THE EFFECTS OF COOLING FAN SPEED ON COMPUTER PERFORMANCE

Garrett Blosen
MIT Department of Mechanical Engineering
Cambridge, MA, United States

ABSTRACT

Optimizing the cooling of computer chips is essential to maximizing the performance and efficiency of hardware while preserving a low noise level and prolonging the lifespan of fans [1]. The effect of fan speed on computer temperature and performance was measured on a water-cooled desktop PC with an anemometer to measure airflow, a standardized graphics processing unit (GPU) stress test to measure performance and a sensor logging software to track the temperature in each test. While a higher fan speed resulted in increased airflow with lower temperatures, the performance score remained constant. The impact of the GPU temperature on computer performance is insignificant provided that the temperature remains within its specified thermal limit, allowing for lower fan speeds, less fan noise, and longer lifespan.

Keywords: PC air cooling, GPU cooling, PC heat transfer, AIO liquid cooling

1. INTRODUCTION

Optimizing the cooling of computer chips is essential to maximizing the performance and efficiency of hardware from the biggest server racks down to the smallest handheld devices. Transferring heat away from these chips allows them to stay at manageable temperatures to compute and execute commands to satisfy a user's needs. As a result, it is beneficial to discover ways to maximize performance with any given set of cooling hardware.

Consumer-level personal computers (PCs) rely heavily on heat dissipation, and they utilize an arrangement of fans mounted on the case to cycle ambient air past components or through radiators. Analyzing the effects of fan speed on computer temperature and performance can provide insight into effective thresholds of fan speeds to cool these components. The resulting data will not only offer relevant information for larger computing servers but will also assist the average PC consumer in optimizing their product. The evolution of software becoming increasingly demanding along with constant releases of new hardware can create a paywall for many consumers around the world as GPU prices have statistically increased about $200 every year [2]. Optimized cooling can allow users to run their fans at lower speeds to prolong the lifespan of their products. As fans produce bothersome noise at high revolutions per minute (RPM), lower fan speeds can also increase the user's quality of life while at the desktop. With all-in-one (AIO) water-cooling systems being a staple cooling accessory for PC consumers, discovering ideal fan speeds with this type of cooling setup can be beneficial to many users.

The speed of the fans and the temperature of critical PC components were recorded during a series of controlled graphics processing unit (GPU) stress tests which output a performance score based on average graphical frames rendered each second (FPS). Fan speeds and volumetric airflow were verified with a handheld tachometer and anemometer before beginning any testing to ensure that the system was working as expected. Standardized stress tests were run at five distinct fan speeds: 300, 600, 900, 1200, and 1500 RPM. Before each test, the computer and liquid in the water cooler were warmed up to a set temperature which is directly reported from internal temperature sensors inside of the critical PC components. To ensure a controlled test, the tachometer readout of fan speeds and internal reporting of the chip's clock speed were actively monitored. The steady-state temperature measurements of each test were extracted and plotted vs fan speed to fit a curve representing the effect of fan speed. With this analysis, PC users can easily set constraints on their fan speeds to discover the most efficient fan speed threshold.

2. BACKGROUND

AIRFLOW PROPERTIES OF PC FANS

Fans are a critical component in any consumer-level PC, being the main source of airflow to cool components. In the test setup for this experiment, the heat generated by the GPU will be transferred through the water-cooling loop into a radiator where forced convection cooling created by the fans will expel the heat into the environment. The amount of heat transferred with this method of cooling is defined in the equation below.

$$q = h_c A dT \tag{1}$$

From the convection heat transfer equation (1), we see that the heat transferred from the chip per unit time q will increase with a higher temperature difference between surface and fluid dT, larger heat transfer area A, and larger heat coefficient h_c which increases with airflow velocity through the radiator (Fig. 1)

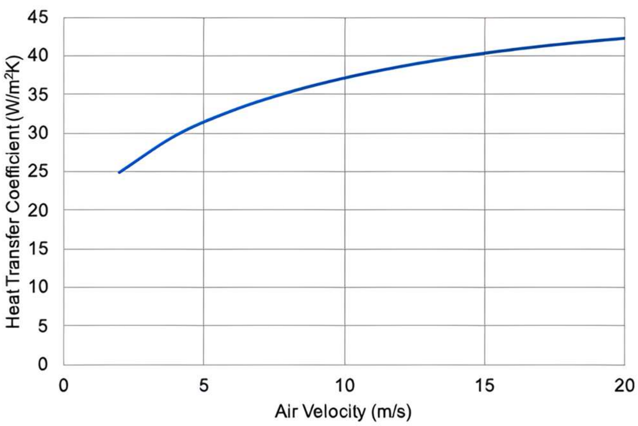

Figure 1: The heat transfer coefficient for forced convection increases nonlinearly with increased air velocity. [3]

In previous studies by tech publication TechPowerUp, different PC fans from the market were put to a performance test to assist consumers in purchasing decisions [4]. A controlled test setup was used to measure important data that would be relevant to consumers. Sound levels were then measured in an anechoic chamber with a sound probe 15.24 centimeters away from the fan. While increased fan speed had a linear correlation with better airflow through the radiator (Fig. 2), fan noise unfortunately increased at higher fan RPM.

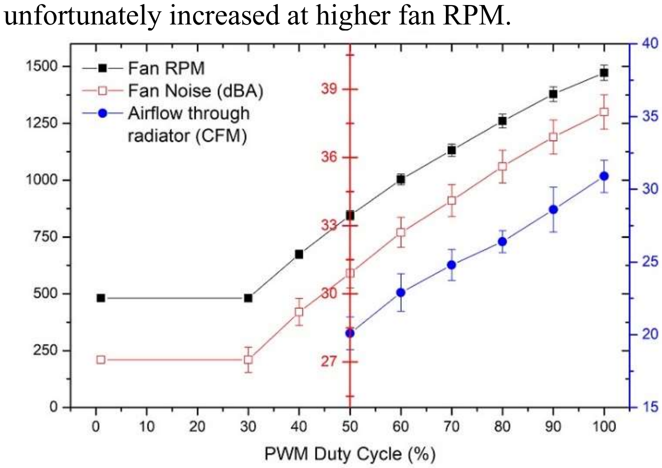

Figure 2: Noise level and airflow data recorded at various fan RPM. Increasing fan speed resulted in a linear increase in ambient fan noise and airflow through the radiator. [5]

EFFECTS OF TEMPERATURE ON COMPUTER CHIPS

When looking at the GPU, temperature has a direct correlation to performance as it affects the physical properties of the chip. Typical consumer GPU chips are made of layers of silicone with metals that form the circuitry inside of the chip. The temperature in the chip builds up over time as power is dissipated along the wires. On the atomic level, increased temperature causes lower electron carrier mobility in non-polar semiconductors such as silicon [6]. As a result, instructions and signals sent across the circuitry inside the chip occur at slower speeds.

While high temperatures may affect performance, they can also shorten the lifespan of components due to thermal stress and eventually cause chip failure. As a result, GPU chip manufacturers implement thermal throttling as a failsafe to prevent their chips from destroying themselves. Once the temperature of the chip reaches a certain temperature, the chip quickly throttles down to a lower clock speed to reduce power consumption and heat generation.

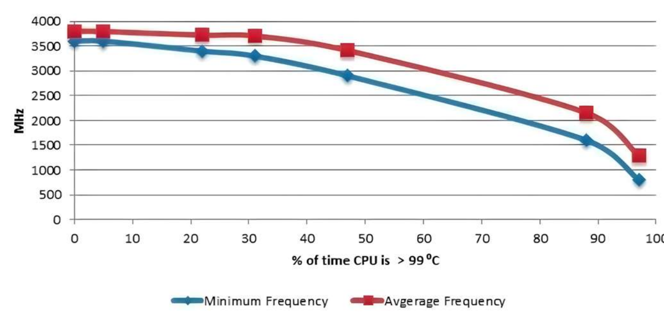

Figure 3: Average computer chip clock speed when subjected to various temperatures during stress tests. As the chip temperature more consistently reaches temperatures above the thermal throttling limit, performance begins to decrease [7].

This decrease in clock speed frequency can be very dramatic (Fig. 3) at higher temperatures which increases the importance of efficient cooling of the chip to maximize clock speed and performance. Not only does the computer chip experience increased wear and lower performance at high temperatures, but the thermal paste interface used to conduct heat away from the chip will degrade at a faster rate. If the condition of the paste is not monitored after use at high temperatures, the heat dissipated by the cooling system will decrease which can cause even higher temperatures that lead to component throttling failure.

3. EXPERIMENTAL DESIGN

To provide relatable measurements to PC users, graphical stress tests were run on a Nvidia RTX 2080 Super which is a well benchmarked and popular graphics card. This GPU has a manufacturer-specified thermal throttling limit of 84℃ and was combined with a standard 240mm x 120mm AIO water cooling system to create a realistic scenario of the chip being cooled under load. To validate consistent airflow through the radiator, a Vernier anemometer was used to measure airflow through the radiator as shown in Figure 4.

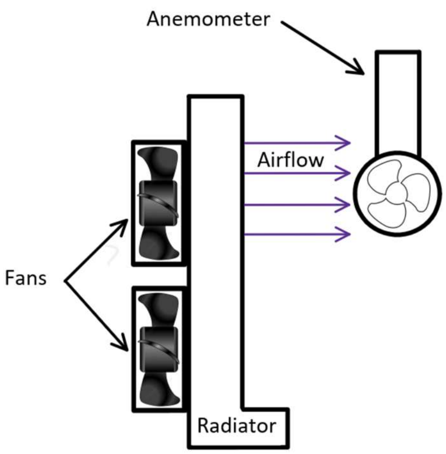

Figure 4: Visualization testing setup and of airflow through the radiator and anemometer for measurement.

Before every test, the PC was put under load to warm up components and the liquid in the water cooler to a steady-state temperature of 53℃ to have a controlled beginning for each trial. To create a repeatable load case for the computer, a 3DMark graphical stress test was conducted to maintain the GPU usage at 100% and output a score based on the performance during the trial. The pump speed was held constant at full power and fan speeds were increased across four different RPM levels.

4. RESULTS AND DISCUSSION

During the initial airflow validation experiment as shown in Figure 4, anemometer measurements were measured and used to validate a linear correlation between fan and air speed.

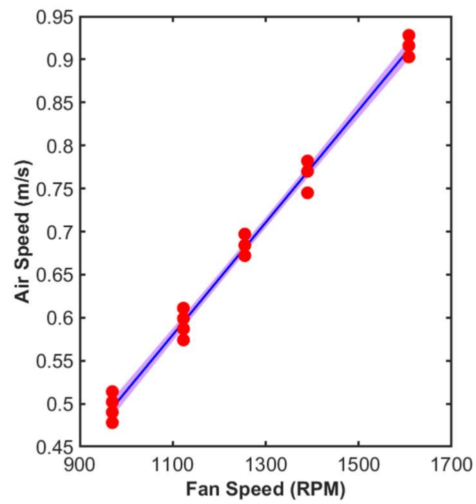

Figure 5: Using data from the setup from the test setup in Figure 4, airflow was observed to linearly increase with higher fan speed at a rate of $0.0007 \pm 0.0001 \ \frac{m/s}{RPM}$.

From each stress test trial, a set of raw data was collected, and the performance score from 3DMark was collected in Table 1 below.

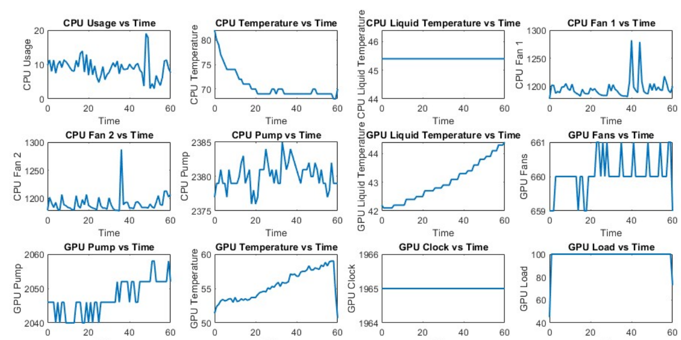

Figure 6: All sensor outputs from the computer were monitored. Measurements from the stress test timespan were plotted and scrutinized to ensure 100% GPU load throughout the trial for consistency. GPU temperature and usage graphs that are boxed in red were used for analysis.

From the raw GPU temperature data, there is a steady increase over time as the liquid in the AIO system absorbs heat from the GPU chip which decreases its cooling capacity. As a result, the metric of temperature delta was extracted to measure its relationship to fan RPM.

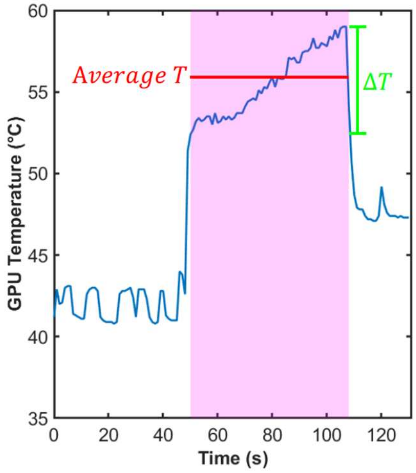

Figure 7: Using the temperature data recorded from the GPU, average temperature and change in temperature metrics were extracted from the highlighted timespan of the stress test.

To see the relationship between temperature delta and fan RPM change in temperature was plotted against fan speed, and a notable trend of strong linear decrease is observed.

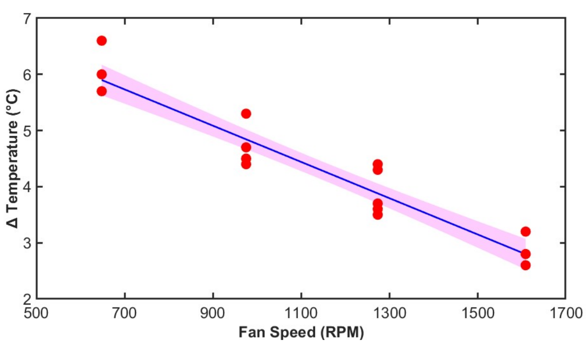

Figure 8: Data points represent distinct graphical stress test trials and their corresponding component temperature and fan RPM. The data has a linear fit of $-0.0032 \pm 0.0005 \ \frac{°C}{RPM}$. The shaded region represents the confidence interval of delta temperature for a given fan RPM.

This observed correlation is influenced by the escalating heat transfer coefficient as the fan RPM increases. Figure 1 illustrates this relationship, where the forced convection heat transfer coefficient is plotted against increasing air velocity. As the heat transfer coefficient increases, there is more heat transferred away from the chip and into the surrounding environment.

Table 1: 3Dmark Performance Scores (Average FPS)

Fan RPM	Test 1	Test 2	Test 3	Test 4	Test 5
648	22.22	22.20	22.27	22.26	22.25
975	22.26	22.21	22.22	22.22	22.22
1273	22.21	22.24	22.20	22.20	22.24
1608	22.27	22.27	22.30	22.23	22.24

Regardless of the increased temperature levels due to higher fan speed cooling down the AIO system, the performance scores remained relatively constant throughout the experiment as recorded in Table 1. The insignificant effect of GPU temperature on performance was likely due to the sufficiently low component temperatures that the chip did not begin to throttle its performance at all. This correlates to the beginning of the graph in Figure 3, where performance remains relatively constant at lower temperatures and drastically falls off at higher temperatures.

The average temperature during a performance test is a well-known standard metric to correlate to performance, so a second metric of average temperature was extracted from the raw data as shown in Figure 7 to plot performance.

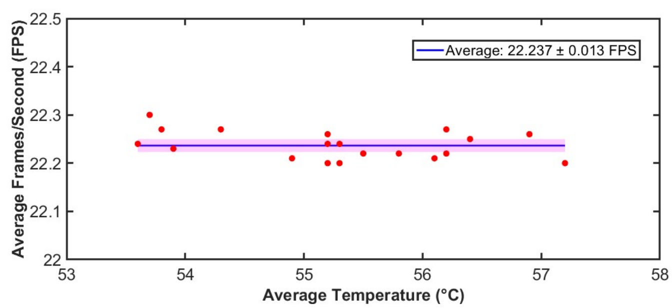

Figure 9: Performance scores during stress tests at various fan RPMs. Data points represent distinct graphical stress test trials. The data has an average of 22.237 ± 0.013 frames per second and the shaded region represents the confidence interval of the score for any given average test temperature.

While analyzing the data in Figure 9, a linear fit would not make sense as it would result in an insignificant slope, so an average value with 95% confidence was calculated. The highest average temperature peaked at just over 57℃ which is much lower than the GPU's thermal limit of 84℃.

These experiments show a direct correlation between higher fan speed and velocity of airflow. This caused a higher rate of forced convection heat transfer from the AIO system into the surrounding environment which lowered the temperature delta during a graphical stress test. It was found from a linear fit that with an increase of 100 RPM in fan speed, there was a 0.32℃ decrease in temperature delta. Even given the small range temperature deltas from the conducted trials, there is a clear significant decrease in temperature with higher fan RPM which can be helpful when tuning the fan curve of a PC especially if components are near throttling temperature.

While the temperature measurements benefited from having a higher fan RPM, the impact of GPU temperature had an insignificant effect on computer performance. This is likely due to sufficiently low temperatures during trials that did not induce thermal throttling. We can conclude that for the range of GPU temperatures measured in this experiment, extra cooling has no effect on performance even as the temperature increases. As a result, the computer can be run with lower fan speeds to help prolong the lifespan of the hardware and decrease noise during usage.

5. CONCLUSIONS

The experiments conducted showed that a higher fan speed and increased airflow significantly lowered component temperatures linearly, yet the performance score remained constant. With the GPU temperature not affecting computer performance, the fans can be run at lower speeds which can decrease ambient noise and increase their lifespan.

Since the temperatures measured were very constrained due to the efficiency of the AIO water cooling system which did not let the chip reach higher temperatures, the scope of this experiment could be increased to include non-liquid cooling systems and higher-power graphics cards. The experiment steps could be repeated under different hardware to examine the performance response once the temperature approaches and passes the thermal throttling threshold of various components.

There have been many studies and trials run on all types of computer components ranging across power consumption, performance, and cooling methods, however, it is difficult for a user to find a benchmark study that directly applies to their personal setup. If a user's PC temperatures under usage remain consistently below the thermal throttling threshold designated by their GPU manufacturer as it did in this study, they can safely reduce fan speeds without concerns about premature fan failure or excessive noise.

ACKNOWLEDGMENTS

The author would like to express gratitude to Dr. Barbara Hughey and Professor Bischofberger for their guidance during the experimentation and analysis phases. The author would also like to thank Ms. Claire Berman for her assistance with discovering databases for background research.

REFERENCES

[1] Black Box (2021) *Server Cabinet Cooling.* (Accessed: 10 May 2024).

[2] Nvidia (2024) Graphics Cards and Desktops. (Accessed: 10 May 2024).

[3] The Engineering ToolBox (2003) *Convective Heat Transfer.* (Accessed: 28 February 2024).

[4] Tech Power Up (2017) *Fan Reviews.* (Accessed: 28 February 2024).

[5] Tech Power Up (2017) *Corsair LL120 RGB Fan Review.* (Accessed: 28 February 2024).

[6] Bart Van Zeghbroeck (2004) Principles of Semiconductor Devices. (Accessed: 29 February 2024).

[7] Puget Systems (2014) *Impact of Temperature on Intel CPU Performance.* (Accessed: 28 February 2024).

[8] https://pngtree.com/freepng/mini-fan-side-view_6952323.html

[9] https://www.vecteezy.com/vector-art/5200348-fan-blades-black-color-icon