i7 outlined in the methods section.  
206 Standard air cooling reached a clock speed of 4.7 GHz before thermal throttling occurred at power of 160  
207 W, which corresponds to a die-level heat flux of  $106 \text{ W/cm}^2$ . For cold plates, clock speeds could be  
208 increased to 5.0 GHz before reaching throttling temperatures at a power of 250 W, which corresponds to  
209 a die-level heat flux of  $168 \text{ W/cm}^2$ , whereas microfluidic cooling could achieve a clock speed up to 5.3  
210 GHz, and no thermal throttling occurred up to a power of 350 W. Overclocking was eventually limited  
211 because the maximum safe operating voltage was reached, limiting any further increase in clock speed.  
212 Extrapolating to a maximum temperature rise of  $105 \text{ }^\circ\text{C}$  gives a die-level heat flux of  $234 \text{ W/cm}^2$ . The total  
213 increase in achievable TDP for microfluidic cooling within safe operating conditions was thus a 219% and  
214 40% increase versus air cooling and cold plates, respectively.

215 All values in Fig 4a are defined at the instance a single core started throttling. Hence, maintaining a  
216 uniform core temperature rise is essential, since a single throttling core defines the peak clock speed for  
217 a given cooling configuration. The variation between thermal resistance based on the peak core  
218 temperature and average core temperature is shown in Fig. 4b, for the chip-aware cooling design on the  
219 i7 in this work, as well as the uniform array of pin fins on the similar chip from the literature. Note that  
220 because of topology optimization, pressure drop is significantly lower on the chip-aware cooling design.  
221 The uniform pin fin design shows a peak core thermal resistance that's more than twice as high as the  
222 average core thermal resistance. This shows that, despite the high average heat transfer coefficient of  
223 embedded microfluidic cooling, it can only translate to increased chip performance if the cooling is  
224 adapted to the local hot-spots on a chip. In contrast, the chip-aware cooling design in Fig. 4b shows a 55%  
225 reduction in peak core thermal resistance at a comparable pressure drop of 1300 mbar, which is only 34%  
226 above the average core temperature and thus a 3x reduction in core temperature spread, demonstrating  
227 the outstanding performance increase that can be achieved using chip-aware microfluidic cooling design

228 A comparison between air cooling, cold plate cooling and microfluidics at maximum achievable stable  
229 operating points under Cinebench (C) and Prime95 (P), with inlet temperatures of  $21 \text{ }^\circ\text{C}$  and  $34 \text{ }^\circ\text{C}$  is shown  
230 in Fig. 4c and Fig 4d, for the i7 and i9, respectively. For cold plates, a single flow rate of 1 LPM was  
231 considered, whereas for microfluidics three flow rates were evaluated. For each all scenarios, microfluidic  
232 cooling consistently achieves the lowest thermal resistance based on both the maximum and the average  
233 core temperature, confirming the potential of microfluidics to outperform conventional cooling systems  
234 available today. For a technology comparison, Fig. 4e shows the maximum TDP that can be achieved for  
235 the i9 based on a  $70 \text{ }^\circ\text{C}$  junction-to-inlet temperature rise, calculated based on thermal resistances using  
236 Prime95 from this work and prior results from the literature that included air cooling, cold plates and two-  
237 phase immersion cooling. The maximum TDP for air cooling was 150 W and cold plates 190W. Immersion  
238 cooling could reach a TDP of 280 W, but it should be noted that a metallic TIM2 was used between the  
239 chip and a boiler plate. Using the same benchmark, microfluidic cooling on the i9 in this work can exceed  
240 a TDP of 420 W, which is a 185% increase compared to standard cold plate cooling, using the exact same  
241 coolant. Breaking down the average core thermal resistance into contributions from TIM1, TIM2, the  
242 integrated heat spreader (IHS), and heat sink (HS) for each solution (Fig. 4f) shows more insights in the  
243 breakdown of the observed temperature rise. Figure 4g shows that for a standard air-cooled or cold plate  
244 configuration, the combined contribution of TIM1, TIM2 and the IHS correspond to  $0.17 \text{ K/W}$ . Immersion  
245 cooling shows a lower thermal resistance, but this can be mostly attributed to the reduction in TIM2, as a  
246 metallic TIM was used. Finally, for microfluidics, the total average core thermal resistance in this work was  
247 reduced to  $0.16 \text{ K/W}$ , which is lower than the fixed contributions due to TIM1, TIM2 and IHS in a default  
248 configuration. This confirms that, even an infinitely good heat sink or cold plate cannot outperform  
249 microfluidic cooling giving the present package and TIM constraints, highlighting the benefits of  
250 embedded microfluidic cooling.

## 251 **Discussion and outlook**

252 In this work, we present a new approach for embedded microfluidic cooling containing biomimetic cooling  
253 architectures adapted to the silicon chip. We show that, by using topology optimization, we can achieve  
254 up to 18 °C lower temperature rise and up to 67% lower pressure drop than a pareto front of canonical  
255 designs such as straight channels and pin fins. We show that these biomimetic cooling architectures are  
256 well suited to be manufactured in silicon by leveraging the design freedom of lithography as a scalable  
257 path to volume manufacturing, but that laser micromachining provides a viable path for rapid prototyping  
258 as well. A new microfluidic packaging solution presented which provided leak-tight operation, qualified at  
259 5 bar of pressure. An extensive experimental overclocking campaign showed that 0.3 GHz increase in clock  
260 speeds can be achieved compared to cold plates without reaching thermal throttling, but instead the  
261 maximum safe operating voltage limited any further increase in clock speed. We further experimentally  
262 validated that maximum temperature rise was reduced by over 55% compared to canonical pin fin  
263 structures that were not tailored to the power map, leading to a 3x improved core temperature  
264 uniformity, which is critical to unlock the full potential of embedded microfluidic cooling. A thermal  
265 resistance network breakdown showed that the total average core-to-inlet thermal resistance of  
266 microfluidic cooling is lower than the thermal resistance of TIM1, TIM2 and the IHS combined. As such,  
267 we find that even an idealized isothermal boundary condition at the heat sink interface cannot outperform  
268 the cooling performance offered by microfluidics, further highlighting the potential of embedded  
269 microfluidic cooling.

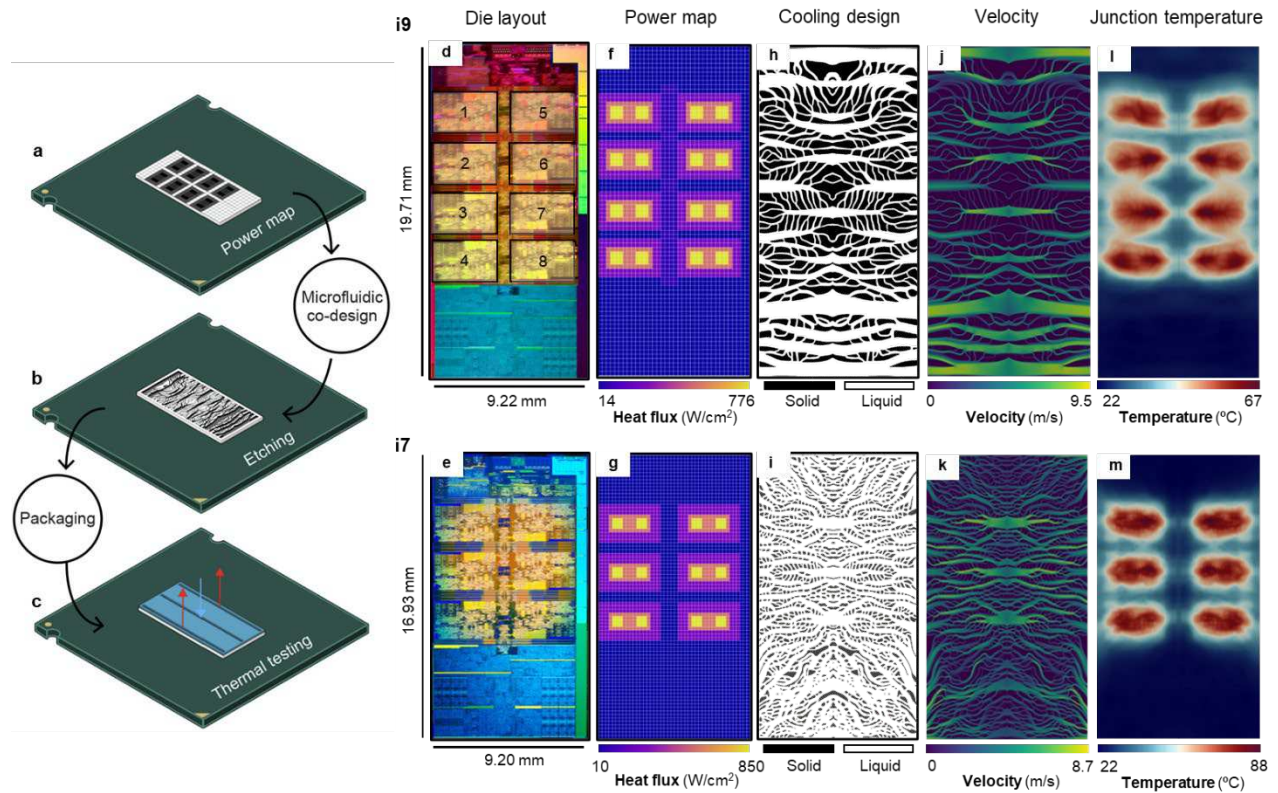
270 We foresee that, as cooling becomes a main limitation on both chip design and data center efficiency, co-  
271 design between chip and cooling will become a critical aspect of high-TDP chips. Given the scalability of  
272 silicon-based wafer-scale manufacturing, the design of micro-scale cooling tailored to the chip can  
273 become part of the tape-out of the chipset in the future. However, this would require industry-wide  
274 standardization of design and simulation tools, a commonly agreed microfluidic ‘process design kit’ (PDK)  
275 between chip designers and foundries, and further investigation on long term reliability, especially on  
276 clogging, erosion and leak-tightness of the packaging during prolonged operation in the field. If these  
277 points can be addressed, it may allow for a further extension of the roadmap of silicon CMOS chips with  
278 increasing power levels, while simultaneously addressing the sustainability concerns at the data center  
279 scale.

## Bibliography

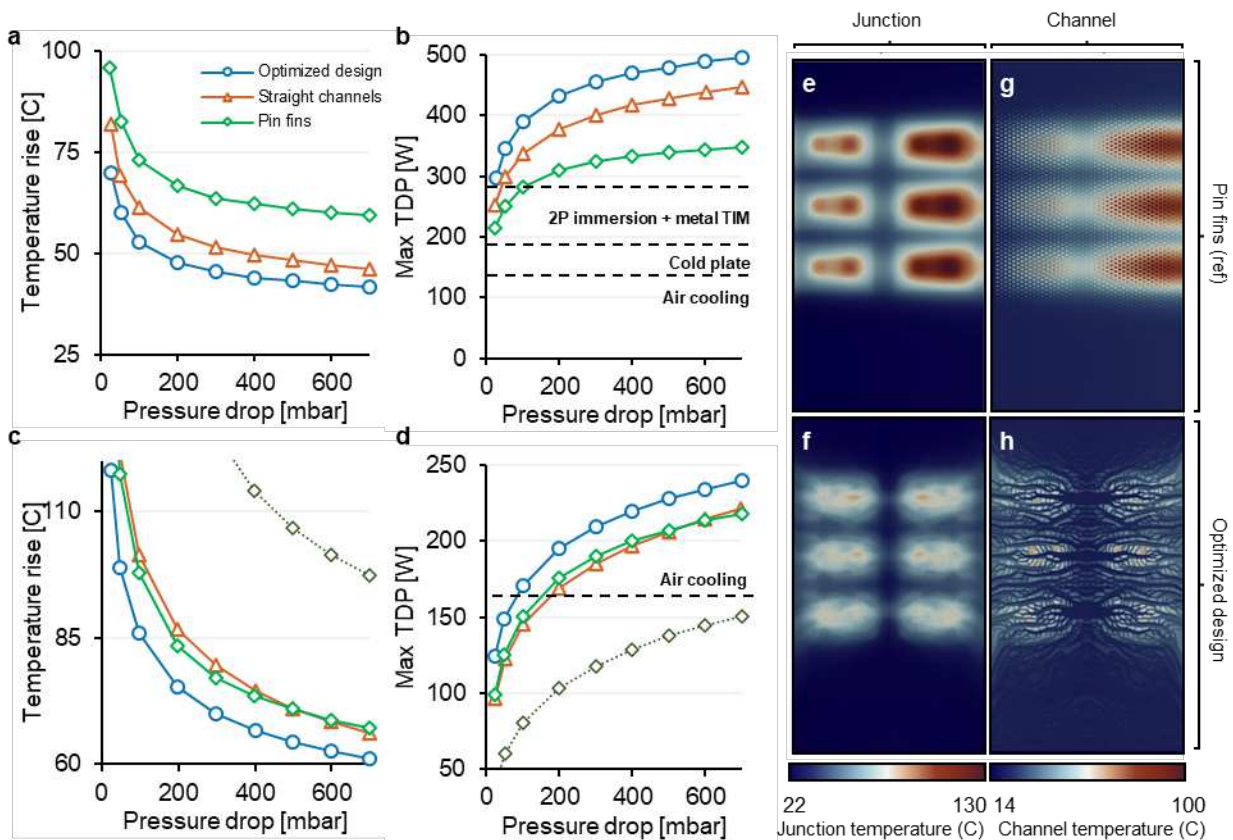
1. Dennard, R. H. *et al.* Design of Ion-Implanted MOSFETs with Very Small Physical Dimensions. *Proceedings of the IEEE* **87**, 668–678 (1999).
2. Rahmani, A. M., Liljeberg, P., Hemani, A., Jantsch, A. & Tenhunen, H. *The Dark Side of Silicon: Energy Efficient Computing in the Dark Silicon Era. The Dark Side of Silicon: Energy Efficient Computing in the Dark Silicon Era* (Springer International Publishing, 2017). doi:10.1007/978-3-319-31596-6.
3. Hardavellas, N., Ferdman, M., Falsafi, B. & Ailamaki, A. Toward dark silicon in servers. *IEEE Micro* **31**, 6–15 (2011).
4. Wesling, P. Heterogeneous Integration Roadmap, 2023 Version.
5. Kim, T. *et al.* Thermal Modeling and Analysis of High Bandwidth Memory in 2.5D Si-interposer Systems. *InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, ITherm 2022-May*, (2022).
6. de Vries, A. The growing energy footprint of artificial intelligence. *Joule* **7**, 2191–2194 (2023).
7. Mytton, D. Data centre water consumption. *npj Clean Water* 2021 4:1 **4**, 1–6 (2021).
8. De Bock, P., Bress, T., Lecoustre, V., Gidwani, A. & Noyes, C. Data Center Energy Reduction by Lowering Chip-to-Supply Thermal Resistance. *InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, ITherm 2023-May*, (2023).
9. Lien, Y. J. *et al.* An Energy-efficient Si-integrated Micro-cooler for High Power and Power-density Computing Applications. *Proceedings - Electronic Components and Technology Conference* 1025–1029 (2024) doi:10.1109/ECTC51529.2024.00164.
10. Zhang, L., Goodson, K. E. & Kenny, T. William. *Silicon Microchannel Heat Sinks*. (Springer, 2004). doi:10.1007/978-3-662-09899-8.
11. Colgan, E. G. *et al.* A practical implementation of silicon microchannel coolers for high power chips. in *Annual IEEE Semiconductor Thermal Measurement and Management Symposium* 1–7 (2005). doi:10.1109/stherm.2005.1412151.
12. van Erp, R., Soleimanzadeh, R., Nela, L., Kampitsis, G. & Matioli, E. Co-designing electronics with microfluidics for more sustainable cooling. *Nature* **585**, 211–216 (2020).
13. Sarvey, T. E. *et al.* Embedded cooling technologies for densely integrated electronic systems. in *Proceedings of the Custom Integrated Circuits Conference* vols 2015-November (Institute of Electrical and Electronics Engineers Inc., 2015).
14. Sarvey, T. E. *et al.* Microfluidic Cooling of a 14-nm 2.5-D FPGA with 3-D Printed Manifolds for High-Density Computing: Design Considerations, Fabrication, and Electrical Characterization. *IEEE Trans Compon Packaging Manuf Technol* **9**, 2393–2403 (2019).

15. Sarvey, T. E. *et al.* Monolithic integration of a micropin-fin heat sink in a 28-nm FPGA. *IEEE Trans Compon Packaging Manuf Technol* **7**, 1617–1624 (2017).
16. Kochupurackal Rajan, S. *et al.* Integrated Silicon Microfluidic Cooling of a High-Power Overclocked CPU for Efficient Thermal Management. *IEEE Access* **10**, 59259–59269 (2022).
17. Yan, S., Wang, F., Hong, J. & Sigmund, O. Topology optimization of microchannel heat sinks using a two-layer model. *Int J Heat Mass Transf* **143**, 118462 (2019).
18. Li, H., Ding, X., Meng, F., Jing, D. & Xiong, M. Optimal design and thermal modelling for liquid-cooled heat sink based on multi-objective topology optimization: An experimental and numerical study. *undefined* **144**, (2019).
19. Kambampati, S. & Kim, H. A. Level set topology optimization of cooling channels using the Darcy flow model. *Structural and Multidisciplinary Optimization* **61**, 1345–1361 (2020).
20. Ozguc, S., Pan, L. & Weibel, J. A. Topology optimization of microchannel heat sinks using a homogenization approach. *Int J Heat Mass Transf* **169**, 120896 (2021).
21. Yaji, K., Yamada, T., Kubo, S., Izui, K. & Nishiwaki, S. A topology optimization method for a coupled thermal–fluid problem using level set boundary expressions. *Int J Heat Mass Transf* **81**, 878–888 (2015).
22. Subramaniam, V., Dbouk, T. & Harion, J. L. Topology optimization of conjugate heat transfer systems: A competition between heat transfer enhancement and pressure drop reduction. *Int J Heat Fluid Flow* **75**, 165–184 (2019).
23. Ramakrishnan, B. *et al.* CPU Overclocking: A Performance Assessment of Air, Cold Plates, and Two-Phase Immersion Cooling. *IEEE Trans Compon Packaging Manuf Technol* **11**, 1703–1715 (2021).
24. Zhang, J., Sadiqbatcha, S., Jin, W. & Tan, S. X. D. Accurate Power Density Map Estimation for Commercial Multi-Core Microprocessors. *Proceedings of the 2020 Design, Automation and Test in Europe Conference and Exhibition, DATE 2020* 1085–1090 (2020) doi:10.23919/DATE48585.2020.9116545.
25. Stolpe, M. & Svanberg, K. An alternative interpolation scheme for minimum compliance topology optimization. *Structural and Multidisciplinary Optimization* **2001 22:2** **22**, 116–124 (2014).
26. Lazarov, B. S. & Sigmund, O. Filters in topology optimization based on Helmholtz-type differential equations. *Int J Numer Methods Eng* **86**, 765–781 (2011).
27. Svanberg, K. The method of moving asymptotes—a new method for structural optimization. *Int J Numer Methods Eng* **24**, 359–373 (1987).

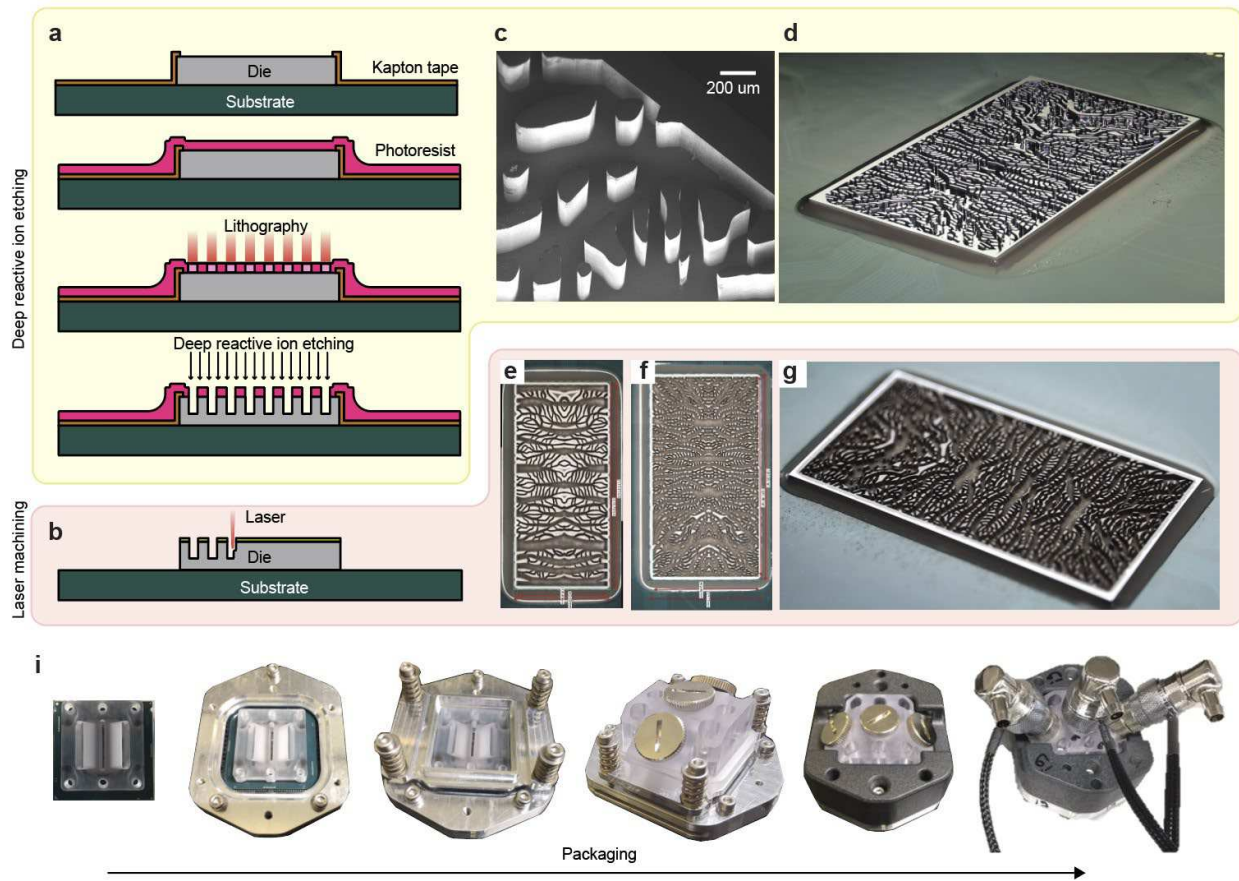
## Figures



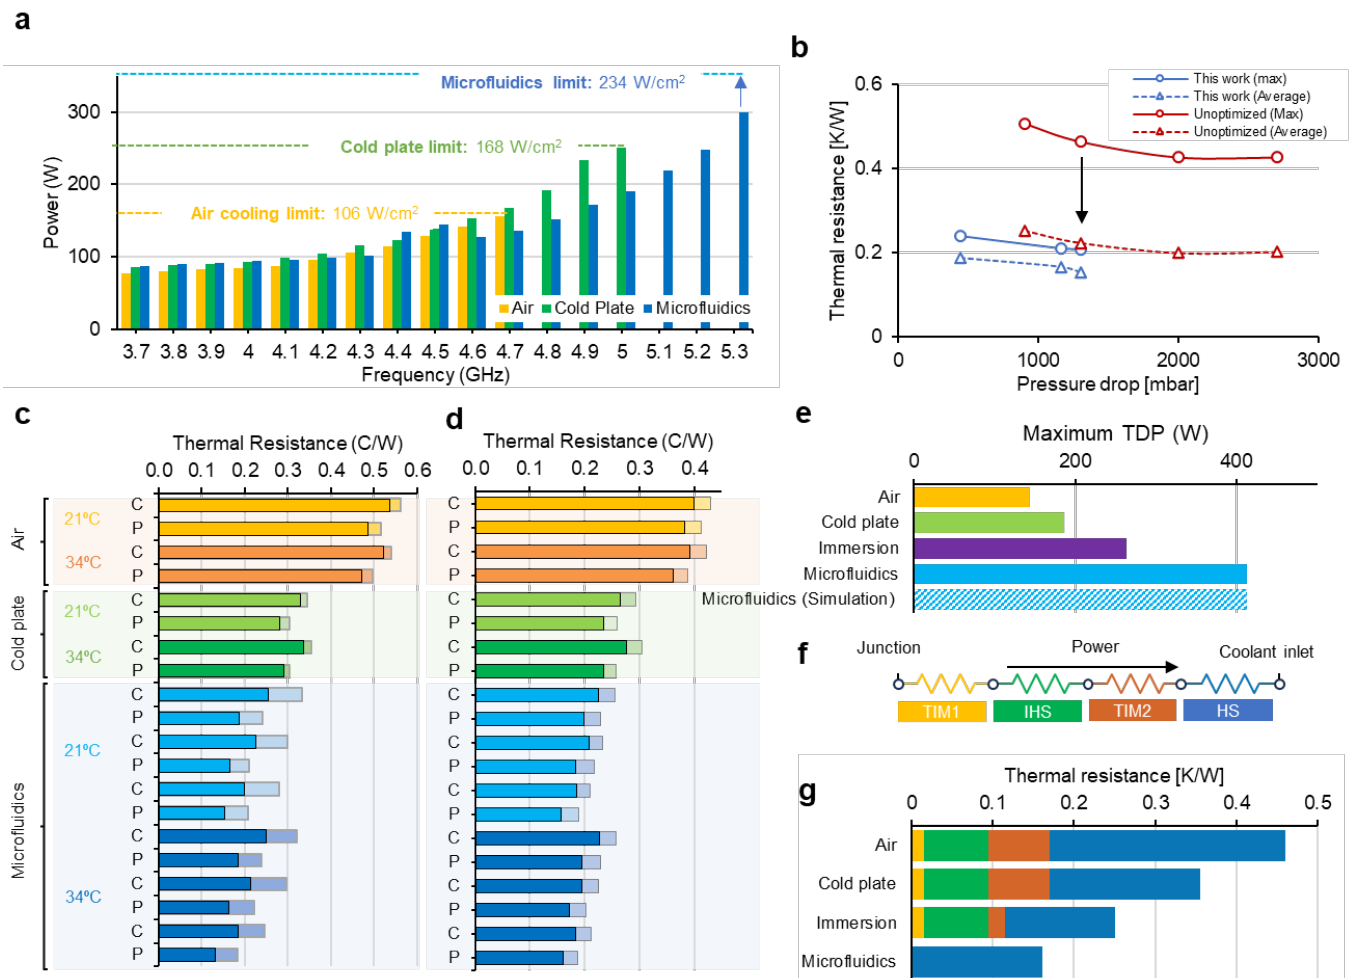
**Figure 1 | Methodology for power-aware microfluidic cooling optimization.** **a**, CPU with heterogeneous power distribution. The power map of this CPU is used to perform topology optimization, which subsequently is etched into **b**, a CPU with integrated microfluidic cooling. **c**, Sealing the microfluidic structure with a well-defined inlet and outlet plenum. **d,f,h,j,l** Intel i9 9900K. **e,g,i,k,m**, Intel i7 8700K. **d,e**, die layout, indication the location of the cores. **f,g**, power maps, with heat flux per cell. **h,i**, Optimized microfluidic design, **j,k**, Simulated fluid velocities, **l,m**, simulated junction temperatures.



**Figure 2 | Numerical results.** **a.** Temperature rise versus pressure drop for the i9. Green diamonds are split-flow pin fins with a diameter of  $109 \mu\text{m}$  at a horizontal pitch of  $200 \mu\text{m}$ . Orange triangles are split-flow straight channels, with an equal channel and fin width of  $50 \mu\text{m}$ . Blue circles are results for the chip-aware microfluidic design. **b.** Max TDP for the i9 for a  $70^\circ\text{C}$  junction-to-inlet temperature rise. The dotted lines refer to experimental results from the literature<sup>23</sup> using air-cooling, cold plate cooling and 2-phase immersion cooling with metallic TIM. **c.** Temperature rise versus pressure drop for the i7. Green diamonds are split-flow pin fins with a diameter of  $99 \mu\text{m}$  at a horizontal pitch of  $200 \mu\text{m}$ . Orange triangles are split-flow straight channels, with an equal channel and fin width of  $50 \mu\text{m}$ . Dark green dotted diamonds are a pin fin design similar to the one presented in the literature<sup>16</sup>. Blue circles are results for the chip-aware microfluidic design. **d.** Max TDP for the i7 for a  $70^\circ\text{C}$  temperature rise. Dashed black line indicate results for air-cooling on the same chip in stock configuration. **e-h:** Comparison of channel temperature and core temperature on the i7 8700K using pin fins (**e,g**) and the optimized design (**f,h**), with a pressure drop of  $380 \text{ mbar}$  and a coolant inlet temperature of  $20^\circ\text{C}$ . **g,h:** channel temperature. **e,f:** junction temperature. The pin fin design is similar to the design used in the literature<sup>16</sup>.



**Figure 3 | Etching microfluidic cooling channels in CPU die and microfluidic packaging design.** **a.** Etching process in a packaged CPU for the i7: First, the substrate is covered with Kapton tape. Photoresist is spin-coated over the chip. The cooling pattern is transferred into the photoresist using photolithography, and channels are finally etched in the die using deep reactive ions etching. **b.** Laser machining of channels in the i9 die. **c.** Close-up SEM image of the surface of the i7 after DRIE showing vertical sidewalls. **d.** Optical picture of the i7 after DRIE. **e,f,** Optical pictures of the i9 and i7, respectively, after laser etching. **g,** optical picture of the i7 after laser etching. **h,** Exploded view of the packaging assembly. **i,** step-by-step pictures of the packaging process. From left to right: The injection head is installed on the die with double sided adhesive and epoxy. The ring is installed over the chip and mounted through the motherboard into the backplate. The pusher is screwed onto the ring with four springs to load the CPU into the socket. The manifold is installed onto the injection head with 3 O-rings, the protective cover is installed, and finally the connectors are mounted on the package.



**Figure 4 | Experimental results.** **a** Overclocking procedure on the i7 using Prime95 and an inlet temperature of 21 C. Bars show the power at stable operating points. Heat flux per cooling configuration is given for a maximum junction-to-inlet temperature rise of 80 C. In case of microfluidics, overclocking was voltage limited and peak heat flux was extrapolated to an 80 C temperature rise. **b**. Comparison of thermal resistance versus pressure drop for chip-aware optimized microfluidics in this work (blue) versus unoptimized pin fins in the literature<sup>16</sup> on the same i7 (red). **c,d** Thermal resistance for all 20 test conditions. Solid bars correspond to thermal resistance based on average core temperature rise, and transparent bars based on maximum core temperature rise. **c**. For the i7, **d**. For the i9. **e**. Maximum TDP for a junction-to-inlet temperature rise of 70 C. **f**. Thermal resistance diagram. **g**. Thermal resistance breakdown for three experimental results on air cooling, cold plate cooling, and two-phase immersion cooling using a metal TIM, from literature<sup>23</sup>, compared with this work on the same chip (i9 9900K).

## Methods

### Chips and power maps

Two chips were selected as part of this study: the Intel i7 8700K (i7) and Intel i9 9900K (i9). The i7 was selected as a comparison point with a prior study<sup>16</sup> that evaluated the cooling performance of embedded micro pin fins. The i9 was selected as a comparison with a prior data set that benchmarked the cooling performance of air-cooling, standard liquid cooling, and immersion cooling<sup>23</sup>. By focusing on these two chips we can both assess the advantage of microfluidic topology optimization, as well as benchmarking these results with respect to other available cooling solutions.

The dimensions of the i7 and i9 chips are given in Extended Data Table 1.  $t$  is the total thickness of the die,  $H$  is the channel height and  $t_b$  is the remaining silicon thickness below the channels ( $t_b = t - H$ ). We consider  $H$  to be the maximum etching depth. One very important parameter in microfluidic cooling is the available silicon thickness on the back of the PCB that can be etched, i.e., the maximum etching depth. This thickness is essentially the channel (or pin fin) height, and it plays an important role in the cooling performance as it can dramatically increase the coolant fluid flow rate for a specific pumping pressure, which in turn can lead to a significant decrease in the resulting case/junction temperature.

**Extended Data Table 1** | Dimensions of the chips.

	Symbol	i7 8700K	i9 9900K
Die thickness	$t$	420 $\mu m$	870 $\mu m$
Die width	$W$	16.9 mm	19.7 mm
Die length	$L$	9.2 mm	9.2 $\mu m$
Channel depth	$H$	230 $\mu m$	500 $\mu m$
Base thickness	$t_b$	190 $\mu m$	370 $\mu m$

Spatial distribution of the power within the die is essential for the simulation and optimization of microfluidic cooling. Since the power maps of the two selected chips are not publicly available, we have employed the techniques described below to reconstruct and estimate power maps based on publicly known information. This technique combines die shots, software readouts, and the fitting of power profiles obtained by infrared measurements. High-resolution pictures, or die-shots, of the i7 and i9 chips, are publicly available and shown in Fig. 1 d,e. Using these die-shots, we can identify the spatial position of the various components within the chip, including 6 cores for the i7, and 8 cores for the i9.

The software tool “Open Hardware Monitor” was used to measure both the total power on the chip (package power) and the power per core (core power) during the operation of the chips. This way, we can distinguish the power per core and the other features on the chip. Using readouts of these software tools from various sources, the power levels for the i7 and the i9 in Extended Data Table 2 were identified. Since we are interested in the highest-power case, we selected the overclocked case (OC) for further analysis.

**Extended Data Table 2** | Power distribution per chip under normal clock speed and at 5.0 GHz.

Chip	Clock speed	Power stressed	Power/core	# cores	Background power
i7 8700K	3.7 GHz	146 W	21.7 W	6	10 W
	5.0 GHz	210 W	31.7 W	6	20 W
i9 9900K	3.6 GHz	168 W	19.8 W	8	10 W
	5.0 GHz	298 W	37.5 W	8	20 W

Based on the die shot and the power levels, a heterogeneous power map can be established by averaging power by the surface area of each component. However, even within the cores, a non-uniform power distribution can be expected. The strategy pursued to model the power distribution inside each core is to derive a typical profile of core-level power distribution and distribute the power per core to match this profile.

Zhang et al.<sup>24</sup> employed infrared temperature measurements to determine temperature profiles on the die during operation and solved an inverse problem to obtain an estimation of the underlying power distribution. The die shot (Fig. 1d,e), and corresponding power maps (Fig. 1f,g) show two distinct peaks per core and a higher heat flux on the ring/interconnect than the rest of the die.

### Numerical model and topology optimization

Topology optimization was used to design the optimal channel network. As such, the design of the cooling structure was posed as a mathematical problem wherein the objective function is the minimization of the temperature at the hotspots and the design variable is the geometry. Other performance goals can be added to the problem as constraint functions such as the maximum pressure drop. Most importantly, the performance metrics need the calculation of the temperature across the entire chip, which is done solving the conjugate heat transfer physics, wherein the Navier-Stokes equations are coupled with an advection-diffusion equation to model the heat transfer.

Modeling a three-dimensional microfluidics network that spans multiple scales would involve tens of millions of degrees of freedom, making the simulation computationally expensive, and the optimization unfeasible. However, the channels in microfluidic cooling structures have a thickness much smaller than the chip width and length, allowing a pre-described profile on the velocity and temperature in the out-of-plane direction, reducing the 3D model to a 2D model with additional terms to account for the thickness effects. Therefore, we chose to employ a 2.5D model<sup>17</sup> for better computational efficiency.

The hydrodynamics are governed by the following 2.5D Navier-Stokes equations where we solve for the velocity field  $\mathbf{u}$  and pressure  $p$ :

$$\begin{aligned} \frac{6}{7}\rho_0\mathbf{u} \cdot \nabla\mathbf{u} &= -\nabla P + \mu\nabla \cdot (\nabla\mathbf{u} + (\nabla\mathbf{u})^T) - \frac{5\mu}{2H_t^2}\mathbf{u} = 0 \\ \nabla \cdot \mathbf{u} &= 0 \end{aligned}$$

$\mu$  and  $\rho_0$  are the coolant viscosity and density respectively.  $H_t$  represents the half-channel thickness and accounts for the flow drag in the dimension orthogonal to the flow direction, i.e. it models the thickness effect. Once the velocity is known, we feed it into the heat transfer equations modeling the top layer, where the channels and fins are located, and the bottom layer, consisting of the remainder solid portion of the chip.

$$\begin{aligned} \frac{2}{3}\rho_0 C (\mathbf{u} \cdot \nabla T_t) - \frac{49}{52}\nabla \cdot (k_t \nabla T_t) - \frac{h}{2H_t} (T_b - T_t) &= 0 \\ -\frac{k_b}{2}\nabla^2 T_b + \frac{h}{2H_b} (T_b - T_t) - \frac{q_0''}{2H_b} &= 0 \end{aligned}$$

Where we solve for the temperature at the top  $T_t$  and bottom  $T_b$  layers. The coolant has a density of  $\rho_0$ , a specific heat of  $C$ . The top layer has a conductivity of  $k_t$  which varies between the coolant thermal conductivity  $k_f$  or the solid  $k_b$  depending on whether it is a channel or a solid fin. Similarly, the heat transfer coefficient  $h$  accounts for the thermal conductivity between the top and bottom layers and depends on whether the top is a channel or a solid fin. Lastly, the chip emits a spatially varying heat flux

represented by  $q_0''$  which is prescribed by the power map of each chip. Material properties used are shown in Extended Data Table 3.

**Extended Data Table 3** | Material properties for numerical model

Material	Property	Symbol	Unit	Value
Water	Density	$\rho$	$kg\ m^{-3}$	998
	Dynamic viscosity	$\mu$	$kg\ m^{-1}s^{-1}$	0.001
	Specific heat	$C$	$J\ Kg^{-1}K^{-1}$	4184
	Thermal conductivity	$k_f$	$Wm^{-1}K^{-1}$	0.598
Silicon	Thermal conductivity	$k_b$	$Wm^{-1}K^{-1}$	148

Prior frameworks for topology optimization of liquid cooling<sup>17</sup> are not directly applicable to real-world examples where Reynolds number are much higher due to an operational pressure drop value up to 1 bar. Because the convergence radius of Newton-Raphson methods is inversely proportional to the Reynolds number (Re), we perform a continuation strategy on Re, whereby solutions for lower Re are used as initial estimation for solving the Navier-Stokes equation for higher Re until we reach the original Re value.

We solve these equations using finite elements with linear basis elements on triangles. To stabilize the convective scales, we use the Streamline Upwind Petrov-Galerkin (SUPG) for the Navier-Stokes and the heat transfer equations and the Pressure Stabilizing Petrov-Galerkin methods (PSPG) for the Navier-Stokes equations to control the instabilities arising from the inf-sup condition. These equations were implemented and solved using an in-house developed finite elements method (FEM) platform.

In this work, we follow the density method, in which the geometry is defined with a continuous field called density that is bounded between 1 (liquid phase) and 0 (solid phase). To favor discrete designs, we used the RAMP penalization method<sup>25</sup> on the thermal conductivity and the channel thickness. The topology optimization of a channel network is an ill-posed problem in the sense that the optimal design is one that contains infinitesimally small features to increase the surface area ad-infinitum. We use a filter operation<sup>26</sup> to ensure no geometric feature has a size smaller than a user-determined threshold. In our case, it is 100 micrometers. Once the optimization problem is set up, we calculate the gradients using the adjoint method and feed them to the method of moving asymptotes (MMA) optimization algorithm<sup>27</sup> to find the optimal channel network.

### Channel fabrication

The final channel designs after optimization are exported as vector files for subsequent integration in the silicon die. First, the integrated heat spreader (IHS) was removed from the chips to gain access to the silicon die. For the i9, heat was applied during this delidding process to ensure the solder-based thermal interface material was melted while removing the IHS, to avoid damaging the silicon die.

Two processes were explored for integrating the microfluidic cooling inside the die, one based on plasma etching for the i7, and one method based on laser machining for the i9. For the plasma etching process, the i7 CPU was installed on a silicon carrier wafer. Due to the wafer-level nature of silicon chip manufacturing, the use of wafer-scale tools to integrate cooling channels in the CPU is a preferred choice for scalability. During a single wafer-scale etching step, hundreds of devices can simultaneously be provided with microfluidic cooling. However, for prototyping, we are limited to packaged chips, which requires an alternative process. Figure 3a outlines the process for deep-reactive ion etching in packaged chips. The CPU was first installed on a 4-inch silicon carrier wafer, to be compatible with wafer-scale equipment. Next, the organic substrate of the chip is covered with Kapton tape. A layer of photoresist is

then spin-coated on the wafer, covering the exposed silicon surface of the die. The cooling pattern was transferred to the photomask using direct laser writing (MLA 150, Heidelberg Instruments), followed by a development step. The carrier wafer with CPU was then placed in an inductively coupled plasma etching chamber (Alcatel AMS 200 SE), to perform deep reactive ion etching (DRIE) using repeating pulses of SF<sub>6</sub> and C<sub>4</sub>F<sub>8</sub> to etch and passivate deep silicon trenches. After etching, the protective kapton film was removed and the photoresist was stripped from the CPU, resulting in a CPU with integrated cooling channels.

Despite the good yield and clear path for scaling up production at wafer-scale, performing DRIE on packaged CPUs with organic substrates is a time-consuming process. In addition, on CPUs with a soldered IHS such as the i9, a thin film of metal is present on the surface of the die after removing the TIM. This metal film is incompatible with the DRIE process. We therefore utilized an alternative method for integrating channels into the i9 using laser micromachining, by performing the following steps: The i9 was positioned in a custom laser micromachining tool using micron-level precision CNC X-Y-Z-theta stages. Next, the desired cooling design was etched into the surface of the silicon using a nanosecond-pulse laser operating at a wavelength of 355 nm. The laser beam was focused onto the surface of the silicon using a f-theta scanning lens and galvanometer scanner. The pattern was machined to a depth of 380 microns in the surface of the silicon, as measured using a Keyence VHX-7000 optical microscope. The laser etching process was optimized to produce well-defined features with no evident cracking or heat-affected zones in the chips. The laser etching was performed in ambient conditions in air with no assist gas, and following etching the chip was lightly air-dusted to remove small particles of debris.

## Packaging

Figure 3h shows an exploded view of the packaging solution designed for this application. From bottom to top: A backplate was installed behind the socket on the motherboard to provide mechanical rigidity. Next, the *Ring* was placed on top of the chip to provide a reference plane for the force required to push the chip into the LGA1151 socket. The Injection head was bonded to the die with double-sided adhesive (3M), to seal the microfluidic channels and provide well-defined inlets and outlets. A ring of epoxy was poured around the injection head as a mechanical lock for the injection head and a second barrier for leaks in case of adhesive failure.

Figure 3i shows the injection head installed on the CPU, followed by the ring installed on the chip. Next, the pusher was installed on the chip with four springs, to provide a balanced distribution of force required to contact the landing grid array (LGA) to the spring-loaded contacts in the socket. A combined force of 300N is required to make electrical contact between the CPU and the LGA1151 socket. The manifold was then screwed onto the injection head, creating a seal using three O-rings. Threads were integrated in the injection head, to separate the compression force on the O-rings from the die. This method prevents directly compressing the die which can lead to a more curved PCB and the creating disconnection. To reduce the bending of the injection head during O-ring compression, contact points were bonded between the injection head and the PCB creating a closed loop for the force path through the injection head. It also absorbs any torque or momentum induced by the package during manipulation. Finally, a 3D-printed SLS in Nylon (PA2) cover was installed to seal and lock the assembly in place, before installing fluidic connectors and temperature sensors to the package.

To ensure leak-tight operation, we performed a pressure decay test after the assembly by pressurizing the system to 500 mbar for 15 minutes. If the drop in pressure does not exceed 100 mbar, the system was validated for the next step, consisting of hydraulic testing with coolant. A hydraulic test for final validation was performed with our custom CDU at a peak pressure of 2 bar.

## Experimental setup

The two microfluidics-cooled chips were compared against corresponding pristine chips subjected to both air and cold plate cooling. The same type of motherboard was employed during all tests since the i7 and i9 share the same socket. The motherboard was installed in a transparent case to create a temperature-controlled environment and to minimize the impact of external irregularities on the airflow during the experiments. Benchmarks were performed at varying inlet temperatures which, for air cooling, was defined as the air inlet temperature measured upstream of the heat sink. For cold plates and microfluidic cooling, inlet temperature was defined as the liquid temperature measured upstream of the chip.

For microfluidic cooling, the assembled package was installed on the motherboard, and hoses were routed out of the case through a PCIe hydraulic pass-through and connected to a custom-built coolant distribution unit (CDU). DI water with corrosion and biological inhibitors was used as a coolant and circulated through the package and heat exchanger using a gear pump (Drifton WT3000-1JA). Two pressure ports were installed before the inlet of the package and after the outlet of the package, to measure the total pressure drop over the package using a differential pressure sensor. A flow meter (Omega FLR1012) was used to monitor the volumetric flow rate. Coolant temperature was measured directly at the inlet and outlet ports of the package using in-line thermocouples. Two parallel 7 micrometer filters were installed upstream of the CPU. To control the coolant inlet temperature at elevated temperature, an inline resistive heater was installed upstream of the CPU, which regulated the inlet temperature of the test section using a PID controller.

For cold plate cooling, an off-the-shelf copper cold plate (EKWB) was used on both i7-8700K and i9-9900K with the same hydraulic loop. Thermal grease (Shin-Etsu X23-7783D) was used as thermal interface material, and the cold plate was tightened onto the chip up to the manufacturer's recommended loading.

For air cooling, a commercially available air-cooled heat sink was used (Intel XT100H). The heat sink was mounted on the CPU using the same thermal grease (Shin-Etsu X23-7783D) as thermal interface material, and tightened onto the chip, up to the manufacturer's recommended loading. For inlet temperatures below ambient, which was around 25C during the experiments, an air conditioning system was connected to the case to stabilize the temperatures, whereas at elevated temperatures a convective heater with relay and PID control was employed to stabilize the inlet temperature.

## Experimental procedure

The performance of microfluidics-cooled chips was experimentally evaluated by performing computational benchmarks using a Windows PC, while gradually overclocking the system to higher clock speeds to stress the CPU. The BIOS was modified according to the settings in Extend Data Table 4 to allow effective overclocking of the system.

**Extended Data Table 4** | BIOS settings for overclocking

<b>Ai Overclock Tuner</b>	XMP I	<b>Long Duration Package Power</b>	4095
<b>SVID Behavior</b>	Best-Case Scenario	<b>Package Power Time Window</b>	127
<b>CPU Core Ratio</b>	Sync All Cores	<b>Short Duration Package Power</b>	4095

<b>DRAM Frequency</b>	DDR4-3000MHz	<b>IA AC Load Line</b>	0.01
<b>DRAM CAS# Latency</b>	16	<b>IA DC Load Line</b>	0.01
<b>DRAM RAS# ACT Time</b>	38	<b>MAX CPU Cache Ratio</b>	40
<b>CPU Load-line</b>	Level 6	<b>CPU Core/Cache Voltage</b>	Adaptive Mode
<b>CPU Current Capability</b>	140%	<b>DRAM Voltage</b>	1.35
<b>CPU C-states</b>	Disabled	<b>CPU Q-Fan Control</b>	Disabled

With the above settings, the clock speed and voltage were controlled during operation with the help of Intel Extreme Tuning Utility software (XTU). The goal of the overclocking study is to achieve the maximum possible core frequency with lowest power as possible while maintaining a stable system. The following iterative steps were followed:

1. CPU clock frequency increased by 0.1GHz in the XTU tool.
2. Run stress test.
3. If stress tests fail due to thermal issue, reduction in the BIOS of the voltage core by increments of 0.01V.
4. If the system fails immediately after booting up again (blue screen) -> XTU -> increase by 0.01V the core voltage.
5. If workloads execute successfully, reboot the system while increasing the frequency by 0.1 GHz. Repeating steps 3 and 4 until success.
6. Making sure that inlet temperature and flow rate are stable before recording test data.
7. As a stability criterion, Prime95 was launched for 5 minutes and Cinebench R20 for three consecutive times with no error.

For all the experiments, the starting voltage for the cores was around 1.2 volts, and maximum voltage was limited to 1.47 volts as a safety precaution. This overclocking sequence was performed for each, for a total of 40 experimental conditions. For both chips, 3 cooling methods were evaluated (air, cold plate and microfluidics). For microfluidics, 3 flow rates were evaluated, namely 15 ml/s, 20 ml/s and 22.5 ml/s. For each combination of cooling and flow rate, two inlet temperatures were tested (21 C and 34C), and finally performance was measured using two benchmarks, Prime95 (P) and Cinebench R20 (C). Extended Data Table 5 and 6 shows a summary of all 40 experimental conditions, on the i7 and i9, respectively.

**Extended Data Table 5** | Experimental conditions for the i7 8700K

<b>Chip</b>	<b>i7 8700K</b>																			
<b>Cooling type</b>	Air				Cold plate				Microfluidics											
<b>Flow rate</b>	-				15 ml/s				15 ml/s				20 ml/s				22.5 ml/s			
<b>Inlet temperature</b>	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C				
<b>Benchmark</b>	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P

**Extended Data Table 6** | Experimental conditions for the i9 9900K

<b>Chip</b>	<b>i9 9900K</b>																			
<b>Cooling type</b>	Air				Cold plate				Microfluidics											
<b>Flow rate</b>	-				15 ml/s				15 ml/s				20 ml/s				22.5 ml/s			
<b>Inlet temperature</b>	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C	21 C	34 C				
<b>Benchmark</b>	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P	C	P

## **Data processing**

During each experiment, all core temperatures were monitored using *Open Hardware Monitor*. The temperature of the hottest core during each experiment is referred to as *hot-spot* temperature, and the average of all cores (either six or eight), is referred to as average core temperature. Junction-to-inlet temperature rise was calculated for each core by subtracting the inlet temperature from the core temperature. Maximum thermal resistance was calculated by dividing the hot-spot temperature rise by the total package power, and average thermal resistance was calculated by dividing the average core temperature by the total package power.

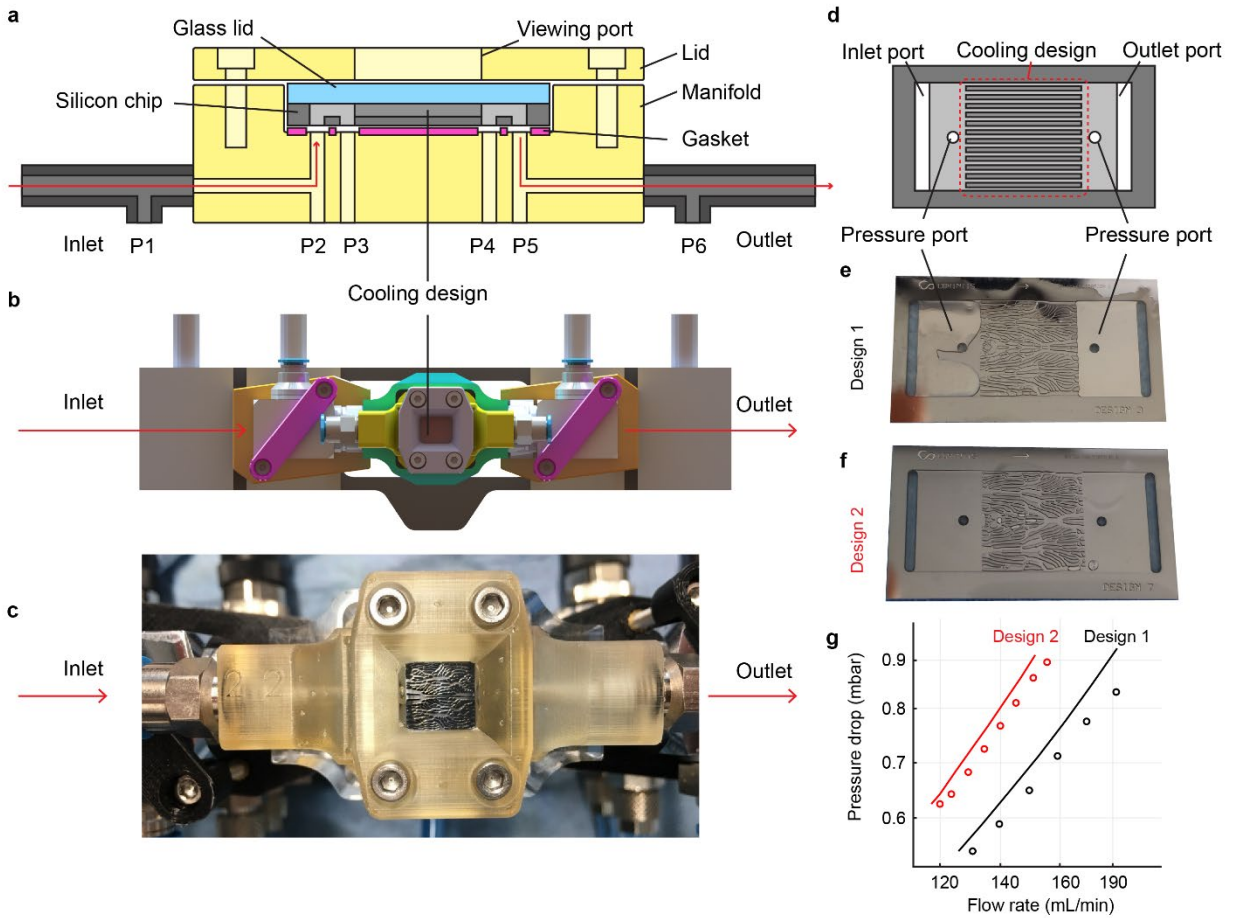
**Acknowledgments** We thank the staff at the Center of Micro and Nano Technology (CMi) for support and advice on the fabrication processes.

**Author contributions**

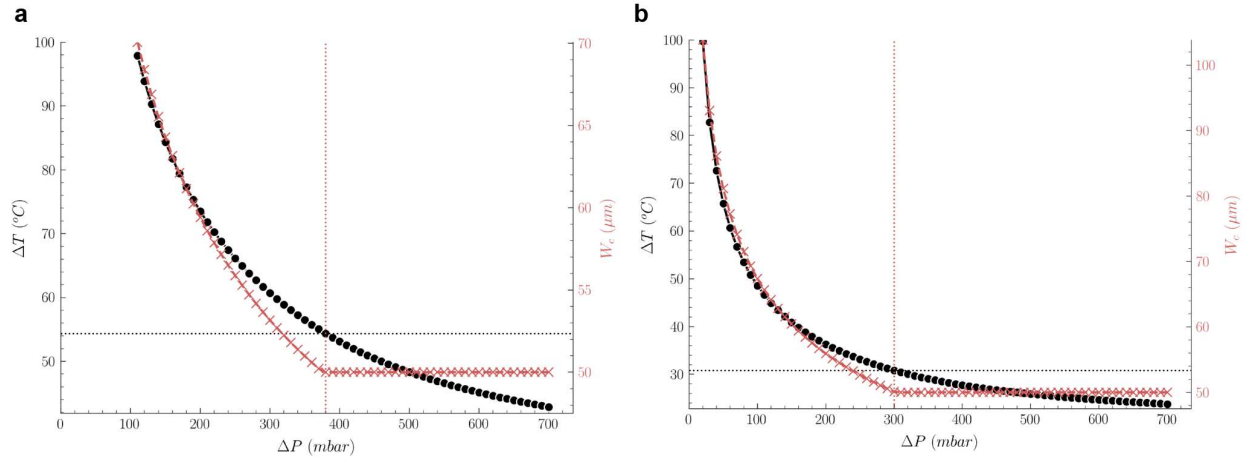
**Competing interest declaration**

**Correspondence and requests for materials** should be addressed to R.v.E.

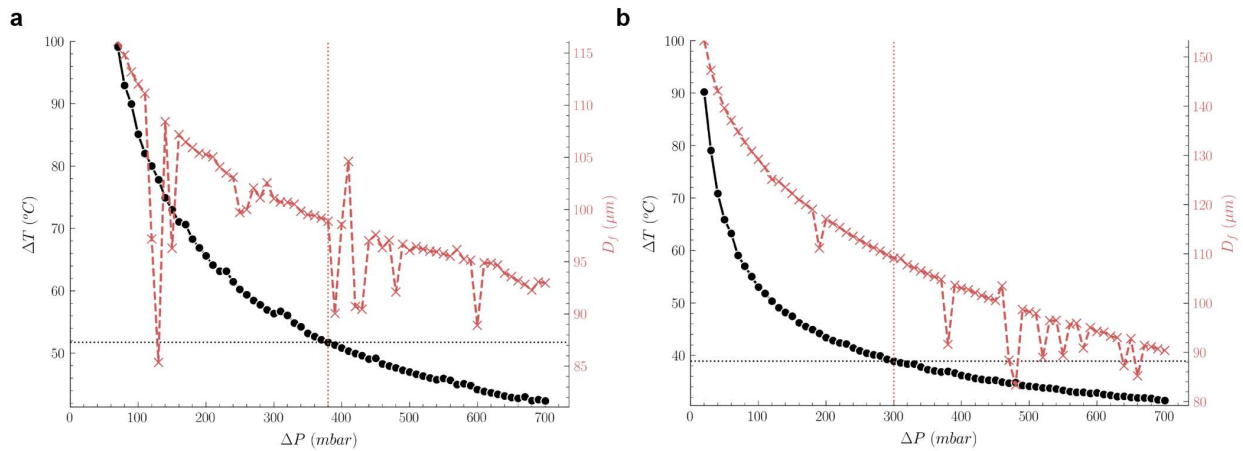
## Extended data




**Extended data Fig 1 | Experimental validation of the numerical model** **a:** Experimental setup: A device under test, consisting of a silicon chip bonded to a glass lid, is connected with an intermediate gasket to a manifold. P1, P2 and P3 are pressure upstream pressure ports, and P4, P5 and P6 are downstream pressure ports. **b:** 3D model of the experimental setup. **c:** Top view of the experimental setup with a sample installed. **d:** Illustration of a device under test, the area marked in red is the cooling design. **e,f:** Picture of sample 1 and sample 2. **g:** Experimental results of the pressure drop measured between the pressure ports embedded inside the device under test for varying flow rates (markers), compared to simulation results (solid lines).



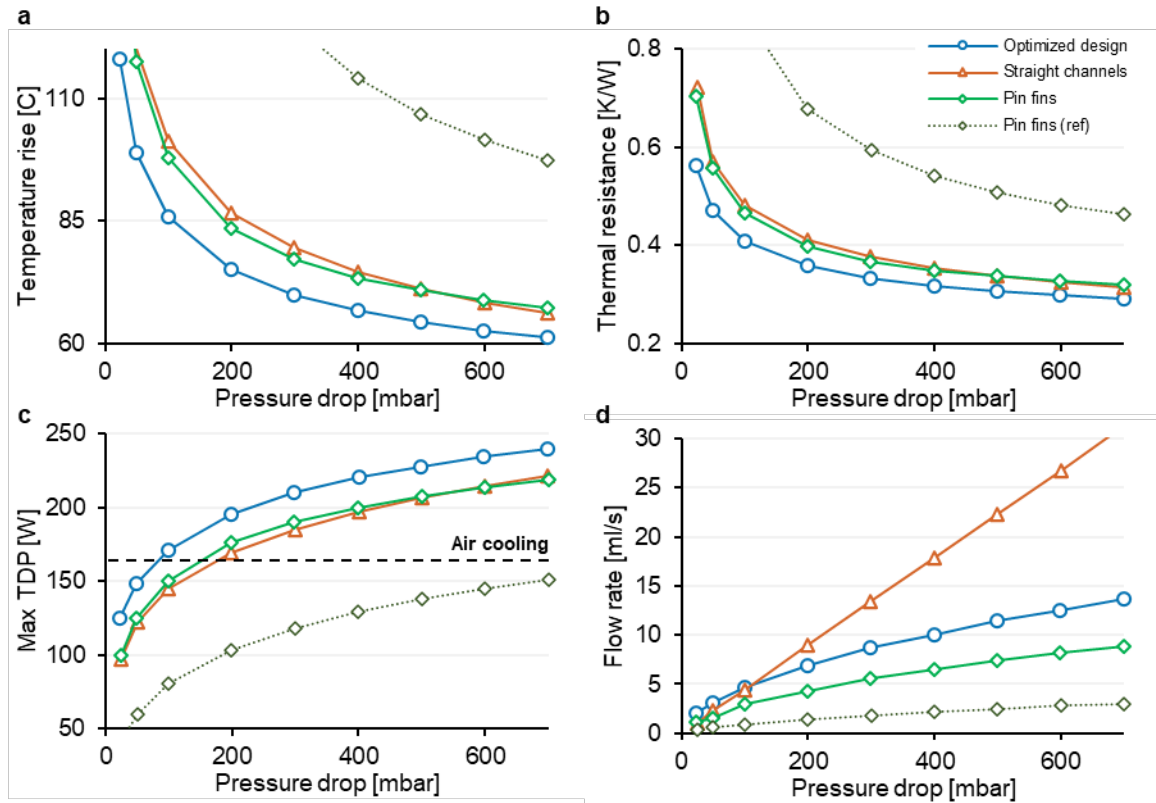
**Extended data Fig 1 | Pareto fronts for straight channel optimization. a:** results for i7. At a target pressure drop of 380 mbar, the minimum allowable 50  $\mu\text{m}$  channel width is the optimum. **b:** Results for the i9, the minimum allowable 50  $\mu\text{m}$  channel width is the optimum at a target pressure drop of 300 mbar.




**Extended data Fig 2 | Pareto fronts for pin fin optimization. a:** Pareto of optimized solutions for the pin fin design for i7 8700k at a power of 210 W. This graph depicts two different quantities for a pressure sweep in [25, 700] mbar (x-axis): the red X markers denote the optimized pin diameter for every pressure drop (right y-axis), while the solid black O markers denote the temperature increase for such a design (left y-axis). **b:** Pareto of optimized solutions for the pin fin design for i9 9900k at a power of 298 W. This graph depicts two different quantities for a pressure sweep in [25, 700] mbar (x-axis): the red X markers denote the optimized pin diameter for every pressure drop (right y-axis), while the solid black O markers denote the temperature increase for such a design (left y-axis). *N.B. The Pareto fronts are discontinuous because there are two other design variables not shown here (longitudinal and transversal pitch), the combination of which leads to a minimum thermal resistance.*



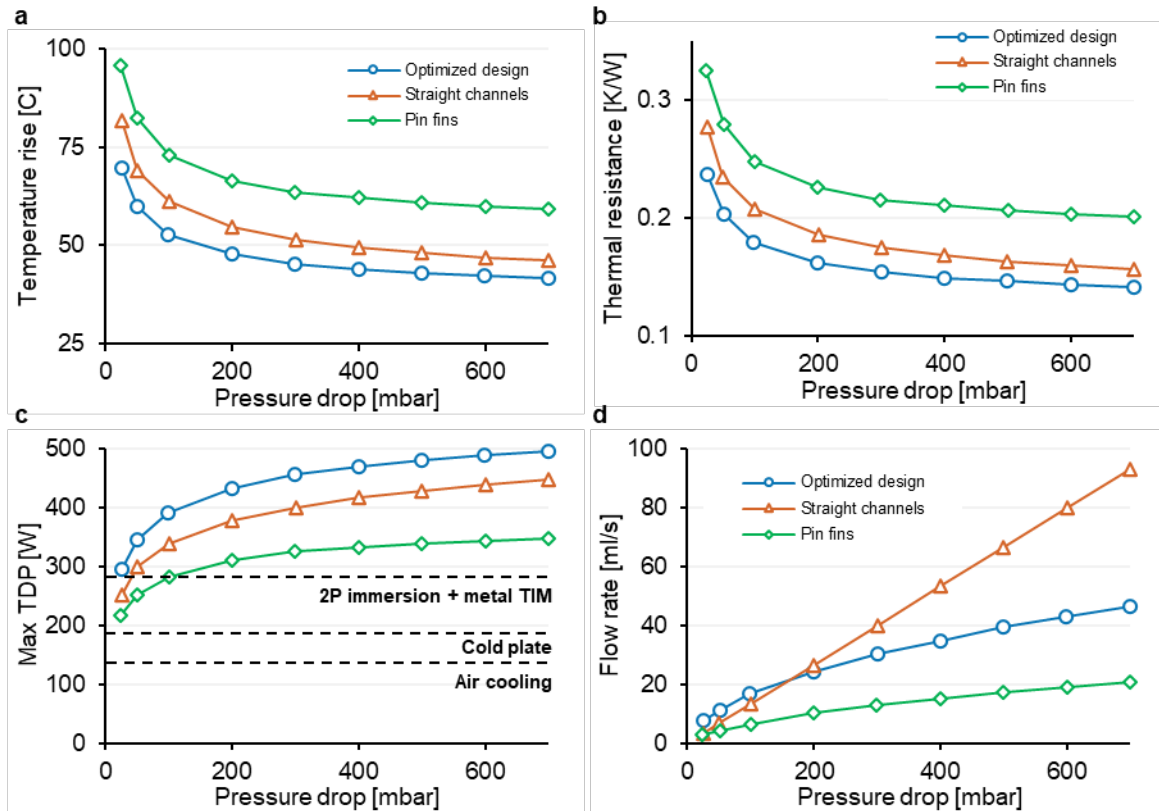
**Extended data Fig 3 | Numerical results for optimized design of the i7: (top) velocity field, (middle) Channel temperature, and (bottom) junction temperature.**




**Extended data Fig 4 | Comparison of canonical designs and optimized microfluidics:** Orange triangular markers: Straight channels. Green diamond markers: Pin fins (split flow). Dark green markers: Pin fins, no split flow, as presented in the literature<sup>16</sup>. Blue circular markers: Optimized cooling design. **a.** Temperature rise versus pressure drop. **b.** Junction-to-inlet thermal resistance. **c.** Max TDP for a 70 °C temperature rise. Dashed black line indicate results for air-cooling on the same chip in stock configuration. **d.** Flow rate versus pressure drop.




**Extended data Fig 5 | Numerical results for optimized design of the i9: (top) velocity field, (middle) Channel temperature, and (bottom) junction temperature.**



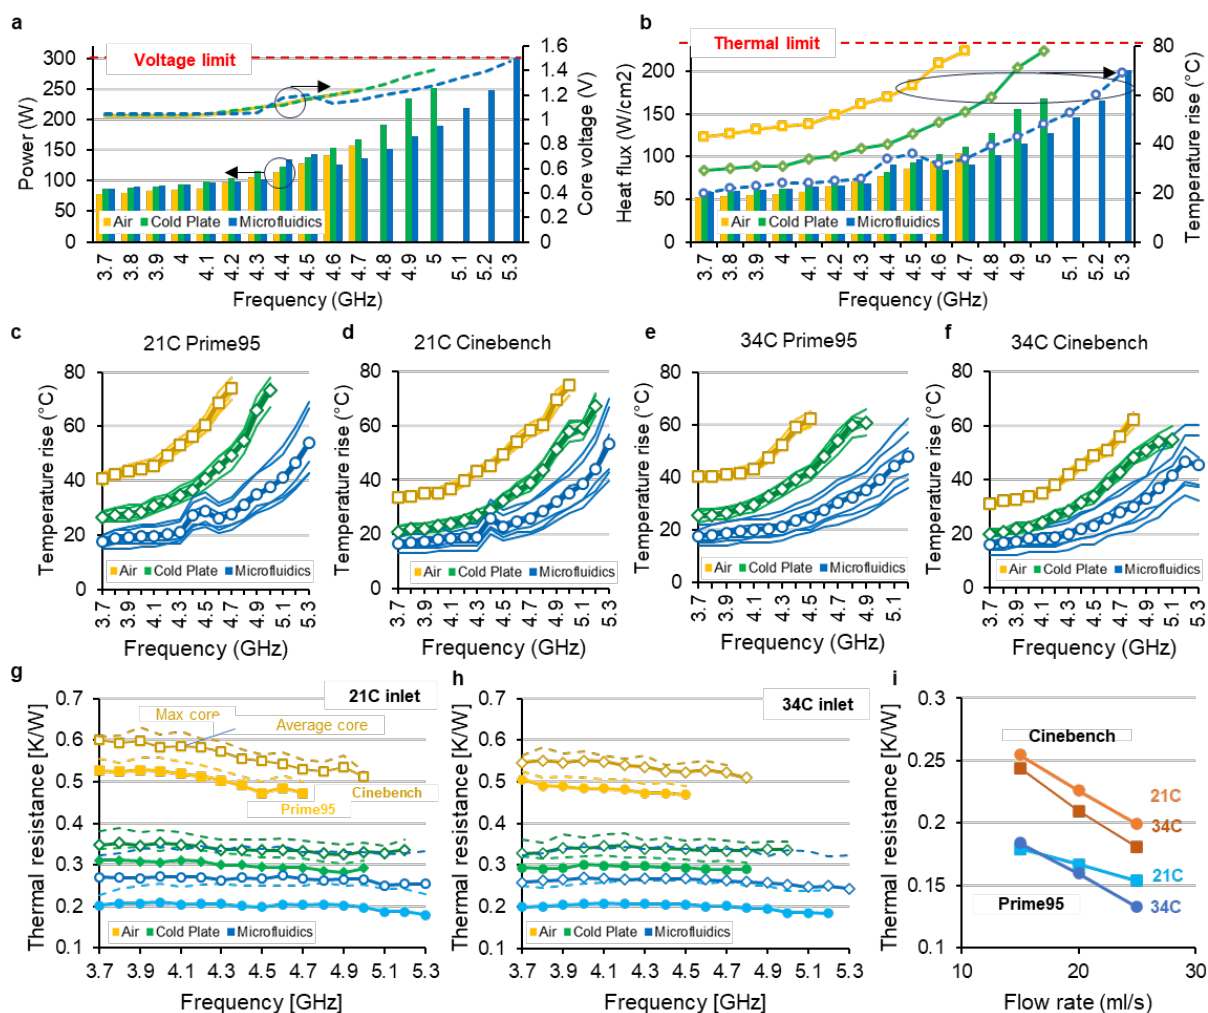
**Extended data Fig 6 | Comparison of canonical designs and optimized microfluidics for the i9:** Orange triangle markers: Straight channels. Green diamond markers: Pin fins. Blue markers: Optimized cooling design. **a.** Temperature rise versus pressure drop. **b.** Thermal resistance versus pressure drop. **c.** Max TDP for a 70 °C junction-to-inlet temperature rise. The dotted lines refer to experimental results from the literature<sup>23</sup> using air-cooling, cold plate cooling and 2-phase immersion cooling with metallic TIM. **d.** Flow rate versus pressure drop



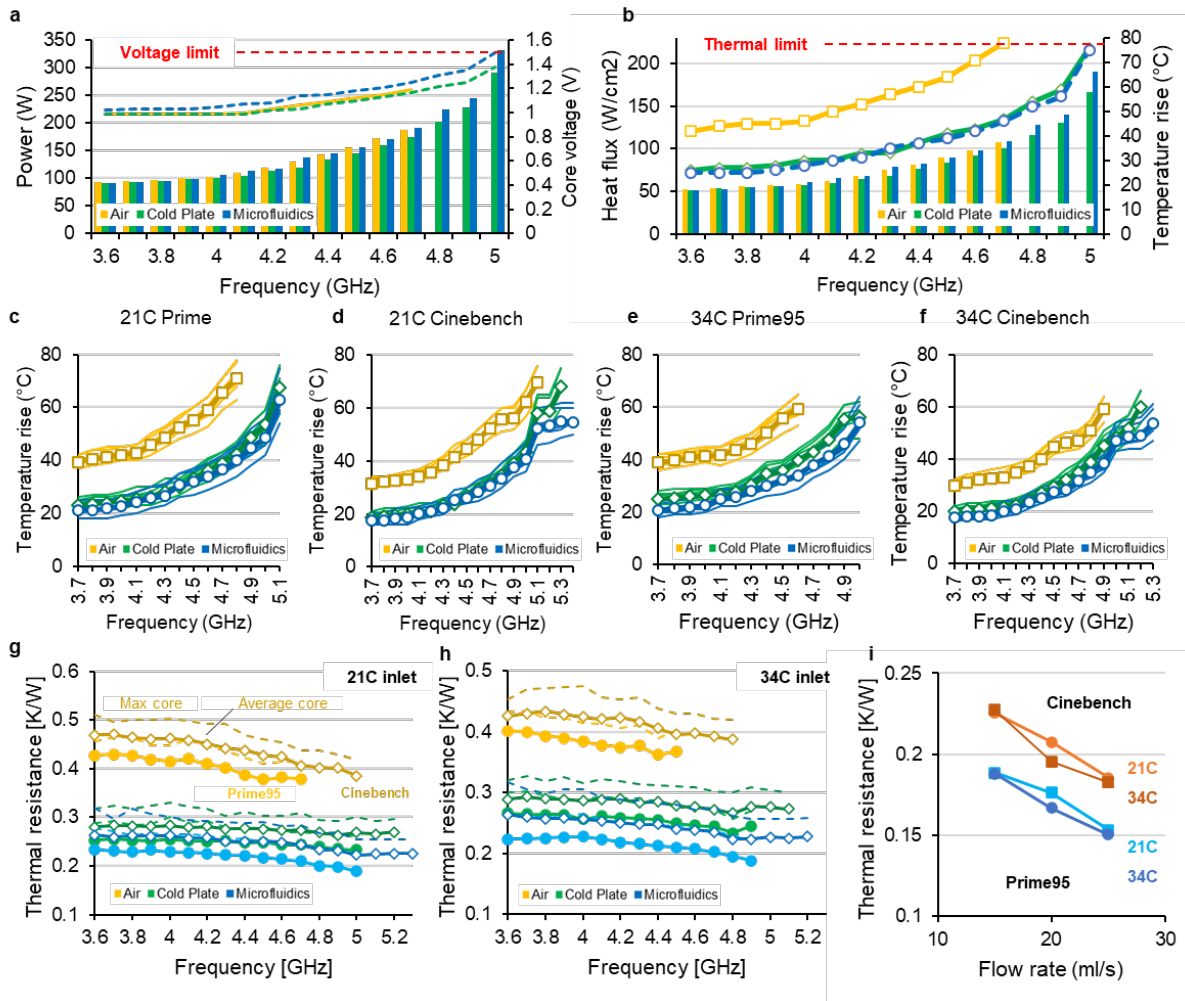
**Extended data Fig 7 | Fabrication process for deep reactive ion etching in the i7. a.** Mounting the CPU on a carrier wafer and covering with Kapton tape. **b.** spin coating photoresist. **c.** Lithography and development. **d.** Etching.



**Extended data Fig 8 | Packaging assembly model and experimental setup for benchmarking: a.** Exploded view of the packaging design. **b.** Cross-section of the microfluidic packaging. **c:** Experimental setup for cold plate cooling. **d:** Experimental setup for microfluidic cooling.



**Extended data Fig 9 | Overclocking process for the i7. a:** Power (bars) and core voltage (lines) for each stable point reached during overclocking. Microfluidics was limited by maximum safe core voltage of 1.5 V at 5.3 GHz as the total power reached 300 W. **b:** die-level heat flux (bars) and maximum core temperature rise (lines) for each stable point reached during the overclocking, indicating that at 5.3 GHz the thermal limit was not reached for microfluidics. **c-f:** Core temperature rise (lines) and average core temperature rise (markers) for each stable point. **c,d:** 21 °C inlet temperature, **e,f:** 34 °C inlet temperature, **c,e:** Prime95 benchmark, **d,f:** Cinebench benchmark. **h,i:** Thermal resistance for each stable overclocking point. Solid markers are results from Prime95 and open markers are Cinebench. Solid lines represent thermal resistance based on average core temperature rise and dashed lines based on maximum core temperature rise. **h:** for an inlet temperature of 21 °C. **i:** for an inlet temperature of 34 °C. **i:** Average thermal resistance for Prime 95 (blue) and Cinebench (Orange). Circular markers are for an inlet temperature of 21 °C, square markers for an inlet temperature of 34 °C.



**Extended data Fig 10 | Overclocking process for the i9.** **a:** Power (bars) and core voltage (lines) for each stable point reached during overclocking. Microfluidics was limited by maximum safe core voltage of 1.5 V at 5.0 GHz as the total power reached 300 W. **b:** die-level heat flux (bars) and maximum core temperature rise (lines) for each stable point reached during the overclocking, indicating that at 5.3 GHz the thermal limit was not reached for microfluidics. **c-f:** Core temperature rise (lines) and average core temperature rise (markers) for each stable point. **c,d:** 21 °C inlet temperature, **e,f:** 34 °C inlet temperature, **c,e:** Prime95 benchmark, **d,f:** Cinebench benchmark. **g,h:** Thermal resistance for each stable overclocking point. Solid markers are results from Prime95 and open markers are Cinebench. Solid lines represent thermal resistance based on average core temperature rise and dashed lines based on maximum core temperature rise. **g:** for an inlet temperature of 21 °C. **h:** for an inlet temperature of 34 °C. **i:** Average thermal resistance for Prime 95 (blue) and Cinebench (Orange). Circular markers are for an inlet temperature of 21 °C, square markers for an inlet temperature of 34 °C.

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Media1.mp4](#)