

# Liquid cooling for hot AI accelerator chips

- Bert Gyselinckx -

The first liquid-cooled computers that grabbed my attention were gaming machines, back in the early '90s. Games like *Prince of Persia* and *Wolfenstein* were just getting into 3D rendering mode, and avid gamers needed all the CPU speed they could get to make these games playable in real time. So, gamers started overclocking their CPUs: they installed faster clocks in their PCs to get additional horsepower for their 3D games. Making the processors run faster also meant making them run hotter. So much so that a traditional air fan provided insufficient cooling. To avoid thermal meltdown of the processors and graphics cards, they turned to liquid cooling.

Many things have changed since then, but the physics and potential of liquid cooling for chips have remained the same. Liquid cooling is more efficient than air cooling because liquid has a higher heat capacity and thermal conductivity than air. This means it can carry away heat more effectively, allowing for better chip cooling.

Despite the clear benefits of liquid cooling over air cooling, it never became mainstream. It was too costly and too cumbersome to maintain and install. That is about to change, big time. Artificial intelligence (AI) and high-performance computing (HPC) drive this change.

The advent of generative AI based on trillion-parameter large language models requires massive amounts of computation and data storage and transfer. To satisfy these demands, CPU and GPU companies are cranking out ever more performing chips. The epitome of this movement is the [NVIDIA Blackwell](#) architecture, announced in March 2024. Chips using this architecture will be integrated into both air-cooled and liquid-cooled servers. The obvious speculation is that the air-cooled systems are designed to satisfy the current legacy data center infrastructure managers. Most of them are managing an infrastructure designed for air cooling and will need time to upgrade their infrastructure to become liquid-cooling-ready. Till then, they can opt for the air-cooled Blackwell systems. The future, however, is in the liquid-cooled systems. They offer the best performance and efficiency. Dell also recently announced liquid-cooled PowerEdge servers that make use of direct-to-chip cooling from startup [JetCool](#), which recently got acquired by [Flex](#). New advanced AI and HPC data centers are expected to move directly to these new liquid-cooled compute marvels. It is only a matter of time before the remaining data centers switch to support liquid cooling for all their AI and HPC servers.

In this paper, I will investigate market and technology trends and depict potential options for liquid-cooled AI and HPC systems. I believe that the market for liquid cooling has reached an inflection point and that the chip cooling ecosystem is about to be disrupted.

## The market for liquid cooling has reached an inflection point

Current AI accelerator chips like the AMD MI300 or the NVIDIA H100 have a thermal design power (TDP - a metric for how much power needs to be drawn away from the chip to remain operational) of around 500W. This is about the limit of what very advanced air cooling can deliver. Moreover, such chips are usually assembled on larger blades that go into a server rack. Such racks now have a total power consumption of several 10s of kW. Again, at the limit of what air cooling can supply. So, either designers find a way to make those AI chips more power efficient and run cooler, or we need to find a better way to cool the next generation of more performant and power-hungry AI chips and server racks.

Many attempts are ongoing to design more power-efficient AI chips, ranging from in-memory computing via analog to optical computing. This topic justifies a separate analysis. In this article, we will take a closer look at the different cooling options.

With a TDP of 1000W, the recently announced NVIDIA GB200 will need liquid cooling to run full throttle. The racks in which these chips are mounted will also need liquid cooling to shun 100kW power per rack. I expect AMD and Intel to follow suit with announcements on AI chips with similar heat dissipation requirements.

Liquid cooling comes in many guises. The most common technique used today is a so-called cold plate. These are direct replacements for air heatsinks, those metal fins you see when opening your computer, sometimes also containing a fan. Replace the fan with a pump and the air with a liquid, and you will have a single-phase liquid cold plate. Obviously, the liquid must be carefully contained inside the heatsink to not damage the electronics. A similar scenario applies when the liquid effectively evaporates under the chip's heat. Such a system is known as a two-phase cold plate.

Liquid cold plates provide a lower thermal resistance to heat transfer than a traditional air fan. This means they can dissipate more heat with minimum rise in the chip's temperature. Let's take a closer look at the operation of these cold plates<sup>1</sup>.

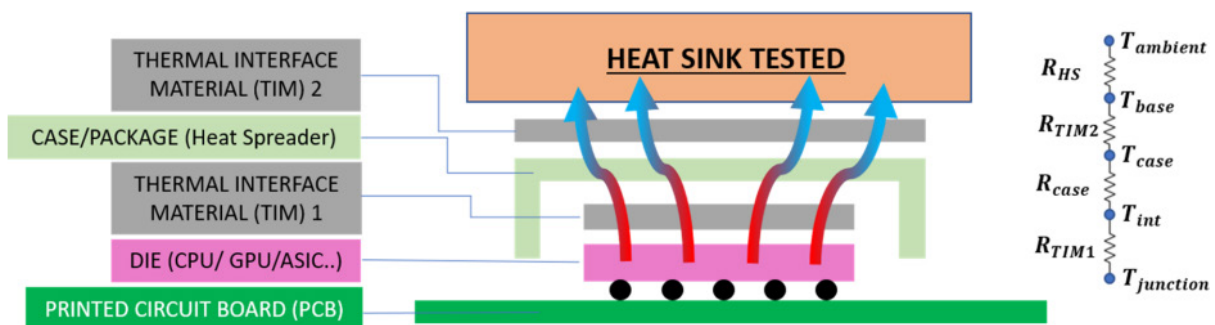


Figure 1 - Illustration of the heat flow from the hot die via the thermal interface materials and the chip package into the heat sink

The heat from the die must find its way through the thermal interface material 1 (TIM) to the package, then through the package and onto the thermal interface material 2, going to the cold plate to finally dissipate in the liquid. The design of such TIMs has become an art in itself, attracting startups with very creative solutions. Startup Carbice is proposing a TIM consisting of vertically-aligned carbon nanotubes, while Proud Technology has come up with a diamond heat spreader.

Ultimately, the best solution is to eliminate the thermal barriers and bring the coolants closer to the chip. A first option is to directly spray a liquid on the backside of the chip as demonstrated by imec<sup>2</sup>. These so-called jet impingement coolers spray cold liquid directly onto the backside of the chip, with a higher flow rate for more cooling, targeting the hottest parts of the chip. The coolants are in direct contact with the chip, providing more effective cooling than traditional cold plates. Disadvantages are that care has to be taken so that the liquids do not interfere with the electrical operation of the chip and that thermal warping can occur, leading to mechanical yield issues.

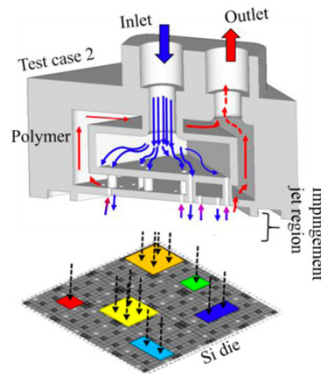


Figure 2 - Illustration of a jet impingement cooler

A cold coolant comes in through the inlet, is distributed in the polymer cold plate and directly jetted onto the backside of the chip hot spots. The hot liquid is discarded through the outlet port.

An alternative, potentially even more effective approach, is to bring the liquids directly into the chip substrate as demonstrated at EPFL<sup>3</sup> and at TSMC. With this approach, microchannels are etched into the substrate. These provide very close access for the liquids to the hot transistor junctions, reducing the thermal resistance and increasing the efficiency of this cooling method. A drawback is that such channels may impact the electrical performance of the transistors as well as the mechanical stability of the die.

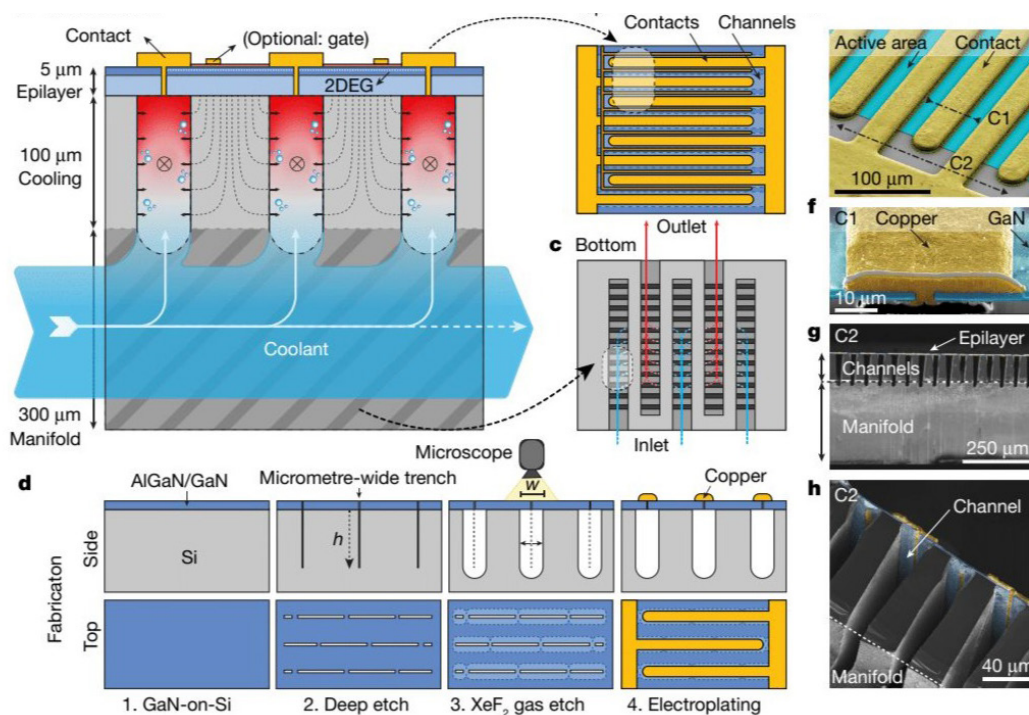


Figure 3 - Illustration of a liquid cooling system integrated into the substrate of a chip  
A coolant is pumped through microchannels etched in the chip substrate, bringing cooling very close to the hot transistor junctions.

An alternative method involves submerging the entire server in a dedicated dielectric liquid<sup>4</sup>. This technology is known as immersion cooling. Such systems can use a single-phase or two-phase cooling loop, depending on whether the server's heat transfer changes the phase of the dielectric liquid. Because of the direct contact of the liquids with the chips, these systems can be very effective coolers. Their disadvantage is the size and weight of the liquid baths containing the servers, which requires drastic design changes in data centers. Maintenance and uptime are also often cited as concerns when it comes to immersion cooling.

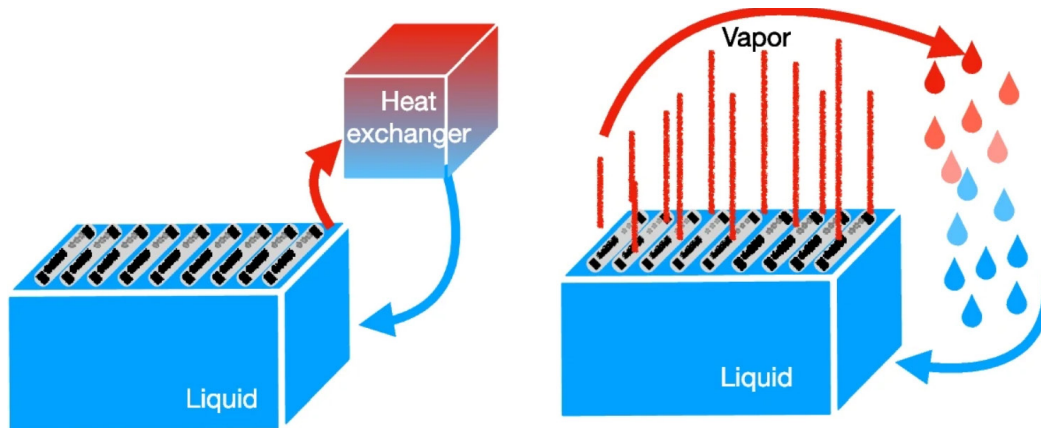


Figure 4 - Illustration of single (left) and two-phase (right) immersion cooling systems  
The server blades are fully immersed in a dielectric liquid.

Recently a cross-over promising the best of cold plates and immersion systems is emerging. Companies like Ferret and Iceotope have come up with solutions in which compute blades are individually submerged in stand-alone liquid-cooled chassis. These systems combine the efficiency of immersion cooling with the modularity of cold plates.

## The ecosystem for chip cooling is about to get disrupted

The liquid cooling market for AI data centers is emerging. The market is highly fragmented, as depicted in figure 5, with primarily players developing subsystems of the overall cooling solution, such as electronic design automation, direct-to-chip cooling solutions, immersion cooling, thermal interface materials and metal micro-machined parts.

### Electronic design automation



### Direct-to-chip cooling



### Immersion cooling



### 3D metal micromachining

FABRIC8LABS

### Chassis immersion cooling



### Thermal interface materials



Figure 5 - Market map of the chip cooling industry

Market research firm [Dell'Oro](#) predicts that “between 2023 and 2028, the overall thermal management market (comprising both air- and liquid-cooling systems) will expand at a 14% CAGR to reach \$12 billion. Of that \$12 billion, liquid-cooling solutions - including direct, immersion, and rear-door heat exchanger systems - are expected to account for \$3.5 billion by 2028. The need for air cooling is not going away, but liquid cooling is what's going to rapidly grow in the market, at a CAGR over 40%.”

I expect to first see more advanced cold plate products from vendors such as JetCool and Corintis. These cold plates are modeled and designed with the help of generative AI software from vendors such as Diabatix, ToffeeX, and Ansys. The intricate channel structures needed for advanced modern cold plates will require advanced 3D metal printing by highly specialized companies. Fabric8Labs - with its electrochemical additive manufacturing process - is in pole position to claim this market. The chassis-based immersion cooling solutions from the likes of Ferretet and Iceotope may also see adoption for dedicated on-prem servers with a very high

compute load. All of these systems will require limited additional data center infrastructure in the form of coolant distribution units and a way to chill the coolants. These cooling systems will command a premium price of several \$100 - maybe even over \$1000 - per chip and will, therefore, target the more advanced chips with the most significant cooling needs.

Once the market matures, I believe cooling will become an integral part of the chip package and be supplied by the OSATs (Outsourced Semiconductor Assembly and Test) such as ASI and Amkor. Today, we already see OSATs delivering chips in lidless packages to allow better contact with the die for better cooling. By 2028, we expect the more advanced AI accelerator chips to come with chip packages that include liquid cooling. These packages will offer better performance than add-on cold plates because they bring the cooling closer to the chip. Moreover, these packages will be more affordable and reliable than cold plates because they are mass-produced in conjunction with the underlying AI chips at an OSAT. Since such packages can bring coolants in direct contact with the chip, they can offer a similar performance as immersion cooling without the significant infrastructure changes immersion requires.

The next frontier could be to add cold micro-channels into the chips' substrates. First experiments have shown these to be extremely efficient while they are also scalable because designed at wafer scale. Further research and development will be needed to make such systems reliable and cost-efficient. The earliest I would expect this in mass-volume production ICs would be by 2030, and then only if the OSAT-provided chip-cooled packages no longer provide sufficient cooling for the power requirements of some of the high-value, high-power chips.

## Conclusion

AI is hot. And the chips delivering it even hotter. Liquid cooling will be a big part of future AI chipsets and data centers alike. A well-designed cooling system will directly impact both the cost and performance of future AI systems. I believe that at first, we will see various types of custom cold plates entering the market and getting integrated by rack integrators. We see early examples of such collaborations between Gigabyte and CoolIt, as well as between Dell and Jetcool. As the market grows, we will see further standardization of chip packages with built-in liquid cooling. This will reduce the cost of such systems, increase their reliability, and ultimately lead to mass adoption.

## About the author

**Bert Gyselinckx** is Partner at imec.xpand. He is a pioneer by birth and an electrical engineer by training.

Bert was responsible for imec's nano-electronics design center in the US. He specializes in DeepTech and is always on the lookout for companies with big, hairy ideas involving semiconductor or photonics technology.



## About imec.xpand

**imec.xpand** is one of the world's largest independent venture capital funds dedicated to early-stage semiconductor innovation. It targets ambitious start-ups where the knowledge, expertise and infrastructure of imec, the world-renowned semiconductor and nano-technology R&D center, can play a determining role in their growth.

imec.xpand has an outspoken international mindset towards building disruptive global companies and strongly believes that sufficient funding from the start is key to future success.

## Endnotes

- 1 CPU Overclocking: A Performance Assessment of Air, Cold Plates, and Two-Phase Immersion Cooling, by Bharath Ramakrishnan et al in IEEE Transactions on Components, Packaging and Manufacturing Technology, Vol 11, No 10, October 2021.
- 2 Low-cost Energy-Efficient On-Chip Hotspot Targeted Microjet Cooling for High-Power Electronics, Tiwei Wei et al, IEEE Transactions on Components, Packaging and Manufacturing Technology, Vol 10, No 4, April 2020.
- 3 Co-designing electronics with microfluidics for more sustainable cooling, R. van Erp et al, Nature, Vol 585, 10 September 2020.
- 4 Enough hot air: the role of immersion cooling, K. Haghshenas et al, Springer Open, Energy Informatics, August 2023.