

Article

Predictive Thermal Management for Dual PWM Fans in High-Power Audio Amplifiers

Andrei Militaru, Emanuel-Valentin Buica * and Horia Andrei 

Doctoral School of Electrical Engineering, Valahia University of Targoviste, 13 Aleea Sinaia, 130004 Targoviste, Romania; andrei.militaru@valahia.ro (A.M.); hr_andrei@yahoo.com (H.A.)

* Correspondence: emmanuel.buica@valahia.ro

Abstract

This paper presents the design and implementation of a low-cost microcontroller-based dual-channel fan controller optimized for high-power audio amplifiers, yet adaptable to power supplies, electronic loads, and other thermally intensive systems. Unlike conventional designs that drive all fans uniformly, the proposed solution provides fully independent cooling via dual I²C temperature sensors, predictive trend analysis, and multi-stage hysteresis. The controller incorporates advanced features including an anti-dust startup sequence, predictive boost with latching, active cross-cooling, anti-heat-soak protection, and stall detection via tachometer monitoring, complemented by LED-based fault signaling and automatic channel muting during overheating or fan failure. Hardware support for 12 V and 24 V fans, dual power-input options, and a compact PCB layout enhance integration flexibility. The firmware employs temperature-driven PWM mapping with EMA filtering and multi-level hysteresis. The experimental results confirm that all implemented features operate as intended, with each function demonstrating clear practical relevance, whether in improving responsiveness, preventing heat accumulation, or enhancing system reliability under a wide range of operating conditions.

Keywords: PWM fan; temperature control; audio amplifier cooling; microcontroller; I²C sensors; stall detection; low-cost design

1. Introduction

1.1. High-Power Amplifiers and Thermal Challenges

At its core, a high-power audio amplifier is an energy conversion device, tasked with transforming electrical energy from its power supply into a powerful audio signal to drive loudspeakers. However, this process is far from perfectly efficient, particularly in widely used linear amplifier designs like Class-AB amplifiers. A significant portion of the electrical power drawn from the wall is not delivered to the speakers; instead, it is dissipated as waste heat [1].

Thermal management is a critical concern in high-power electronic systems such as audio amplifiers, switched-mode power supplies (SMPSs), or programmable electronic loads [2]. Excessive heat can degrade component performance, reduce system lifetime, and compromise reliability [3]. Traditionally, cooling is achieved using brushless DC (BLDC) fans, with fan speed controlled by analog temperature sensors that trigger control relays or transistors. While effective, these approaches have limitations: most existing controllers regulate all fans at the same speed regardless of localized thermal conditions, do not



Academic Editor: Teen-Hang Meen

Received: 19 January 2026

Revised: 24 February 2026

Accepted: 24 February 2026

Published: 26 February 2026

Copyright: © 2026 by the authors.

Published by MDPI on behalf of the International Institute of Knowledge Innovation and Invention. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

monitor or react to fan stall conditions, and lack mechanisms for cross-assisted cooling, where one fan supports another during asymmetric thermal loads [4,5].

In high-power audio amplifiers, channel loading and thermal conditions can be asymmetric [6]. A single channel may generate substantially more heat depending on the audio program, load characteristics, or elevated ambient temperature. Using a shared fan-speed control loop in such cases leads to inefficient cooling: fans may run at unnecessarily high speed [7], reducing overall amplifier efficiency, or may trigger premature muting despite other channels remaining within safe limits.

The effectiveness of a cooling system directly influences the thermal headroom of an audio amplifier. A fast and responsive airflow adjustment helps maintain lower device junction temperatures during transient high-power bursts, preventing early activation of thermal limiters and improving both short-term output capability and long-term reliability. These principles are well documented in the power electronics literature, where reduced temperature rise rates and minimized thermal cycling correlate with extended device lifetime [8–10].

1.2. The Four-Wire PWM Fan Standard (Technical Background)

The design of a modern, intelligent fan controller relies on the use of four-pin PWM fans, standardized by Intel Corporation (Santa Clara, CA, USA) [11]. This standard is a notable improvement over older two- and three-wire fan designs, offering a more sophisticated and reliable method for speed control [12]. The four wires are: power (+12 V), ground, tachometer, and PWM input. In a two-wire fan, speed is controlled by varying the supply voltage or by using a low-frequency PWM signal to chop the power line. This is an “open-loop” control method [13], as there is no feedback on the actual fan speed, and it is impossible to know if the fan is even spinning [14]. The three-wire fan adds a tachometer output, providing a feedback signal proportional to the fan’s rotational speed [15], which allows for closed-loop control. However, when using a low-frequency PWM signal, the tachometer output is only available during the “on” phase of the pulse, as the fan’s internal circuitry is unpowered during the “off” phase [14]. Four-wire fans solve this problem. The fan’s internal control circuitry is continuously powered by the dedicated +12 V line, leaving the PWM input solely to control the speed of the motor’s drive coils. This design ensures that the tachometer signal (open collector, two pulses/rev), which provides crucial feedback on the fan’s rotational speed, is always available, regardless of the PWM duty cycle [14]. This continuous feedback is vital for a robust control system, as it allows the microcontroller to detect fan faults, such as a stalled fan, and to precisely adjust the PWM signal to achieve a specific target RPM. However, the internal transistor switching and winding commutation circuits inside the fan act as multiple electromagnetic radiation sources, which can introduce electromagnetic interference (EMI) during PWM speed regulation [16].

1.3. Heatsinks and Heat Dissipation in High-Power Audio Amplifiers

High-power audio amplifiers can dissipate substantial amounts of heat, both in idle and under load. Typical Class-AB stages may generate 50–150 W of thermal dissipation per channel even at moderate output levels, while sustained high-power operation can raise this figure significantly as output devices approach their safe operating limits. Such thermal loads necessitate an efficient heatsink design and reliable forced-air cooling, particularly in compact 1 U and 2 U professional enclosures. In professional audio equipment, rack-mount enclosures are specified in rack units (U), where 1 U corresponds to a height of 44.45 mm (1.75 in) and 2 U provides 88.9 mm (3.5 in) of vertical clearance. Such low-profile formats impose strict constraints on heatsink volume, airflow path geometry, and fan placement, significantly increasing the thermal management challenge. As a result, forced-air cooling

becomes mandatory, with fans pulling in cooler ambient air and expelling heated air through the chassis vents [17]. High-quality amplifiers rely on aluminum heatsinks with high thermal conductivity, which rapidly transfer heat from the power transistors to the fins of the radiator [18,19]. In rare cases, some manufacturers employ copper heatsinks; copper offers even higher thermal conductivity than aluminum [20], allowing faster heat transfer and more uniform temperature distribution across the fins [21], though at the expense of increased cost and weight.

Certain heatsinks are specifically engineered for forced-air cooling and exhibit poor thermal performance when operated without adequate airflow. Heat dissipation in such systems occurs through three primary mechanisms. First, conduction transfers heat from the semiconductor die to its package and subsequently through thermal interface materials into the heatsink, where it is further spread across the base and fins. Second, convection most effectively achieved through forced airflow removes heat from the heatsink surfaces into the surrounding environment [22]. Finally, thermal radiation contributes to heat transfer by emitting infrared energy from hot surfaces to cooler surroundings. Experimental studies show that under forced convection, radiation can contribute around 10–12% of the total heat dissipation in heatsinks with surface treatments that increase emissivity (e.g., oxidized or chemically treated surfaces) [23]. Under natural convection, the fraction from radiation can rise substantially, in some configurations reaching 20–50% of the total heat dissipation, depending on geometry and temperature differences [24].

Not only the heatsink design but also the placement of the fan has a critical impact on overall cooling performance. Investigations on impinging air cooling with elliptical fin heatsinks have shown that optimized fan positioning can enhance heat transfer, with Nusselt number [25] analysis confirming superior performance when the fan is embedded in a cut-out configuration [26].

1.4. Limitations of Existing Control Approaches and Contribution of This Work

Several commercial integrated circuit fan controllers exist (e.g., MAX6650/MAX6651 Maxim Integrated Products, Inc., San Jose, USA [27], ST's STLM75, STMicroelectronics, Plan-les-Ouates, Switzerland [28]-based solutions), but these devices are primarily intended for PC/server environments and typically provide only group fan control. While functional, they lack the flexibility required for audio amplifier applications [29], where independent channel regulation and dedicated fault handling are critical. Moreover, such controllers usually require an additional MCU to supervise and fully configure their operation, which increases design complexity and overall cost.

This paper presents a standalone active thermal control [8] using a module designed specifically for high-power audio amplifiers. The system independently regulates two four-wire PWM fans, each associated with its own digital I²C temperature sensor, thereby enabling per-channel thermal management. Fan tachometer signals are continuously monitored to measure RPM and detect potential stalls, ensuring reliable operation. In the event of overheating or fan failure, the controller asserts a mute signal to the affected amplifier channel, protecting the system from damage. System status feedback is provided through onboard LEDs as well as an external front-panel indicator.

2. System Architecture and Implementation

2.1. Hardware Design—System Overview

The proposed controller was implemented around a low-cost ATmega328P 8-bit microcontroller (Microchip Technology Inc., Chandler, AZ, USA), chosen for its wide availability, robust peripheral set (including multiple timers, PWM outputs, I²C interface, and analog/digital I/O), and proven reliability in embedded control applications [30].

Other scientific papers demonstrated that the ATmega328P is capable of implementing temperature-based fan speed control systems: for example, Amadi Christopher Chidi designed a smart controller with this MCU to regulate fan speed using environmental temperature readings and an LCD display [31].

Similarly, a temperature-controlled DC fan was regulated using ATmega328, confirming that the MCU's PWM and sensor-reading capabilities are sufficient for precision control in such systems [32]. Additionally, ATmega328P has been used in industrial automation contexts requiring stable I/O and reliable operation under variable loads [33]. Finally, its ability to implement PWM schemes with consideration for EMI is supported by studies using the ATmega328P in power electronics switching [34] (e.g., EMI reduction in boost converters).

All control features presented in this work were experimentally validated using two independent programmable electronic loads configured to emulate the thermal behavior of a dual-channel audio amplifier. This controlled setup ensures reproducible thermal profiles while accurately representing the conditions encountered in practical high-power audio systems.

2.2. Temperature Sensing and Sensor Selection

The analog temperature sensors—such as thermistors, resistance temperature detectors (RTDs), or thermocouples—measure temperature via an output that is a continuous variable (voltage or resistance) that must be converted, scaled, linearized, and often conditioned with external circuitry such as amplifier stages, bias networks, and analog-to-digital converters (ADCs). While analog sensors can offer high sensitivity, their accuracy in a full system depends heavily on the precision of these external components, the linearity of the analog stage, the stability and noise of the ADC reference, lead wire resistance [35], external thermal effects, and calibration. In contrast, digital temperature sensors integrate the sensing element, calibration, linearization, and conversion internally; they communicate in a digital form (e.g., via I²C or SPI), reducing the number of external analog components and the burden of calibration, and often improving repeatability and ease of integration. The authors of [36] provided an example of a direct comparison between resistance thermistor systems and digital bandgap sensors in long-term field conditions. Although thermistors had somewhat lower measurement noise under certain conditions, digital sensors showed good agreement in temperature readings (± 0.1 °C in many cases), with acceptable delays and acceptable noise levels for many monitoring purposes, while being less sensitive to drift, moisture, and mechanical strain.

In our proposed controller system, temperature monitoring is achieved using two Microchip TC74 (Microchip Technology Inc., Chandler, AZ, USA) [37] digital temperature sensors connected via the I²C bus. Each amplifier channel is assigned a dedicated sensor, thereby allowing independent thermal supervision and control. This particular sensor model was selected for its TO-220 package, which closely resembles the form factor of a power transistor, making it exceptionally easy to mount directly onto the heatsink with reliable thermal contact and minimal mechanical adaptation. Given the confined and potentially electrically noisy setting of high-power audio amplifiers, digital sensors are preferred because they offer greater robustness. They resist drift caused by temperature cycling, humidity, and mechanical stress better than analog alternatives.

For connectivity, JST plug-and-play connectors were selected to facilitate rapid assembly and reliable integration. Two connectors are provided for the PWM fans, while additional JST headers allow external attachment of temperature sensors. This modular approach simplifies installation inside the amplifier chassis, ensuring that the controller can be replaced or upgraded without soldering or rewiring.

The power input stage supports dual configurations, reflecting the diversity of integration scenarios in professional audio equipment. A right-angled 90° pin connector enables direct slot-in connection to the main power supply unit (PSU) of the amplifier, minimizing cabling and optimizing mechanical stability. Alternatively, a standard two-pin JST header is available for standalone operation or use in other equipment, such as laboratory electronic loads or power supplies. To accommodate a wide range of fan types, the board incorporates an onboard voltage regulator that provides a stable MCU supply while supporting both 12 V and 24 V input rails. This allows the controller to be used interchangeably with either 12 V or 24 V PWM fans without additional hardware modification. The dual-input topology, together with adaptive supply support, greatly enhances the versatility of the module across different applications.

For diagnostic and operational feedback, two onboard LEDs are included. A permanent 1 Hz “heartbeat” LED provides continuous indication of MCU activity, confirming proper operation of the firmware, while a dedicated “fault” LED signals thermal or fan-related errors. An additional header is provided for an external LED, which may be mounted on the amplifier’s front panel, ensuring that fault information is visible to the end user.

The overall hardware design emphasizes compactness, low EMI generation, and cost efficiency, with a target bill of materials (BOM) cost of approximately €10. These constraints were considered essential to ensure seamless adoption in high-power audio amplifiers, where space is limited and reliability is critical.

2.3. Block Diagram

Figure 1 illustrates the block diagram of the controller. The system is built around a central microcontroller (MCU) that manages all thermal monitoring, fan control, and protection functions. Two temperature sensors (one per amplifier channel) feed real-time thermal data to the MCU via a digital interface. Based on this information, the MCU generates PWM signals to control the speed of Fan 1 and Fan 2, while simultaneously reading their tachometer outputs to monitor RPM and detect stalls or failures. Power is supplied through a 12/24 V input, later fed into a fixed 5 V voltage regulator to provide stable operating voltage. The MCU also drives status LEDs for user feedback and controls the mute amplifier lines for both channels, allowing it to protect the system under overheating conditions.




Figure 1. Controller block diagram.

2.4. Sensors Placement on Heatsinks

Sensor placement on the heatsink is not arbitrary: optimal locations depend on the heatsink geometry, the airflow direction imposed by the fans, and the thermal time constants of the cooled components [38]. In rack amplifiers, the fans are commonly collocated with the chassis intake/exhaust slots, so the sensors must be chosen and positioned with respect to the prevailing flow direction. When a fan is mounted at the front (intake) and draws cool ambient air through a tunnel of radiator fins, the air warms as it traverses the fin array, and the highest temperatures are typically found at the downstream (exhaust) end; placing the sensor near this downstream hotspot yields the earliest reliable indication of inadequate cooling for that channel. Conversely, if the fan is mounted at the end of the heatsink and pulls air through the fin tunnel to exhaust directly outside the chassis, the hottest region migrates toward the fan side, and the sensor should be placed accordingly. Figure 2 illustrates the expected temperature distribution along the heatsink depending on airflow direction, highlighting the regions where thermal gradients are most pronounced. For maximum convective efficiency, the fan dimension should closely match the heatsink footprint (many amplifier heatsinks are designed around common fan sizes such as 80–90 mm), and some high-power designs even use two fans per channel (one intake and one exhaust) to create higher mass flow. For example, the air entering the heatsink (near the intake fan) will be cooler, and the air exiting the heatsink (near the exhaust fan) will be hottest because it just absorbed all the heat.




Figure 2. Heat accumulation relative to fan position, (a) intake; (b) outtake.

Several prior works address optimal thermal sensor allocation and placement, emphasizing that accurate hotspot detection can be achieved even with a limited number of sensors if they are positioned in regions with steep thermal gradients or anticipated failure locations [38]. Heuristic methods, such as genetic algorithms [39] and dual clustering approaches [40], further confirm that placing sensors where thermal gradients are greatest—rather than distributing them uniformly—improves overheating detection. Experimental studies additionally show that model-based estimation techniques, including Kalman filtering [41], can reconstruct the underlying temperature field with high fidelity when combined with strategically placed sensors, while also quantifying the trade-off between the number of sensors and achievable accuracy. Finally, practical considerations such as sensor attachment (clips, epoxy, bolted contacts) and thermal coupling (e.g., high-quality thermal grease) influence measurement fidelity and must be considered in real-world heatsink monitoring.

2.5. PCB Design

The controller is implemented on a compact 45 × 35 mm, 1.6 mm thick two-layer PCB, optimized for ease of integration inside amplifier power supply units. Additional, mechanical stability is ensured by two 2.2 mm diameter screw mounting holes, allowing secure attachment to the chassis and minimizing vibration-induced stress on the connectors. A six-pin in-circuit programming header is included for firmware flashing and debugging, providing flexibility for future software upgrades or parameter adjustments.

Power distribution is carefully structured to support the dual-input topology described previously. The incoming supply, whether provided through the amplifier PSU's 90° slot-in connector or the auxiliary two-pin JST XHheader (JST Mfg. Co., Ltd., Osaka, Japan), is routed directly to the fan connectors and simultaneously to the input of the onboard voltage regulator. This arrangement ensures that the fans always operate from the raw 12 V or 24 V rail, while the microcontroller and associated logic consistently receive a regulated, fixed supply voltage. Since the current drawn by the system is dominated by the cooling fans, the required power input capacity depends directly on the fan specifications. For example, with two high-performance PWM fans rated at approximately 1.2 A each, the amplifier PSU must be able to provide at least 3 A of headroom (2.4 A for the fans plus a small margin for the regulator and MCU). Such an architecture not only maximizes flexibility but also guarantees robust operation across a wider input range.

To maintain stable MCU operation, the onboard voltage regulator is supplemented by local decoupling capacitors placed very close to the MCU's power pins. Such decoupling capacitors play a critical role in suppressing high-frequency transients and switching noise originating from the fan drivers or the amplifier's main PSU [42]. Studies have shown that improper decoupling or placing capacitors far from the device pins can lead to increased radiated emissions as ground/power loops expand, degrading both EMC performance and signal integrity [43]. Additionally, the ground layout is arranged with separation of the power and logic domains; sensitive logic and sensor circuits are physically isolated from high-current fan return paths. This minimizes ground bounce and prevents coupling of noise into the analog/digital front ends, a technique supported by the PDN (power distribution network) design literature showing that properly placed decoupling and plane segmentation significantly reduce power-supply-induced noise [44].

Although the tachometer output of a four-wire fan is not directly affected by the PWM control signal, since the PWM input drives the internal motor controller while the tach line is generated by a separate open-collector stage, high-frequency switching inside the fan can still introduce electrical noise on the tach signal. To ensure reliable RPM measurement, the tach line is therefore biased with a pull-up resistor and filtered with a low-pass RC network that attenuates high-frequency disturbances while preserving the pulse train corresponding to the actual rotational speed.

The tachometer output of a four-wire fan produces a pulse train proportional to rotational speed, with the frequency expressed as:

$$f_{tach} = \frac{\text{RPM}}{60} \times 2 \quad (1)$$

where the factor of two reflects the industry-standard behavior of providing two pulses per revolution. In high-power audio amplifiers, the use of high-RPM fans is comparable to those found in server equipment, which leads to significantly higher tachometer frequencies than those encountered with standard PC cooling fans. For example, a fan spinning at 10,000 RPM yields a tach frequency of approximately 333 Hz, while at 20,000 RPM the frequency increases to about 667 Hz. Thus, while consumer-grade fans typically generate signals in the range of 50–200 Hz, server-grade devices may reach 300–700 Hz.

To preserve the tach information while suppressing interference, the filter cutoff must be carefully chosen. In this design, the target cutoff frequency was set to approximately 2 kHz [45], which is high enough to pass tach pulses up to 700 Hz without distortion, yet low enough to attenuate switching noise. The cutoff frequency is defined by the time constant of the RC network:

$$RC = \frac{1}{2\pi f_c} = \frac{1}{2\pi \times 2000} \approx 80 \mu s \quad (2)$$

Equation (2) represents a target time constant of approximately 80 μs for a 2 kHz cutoff frequency. In the implemented prototype, standard component values (3.3 k Ω and 22 nF) were selected, resulting in an effective time constant of approximately 73 μs , corresponding to a cutoff frequency of about 2.19 kHz. This provides a sufficient margin above the maximum tachometer pulse frequency while maintaining effective attenuation of high-frequency electromagnetic interference.

In practice, this configuration ensures that the microcontroller receives a clean, logic-level tach signal with minimal jitter, enabling precise RPM measurement and reliable stall detection [46]. Similar filtering and observer-based techniques have been shown in the related motor and fan speed estimation literature to significantly reduce measurement error in noisy environments. For example, Ivan Petrić in [47] demonstrate that low-pass filtering of tach or encoder-derived signals improves speed estimation accuracy for induction machines under high switching noise.

2.6. Firmware Design and Control Logic

The firmware periodically samples temperature data from the two TC74 digital I²C sensors at one-second intervals, a rate that offers sufficient responsiveness for the thermal dynamics of power amplifier heatsinks while avoiding unnecessary microcontroller load and sensor noise accumulation. Each channel operates independently, receiving its own temperature measurement and adjusting its corresponding fan speed accordingly.

2.6.1. Linear Mapping with Deadband Hysteresis

Fan speed control is implemented through linear interpolation between predefined temperature limits rather than a closed-loop proportional–integral–derivative (PID) regulator. A closed-loop PID regulator was intentionally not used for fan control, despite its common application in temperature regulation, because the thermal dynamics of amplifier heatsinks are inherently slow and dominated by large time constants. Under such conditions, a PID loop offers no practical advantage over a simpler open-loop mapping scheme, as the system's thermal inertia prevents rapid corrective action and naturally suppresses high-frequency disturbances. Moreover, implementing a stable PID controller for a slow-response plant requires careful gain tuning and anti-windup protection, increasing firmware complexity and computational load without improving practical performance. For these reasons, linear temperature-to-duty-cycle mapping with hysteresis provides a deterministic, low-noise response that is fully adequate for heatsink cooling, while avoiding the oscillations and overshooting that poorly tuned PID controllers can introduce in thermally sluggish systems.

Beyond the main control law, the firmware also incorporates a set of well-defined thermal thresholds that govern the system's protection behavior. At 50 °C, if the temperature remains elevated for several minutes, the fan is forced to operate at 100% duty cycle to prevent long-term heat accumulation within the enclosure (anti-heat-soak protection). When the temperature reaches 70 °C, the system enters a warning state and activates the fault indicator, signaling increased thermal stress while allowing normal operation to continue. If the temperature rises to 90 °C, the controller asserts the MUTE line for

the affected amplifier channel, restoring output only after the temperature falls below the 70 °C warning level. Together, these thresholds establish a clear and hierarchical protection scheme that ensures stable, deterministic thermal management under all operating conditions. Below T_{min} , the fan remains off, while above T_{max} , the fan operates at full duty cycle. For intermediate values, the fan PWM duty cycle $D(T)$ is obtained by a linear mapping function:

$$D(T) = \begin{cases} 0, & T \leq T_{off} \\ D_{min} + \frac{(T-T_{min})}{T_{max}-T_{min}} \times (D_{max} - D_{min}), & T_{min} < T < T_{max} \\ D_{max}, & T \geq T_{max} \end{cases} \quad (3)$$

where $D_{min} = 5\%$ and $D_{max} = 100\%$.

To prevent rapid on–off cycling when the measured temperature oscillates near the activation threshold, a 3 °C deadband hysteresis is applied: the fans engage when the temperature reaches 30 °C and deactivate only once it falls below 27 °C. A secondary hysteresis regulates duty-cycle updates to avoid unnecessary micro-adjustments—fan speed is recalculated only when the sensed temperature changes by more than 2 °C from the last update point. This dual hysteresis prevents rapid on–off oscillations that may occur when minor temperature fluctuations or sensor quantization produce frequent threshold crossings.

To further improve stability, an Exponential Moving Average (EMA) filter is applied to the raw temperature data before they reach the control logic. Unlike a simple arithmetic moving average, which gives equal weight to all recent samples and requires storing multiple readings, the EMA applies progressively smaller weights to older values. This makes it more memory-efficient, computationally lighter, and faster to adapt to real thermal trends while still filtering out short-term noise [48].

The smoothing factor of the Exponential Moving Average was selected based on objective performance criteria rather than arbitrary tuning. Three candidate values ($\alpha = 0.1, 0.2$, and 0.5) were evaluated against identical thermal profiles, and the resulting filtered signals were compared according to three metrics: (1) response time to a genuine temperature increase, (2) suppression of short-term sensor noise, and (3) avoidance of unnecessary fan-speed adjustments. An α of 0.1 produced excessive lag, delaying the controller’s reaction to rapid thermal transitions, whereas $\alpha = 0.5$ amplified minor sensor fluctuations and generated frequent duty-cycle changes. The intermediate value $\alpha = 0.2$ provided the best balance, reducing noise while maintaining sufficient responsiveness to real heating events. For this reason, $\alpha = 0.2$ was adopted as the operational smoothing factor in all subsequent tests.

Finally, the filtered and stabilized temperature data are processed through a linear interpolation function to determine the PWM duty cycle. This open-loop mapping approach, as it is outlined in the manufacturer application literature from STMicroelectronics [28], ensures a predictable and consistent thermal response without the complexity of feedback control. Combined with hysteresis and EMA filtering, it achieves a balance between simplicity, responsiveness, and reliability, delivering smooth fan transitions, stable operation, and efficient thermal management even under dynamically varying workloads and ambient conditions [49].

The overall control algorithm is summarized in Figure 3. The system operates within a main loop architecture, where temperature acquisition and filtering are performed continuously, while the control logic is evaluated at fixed time intervals. The algorithm first computes the base PWM duty cycle from the filtered temperature values, after which cooperative fan assistance and optional anti-heat-soak mechanisms are applied. Safety-related checks, including tachometer-based stall detection and predictive latch activation,

have higher priority and may override the computed PWM values by enforcing maximum duty. Thermal muting is handled independently per channel with hysteresis to prevent oscillatory behavior. This layered decision structure ensures stable operation under normal conditions while maintaining deterministic high airflow demand response capability.




Figure 3. Control algorithm architecture (flowchart).

2.6.2. Fan Monitoring, Fault Detection, and Thermal Protection

While the fans are operating, the tachometer signal from each unit is sampled once per second. The microcontroller counts the number of tach pulses accumulated during this 1 s interval and converts the result into rotational speed according to relation:

$$N_{fan} = \frac{60}{n_p} \times f_{tach} \quad (4)$$

where $n_p = 2$ represents the two tach pulses generated per revolution by standard four-wire PWM fans. This pulse-counting method provides a stable and low-noise estimate of fan speed with minimal computational overhead. If the resulting RPM remains below the stall threshold (typically around 1000 RPM) for more than two consecutive measurement intervals, a fault condition is declared. In this state, the remaining fan for that thermal zone is automatically commanded to full duty cycle to maintain adequate airflow. This scenario is illustrated in the experimental results section.

Thermal protection logic complements the control system. When the measured temperature exceeds 70 °C, the system enters a warning state indicated by the fault LED. If the temperature surpasses 90 °C, the controller asserts the mute signal for the corresponding amplifier channel, ensuring thermal protection of the power stage.

2.6.3. Proactive Control

As already stated, most systems regulate fan speed reactively: the current temperature directly dictates the fan's output. However, a smarter approach involves extending this logic toward a model predictive control (MPC) concept, where the system not only reacts but also anticipates thermal dynamics. Building on this principle, several intelligent features have been introduced to improve reliability, thermal stability, and long-term performance: anti-dust startup sequence, predictive boost, active cross-cooling, and anti-heat-soak protection.

2.6.4. Anti-Dust Startup Sequence

Inside an amplifier's chassis, the presence of constantly running fans leads, over time, to the accumulation of superficial, un-embedded dust particles (often called floating dust) that do not adhere firmly but stay suspended or loosely settled on surfaces. The purpose of the anti-dust startup sequence is to expel this loose dust by briefly running the fans at a higher speed for a short duration—5 s immediately upon startup—creating a burst of airflow that helps clear the internal air space before the system enters its normal thermal control mode. For instance, a patent demonstrates using a device for cleaning dust in computers cases that relies on fan-reversal or a high-speed boost sequence to remove dust build-up in electronics [50]. Furthermore, studies of dust removal from surfaces via elevated airflow velocities confirm that higher-speed air bursts significantly enhance the removal of loosely adhered particles [51]. Also, during the startup sequence the firmware logic can check if any of the fans are not already stalled or defective (an RPM threshold is defined in firmware to verify whether a fan's speed falls below the expected range, lower than 1000 RPMs for chosen fans, which would indicate early-stage failure), immediately prompting the user with a visual warning by the on-board and front panel LEDs.

2.6.5. Predictive Boost

The predictive boost feature introduces a proactive control layer inspired by model predictive control principles, allowing the system to anticipate rather than merely react to thermal changes [52]. The controller continuously monitors recent temperature samples and maintains a short-term history for each sensor. A big heatsink, as is typically used in high-power audio amplifiers, comes with big inertial thermal mass, and because of this it takes time to heat it or cool it down. The predictive boost method aims to move cool air before the entire heatsink can become saturated with heat. This keeps the core temperature of the device (the transistor die itself) under control before it spikes up and rather than applying powerful cooling when it is already too late. By aggressively pre-cooling, we slow down the rate of temperature rise during a peak load and extend the time the device can operate at its maximum capacity before it has to thermal throttle, as demonstrated in the experimental validation..

2.6.6. Active Cross-Cooling

The active cross-cooling feature ensures balanced cooling across the system when thermal conditions between the two heatsinks become uneven. When one temperature sensor registers a value more than 15 °C higher than the other, the cooler zone's fan temporarily increases its speed to 50% of the hotter fan's duty cycle. This cooperative response helps equalize airflow distribution, accelerating heat dissipation from the hotter area and preventing localized thermal buildup. To maintain stability, the assist feature deactivates once the temperature difference between the two zones drops below 10 °C, creating a 5 °C hysteresis window. The experimental validation of this mechanism is presented in results section.

2.6.7. Anti-Heat Soak

Sometimes, a heavy load is not sharp and brief, it is sustained thermal stress that keeps temperatures high but stable for several minutes or even hours. During such periods, while the fan may successfully maintain that steady temperature, and the entire chassis gradually heat-soaks, turning into a large, warm radiator that stores thermal energy.

The anti-heat-soak feature is specifically designed to prevent this scenario. When a temperature of 50 °C or higher is detected for a prolonged duration (approximately 5 min), the corresponding fan is forced to operate at a 100% duty cycle for as long as the

temperature remains at or above the threshold. This response helps purge accumulated heat from the heatsink and enclosure, ensuring that internal components do not remain thermally saturated. By clearing residual heat buildup early, the system maintains better cooling efficiency, shortens post-load recovery time, and enhances long-term reliability.

2.7. Experimental Setup and Validation Procedure

To ensure consistent and repeatable thermal conditions throughout testing, two independent programmable electronic loads rated for 150 V and 20 A per channel were used in place of a complete audio power amplifier. Each channel was operated in constant power (CP) mode, allowing precise and reproducible emulation of controlled thermal dissipation profiles representative of high-power amplifier operation. This configuration accurately simulates a dual-channel amplifier while enabling rapid power-step transitions, gradual heating ramps, and sustained high-power operation under controlled conditions.

Although validation on a complete audio amplifier would represent a realistic application scenario, controlled thermal characterization under such conditions presents significant practical limitations. Achieving comparable and repeatable thermal stress would require high-power dummy loads and sustained near-maximum sinusoidal excitation, operating the amplifier under extreme and potentially damaging stress conditions that do not reflect typical music-driven signals, which are inherently dynamic and non-stationary. In contrast, programmable electronic loads provide deterministic power dissipation independent of signal content, enabling safe, accelerated, and fully repeatable evaluation of the controller across its entire thermal operating range.

For accurate thermal feedback, each temperature sensor was mounted directly above the power transistor and secured with thermally conductive grease to minimize measurement lag and improve thermal contact. Cooling was provided by two high-performance QUR0812SH fans (Delta Electronics, Inc., Taipei, Taiwan) (12 V, 0.5 A, 5500 RPM), selected for their high RPM and cubic feet per minute (CFM). In the context of high-power audio amplifiers, fan noise is not a limiting factor, as the acoustic output of the system far exceeds any audible contribution from the cooling fans; however, for environments where acoustic noise is critical, the controller's adaptive duty-cycle regulation ensures reduced fan speed under low thermal load conditions, thereby minimizing audible noise emissions [53].

Temperatures, fan RPM, and function states data were logged via a separate Arduino-based data acquisition unit, using ArduSpreadsheet v1.1 for real-time monitoring and structured data export. The main controller board was powered by a battery-driven switch mode supply.

This setup provided a controlled, repeatable, and instrumented environment for comprehensive evaluation of the fan control system across all implemented features.

Figure 4 illustrates the mechanical integration of the TC74 temperature sensor, secured directly onto the package of the power transistor within the electronic load. This arrangement ensures that the sensor measures the actual device case temperature. A thin layer of thermal grease was applied between the sensor and the transistor to improve thermal coupling and enhance the accuracy of temperature acquisition under dynamic load conditions.

Figure 5 presents the complete test configuration of the controller, with both PWM fans and TC74 temperature sensors connected to the board through the designated JST headers. A white TX lead (used solely for tests during development) links the microcontroller to the external data-acquisition module, while the orange wire provides a shared ground reference between the two systems. The module is powered through the red-black supply pair, and the illuminated status LED confirms proper firmware operation during real-time thermal monitoring and control experiments.




Figure 4. Mounting method of the TC74 digital temperature sensor.




Figure 5. Fully assembled fan controller PCB connected to the PWM fans, temperature sensors, data acquisition interface, and power supply.

Figure 6 shows the full laboratory configuration employed to evaluate the controller's behavior under reproducible thermal loads. Two programmable electronic loads operate independently to emulate the dissipation profiles of a dual-channel amplifier, supplied by a laboratory power source positioned on the left. On the right, a portable battery-powered supply feeds the fan controller module, ensuring electrical isolation and eliminating PSU-induced noise during measurement. The laptop in the background runs ArduSpreadsheet, receiving real-time data from the acquisition interface and enabling synchronized logging of temperature, RPM, and functional state transitions throughout the tests.




Figure 6. Full experimental setup used for validating the dual-channel thermal management system under controlled power-dissipation conditions.

3. Results

This section presents the experimental validation of the proposed dual-channel thermal management system. All functional mechanisms described in Section 2 were evaluated under controlled and repeatable thermal loading conditions. The following subsections analyze the controller behavior under fault conditions, overheating scenarios, predictive operation, and sustained thermal stress.

3.1. Experimental Thermal Performance Evaluation

Figure 7 illustrates the fan failure scenario. When Fan 1, responsible for cooling the active high-load channel, was intentionally disconnected, Fan 2 automatically ramped up to full speed to compensate for the loss of airflow and maintain adequate cooling for the system. At this time, the fault LED flashes at 4 Hz. After some time, Fan 1 is reconnected, the fault condition is removed, and Fan 2 ramps down.




Figure 7. System reaction to a fan failure.

Figure 8 illustrates the system's behavior under an intentional overheating condition. A sustained high load was applied until the temperature reached the 90 °C overheating threshold, at which point the MCU's mute pin was driven low, as shown by the green trace in the graph (scaled $\times 10$ for visibility). This action immediately muted the affected channel, allowing the temperature to gradually decrease. Once it cooled below 70 °C, the system automatically re-enabled the output, after which the temperature began to rise again as the load was reconnected, demonstrating readiness for subsequent thermal protection cycles.




Figure 8. System response to overheating conditions.

The graph in Figure 9 illustrates the operation of the predictive boost mechanism. Activation occurs when the predictive delta threshold exceeds 7 °C, calculated as the difference between the current temperature and the value recorded 10 s earlier. When this condition is met, the controller immediately engages both cooling fans at 100% duty cycle for a fixed duration—60 s in testing—rapidly evacuating heat before it can saturate the heatsink.




Figure 9. Experimental validation of the predictive boost control strategy.

As shown, once the function was triggered, the temperature increased from 25 °C to 60 °C within 119 s, after which it stabilized at a more thermally steady equilibrium. If the temperature continues to rise during the active period, the latch automatically extends the boost interval, sustaining maximum fan output until conditions normalize.

In essence, predictive boost acts as an emergency, forward-looking safeguard that rapidly intervenes during fast heating events, allowing the system to recover to safe operating levels or at least slow down as much as possible when a heavy transient load causes a temperature spike in the power devices.

This anticipatory behavior represents a simplified form of model predictive control, where the system reacts to thermal trends rather than static thresholds, effectively looking ahead to counteract imminent overheating before it occurs. In order to prove the predictive behavior, a second set of data was acquired with the function turned off, the fan speed being adjusted through normal linear interpolation only. Such behavior is displayed in Figure 10.




Figure 10. Thermal behavior without predictive boost.

As illustrated in Figure 10, the temperature increased from 20 °C to 60 °C within approximately 83 s, after which it began to stabilize.

To compare the effectiveness of the predictive boost mechanism, identical thermal load profiles were applied under two configurations: baseline linear temperature–PWM mapping (Predictive OFF) and predictive control enabled (Predictive ON). Starting from approximately 20 °C, the system response was evaluated in terms of transient heating behavior, peak temperature, and steady-state performance. The time required to reach 60 °C increased from 83s in the baseline configuration to 119s with predictive boost enabled, corresponding to a 43% delay in thermal rise. The maximum recorded temperature decreased from 69 °C to 62 °C, while the steady-state mean temperature was reduced from 68 °C to 61 °C. The initial heating rate decreased slightly from ~0.354 °C/s to ~0.336 °C/s, indicating that the primary benefit of predictive control lies in early high-duty airflow activation rather than drastic slope suppression alone. As expected, the average PWM duty cycle increased from 63% to 70%, reflecting the proactive cooling strategy that trades moderate additional airflow for a lower peak and steady-state temperatures.

Overall, these quantitative results, also represented in Table 1, confirm that the predictive boost mechanism extends the safe thermal operating window and reduces maximum thermal stress without altering the fundamental steady-state physics of the cooling system.

Because the two electronic loads used in the experimental setup were housed in separate enclosures, the corresponding heatsinks were not thermally coupled. As a result, the activation of the cross-cooling assistance did not produce a measurable temperature change on the adjacent heatsink. Nevertheless, the test represented by Figure 11 confirms that the control logic functions correctly: the auxiliary fan responds deterministically to inter-channel temperature differences, transitions cleanly into and out of the assist

state, and follows the defined hysteresis rules. Although thermal equalization could not be demonstrated in this isolated configuration, the experiment validates the intended algorithmic behavior and its readiness for operation in a shared-enclosure amplifier where the two thermal zones interact physically.

Table 1. Comparison between baseline control and predictive boost configuration.

Metric	Baseline (Predictive OFF)	Predictive ON	Relative Change
Time to 60 °C (s)	83	119	+43%
Initial heating rate (°C/s)	0.354	0.336	−5%
Maximum temperature (°C)	69	62	−10%
Steady-state mean temperature (°C)	69	62	−10%
Average PWM duty cycle (%)	63	70	12%




Figure 11. Experimental behavior of the active cross-cooling mechanism under asymmetric thermal conditions.

Figure 12 illustrates the activation cycle of the anti-heat-soak feature. Once the function engages, fan speed and exhaust airflow increase sharply, expelling accumulated hot air from within the amplifier enclosure. As a result, the internal temperature begins to gradually decline, demonstrating how the intensified airflow effectively removes trapped heat and restores the system toward a cooler, steady-state operating condition. With the load remaining active, the function will be triggered again for the next cycles.

Table 2 summarizes the dynamic response of the anti-heat-soak mechanism under sustained thermal load. The Start of max-duty (PWM 100%) indicates the moment at which the controller forced both fans to maximum-duty cycle following prolonged exposure above the defined thermal threshold. The Duration of max-duty activation phase represents the total time interval during which the maximum PWM override remained active. Time to reach 50 °C threshold denotes the recovery time measured from the onset of the max-duty phase until the temperature decreased below 50 °C. The Temperature at activation phase corresponds to the measured temperature at the instant of override initiation, while the Minimum temperature during max-duty phase reflects the lowest temperature reached before normal control resumed. Finally, Mean RPM during max-duty phase indicates the average fan rotational speed during the forced maximum-duty interval, confirming full-speed operation and enhanced airflow.




Figure 12. Anti-heat-soak.

Table 2. Anti-heat-soak behavior under sustained load.

Metric	Value
Start of max-duty (PWM 100%)	170 s
Duration of max-duty activation phase	210 s
Time to reach 50 °C threshold	180 s
Temperature at activation phase	57 °C
Minimum temperature during max-duty phase	48 °C
Mean RPM during max-duty phase	~5000 RPM

3.2. Comparison with Existing Fan Control Approaches

To position the proposed system within the current state of the art, a comparison was performed against representative fan control approaches reported in the literature and commercially available integrated solutions. Prior work on server-oriented cooling systems primarily focuses on reactive PID or optimized linear fan control strategies aimed at reducing energy consumption in data center environments [49,54], without incorporating predictive multi-zone cooperation tailored for asymmetric thermal loads. Conventional temperature-based fan controllers typically rely on linear or PID-based temperature-to-duty mapping and are often designed for general-purpose electronics or server environments. While such systems may incorporate tachometer feedback and closed-loop RPM regulation, they usually operate on a shared control loop without independent thermal supervision per channel.

Commercial integrated controllers (e.g., MAX6650/MAX6651 [27], ST-based digital sensor + MCU combinations) provide reliable fan speed regulation and fault monitoring; however, they are primarily optimized for PC/server architectures and generally lack application-specific features such as predictive thermal intervention, cross-assisted cooling between asymmetric thermal zones, or anti-heat-soak protection tailored for high-power audio amplifiers.

In contrast, the proposed controller integrates independent dual-channel temperature supervision, predictive boost logic based on short-term thermal trend analysis, active cross-cooling, anti-heat-soak enforcement, stall detection with automatic compensation, and channel-specific mute protection in a low-cost standalone architecture. Table 3 summarizes the key functional differences.

Table 3. Functional comparison between representative fan control approaches and the proposed system.

Feature	Conventional Linear/ PID Fan Control	Commercial IC Fan Controllers	Proposed System
Independent dual-channel control	✗	Partial ¹	✓
Tachometer-based stall detection	Partial ²	✓	✓
Predictive thermal intervention	✗	✗	✓
Cross-assisted cooling between zones	✗	✗	✓
Anti-heat-soak enforcement	✗	✗	✓
Channel-specific mute protection	✗	✗	✓
Application-specific tuning for audio amplifiers	✗	✗	✓

✓—feature fully supported; ✗—feature not supported; Partial—limited or conditional support. ¹ In certain multi-fan implementations, independent control ICs are used for each fan channel. While this allows separate speed regulation, the controllers typically operate without inter-channel communication, preventing coordinated thermal assistance between asymmetric zones. ² Fan stall detection in conventional analog designs may require additional hardware such as current shunt resistors and comparator circuits, which increases component count, PCB area, and overall bill of materials (BOM) cost.

4. Conclusions

This study introduced and experimentally validated a hierarchical multi-layer fan control architecture for thermally demanding embedded systems. Unlike conventional temperature-to-PWM mapping strategies, the proposed controller integrates predictive trend evaluation, cooperative cross-channel assistance, staged hysteresis, and anti-heat-soak logic within a structured priority-based override framework.

Experimental evaluation confirmed measurable improvements across all implemented mechanisms. The predictive boost reduced peak temperature under identical load conditions and improved thermal response during rapid transients. Cooperative cross-cooling enhanced multi-zone stability by dynamically assisting the thermally stressed channel, limiting inter-channel temperature divergence. Under sustained load, the anti-heat-soak function actively mitigated long-term thermal accumulation, enabling temperature recovery below 50 °C within 180 s while maintaining full cooling capacity during the override interval.

By combining independent sensing, tachometer-based stall detection, and deterministic override layers, the system ensures robust behavior under both nominal and fault conditions. The proposed architecture demonstrates that intelligent coordination between predictive, cooperative, and protective control layers can improve thermal stability without increasing hardware complexity.

The presented framework establishes a scalable methodology for advanced airflow control in high-power electronic systems and provides a foundation for future adaptive and workload-aware thermal management strategies.

Author Contributions: Conceptualization, A.M., E.-V.B., and H.A.; methodology, A.M.; software, A.M. and E.-V.B.; validation, A.M., E.-V.B., and H.A.; resources, A.M.; data curation, A.M. and E.-V.B.; writing—original draft preparation, A.M.; writing—review and editing, A.M.; visualization, E.-V.B.; supervision, H.A.; funding acquisition, A.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data are unavailable due to privacy restrictions.

Acknowledgments: During the preparation of this manuscript/study, the author(s) used Nano Banana 3 for the purposes of creating an illustrative thermal gradient overlay on a CAD-modeled heatsink figure. In addition, GenAi (Gemini 3 Flash Preview) was used only to support language refinement and academic expression, while all technical content, analysis, and conclusions remain entirely the responsibility of the authors. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

SMPS	Switched-Mode Power Supply
BLDC	Brushless Direct Current
PWM	Pulse-Width Modulation
RPM	Revolutions Per Minute
EMI	Electromagnetic Interference
U	Rack Unit (standardized height unit for rack equipment: 1 U, 2 U, etc.)
CFM	Cubic Feet per Minute
I ² C	Inter-Integrated Circuit
MCU	Microcontroller Unit
PSU	Power Supply Unit
BOM	Bill of Materials
RC	Resistor-Capacitor (filter topology)
Hz	Hertz (unit of frequency)
PID	Proportional-Integral-Derivative
EMA	Exponential Moving Average
MPC	Model Predictive Control
LCD	Liquid Crystal Display

References

1. Militaru, A.; Buica, E.V.; Horia, L.A. Technical Review of Class TD Audio Power Amplifiers. In Proceedings of the International Conference on Electronics, Computers and Artificial Intelligence, Targoviste, Romania, 26–27 June 2025.
2. Ronak, I. An Analysis of Heat Dissipation Techniques in Power Electronics. *J. Eng. Appl. Sci. Technol.* **2023**, *5*, 1–6.
3. Luke, M.; Masud, B.; Graham, M. An Experimental and Computational Study of Heat Transfer in High Power Amplifiers. *Heat Transf. Eng.* **2007**, *26*, 81–92.
4. Wen-Xiao, C.; Yun-Hsuan, L.; Kuei-Ru, H.; Chi-Chuan, W. Energy saving of fans in air-cooled server via deep reinforcement learning algorithm. *Energy Rep.* **2021**, *7*, 3437–3448. [\[CrossRef\]](#)
5. George, L.; Sergey, K.; Ilya, K. Regulation of the cooling fan rotation speed, as an effective instrument for reducing mechanical consumption of energy. In Proceedings of the International Conference on Modern Trends in Manufacturing Technologies and Equipment: Mechanical Engineering and Materials Science (ICMTMTE 2019), Sevastopol, Russia, 9–13 September 2019.
6. Ivan, M.; Ivan, Š. Improving Efficiency of AB-Class Audio Power Amplifier Using Thermoelectric Generators. *Electronics* **2025**, *14*, 34.
7. Porya, S.H.P.; Abbas, R.; Mojtaba, H.T. Temperature and performance evaluation of multiprocessors chips by optimal control method. *Bull. Electr. Eng. Inform.* **2022**, *12*, 749–759.
8. Anas, I.; Mohamed, S.; Mohamad, K.M.J.; Maria, T.D. Review of Active Thermal Control for Power Electronics: Potentials, Limitations, and Future Trends. *IEEE Open J. Power Electron.* **2024**, *5*, 414–435. [\[CrossRef\]](#)
9. Markus, A.; Giampaolo, B.; Marco, L. Study of reliability-efficiency tradeoff of active thermal control for power electronic systems. *Microelectron. Reliab.* **2016**, *58*, 119–125. [\[CrossRef\]](#)
10. Gagner, T. PWM Fan Controller Benefits for Enhanced Cooling and Efficiency. Innovations in Cooling Technologies. 2025. Available online: <https://gagner-toomey.com/pwm-fan-controller-benefits-for-enhanced-cooling-and-efficiency/> (accessed on 17 February 2026).
11. Intel Corporation. *4-Wire Pulse Width Modulation (PWM) Controlled Fans*; Intel Corporation: Santa Clara, CA, USA, 2005.
12. ACDCEFan. What Is PWM Fan and Why It's Crucial for Efficient Cooling Systems. Available online: <https://www.acdcefan.com/what-is-pwm-fan/> (accessed on 25 September 2025).

13. Analog Devices. Fan Speed Control Is Cool. Technical Articles. 2002. Available online: <https://www.analog.com/en/resources/technical-articles/fanspeed-regulators.html> (accessed on 24 September 2025).
14. Analog Devices. How to Control Fan Speed. Available online: <https://www.analog.com/en/resources/analog-dialogue/articles/how-to-control-fan-speed.html> (accessed on 24 September 2025).
15. Rahul, H.; Hua, H. Motor Rotation Speed Estimation based on Magnetic Inductive Sensing. *arXiv* **2025**, arXiv:2503.10885. [CrossRef]
16. Jinsheng, Y.; Xuan, Z.; Baolin, N. An Equivalent Modeling Method for Electromagnetic Radiation of PWM Fans with Multiple Radiation Sources. *Appl. Sci.* **2025**, *15*, 2887. [CrossRef]
17. Zhihao, Z.; Xuehui, W.; Yuying, Y. A review of the state-of-the-art in electronic cooling. *e-Prime-Adv. Electr. Eng. Electron. Energy* **2021**, *1*, 100009.
18. Imran, Z.; Muhammad, F.; Muhammad, F.; Muhammad, U.; Adnan, Q.; Muhammad, I.; Mejdal, A.A.; Saqib, A.; Muhammad, S.; Muhammad, Y.J. Thermal Performance Analysis of Various Heat Sinks Based on Alumina NePCM for Passive Cooling of Electronic Components: An Experimental Study. *Energies* **2022**, *15*, 22. [CrossRef]
19. Ting, K.; Yuxin, Y.; Yuncong, J.; Yanmei, K.; Binbin, J. Enhanced Thermal Management of GaN Power Amplifier Electronics with Micro-Pin Fin Heat Sinks. *Electronics* **2022**, *9*, 11.
20. Mahmodul, A.M.; Rizwanur, R.; Ahmad, K.A.K.; Hamza, A. Numerical Investigation of Heat Transfer Characteristics of Aluminium and Copper Heat Spreader. In Proceedings of the International Conference on Mechanical Engineering and Renewable Energy, Chattogram, Bangladesh, 12–14 December 2021.
21. Ramesh, T.; Ramesh, V.; Karthik, K. Design of Heat Sink and Thermal Analysis for Electronics Applications. *Int. J. Innov. Technol. Explor. Eng.* **2019**, *8*, 4091–4094. [CrossRef]
22. Oisin, M.; Rajesh, N.; Syed, M.A.; Tim, P. A Parametric Design Study of Natural-Convection-Cooled Heat Sinks. *Fluids* **2023**, *8*, 234.
23. Yuxin, Y.; Beibei, Z.; Sulian, T.; Zihui, L.; Biao, T.; Rui, Z.; Dong, Y. Effect of Surface Microstructure on the Heat Dissipation Performance of Heat Sinks Used in Electronic Devices. *Micromachines* **2021**, *12*, 265. [CrossRef]
24. Wen-Xiao, C.; Yu-Chieh, L.; Ching-Yao, C.; Chi-Chuan, W. Experimental and numerical study on the performance of passive heat sink having alternating layout. *Int. J. Heat Mass Transf.* **2019**, *135*, 822–836. [CrossRef]
25. Wikipedia. Nusselt number. Available online: https://en.wikipedia.org/wiki/Nusselt_number (accessed on 26 September 2025).
26. Adil, A.M.; Saad, A.R. Effect of air fan position on heat transfer performance of elliptical pin fin heat sink subjected to impinging air flow. *J. Therm. Eng.* **2021**, *7*, 1406–1416. [CrossRef]
27. Analog Devices. MAX6650. Available online: <https://www.analog.com/en/products/max6650.html> (accessed on 26 September 2025).
28. STMicroelectronics. Fan Speed Controller Based on STDS75 or STLM75 Digital Temperature Sensor and ST72651AR6 MCU. Available online: https://www.st.com/resource/en/application_note/an2680-fan-speed-controller-based-on-stds75-or-stlm75-digital-temperature-sensor-and-st72651ar6-mcu-stmicroelectronics.pdf (accessed on 26 September 2025).
29. William, B.T. PWM/PID Fan Controller. Available online: <https://www.brodie-tyrrell.org/fancontrol/> (accessed on 26 September 2025).
30. Tareq, S.E.H.; Mohammad, A.; Wasif, S.; Ahmad, A.J. Arduino and Labview Based Control for Efficient Drive of Cooling Fan System. *Res. J. Appl. Sci. Eng. Technol.* **2016**, *13*, 771–780.
31. Amadi, C.C.; Nnamdi, H.I. Design and Construction of a Smart Microcontroller Temperature-Based Fan Speed Controller Using ATmega328P Microcontroller and Liquid Crystal Display (LCD). *J. Inovasi Teknol. Eduk. Tek.* **2025**, *5*.
32. Joshua, D.B.; Emmanuel, S.S.; Peter, O.I.; Haonat, O.A.; Bakare, A.A. Microcontroller-Based Temperature Controlled DC Fan Using Atmega328. *Eng. Technol. J.* **2022**, *6*.
33. Vishal, C.; Atma, R.G. Industrial Automation Using ATmega328p-Based Programmable Logic Controller. In *Smart Energy and Advancement in Power Technologies*; Springer: Singapore, 2022; pp. 17–26.
34. Dhia, C.; Lassâad, S. Implementation of different randomized PWM schemes using Atmega328p microcontroller for EMI reduction in boost converter. *Electr. Eng.* **2020**, *102*, 1063–1071. [CrossRef]
35. Apinai, R.; Supatsorn, P.; Thawatchai, K.; Vanchai, R.; Wandee, P. A Procedure for Precise Determination and Compensation of Lead-Wire Resistance of a Two-Wire Resistance Temperature Detector. *Sensors* **2022**, *22*, 4176.
36. Lars, W.; Marcia, P.; Chasper, B. Brief communication: Comparison of the performance of thermistors and digital temperature sensors in a mountain permafrost borehole. *Cryosphere* **2023**, *17*, 10. [CrossRef]
37. Microchip. TC74 Tiny Serial Digital Thermal Sensor. Available online: <https://ww1.microchip.com/downloads/en/DeviceDoc/21462D.pdf> (accessed on 30 September 2025).
38. Xin, L.; Wei, Z.; Wen, J. Rigid sensor allocation and placement technique for reducing the number of sensors in thermal monitoring. *J. Shanghai Jiaotong Univ.* **2017**, *22*, 481–492.
39. Xin, L.; Xueting, W.; Wei, Z. Heuristic thermal sensor allocation methods for overheating detection of real microprocessors. *IET Circuits Devices Syst.* **2017**, *11*, 559–567.

40. Xin, L.; Mengatian, R.; Tao, L.; Liang, Z. Research of thermal sensor allocation and placement based on dual clustering for microprocessors. *SpringerPlus* **2013**, *2*, 253. [[CrossRef](#)]
41. Xin, L.; Xingtao, O.; Zhi, L.; Henglu, W.; Zhemin, D. On-Line Temperature Estimation for Noisy Thermal Sensors Using a Smoothing Filter-Based Kalman Predictor. *Sensors* **2018**, *18*, 433.
42. Rajesh, R.; Gajanan, K.; Jayant, C. Analysis of Decoupling Capacitor Performance in Improving Power Integrity in Two Layer/Three Layer Printed Circuit Boards. *Eur. J. Electr. Eng. Comput. Sci.* **2023**, *7*, 34–37. [[CrossRef](#)]
43. Deutschmann, B.; Juch, N. Influence of the IC power supply and decoupling capacitor arrangement on the electromagnetic emission. *Elektrotech. Informationstech.* **2023**, *141*, 66–75. [[CrossRef](#)]
44. Obaidul, H.M.; Yang, Z.; Hesam, F.M.; Yue, Z.; Michael, D.; Muhannad, S.B. Design Space Exploration of Power Delivery for Advanced Packaging Technologies. *arXiv* **2020**, arXiv:2008.03124. [[CrossRef](#)]
45. Hongze, F.; Yunting, L.; Qingyang, Z. Research on the Passive and Active Low Pass Filters. *Highlights Sci. Eng. Technol.* **2025**, *46*, 260–265.
46. Mehmet, C.O.; Gungor, B.; Durmus, U. Design and application of a novel high precision and low cost electronic tachogenerator for sensor-based brushless direct current motor drivers. *Rev. Sci. Instrum.* **2017**, *88*, 35005.
47. Ivan, P.; Slobodan, V. High-performance speed estimation of induction machines based on adaptive filtering. *IET Electr. Power Appl.* **2020**, *14*, 695–704.
48. Jim, E. The exponentially weighted moving average applied to the control and monitoring of varying sample sizes. *WIT Trans. Model. Simul.* **2011**, *51*, 3–13. [[CrossRef](#)]
49. Jeffrey, S.; Rickard, B.; Jonas, G.; Jon, S. Experimental Analysis of Server Fan Control Strategies for Improved Data Center Air-based Thermal Management. In Proceedings of the 19th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 21–23 July 2020.
50. Shih-Chi, C. Method of Cleaning Dust Using a Cooling Fan of a Computer and Related Dust-Cleaning Device. U.S. Patent US20070095522A1, 5 October 2010.
51. Zheng, Z.; Lu, H.; Zhao, W. Research of dust removal performance and power output characteristics on photovoltaic panels by longitudinal high-speed airflow. *Ideas* **2024**, *304*, 132202. [[CrossRef](#)]
52. Sánchez-Amores, A.; Ferramosca, A.; Maestre, J.M.; Camacho, E.F. Partially cooperative robust distributed model predictive control for tracking. *Automatica* **2025**, *182*, 112547. [[CrossRef](#)]
53. Jianhu, G. The Design and Implementation of a Noise Reduction System for a Microcomputer Cooling Fan. *J. Phys. Conf. Ser.* **2021**, *1852*, 022060.
54. Chengming, L.; Rongshun, C. Optimal Self-Tuning PID Controller Based on Low Power Consumption for a Server Fan Cooling System. *Sensors* **2015**, *15*, 11685–11700. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.