

Quest for Cooling: Achieving Optimal Steady-State Temperature Equilibrium in Microprocessors

A. C. Li

Abstract—This paper introduces an efficient Poisson’s equation solver used to study the steady-state temperature for different CPU systems. Experiments are conducted to investigate the effects of heat sinks and convection on the steady-state temperature distribution. Experimental results show that incorporating heat sinks and forced convection can reduce microprocessor temperatures by up to 20-fold. Validation methods including energy conservation and the Biot number confirm the robustness of the solver.

I. INTRODUCTION

THE central process unit (CPU), also called a processor, is the active part of the computer[2]. As technology advances, microprocessors have evolved to become increasingly smaller and powerful. Modern CPUs can generate an incredible amount of heat while working. This paper provides a study of steady-state heat dissipation in microprocessors and investigates how the implementation of different heat sink architectures and forced convection can greatly reduce the steady-state temperature of microprocessors. This investigation holds strong significance as it demonstrate the importance of thermal management for microprocessors.

A. Background and Context

The goal of this investigation is to analyze and optimize the steady-state temperature distribution of personal computer CPUs. The problem is simplified to 2D such that the length of the components can be ignored[1]. Fig. 1 illustrates the placement of different components. The microprocessor has a width of $14mm$ and a thickness of $1mm$ and a ceramic casing with a width of $20mm$ and a thickness of $2mm$ is attached on top of it. Aluminum Heat sinks of various sizes may be attached on top of the casing. The heat sink, ceramic casing, and the microprocessor have thermal conductivities of 250 , 230 and $150Wm^{-1}K^{-1}$ respectively. The ambient surroundings is assumed to have a fixed temperature

of $293K$. The microprocessor produces $0.5Wmm^{-3}$ of thermal power.




Fig. 1: An illustration of different components. The heat sink is detachable and the number of fins, separation between fins, length of fins are variable.

II. THEORY

A. Heat Transfer

Elliptic partial-differential equations(PDEs) have two standard forms when the problem is in 2D[3]:

$$c\nabla^2 u = c\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right) = 0 \quad (1)$$

$$c\nabla^2 u = c\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right) = f(x, y) \quad (2)$$

The first equation is Laplace’s equation and the second is Poisson’s equation. We focus on Poisson’s equation where u is temperature, T ; c is negative thermal conductivity, k ; and $f(x, y)$ is the power density, $q(x, y)$.

As heat dissipates into the ambient surroundings from the boundaries of our components through convection, we introduce Newton’s Law of Cooling[1]:

$$\phi_s = h(T_s - T_a) \quad (3)$$

where ϕ_s is the surface heat flux and h is the heat transfer coefficient at the interface which can take different values for natural and forced convection. To simplify the equations and increase numerical

stability, the equations used in the solver program are dimensionless, and they are given by:

$$\frac{\partial^2 \hat{T}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{T}}{\partial \hat{y}^2} = \frac{qr_0^2}{-kT_0} \quad (4)$$

$$\frac{\phi_s}{1.31T_0^{4/3}} = (\hat{T}_s - \hat{T}_a)^{4/3} \quad (5)$$

$$\frac{\phi_s}{(11.4 + 5.7v)T_0} = (\hat{T}_s - \hat{T}_a) \quad (6)$$

Eqn. 4 is the dimensionless Poisson's equation, variables with hats indicate that they are dimensionless, r_0 is the length scale and T_0 is the temperature scale. Eqn. 5 and Eqn. 6 are the dimensionless Newton's law of cooling for natural and forced convection. v is the wind speed.

III. METHOD

Numerical methods will be used to solve Eqn. 4. Unlike analytical solutions, numerical solutions only allows the determination of temperature on discrete points. But it can be easily extended to solve 3D problems.

A. Finite-Difference Method

To find the points of interest, the medium is subdivided into smaller squares. At the center of each square, we define the nodal point, where its value represents the average temperature within the region. The ensemble of the nodal points is called a mesh where i and j indices are used to describe the x and y locations. By applying energy conservation to an area around a node, we can derive the finite-difference equation for the node. Fig. 2 illustrates that the energy flowing into the control volume will be equal to the energy flowing out and we can express this relationship by[4]:

$$\sum_{i=1}^4 \xi \phi_{(n) \rightarrow (i,j)} + q\xi^2 = 0 \quad (7)$$

where ξ is the distance between each node and $\xi \phi_{(n) \rightarrow (i,j)}$ represents the energy flow from neighbouring nodes. Since we are working in 2D, the overall unit of Eqn. 7 is Wm^{-1} . This equation can be expanded to give:

$$\sum_{i=1}^4 T_{(n) \rightarrow (i,j)} + \frac{q\xi^2}{k} - 4T_{i,j} = 0 \quad (8)$$




Fig. 2: Conduction to an interior node from its neighbouring points

For the ceramic casing and the heat sink, where there is no power generation, this expression reduces to:

$$T_{i,j+1} + T_{i,j-1} + T_{i+1,j} + T_{i-1,j} - 4T_{i,j} = 0 \quad (9)$$

At the interface of two different materials, a slightly different finite difference equation needs to be applied. Fig. 3 shows two nodes at the aluminium ceramic interface with a thermal contact resistance, where the rate of heat transfer can be expressed by:

$$\xi \phi_{(i,j-1) \rightarrow (i,j)} = \frac{T_{i,j-1} - T_{i,j}}{R_{tot}} \quad (10)$$

$$R_{tot} = \frac{1}{2k_s} + \frac{1}{2k_c} + \frac{R''_{t,c}}{\xi} \quad (11)$$

Since we assume perfect contact, $R''_{t,c}$ equals zero and we obtain:

$$\xi \phi_{(i,j-1) \rightarrow (i,j)} = \frac{2k_s k_c}{k_s + k_c} (T_{i,j-1} - T_{i,j}) \quad (12)$$

After implementing this change in heat transfer and a pictorial operator to the finite-difference equation, the equation for the microprocessor temperature at the microprocessor-ceramic boundary can be expressed as:

$$T_{i,j} = \alpha \left(\begin{bmatrix} 1/R_{tot} & & \\ k_m & & k_m \\ & k_m & \end{bmatrix} T_{i,j} + \xi^2 q \right) \quad (13)$$

$$R_{tot} = \frac{2k_m k_c}{k_m + k_c}, \alpha = \frac{1}{1/R_{tot} + 3k_m} \quad (14)$$

The other boundary equations can be derived similarly. For the heat sink ceramic boundary, q equals zero. Since the linear system that we try to solve




Fig. 3: Conduction between different materials at the interface

is derived from the discrete Poisson equation, we can find the truncation error in the approximation. The neighbouring points of $T_{i,j}$ in Eqn. 8 can be approximated by the Taylor series expansion of the central point, after simplification the equation becomes:

$$\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} + \frac{\xi^2}{12} \left(\frac{\partial^4 T}{\partial x^4} + \frac{\partial^4 T}{\partial y^4} \right) = \frac{q}{-k} \quad (15)$$

which has a truncation error proportional to ξ^2 , suggesting that the solution is accurate up to ξ^2 [5].

Algorithm 1 System with a heat sink

Input: number of iterations, number of fins, fin separation, fin height, microprocessor grid dimension, forced convection or not, convergence criteria

Output: Plot of steady-state temperature distribution for the system

Generate fins, heat sink base, ceramic casing, and microprocessor grids from microprocessor grid dimension, calculate step size

for each iteration **do**

 Update ghost points for all components

 Calculate power out across all boundaries

 Apply pictorial operator except interface nodes

 Apply modified pictorial operator to interface nodes

if convergence reached **then**

 Break

end if

 Combine all components and create plot

end for

At the boundaries of the system, a layer of ghost points outside the mesh are calculated from Newton's Law of Cooling using central difference,

such that the same pictorial operator can also be applied to the boundary points. In the solver, the microprocessor, ceramic casing, and the rectangular base of the heat sink are each stored separately with their ghost points. Each rectangular heat sink fin is stored separately with their ghost points. Algorithm 1 shows the program used to solve the steady-state temperature of the system with a heat sink, where the Jacobi method is used. Convergence criteria for both temperature and energy conservation are used to ensure accurate result and efficient computation. The heat in and out of the system both have units Wm^{-1} due to the 2D nature of the problem. The heat supplied into the system are solely from the microprocessor and can be calculated by multiplying the power density by the area of the microprocessor, and the heat that leaves the system through the boundaries can be calculated by a similar approach described in Eqn. 7.

IV. EXPERIMENTS

In this section, an analysis of the computational efficiency of the program will first be conducted, followed by a grid study of steady-state temperature for fine grids against course grids. After that, temperature distribution results for various systems will be analyzed including the steady-state temperature distribution for different heat sink fin architectures. Lastly, different validation methods will be discussed.

A. Computational Efficiency

Table 1 shows the time taken to run the program under different convection, steps of iteration, and step size. The program runs very fast due to the implementation of vectorized operations. It can be seen from the table that the program tends to run faster for forced convection than Natural convection which is expected due to simpler equations.

B. Grid studies

The grid study is conducted by varying the step size between the nodes. In this investigation, we consider only the microprocessor and the ceramic casing. The initial temperature is set to $293K$. Convergence criteria of 0.01% is used, meaning that the program stops only when both the percentage difference between the heat-out and heat-in and

Convection	Iteration	Step size(mm)	Time(s)
Natural	1e3	0.1	0.3154
Natural	1e3	0.01	0.3784
Natural	1e5	0.1	17.29
Natural	1e5	0.01	23.47
Forced	1e3	0.1	0.2887
Forced	1e3	0.01	0.3500
Forced	1e5	0.1	12.05
Forced	1e5	0.01	22.36

TABLE 1: The performance of the steady-state temperature distribution solver for microprocessor and ceramic casing only on an Apple M1 Max chip. Step size is equivalent to ξ described in the previous section

the percentage difference between the current mean microprocessor temperature and that of the last iteration are both less than 0.01%. In fig. 4, the data points show a quadratic decrease as the step size decreases with a subsequent rise at around 0.2mm. A best fit line is plotted and the quadratic form of the best fit line suggesting a truncation error of order ξ^2 . Moreover, the rise suggests that the round-off error exceeds the truncation error as the step size approaches zero. This is expected as the number of iterations increase since the step size decreases which can lead to accumulation of errors that eventually overpower the truncation error.

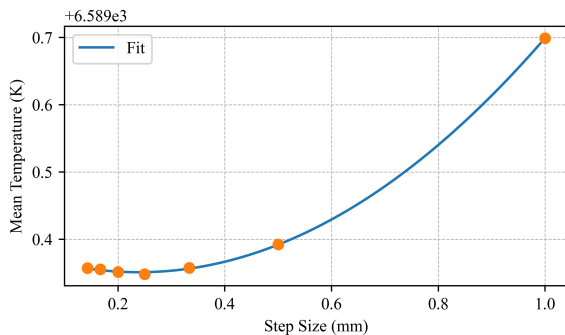


Fig. 4: Plot of mean steady-state microprocessor temperature against step size using the microprocessor and ceramic solver. The convergence criteria is 0.01%, and the system is under natural convection.

C. No Heat Sink

Fig. 5 shows the steady-state solution for the ceramic microprocessor system under natural convection. The temperature distribution is continuous

throughout the system which agrees with the theory. The figure shows that only using the ceramic casing to cool the microprocessor is not viable, as the mean steady-state temperature of the system is higher than the melting points of all metals. Moreover, if there were no convection, heat will not flow out, so there will never be a steady-state temperature as temperature is always increasing.




Fig. 5: steady-state temperature distribution for the ceramic casing and microprocessor under natural convection. A step size of 0.05mm and convergence criteria of 0.1% are used

D. With Heat Sink

Fig. 6 shows how the mean steady-state temperature across the microprocessor can vary with both the number and height of fins. The number of the fins range from 10 to 25 and the length of the fins range 1mm to 40mm. Both plots show a diminishing rate of decrease in temperature, suggesting that increasing the number and height of the fins beyond the plotted range gives less significant returns. Since there is a dimensional constraint for the CPU to fit into a computer, it is impossible to keep the average steady-state temperature to the working temperature of 373K without forced convection.

E. Forced Convection

When a constant wind speed of $20m/s^{-1}$ is introduced, the mean steady-state temperature for the microprocessor is greatly reduced such that it can work at a temperature of around 353K ($80^{\circ}C$). The width of the heat sink is restricted to under 70mm and the height of the fins are restricted to under 40mm. Fig. 7 and Fig. 8 shows two heat sink configurations optimising the fin height and number of fins under this constraint. Both plots show similar trends in temperature distribution and the mean temperature across the microprocessor are under 353K for both.

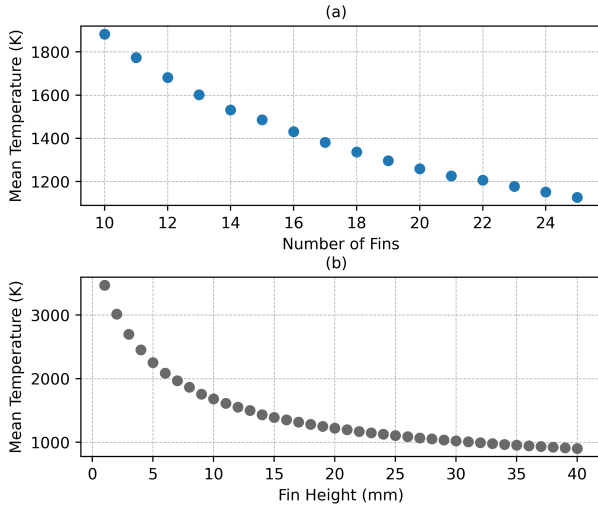


Fig. 6: Mean Microprocessor temperature against fin height and number of fins on the heat sink with a constant fin separation of 2mm and convergence criteria of 1%




Fig. 7: steady-state temperature distribution under 20ms^{-1} forced convection. The heat sink has 35 fins and the fin height is 25mm . A step size of 0.2mm and convergence criteria of 1% are used. The mean temperature across the microprocessor is 352.5K

F. Validation Method

As discussed in the previous section, the main method of validation used in the solver is temperature convergence and energy conservation. An interesting observation from the previous experiments is that energy conservation is almost perfect at large iterations even when the grids are extremely coarse. This can be explained by the Lumped Capacitance approximation[4], which is the process of reducing a system to a number of discrete lumps and assuming that temperature is uniform inside each lump. In our problem, each cell represented by a node acts as a capacitive reservoir that absorbs heat until steady-state is reached. For the entire microprocessor, we




Fig. 8: steady-state temperature distribution under 20ms^{-1} forced convection. The heat sink has 25 fins and the fin height is 40mm . A step size of 0.2mm and convergence criteria of 1% are used. The mean temperature across the microprocessor is 349.6K

can calculate the Biot number, given by:

$$Bi = \frac{hL_c}{k} \quad (16)$$

where h is the heat transfer coefficient and the characteristic length L_c is reduced to the thickness for a plane wall. The Biot number for the microprocessor is therefore 0.0117 , smaller than the threshold of 0.1 for the Biot approximation. The temperature distribution across the thickness of the microprocessor is illustrated in Fig. 9. The difference between the maximum temperature at the center and the minimum temperature on the edge is 0.10670% , which matches the Lumped Capacitance approximation.




Fig. 9: steady-state temperature distribution across the center of the microprocessor under 20ms^{-1} forced convection. A step size of 0.2mm and convergence criteria of 0.1% are used.

Another way of verifying the result is by checking the vertical temperature distribution across two com-

ponents. Fig. 10 shows how the temperature ranges from the ceramic plate to the microprocessor along the x axis. The temperature is continuous along the system but there is a sudden change in gradient at the interface at $2mm$, caused by the jump in thermal conductivity.




Fig. 10: steady-state temperature distribution across the vertical center of the microprocessor and heat sink under $20m/s$ forced convection. A step size of $0.05mm$ and convergence criteria of 1% are used.

G. Maximum cooling

Fig. 11 shows how the mean steady-state temperature changes as the wind speed increases. There is a significant drop in steady-state temperature when the wind speed is doubled from $20m/s$ to $40m/s$. We can implement the extra cooling together with the best performing heat sink under our constraints to achieve effective cooling. Fig. 12 shows the steady-state temperature distribution for such system, where the microprocessor's mean temperature is only $334.3K$ ($61.3^\circ C$) which is a huge improvement from the mean temperature of over $6000K$ that we started with.




Fig. 11: Mean steady-state temperature against wind speed. The heat sink has 12 fins with a fin height of $10mm$ and fin separation of $1mm$. A step size of $1mm$ and convergence criteria of 1% are used.




Fig. 12: steady-state temperature distribution under $40m/s$ forced convection. The heat sink has 25 fins and the fin height is $40mm$. A step size of $0.2mm$ and convergence criteria of 0.1% are used. The mean temperature across the microprocessor is $334.3K$

V. CONCLUSION

In conclusion, this paper introduces a Poisson's equation solver that utilizes the finite difference method to define the mesh. The finite difference method gives the equations for the pictorial operators used in the Jacobi method to solve for the steady temperature distribution in various systems through an iteration scheme. The experiments shows the optimal architecture for heat sinks and when incorporated with forced convection can reduce the microprocessor temperature by up to 20 times. Lastly, different validation methods like the Biot number and temperature continuity are introduced to show the robustness of the proposed solver.

REFERENCES

- [1] M. Scott and J. Owen, "Project 4: Heat dissipation in microprocessors," 2023.
- [2] D. A. Hennessy, *COMPUTER ORGANIZATION AND DESIGN RISC-V EDITION : the hardware software interface*. S.L.: Morgan Kaufmann Publisher, 2021.
- [3] C. F. Gerald and P. O. Wheatley, *Applied numerical analysis*. Boston: Pearson/Addison-Wesley, 2004.
- [4] T. L. Bergman, *Fundamentals of heat and mass transfer*, 7th ed. Hoboken: J. Wiley Sons, Cop, 2011.
- [5] G. D. Smith, *Numerical solution of partial differential equations : finite difference methods*. Oxford: Oxford University Press, Cop, 1985.