SECTION: MECHANICS AND ELECTRICAL ENGINEERING

DOI 10.70286/ISU-10.12.2025.010

ANALYSIS OF THE THERMAL STATE OF PROCESSORS USING THREE-DIMENSIONAL MODELING

Taboranskiy Serhii
Ph.D., Head of the Department of Information Technologies
Voronin Stepan
Student of Information Technologies
Khomenko Violetta
Student of Information Technologies
Private Higher Educational Institution Kharkiv University of Technology "STEP":
Kharkiv, UA.

There are several main types of CPU cooling systems, varying in efficiency, noise level, and cost. The simplest and cheapest is passive cooling: a massive aluminum or copper heatsink without a fan. It is completely silent and reliable, but is only suitable for low-power processors (TDP up to ~65 W). Active air cooling is more common: a heatsink and one or more fans.

Budget models ("stock" coolers) handle TDPs of ≈65–90 W fairly well (≈95–110 W is already hot, close to throttling, and >120 W almost always means overheating or throttling). Most of the best dual-tower coolers can temporarily handle around 220–300 W with their fans at full load (maximum RPM, good case ventilation). At constant full load, closer to the upper limit, they will produce high temperatures (≈90–100 °C) and loud noise. The most effective cooling solution today is liquid cooling (AIO and custom loops) [1]. Pre-built, maintenance-free systems with 240 mm radiators (stable with ~150–220 W) and 360 mm radiators can easily handle flagship CPUs with TDPs of 200–350 W [2]. Custom cooling systems with multiple radiators and powerful pumps can dissipate over 500–700 W (HEDT/server-grade CPUs, such as the Xeon Sapphire Rapids/Emerald Rapids, AMD EPYC 9004 “Genoa”/“Bergamo,” AMD Threadripper PRO 7000/5000, or extremely overclocked desktop CPUs, such as the Intel Core i9–13900K/14900K or AMD Ryzen 9 7950X/7950X3D).

It's important to consider that one of the main challenges in dissipating heat from a processor is heat flux density (W/mm²). The desktop Intel i9 Raptor Lake has a die size of approximately 200 mm², and its boost power is 250–300 W. Its heat flux at a given power level is 1.2–1.5+ W/mm², which is very high. EPYC Genoa, for example, has a CCD die area of ≈72 mm², but there are 12 of them, meaning the actual distributed heat transfer area is significantly larger. With a total power of ~360–400 W, heat dissipation occurs over a larger total area. Accordingly, the heat flux per unit area is 0.3–0.5 W/mm², which is several times lower than that of a desktop processor.

Accurate thermal calculations are complicated by the fact that a processor's actual heat dissipation depends heavily not only on its stated TDP, but also on the specific workload (AVX-512, Prime95, gaming, rendering), voltage, frequency, and even the chip's stepping version. For example, 13th-14th generation Intel processors can consume 50–100% more than their stated PL2 power when running heavy AVX tasks. Furthermore, thermal conductivity between the die and the processor cover (especially on delinked Intel LGA1700 processors), the quality of thermal paste or liquid metal application, cooler clamping pressure, room temperature, and airflow velocity inside the case are all important factors [3]. All this makes "universal" recommendations (for example, "240 mm of liquid cooling is sufficient for any CPU") highly conditional: in a poorly ventilated case with high air temperatures, even a 360 mm system may not be able to handle an overclocked 250-watt processor. Therefore, the final choice of a cooling system always requires testing under real-world operating conditions.

The finite element method significantly outperforms analytical calculations in thermal modeling of modern processors, as the latter are inapplicable to complex 3D geometries and highly heterogeneous materials. Under the heat-spreading lid are a multilayer die with varying transistor densities, zones with local heat dissipation of up to 1000 W/cm², solder or polymer TIM, a copper lid of variable thickness—all with different and often anisotropic thermal conductivity coefficients, which cannot be described analytically. FEM allows for the generation of a realistic 3D model and an accurate power map with micron resolution, accounting for hot spots that cause throttling long before the average TDP is reached. Furthermore, analytical solutions are limited to the simplest boundary conditions and cannot account for uneven cooler pressure, contact resistances, variable convective heat transfer, thermal radiation, the nonlinear dependence of material properties on temperature, and transient processes under sudden load increases. As a result, analytical estimates yield an error of 30-70% or more, while a properly performed finite element analysis provides an accuracy of 3-7%. This is why Intel, AMD, and all leading cooling system manufacturers use FEM for design and validation. Without it, it is impossible to accurately predict hot spot temperatures, delinking efficiency, liquid metal temperature, or optimal cooler clamping pressure.

Let's consider a practical thermal analysis of a processor model with a 30 W thermal package cooled by a heatsink with both passive and active cooling systems.

Figure 1 shows a schematic of the initial conditions problem.

The initial air temperature is 20°C. The heat transfer coefficient is 20W\K·m^2 for velocities of 3 m/s and 20W\K·m^2 for velocities 6 m/s.

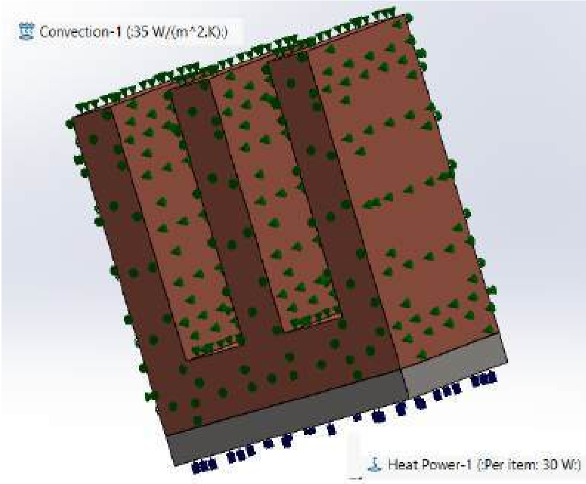

Fig. 1. Initial conditions

Figure 2 shows a tetrahedral finite element mesh.Fig. 2 shows a sketch of the shaft design of a synchronous motor in the basic version.

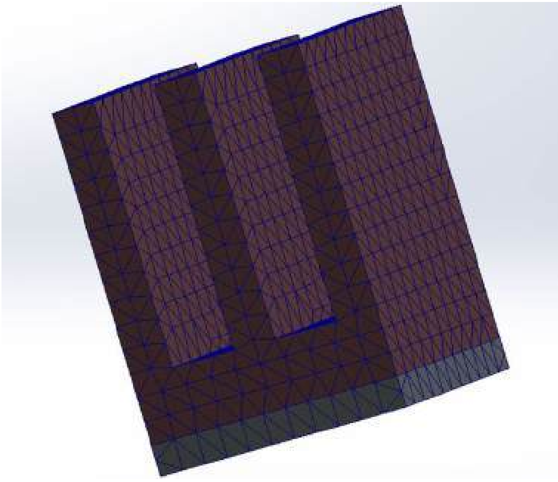

Fig. 2. Finite element mesh

In Figures 3 and 4, two different heat modes are shown above

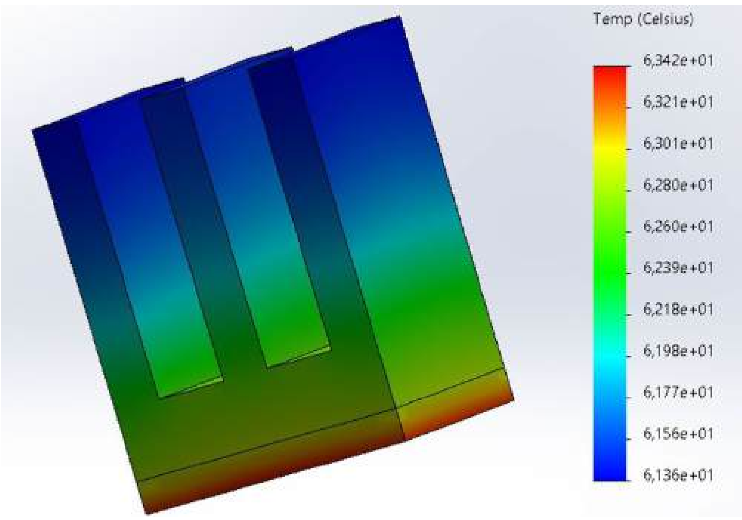

Fig. 3. Temperature field at a cooling air velocity of 3 m/s.

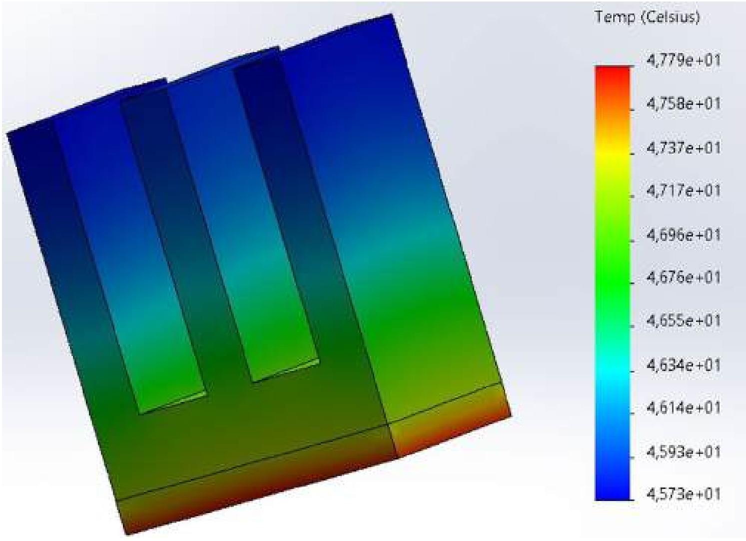

Fig. 4. Temperature field at a cooling air velocity of 6 m/s.

As you can see from the results of the heat calculation, increasing the cooling air velocity by 2 times allows to reduce the temperature of the processor by more than 10 degrees. Therefore, the further direction of improvement of cooling systems may be related to improving the efficiency of cooling systems by increasing water consumption provided that air velocities are increased.

References

1. Gullbrand, J., Luckeroth, M. J., Sprenger, M. E., and Winkel, C. (March 1, 2019). "Liquid Cooling of Compute System." ASME. J. Electron. Packag. March 2019; 141(1): 010802. https://doi.org/10.1115/1.4042802
2. S. Kang and D. Miller (2007). Closed Loop Liquid Cooling for High Performance Computer Systems. ASME 2007 InterPACK Conference collocated with the ASME/JSME 2007 Thermal Engineering Heat Transfer Summer Conference. https://doi.org/10.1115/IPACK2007-33870
3. Yuan Z., Coskun A.K. (2021). Neural network-based cooling design for high-performance processors. iScience. 2021 Dec 9;25(1):103582. https://doi.org/10.1016/j.isci.2021.103582