

Evaluating Active Thermal Regulation in Polycrystalline Silicon Solar Cells Under Thermal Stress and Constant Irradiance

Naavya Garg
naavyagarg4@gmail.com

ABSTRACT

This study examined whether a 3D-printed, microcontroller-controlled cooling system could recover the efficiency that polycrystalline silicon solar cells lose under thermal stress, while keeping irradiance fixed throughout. A single poly-Si cell was tested under three conditions: no cooling, a passive 3D-printed heat sink, and the same heat sink paired with an Arduino-controlled fan that switched on automatically once surface temperature reached 45°C. The study has demonstrated a relationship that is well established in theory but rarely shown directly with low-cost, accessible hardware: as a solar cell heats up, its voltage and power output fall predictably, and a simple active cooling system can partially counteract that fall.

Active cooling produced no statistically detectable advantage at moderate temperatures (60-70°C), however a significant one at the highest temperatures tested. At 70-80°C, the active-cooling condition delivered a mean power output 1.36 mW (roughly 3.9%) higher than passive cooling alone (Welch's t-test, $p = 0.034$). A corresponding voltage advantage of 0.26 V was observed in the same band. Within-trial autocorrelation means this voltage difference is reported descriptively rather than as a statistically confirmed effect. Therefore, the benefit appears where thermal stress is most severe. However, the fan's simple ON/OFF control could not bring temperature back below its activation threshold once heat was applied continuously, indicating that the system slowed overheating rather than reversing it.

At the single-cell scale tested, the fan's own power draw exceeds the power it recovers, so net energy gain is negative, and the case for practical viability rests on module-scale deployment. These results illustrate that the gap between a solar panel's rated performance and its output in hot climates does not require expensive infrastructure to begin closing. It is directly relevant to regions, such as South Asia and the Middle East, where low-cost polycrystalline silicon dominates.

Keywords: Polycrystalline Silicon; Solar Photovoltaic; Active Cooling; Thermal Regulation; Temperature Coefficient; 3D Printing; Microcontroller; Thermal Stress.

INTRODUCTION

The Popularity of Photovoltaic Technology (PV)

The last decade has confirmed that solar photovoltaic (PV) technology is no longer a supplement to the global energy mix: it is increasingly becoming the centre of it. Global installed solar PV capacity surpassed 2.25 TW by the end of 2024, up from 1.65 TW in 2023, with annual additions reaching a record 602 GW. This is more capacity in a single year than the entire global solar fleet had accumulated before 2019 (REN21, 2025). Investment highlights the same story: in 2023, solar PV attracted approximately USD 480 billion globally, exceeding investment in all other power generation technologies combined (IEA, 2024). These figures establish two conditions that frame this study. First, the installed base is large and compounding, which means that any proportional loss in real-world output translates into an absolute loss that grows year on year. Second, because capital has effectively settled on solar as the default new generation source, improvements in how much of a panel's rated capacity is actually delivered now carry more weight compared to further increases in deployment volume. A performance loss of 0.4-0.5% per degree Celsius above standard test conditions is negligible for a single panel. Applied across 2.25 TW of installed capacity, however, it becomes a systematic and largely unrecovered shortfall.

Connection Between Photovoltaic Technology and Crystalline Silicon

Photovoltaic (PV) technology is built on the ability of semiconductor materials to convert sunlight into electricity via the photovoltaic effect. Silicon has held the dominant position among these materials since the beginning of practical solar development. The first silicon p-n junction photocell efficient enough to be useful as a power source was demonstrated at Bell Laboratories in 1954, achieving approximately 6% conversion efficiency (Chapin, Fuller, & Pearson, 1954). Silicon cells provided the power for the first solar-powered satellite, Vanguard 1, in 1958. Two properties can explain this durability. First, silicon is the second most abundant element in the Earth's crust after oxygen, comprising roughly 28% by mass (Wedepohl, 1995), which makes large-scale manufacture economically feasible in a way that rarer semiconductors are not. Second, its bandgap of approximately 1.12 eV sits close to the optimum for the terrestrial solar spectrum, allowing it to absorb a large proportion of incident photons whilst still producing a usable output voltage (Green, 1982; Goswami, 2023). More than 70 years after that first cell, crystalline silicon still accounts for approximately 98% of the global PV market (REN21, 2025).

Crystalline silicon (c-Si) is available in two main forms:

- **Monocrystalline silicon (mono-Si):** grown from a single crystal lattice, offering higher efficiency but at greater manufacturing cost and energy input.
- **Polycrystalline silicon (poly-Si):** cast from multiple crystal grains, more cost-effective to produce, and the backbone of mass-market solar deployment.

The relationship between PV and crystalline silicon is one of co-evolution. As the solar industry scaled up, it was crystalline silicon, and polycrystalline silicon in particular, that made grid-scale and residential deployment economically viable across diverse markets. Module costs fell by roughly 99% between 1980

and 2012, driven primarily by economies of scale in manufacturing and by incremental gains in cell efficiency and materials use (Kavlak et al., 2018), reductions achieved within the crystalline silicon platform rather than by displacing it.

Understanding Polycrystalline Silicon Solar Cells

Polycrystalline silicon solar cells are produced by melting and casting silicon into ingots composed of multiple randomly oriented crystal grains. Unlike monocrystalline cells, which are grown from a single uniform lattice at greater energy cost, poly-Si cells are cast from metallurgical-grade feedstocks into multicrystalline blocks (Goswami, 2023). When photons strike the cell, they excite electrons in the silicon lattice above the conduction band threshold, generating electron-hole pairs that the built-in electric field of the p-n junction sweeps toward their respective contacts to produce a direct current (Goswami, 2023). The minimum photon energy required is set by silicon's bandgap of approximately 1.12 eV. Photons below this energy pass through unabsorbed, which is one intrinsic reason no silicon cell converts all incident sunlight (Green, 1982; Green et al., 2021). The theoretical conversion limit for single-junction silicon under AM1.5 illumination is around 29%. Yet, real-world poly-Si modules can achieve 14–18% under standard test conditions of 1,000 W/m² irradiance, AM1.5G spectrum, and a cell temperature of 25°C (Green et al., 2021; Goswami, 2023). The final figure of 25°C is where the problem lies.

Active Thermal Regulation

Despite this rapid growth, a persistent and underaddressed challenge is solar PV systems: heat. As module operating temperatures rise above the standard test condition of 25°C, efficiency degrades measurably and continuously. For polycrystalline silicon cells, the temperature coefficient of power typically falls in the range of −0.4% to −0.5% per degree Celsius (Skoplaki & Palyvos, 2009). Under field conditions, module surface temperatures routinely reach 50°C to 75°C, particularly in high-irradiance, low-wind environments, representing a consistent and largely unrecovered efficiency loss (Skoplaki & Palyvos, 2009; Dubey et al., 2013).

Active thermal regulation addresses this by introducing a controlled mechanism: a forced airflow or cooling that removes heat from the cell surface in real time. Unlike passive approaches, such as heat sinks or thermally conductive backsheets, active systems are capable of responding dynamically to changing thermal loads. Thus, maintaining cell temperature within a narrower and more performance-optimal range (Bahaidarah et al., 2013; Chandel & Agarwal, 2017).

3D-Printed Microcontroller-Integrated Cooling System

The practical implementation of active thermal regulation has been constrained by cost, complexity, and the difficulty of integrating sensing and actuation into compact form factors. Advances in additive manufacturing and embedded systems now offer a route past these barriers. A 3D-printed cooling system allows for the rapid prototyping of precise channel structures and airflow housings that can be tailored directly to the module geometry, without the tooling costs associated with conventional fabrication (Niknam et al., 2021).

When integrated with a microcontroller, such as an Arduino or equivalent platform, the system gains the capacity for closed-loop thermal management. Therefore, onboard temperature sensors feed real-time data to the controller, which modulates cooling output in proportion to the measured thermal load. This approach is well suited to experimental investigation of PV thermal performance under controlled conditions.

The Meaning of Constant Irradiance

To isolate the effect of temperature on solar cell performance, irradiance needs to be held constant throughout experimentation. Irradiance, measured in watts per square metre (W/m^2), describes the incident solar power per unit area striking the cell surface. It is one of the two primary drivers of PV output. The other one is temperature, and the two variables are not independent under natural outdoor conditions, where cloud cover and angle of incidence cause both to fluctuate simultaneously (Duffie & Beckman, 2013).

In a controlled experimental setting, fixing irradiance at the standard test condition value of $1,000 \text{ W/m}^2$ using an artificial light source eliminates this variability. Therefore, any observed changes in open-circuit voltage, current or maximum power output can be attributed to thermal effects rather than fluctuations in the incoming radiation. Thus, constant irradiance is a prerequisite for drawing valid conclusions about the thermal behaviour of the cell.

Identifying Thermal Stress

Thermal stress in photovoltaic cells refers to the condition in which elevated operating temperatures produce measurable and cumulative degradation in electrical performance and material integrity. At the electrical level, increasing temperature reduces the bandgap of silicon. This in turn lowers the open-circuit voltage (V_{oc}) while producing only a marginal increase in short-circuit current (I_{sc}), resulting in a net reduction in maximum power output (Green, 1982; Goswami, 2023). Silicon's bandgap of approximately 1.12 eV (Green, 1982) sets both the cell's photon absorption threshold and its temperature sensitivity. As temperature rises, the bandgap narrows, directly lowering the open-circuit voltage the cell can sustain (Goswami, 2023).

Exposure to repeated thermal cycles induces mechanical stress arising from different coefficients of thermal expansion at the interfaces between the materials that make up a module (Köntges et al., 2014). Silicon forms the active semiconductor layer, and the encapsulant is the polymer layer that binds and protects the cell. The backsheet is the rear-facing protective film that provides electrical insulation and moisture resistance, and the solder joints are the metallic connections linking the cell to its electrical contacts and interconnecting ribbons. Because each of these materials expands and contracts at a different rate under heating and cooling, the interfaces between them experience cumulative mechanical strain with each thermal cycle. Over time, this leads to micro-crack propagation and solder fatigue, both of which contribute to the long-term degradation of module output (Köntges et al., 2014).

Identifying thermal stress requires two parallel forms of measurement. The first is real-time electrical characterisation through current-voltage (I-V) curve analysis, which tracks how output power responds as

temperature rises (Dubey et al., 2013). The second is controlled temperature monitoring, which pinpoints when heat accumulation becomes significant (Skoplaki & Palyvos, 2009). These approaches allow the threshold at which thermal effects begin to meaningfully limit performance to be precisely identified.

Importance of the Study

The significance of this research operates at several levels. For the engineering research community, it provides an experimentally validated performance dataset for a novel integrated cooling configuration which combines factors that existing studies have treated separately. In terms of policy and deployment contexts, the findings are relevant to solar energy planners in high-irradiance regions, such as South Asia and the Middle East, including India, Pakistan and Bangladesh, where ambient temperatures routinely exceed 40°C and polycrystalline silicon dominates installed fleets due to cost constraints. In these regions, thermal degradation of performance is a persistent and poorly compensated factor in energy yield forecasting.

Multiplied across gigawatts of installed capacity, this represents a systematic overestimation of solar energy output, especially for governments and national energy planners. Addressing this gap has implications for the accuracy of national renewable energy targets and the long-term financial viability of solar investment in the Global South. For manufacturers and system integrators, a 3D-printable cooling architecture represents a low-tooling-cost path that does not depend on global supply chains for specialised heat exchange components, further lowering import costs.

Aims and Goals of the Paper

This paper investigates whether active thermal regulation delivered through a purpose-built 3D-printed and microcontroller-integrated cooling system can meaningfully recover the efficiency losses that polycrystalline silicon solar cells sustain under real-world thermal stress conditions. The central argument is that the gap between rated and actual PV performance in high-irradiance environments is not a limitation of the technology. It is a solvable engineering problem, and low-cost, accessible solutions are viable and necessary. This is particularly for markets where thermal degradation is highest and resources for commercial cooling infrastructure are most constrained.

The literature review sub-section will analyse the relevant literature on polycrystalline silicon PV performance degradation under thermal stress, existing active and passive cooling methodologies, microcontroller-based thermal management systems, and the application of 3D printing in solar-adjacent hardware. It will aim to establish the body of work upon which this study builds, and address the gaps in current literature.

In the Methods section, the study will describe the design rationale and fabrication process for the 3D-printed cooling assembly, the microcontroller logic and sensor integration, the test cell characterisation protocol, and the thermal stress conditions applied. In the Results section, the data will be presented. This includes descriptive statistics for each condition, the measured relationship between temperature and open-circuit voltage, band-matched comparisons of voltage and power output between the passively and actively cooled conditions with accompanying significance tests, and an accounting of

the cooling system's own power consumption. In the Discussion section, the findings will be interpreted in relation to prior literature, the limitations of the design are examined, and directions for subsequent work are proposed. The Conclusion will summarise what the study establishes and what remains unresolved.

Hypothesis

Based on the negative temperature coefficient of polycrystalline silicon and the theoretical advantage of forced convection over passive heat dissipation, this study hypothesised that active cooling would produce a statistically measurable improvement in voltage and power output relative to passive cooling alone, and that this advantage would increase with rising cell temperature as the passive heat sink's effectiveness diminished.

LITERATURE REVIEW

The Concept of Polycrystalline Silicon (Poly-Si)

Standard test conditions exist as a controlled reference and are rarely the case in the real world. Constant irradiance maintained at $1,000 \text{ W/m}^2$ throughout a cell's operational life is merely an assumption. In practice, a poly-Si module installed in an arid or tropical environment operates under sustained high irradiance, often above 800 W/m^2 (Duffie & Beckman, 2013), whilst simultaneously accumulating heat from absorbed radiation that cannot be converted into electricity. This is because polycrystalline silicon cells convert only a fraction of incident solar energy into electrical power, typically between 14% and 18% under ideal conditions, with the remainder absorbed as heat (Green et al., 2021; Goswami, 2023). Therefore, the gap between what the cell absorbs and what it converts becomes thermal energy trapped within the module. Cell temperatures can exceed 70°C in such conditions (Skoplaki & Palyvos, 2009; Dubey et al., 2013), and the relationship between rising temperature and falling performance becomes direct.

As cell temperature increases beyond 25°C , open-circuit voltage declines at approximately 2.3 mV per degree Celsius, and overall power output falls by roughly 0.4% to 0.5% per additional degree (Skoplaki & Palyvos, 2009; Hudişteanu et al., 2024). This relationship is not uniform across silicon cell types. Baghel and Chander (2022) found that under the same tropical irradiance conditions, poly-Si panels recorded a maximum performance ratio of 0.86 compared to 0.89 for monocrystalline panels in December, with thermal losses consistently more pronounced in poly-Si modules. This is attributable to the higher density of grain boundaries in polycrystalline silicon, which act as recombination sites and amplify carrier losses as thermal agitation increases. This effect is less present in the more structurally uniform monocrystalline lattice (Green, 1982). Over longer timescales, the cumulative consequence is considerable. Long-term field studies have reported median degradation rates of around 0.5% per year for crystalline silicon modules, with thermal stress among the primary degradation mechanisms (Jordan & Kurtz, 2013).

Thus, this means a polycrystalline silicon panel is made to operate under conditions its rated specifications never assumed. A 300 W panel installed in Mumbai, Riyadh, Lagos, Karachi, Jakarta

and/or Nairobi, where ambient temperatures regularly exceed 40°C and module surface temperatures can reach 65°C to 70°C, will produce 15% to 20% less power than its nameplate rating would suggest (Dubey et al., 2013).

Passive and Active Thermal Regulation

The core problem is that polycrystalline silicon modules deployed in high-irradiance environments operate persistently above the 25°C standard test condition (Dubey et al., 2013). Without intervention, this excess heat has nowhere to go. The module structure is largely sealed, designed for weather resistance and electrical insulation rather than thermal dissipation, so the packaging materials themselves substantially determine operating temperature (Pavgi et al., 2021). Absorbed radiation that cannot be converted into electricity therefore accumulates within the laminate as heat. As cell temperature increases, efficiency falls continuously and the losses compound across operational hours (Skoplaki & Palyvos, 2009). If left unaddressed, this thermal burden represents a permanent and growing gap between rated and actual energy output, and it has been identified across the literature as a foundational constraint on real-world PV performance (Skoplaki & Palyvos, 2009; Dubey et al., 2013; Chandel & Agarwal, 2017). Consequently two broad families of solutions have emerged: **passive and active thermal regulation**.

Passive approaches require no external power input because they rely entirely on inherent material properties or fixed physical configurations to dissipate heat without any mechanical or electronic actuation. For instance, heat sinks are thermally conductive structures, typically aluminium fins, bonded to the rear surface of a module. They work by increasing the effective surface area available for convective and radiative heat transfer. Thus, allowing excess thermal energy to dissipate into the surrounding air through natural convection (Chandrasekar & Senthilkumar, 2015). Thermally conductive backsheets function similarly. They replace the standard polymer layer of a module with a material of higher thermal conductivity. Therefore, enabling heat to conduct away from the cell more readily without any moving components. This is an approach that is demonstrated in field trials to reduce daily average module temperature by up to 3°C relative to conventional backsheets (Pavgi et al., 2021). Radiative coatings operate on a different principle: they are applied to module surfaces to increase emissivity, allowing the module to shed heat as longwave radiation directly to the sky, even under direct sunlight (Zhao et al., 2019).

Heat sink configurations can include their own limitations. First, their effectiveness is dependent on the ambient temperature and the availability of natural convection. These are least favourable in the high-temperature and low-wind environments where cooling is most needed (Chandrasekar & Senthilkumar, 2015). Second, heat sinks dissipate heat through the temperature difference between the module surface and the surrounding air. As ambient temperature rises, this gap narrows. A smaller gap means less driving force for heat to move from the module into the air. Thus, the rate at which the heat sink can shed thermal energy falls on a hot day, when the module needs cooling the most (Chandrasekar & Senthilkumar, 2015). Third, in densely packed rooftop arrays, warmed air rising from one module can be drawn across its neighbours. Thereby, further reducing the temperature gradient available for passive dissipation. Thermally conductive backsheets and radiative coatings share a further common constraint - neither can adapt to changing thermal conditions. Both are fixed material solutions whose performance is

determined at manufacture, offering no capacity to intensify heat removal as irradiance and cell temperature climb through peak operating hours (Zhao et al., 2019).

PCM-based systems absorb excess heat as latent energy during phase transition, buffering temperature without mechanical intervention. As the PCM melts from solid to liquid, it absorbs thermal energy at a near-constant temperature, acting as a thermal buffer between the cell and peak irradiance conditions; experimental work on docosane PCM in crystalline silicon modules has demonstrated measurable reductions in cell temperature through this mechanism (Amalu & Fabunmi, 2022). The constraint, however, is saturation. Once the material has fully transitioned to its liquid phase, its capacity to absorb further heat is exhausted, and no additional buffering is available during the sustained peak irradiance hours when thermal regulation is most critical (Chandel & Agarwal, 2017). At that point, cell temperature rises unchecked and the efficiency losses described above continue. This is because PCMs are finite thermal reservoirs: their buffering capacity is fixed by the mass and latent heat of the material used, so a PCM can only absorb heat up to the total energy its phase transition can accommodate. Unlike active systems, which remove heat from the module continuously, a saturated PCM behaves as a passive thermal mass with no further cooling function, offering no protection precisely during the longest and most thermally demanding portion of the day.

Active cooling approaches address the saturation and gradient problems directly by continuously removing heat through forced fluid circulation. Water-channel cooling circulates water through pipes or channels bonded to the rear surface of the module (Zondag, 2008). As cooler water flows through the channels, it absorbs heat conductively from the cell surface and carries it away from the module, with the warmed water subsequently recovered for thermal use elsewhere (Zondag, 2008). Furthermore, forced-air systems use dedicated fans to drive airflow through ducted channels behind the module. Thereby, enhancing convective heat removal beyond what natural convection can achieve alone, the fan speed and duct geometry determine the rate of heat extraction (Sarhaddi et al., 2010). Spray cooling applies fine water droplets directly to the module surface, removing heat through a combination of convective cooling and evaporative latent heat loss as the water film vaporises. This can achieve rapid surface temperature reductions with relatively low water volumes (Moharram et al., 2013).

The results of active cooling are substantially more robust than passive alternatives because heat removal is continuous and responsive rather than finite or gradient-dependent. Zanlorenzi et al. (2018) demonstrated that a water-based hybrid PV/T collector installed in the subtropical state of Paraná, Brazil, reduced average module temperature and improved electrical efficiency. The study reported a reduction in maximum effective cell temperature of 8.83°C and an improvement in mean electrical efficiency from 13.95% to 15.20%. Another study conducted in Saharan conditions reported efficiency gains from 17.58% to 21.92%, representing a maximum improvement of 24.68%. This was achieved through intermittent water flow activated at temperature set points (Chaich et al., 2025). Furthermore, Hamed et al. (2025) explained that an Arduino Uno microcontroller-driven ON/OFF water cooling circuit achieved temperature reductions of up to 31.5°C on a 40 W PV panel, with set-point thresholds determining cooling onset.

It is important to note that these results are encouraging. However, they come with limitations. Most active cooling studies test commercially fabricated and/or generic pipe configurations which are designed for laboratory convenience rather than scalability. Thus, leaving channel geometry, such as width, pitch and contact area as an underexplored variable relative to flow rate and coolant type. Microcontroller-based implementations similarly tend toward simple ON/OFF threshold logic, without feedback responsive to real-time irradiance variation (Kusuma et al., 2023; Raju et al., 2024).

Gaps in the Literature

The gap is specific and addressable. Prior work has established each component of this problem separately. Active water cooling has been demonstrated to improve module performance across climates, from subtropical Brazil (Zanlorenzi et al., 2018) to Saharan conditions (Chaich et al., 2025). Collector and channel configuration is known to influence heat extraction, with flow geometry treated as a design variable in PV/thermal systems (Zondag, 2008; Sarhaddi et al., 2010). Microcontroller feedback has been used to automate thermal regulation through temperature set-point control (Kusuma et al., 2023; Hamed et al., 2025; Raju et al., 2024). Additive manufacturing has been shown to enable complex heat-exchange geometries at low tooling cost (Niknam et al., 2021). What remains unproven is the integrated system: a 3D-printed cooling structure designed for contact coverage across the cell's rear surface, driven by microcontroller-based temperature control, applied to a polycrystalline silicon cell under controlled and reproducible thermal stress, and evaluated systematically against both an uncooled and a passively cooled reference.

This study explores the missing gap. It investigates whether active thermal regulation through a 3D-printed, microcontroller-integrated cooling system produces statistically significant improvements in the electrical performance of polycrystalline silicon solar cells subjected to thermal stress. The conceptual framework is based on the negative temperature coefficient effect in semiconductors. Therefore, the concept that the open-circuit voltage and efficiency of polysilicon solar cells reduce as the temperature increases due to certain established thermodynamic effects. The other component of the framework is heat transfer by forced convection.

The research aims to:

1. Design and fabricate a 3D-printed, microcontroller-integrated active cooling system for polycrystalline silicon solar cells, with fin geometry configured for thermal contact coverage across the cell's rear surface.
2. Characterise the electrical performance of a poly-Si cell under defined thermal stress conditions, with and without cooling. This incorporates open-circuit voltage, current and power output.
3. Quantify the temperature differential achieved by the active cooling system and establish the correlation between temperature reduction and efficiency recovery.
4. Evaluate the net energy gain accounting for parasitic power consumed by the cooling system in order to assess practical viability.

METHODS

Study Design

This study employed a controlled and comparative experimental design to evaluate whether active thermal regulation through a fan-assisted, microcontroller-integrated heat sink, improved the electrical performance of a polycrystalline silicon solar cell relative to uncooled and passively cooled states.

Three conditions were tested: a cell with no heat sink attached, the same cell type fitted with a 3D-printed heat sink with the fan switched off, and the heat sink-equipped cell with the fan active under microcontroller threshold control. This design was selected because it isolates the contribution of each cooling mechanism in sequence - from no intervention to passive and to active forced convection. Therefore, allowing the incremental benefit of each stage to be attributed specifically to the cooling configuration. All testing was conducted in a single laboratory setting over a continuous data collection period in June 2026. All procedures were conducted in accordance with standard laboratory safety practice for low-voltage electronic apparatus.

Apparatus

The experimental apparatus consisted of a polycrystalline silicon solar cell. It was selected as the test subject because poly-Si remains the dominant commercial PV technology and the cell type whose thermal degradation behaviour this study sought to characterise. A 3D-printed pin-fin heat sink fabricated in Polyethylene Terephthalate Glycol (PETG) filament was used as additive manufacturing allowed the fin geometry to be produced quickly and at low cost without the tooling investment required by conventional metal fabrication. PETG was chosen over more common filaments, such as PLA, for its higher heat deflection temperature under sustained thermal load. Thermal paste was used to bond the heat sink to the rear surface of the cell because it minimises air gaps at the interface and reduces contact thermal resistance. Thereby, ensuring that heat generated at the cell was conducted into the heat sink efficiently rather than lost at the bonding layer.

A DS18B20 digital temperature sensor was affixed to the cell surface to provide continuous temperature readings. This sensor was selected for its digital output, which avoids the calibration drift associated with analogue thermistors, and for its ease of integration with the Arduino platform via a single-wire communication protocol. The DS18B20 has a manufacturer-stated accuracy of $\pm 0.5^{\circ}\text{C}$ over -10°C to $+85^{\circ}\text{C}$ and a 12-bit resolution of 0.0625°C (Maxim Integrated, 2019). This is adequate for the 10°C analysis bands that are used in this study. An INA219 sensor module was connected in series with the cell output to record voltage and current because it provides simultaneous, high-resolution measurement of both parameters over a single I2C interface. Therefore, reducing wiring complexity relative to separate voltage and current sensing circuits.

The INA219 employs a 12-bit ADC with a bus-voltage resolution of 4 mV and a maximum gain error of $\pm 1\%$ (Texas Instruments, 2015). Under the standard 32 V / 2 A calibration used in this study's firmware, logged current resolves in 0.1 mA steps and computed power in 2 mW steps, which is confirmed by the quantisation visible in the raw logs. The implications of these resolution limits for the effect sizes reported

are addressed directly in the Results. An Arduino Uno microcontroller processed sensor input and controlled a 5 V DC fan via a Metal-Oxide-Semiconductor Field-Effect Transistor (MOSFET) switching circuit. This activated forced-air cooling when the measured surface temperature reached a predefined threshold of 45 degrees Celsius and deactivated it once temperature fell below this value.

The Arduino Uno was selected for its low cost yet sufficient processing capacity for the threshold logic required. Whilst the MOSFET switching stage was necessary because the fan's current draw exceeds what the microcontroller's output pins can safely supply directly. The 45°C threshold was chosen because it sits comfortably above both standard room temperature and the 25°C reference condition, while remaining below the range at which material degradation, including solder fatigue, becomes a significant concern.

Thermal stress was applied by using a heat gun held at a fixed distance from the cell surface. It was selected over solar or halogen lamp exposure because it provided a stable and repeatable thermal input independent of ambient irradiance fluctuation. Thereby, allowing the effect of temperature on electrical output to be examined in isolation. A lux meter was used to confirm that ambient light levels at the cell surface remained stable across conditions. This ruled out incidental irradiance variation as a contributing factor to differences in electrical output. This step was included specifically to safeguard the assumption of constant irradiance established earlier as a prerequisite for valid thermal analysis.

All sensor output, including temperature, voltage, current and fan state, was logged at a fixed interval of one second and transmitted via serial connection to a laptop running a Python data logging script. A one-second logging interval was selected for two reasons. First, it was fast enough to capture the brief electrical changes that occurred right as the fan switched on or off. Second, it was slow enough that the resulting dataset stayed a manageable size given how short each trial was. The complete apparatus is shown in Figure 1.




Figure 1. Experimental Setup Showing the Cell, Heat Sink, Sensors, Arduino/MOSFET/Fan Circuit, Heat Gun, and Lux Meter, with Serial Connection to the Logging Laptop.

Procedure

For each of the three conditions, the cell was allowed a cooling-down interval between trials to limit heat carry-over. Data logging commenced once the heat gun was applied, so the first logged readings reflect an already-elevated surface temperature (38.9°C, 45.8°C, and 45.4°C for the three retained trials respectively) rather than the ~25°C room-temperature reference. Therefore, baseline open-circuit voltage and power values quoted for each condition refer to the first logged readings of each trial, and the differing starting temperatures across trials are considered explicitly in the Limitations.

The three conditions were tested in a fixed order, no heat sink first, then heat sink with fan off, then heat sink with fan on, rather than a randomised order. Randomisation was not practicable because a single physical cell was used throughout to eliminate cell-to-cell manufacturing variation. Additionally, attaching and detaching the thermally bonded heat sink between interleaved trials would have risked damaging the cell's rear surface and altering the thermal interface between conditions. The consequence of this fixed sequential design, the possibility of cumulative thermal history affecting later trials is examined in the Limitations. Next, the heat gun was applied at a constant distance and held in place for the duration of the trial while the DS18B20 and INA219 sensors logged surface temperature, voltage and current at one-second intervals. Figure 2 summarises the procedure and Figure 3 illustrates the three test

conditions.

In the no heat sink condition, the bare cell was tested against thermal stress, establishing a baseline of thermal and electrical response. In the heat sink, fan off condition, the same heat gun was applied to the cell fitted with the 3D-printed heat sink, with the fan disabled throughout. Thus, isolating the passive method of the heatsink's fin geometry alone. In the heat sink, fan on condition, the Arduino-controlled fan was enabled. Thereby, activating automatically once the 45 degree Celsius threshold was reached and deactivating once temperature fell below it. This was done to assess the additional benefit of active forced-air cooling. Each condition was run three times to verify that the apparatus behaved reproducibly. However, complete serial logs were retained for only one continuous trial per condition, and the repeat runs were used solely as qualitative checks of protocol stability. All quantitative analysis in this paper is based on a single logged trial per condition, with logged datasets of 206 readings, approximately 3.43 minutes for the no heat sink condition; 122 readings, approximately 2.03 minutes, for the heat sink, fan off condition; and 208 readings, approximately 3.47 minutes, for the heat sink, fan on condition. Prior to data collection, the Arduino firmware was reviewed and tested to confirm that relay switching, sensor reading and serial logging functioned correctly, and no timing errors or sensor dropouts were present in the output stream.




Figure 2. Schematic of the Experimental Procedure and the Three Test Conditions (No Heat Sink; Heat Sink, Fan Off; Heat Sink, Fan On at 45°C).

The left panel shows the physical arrangement of cell, sensors, and heat gun. The centre panel gives the procedural sequence common to all three trials. The right panel shows the three cooling configurations, with the fan-on condition switching at the 45°C threshold.




Figure 3. The Three Testing Conditions Applied to the Polycrystalline Silicon Cell: (1) No Heat Sink, (2) Heat Sink with Fan Off, (3) Heat Sink with Fan On Above the 45°C Threshold.

Outcome Measures

The primary outcome measure was electrical power output. It was calculated as the product of voltage and current at each logged interval, and evaluated against the corresponding surface temperature to characterise the relationship between thermal load and power degradation under each condition. Secondary outcome measures included voltage decline as a function of rising temperature and the rate of temperature increase over time. These indicated the thermal responsiveness of each cooling configuration. The measures were selected because they directly reflect the temperature-dependent behaviour which is described by the negative temperature coefficient of silicon photovoltaic cells. In addition, they permit a like-for-like comparison of cooling effectiveness across the three conditions by using identical instrumentation and logging protocol.

Fill Factor (FF) and Why It Could Not Be Measured

One standard photovoltaic performance metric is deliberately absent from the outcome measures above, and the reason is methodological rather than incidental. The fill factor measures how "square" the I-V curve is:

- **Short-circuit current** (I_{sc}): the maximum current the cell produces, measured when the terminals are directly connected and voltage is zero.
- **Open-circuit voltage** (V_{oc}), the maximum voltage the cell produces, measured when no current is flowing.
- **Maximum power point** (V_{mp} , I_{mp}), the specific combination of current and voltage on the curve where their product (power = $V \times I$) is greatest.

$$FF = (V_{mp} \times I_{mp}) / (V_{oc} \times I_{sc})$$

Geometrically, it is the ratio of the area under the real I-V curve to the area of the rectangle formed by V_{oc} and I_{sc} . A perfect cell would have $FF = 1$, though no real cell reaches this because of internal resistance and recombination losses.

FF is a useful variable to track because it isolates how *degraded* the shape of the I-V curve is, separately from the cell's absolute voltage or current values. This means it can show internal losses even when temperature shifts V_{oc} up or down. However, calculating FF requires a full I-V sweep at each fixed temperature point, capturing current across the entire range from short-circuit to open-circuit. This study's logging setup recorded continuous voltage and current readings at a single operating point per second while temperature changed in the background rather than performing discrete sweeps, so FF could not be calculated from the resulting data. This is addressed further in the Discussion and listed as a direction for future work. For commercial polycrystalline silicon under standard test conditions (25°C, 1000 W/m²), FF typically falls between 0.70 and 0.78.

Data Processing and Analysis

Logged voltage, current, power and temperature values were exported from the serial output and processed in Python. They were selected for consistency with the existing data logging pipeline and its established libraries for analysis, regression and visualisation. For each condition, descriptive statistics, including mean and standard deviation, were calculated for power output and surface temperature across the logged duration.

Power output was plotted against temperature and against elapsed time for each condition to visualise the rate and extent of performance degradation. Where appropriate, the relationship between temperature and power was assessed by using linear regression to quantify the temperature coefficient of power observed

under each condition. Thus, enabling direct comparison with coefficients reported in the existing literature. To test whether differences between the fan-off and cooling-active fan-on conditions were statistically distinguishable from measurement noise, Welch's two-sample t-tests ($\alpha = 0.05$) were applied to voltage and power readings within matched 10°C temperature bands with 95% confidence intervals reported for each difference. As consecutive one-second readings within a trial are not independent, a slowly drifting quantity produces serially correlated readings, lag-1 autocorrelation was computed for each series to assess the validity of these tests.

Band-level power readings showed negligible autocorrelation (lag-1 r between -0.27 and $+0.12$), so t-tests on power are treated as approximately valid. Voltage readings were strongly autocorrelated (lag-1 $r \geq 0.93$), so voltage differences are reported descriptively with the caveat that their nominal p-values overstate significance. This distinction directly limits which of the study's comparisons can be claimed as statistically established. All raw data files were retained in their original logged format to preserve an unmodified record for verification.

RESULTS

The three experimental conditions generated 206, 122 and 208 logged readings respectively. Table 1 presents descriptive statistics across all conditions. To isolate the forced-convection effect cleanly, only the 81 fan-on readings where surface temperature met or exceeded the 45°C threshold are used, since these are the readings where the fan was actually engaged in active cooling. This 81-reading subset, with a mean surface temperature of 68.7°C, is compared against the heat sink, fan-off condition ($n = 122$, mean temperature 68.6°C). This is because the two groups sit at near-identical mean temperatures and allow a like-for-like assessment of the marginal contribution of forced-air cooling.

Condition	n	Baseline Voc (V)	Baseline Temp (°C)	Mean Temp (°C)	Peak Temp (°C)	Mean Voc (V)	SD Voc (V)	Mean Power (mW)
No heat sink	206	20.54	38.9	56.4	76.8	19.14	0.74	34.9
Heat sink, fan OFF	122	19.86	45.8	68.6	84.3	19.34	0.34	35.9
Heat sink, fan ON (cooling-active only)	81	16.17*	45.4	68.7	82.8	19.33	0.57	35.8

Table 1. Descriptive statistics across experimental conditions, listed in the order tested. Baseline values are the first logged reading of each trial. *The fan-on baseline is transient-depressed by the fan switch-on event and does not represent a settled operating voltage; see Limitations.

Voltage- Temperature Relationship and Temperature Coefficient

The heat sink fan-off trial gives the clearest picture of how voltage responds to temperature. This is because the heat gun pushes temperature up steadily from 45.77°C to 84.31°C, with no interruptions. Running a linear regression of voltage against temperature on this 122-reading dataset gives a slope of -0.0322 ± 0.0010 V/°C (standard error; 95% CI ± 0.0019), with an R^2 of 0.900. That R^2 means temperature explains 90% of the change in voltage in this trial. Relative to the mean voltage of 19.34 V, this slope works out to a temperature coefficient of $-0.167\%/^{\circ}\text{C}$. One caveat applies to this regression: it is fitted to a single continuous time series in which consecutive readings are autocorrelated, so the 122 readings do not represent 122 independent observations. Therefore, the R^2 of 0.900 describes how tightly voltage tracked temperature within this one trial, and should not be read as significance evidence from independent samples. Across the full fan-off trial, voltage dropped from 19.86 V to 18.81 V, a fall of 1.05 V over a 28.4°C rise in temperature.

The no-heat-sink trial demonstrated a larger total voltage drop of 2.76 V, falling from 20.54 V to 17.78 V as temperature rose from 38.91°C to a peak of 76.81°C. However, the linear fit for this trial, returns an R^2 of only 0.047, far below the 0.900 obtained for the fan-off trial. Unlike the fan-off trial, temperature in this condition did not rise monotonically: it climbed, fell back to 65.32°C as the heat gun was withdrawn, and rose again, while voltage continued to decline briefly before recovering.




Figure 4. Open-Circuit Voltage (V) versus Cell Surface Temperature (°C), Heat Sink with Fan Off (n = 122).

Each point represents one second of logged data. The fitted regression gives $dV/dT = -0.0322 \text{ V/}^\circ\text{C}$ with $R^2 = 0.900$.




Figure 5. Open-Circuit Voltage (V) over Time (s) for All Three Conditions.

Voltage in the no-heat-sink condition declines most steeply and most noisily, tracking that trial's uneven thermal profile. Irradiance was held constant and verified by lux meter throughout. The fan-on trace is dominated early by the fan's startup transient, and only its cooling-active portion is used in the quantitative comparisons.

Figure 4 shows this relationship point by point. Voltage falls from just over 20 V at 45°C to approximately 18.7 V at 84°C, and the points cluster tightly around the trend line, particularly between 45°C and 65°C, where scatter is minimal. The decline steepens slightly above 70°C, and the lowest-temperature and highest-temperature regions show visible clustering where the cell dwelled longest. Because irradiance was held constant and the fan remained off throughout, temperature is the only variable that changed during this trial, making it the clearest single demonstration in this study of the negative temperature coefficient in operation.

Figure 5 places all three trials on a common timeline and shows how differently each behaved. The fan-off trace (blue) declines smoothly and almost linearly with very little noise. The no-heat-sink trace (orange) begins highest, at around 20.5 V, but oscillates sharply throughout, rising and falling several times before dropping to 17.8 V at the end; this reflects the uneven thermal profile of that trial rather than measurement error, since the heat gun was repositioned and withdrawn during the run. The fan-on trace

(green) begins lowest and climbs for most of the trial, the opposite direction to the other two. This is not the cell improving as it heats: it is recovery from the switch-on transient visible as the sharp dip to 12.1 V at around 15 seconds. By approximately 150 seconds all three traces converge near 19 V.

Active Cooling Performance

To assess the marginal benefit of active forced-air cooling over the passive heat-sink baseline, the dataset was split into 10°C temperature bands. This controls for temperature when comparing the two conditions. It matters because the fan-off and fan-on trials each follow a different temperature path over time, so comparing them band by band is the only fair way to isolate the fan's effect. Results are presented in Table 2, with the corresponding significance tests in Table 3 and the comparison plotted in Figure 6.

Temp Band	Fan OFF Voc (V)	Fan ON Voc (V)	Δ Voc (V)	Fan OFF Power (mW)	Fan ON Power (mW)	Δ Power (mW)
50-60°C	19.785	19.139	-0.646	39.57	35.33	-4.24
60-70°C	19.538	19.623	+0.085	36.46	37.11	+0.65
70-80°C	19.066	19.328	+0.262	34.48	35.84	+1.36

Table 2. Mean voltage and power by temperature band: fan-off versus cooling-active fan-on periods

Band	Δ Power (mW)	Welch's t	p	95% CI (mW)
60-70°C	+0.65	t(28) = 0.59	0.56	-1.60 to +2.90
70-80°C	+1.36	t(79) = 2.16	0.034	+0.11 to +2.62

Table 3. Welch's *t*-test results for the power difference between fan-off and cooling-active fan-on readings, by temperature band.

Tables 2 and 3 demonstrate a consistent directional pattern with only one statistically supported result. Both voltage and power differences move from negative at 50–60°C, through marginally positive at 60–70°C, to their largest positive values at 70–80°C. The voltage and power columns agree in sign in every band. Only the 70–80°C power difference is statistically significant. However, its confidence interval (+0.11 to +2.62 mW) excludes zero. Whereas the 60–70°C interval (–1.60 to +2.90 mW) spans it, meaning a difference of zero or a difference in the opposite direction, remains consistent with those data. The negative values in the 50–60°C row reflect the post-activation transient rather than a cooling penalty.




Figure 6. Mean Voltage (V) and Power (mW) by Temperature Band, Fan Off versus Cooling-Active Fan On, with 95% Confidence Intervals.

The 70 - 80°C power difference is statistically significant ($p = 0.034$); the 60–70°C differences are not ($p = 0.56$). The 50 - 60°C readings in the fan-on group are an anomaly: these are the first few seconds right after the fan switches on, while the cell is still cooling down from a hotter state earlier in the trial. Voltage had not stabilised at this point, so these readings are not directly comparable with the higher bands.

At 70 - 80°C, where the cell is under the most thermal stress, the active-cooling condition delivered a mean power of 35.84 mW against 34.48 mW for the passive-only condition, a gain of 1.36 mW or approximately 3.9%. Welch's t-test confirms this difference is statistically significant ($t(79) = 2.16$, $p = 0.034$, 95% CI +0.11 to +2.62 mW), and because band-level power readings showed negligible serial correlation, this test can be taken at face value. The accompanying voltage difference of 0.26 V (19.33 V against 19.07 V) points in the same direction and is roughly 65 times the INA219's 4 mV voltage resolution, so it is not a quantisation artefact. However, because the voltage series is strongly autocorrelated within each trial, its nominal significance cannot be relied upon and it is reported as descriptive support rather than independent statistical evidence. By contrast, at 60–70°C, neither the voltage difference (+0.085 V) nor the power difference (+0.65 mW, $p = 0.56$) is statistically distinguishable from measurement noise. Within this band, active cooling produced no detectable benefit over the passive heat sink alone.

Additionally, these effect sizes must be assessed against instrument resolution. The 1.36 mW power difference is smaller than the INA219's single-reading power resolution of 2 mW. It is detectable only because it is a difference of means across 38 and 54 readings whose underlying values span many quantisation steps, and the confidence interval reported above already accounts for the resulting reading-to-reading scatter. The 0.65 mW difference at 60–70°C falls well inside both the quantisation

scale and its own confidence interval, which is why it is reported as having no detectable effect. Voltage differences are on considerably safer ground relative to the 4 mV bus-voltage resolution, and the INA219's gain error of up to $\pm 1\%$ is systematic rather than random. Therefore, shifting all readings from the same device together and largely cancelling in within-device comparisons such as these. The DS18B20's $\pm 0.5^\circ\text{C}$ accuracy is small relative to the 10°C analysis bands used.

Threshold Control Behaviour

Of the 208 readings in the fan-on trial, 81 were recorded while the cell was at or above the 45°C activation threshold with the fan engaged, and 127 while it was below. Thus, the Arduino's threshold logic switched the fan on and off as designed. Mean surface temperature during the cooling-active periods was nonetheless 68.7°C , well above the 45°C trigger point, and temperature continued to climb during these periods from 45.22°C to a peak of 82.77°C with the fan running. The implications of this overshoot for control design are considered in the Discussion.

Net Energy Accounting

Aim 4 asked whether the cooling system produces a net energy gain once its own power consumption is subtracted. At the scale tested here, it does not. The 5 V DC fan used is rated at 0.15 A, corresponding to approximately 0.75 W of draw while running, which is several hundred times the 1.36 mW of additional power the cooling recovered in the $70\text{--}80^\circ\text{C}$ band. On a single test cell producing approximately 36 mW, active cooling represents a substantial net energy loss, and it would be misleading to present the system as energy-positive at this scale.

DISCUSSION

This study set out to determine whether low-cost active thermal regulation produces a measurable recovery of the electrical performance that polycrystalline silicon loses under thermal stress. The results support that proposition in part: a statistically significant power advantage was observed in the highest temperature band tested, no detectable advantage was found at moderate temperatures, and the system's parasitic power draw exceeds its recovery at single-cell scale. The following sections interpret these findings in relation to the established literature, examine what the control architecture reveals about the design, and consider what the results do and do not permit to be claimed.

Interpreting the Temperature Coefficient

The measured coefficient of $-0.167\%/^\circ\text{C}$ is in the same direction as published values for poly-Si, where Skoplaki and Palyvos (2009) reported $-0.30\%/^\circ\text{C}$ to $-0.45\%/^\circ\text{C}$, for maximum power rather than voltage alone. The difference is expected - voltage falls less steeply than power because current rises slightly with temperature, offsetting part of the voltage loss when power is considered instead. In addition, the coefficient is within the range reported by Hudişteanu et al. (2024) and is consistent with the underlying physics described by Green (1982), where rising temperature narrows the silicon bandgap and thereby lowers attainable voltage.

The contrast in fit quality between the two trials is also instructive. The low R^2 of 0.047 in the no-heat-sink condition does not indicate an absence of thermal dependence, nor reflect irradiance variation, as light levels were held constant and verified by lux meter throughout. Instead, it arises from the non-monotonic temperature profile of that trial. For part of the recording, voltage and temperature moved in opposite directions as the cell's electrical response lagged behind the withdrawal of the heat source. Linear regression assumes a consistent directional relationship, and a lagged response violates that assumption regardless of how strong the underlying physical dependence is. Therefore, the fan-off trial, whose temperature rose smoothly and without interruption, provides the cleaner estimate of the coefficient. This is why it forms the basis of the value reported above.

Comparison with Prior Cooling Studies

The pattern observed here, where forced convection contributes measurably only at the highest temperatures, is consistent with the mechanism described by Chandrasekar and Senthilkumar (2015): passive fins lose effectiveness as the temperature gap between the cell surface and the surrounding air narrows. Thus, leaving the largest recoverable margin for active cooling precisely where the cell is hottest. The temperature dependence is directionally consistent with Zanlorenzi et al. (2018), who reported larger cooling gains at higher module temperatures. However, in the present study, the moderate-temperature difference was too small to distinguish from measurement noise, so no gain can be claimed in that band. The 3D-printed PETG heat sink demonstrates that additive manufacturing can produce a functional cooling structure. This is without the tooling costs associated with metal fabrication. Thereby, extending the case made by Niknam et al., (2021) by adding integrated sensing and automatic fan control.

Moreover, these results point in the same direction as Hamed et al. (2025), who recorded temperature drops of up to 31.5°C and corresponding gains in power by using an Arduino-controlled ON/OFF water-cooling setup. The temperature reductions achieved in this study are smaller by comparison, as the fan here worked against a heat gun directed continuously at the cell rather than against natural solar heating that builds and eases over the course of a day. Under real outdoor conditions, where the cell would be exposed to peak heat for only part of the day, the same fan-based system would likely hold the cell below the 45°C threshold for longer, accumulating a larger overall efficiency gain across a full day of operation.

Control Architecture and Thermal Overshoot

The cell that continues to heat during cooling-active periods is a direct consequence of ON/OFF control, which responds only to whether a fixed threshold has been crossed and not to the rate at which heat is accumulating. This is a recognised limitation of microcontroller-based PV cooling implementations. Kusuma et al., (2023) and Raju et al., (2024) noted that simple threshold logic of this kind tends to overshoot, and recommended Proportional, Integral, Derivative (PID) or Pulse Width Modulation (PWM) based fan control where the fan speed adjusts continuously rather than switching fully on or off. The present data supports that assessment. During the cooling-active period, temperature climbed from 45.22°C to a peak of 82.77°C with the fan running, indicating that a single 5 V fan cannot counteract the

heat gun's continuous thermal input. A configuration using multiple fans or a single fan under proportional speed control would be expected to hold cell temperature considerably closer to the 45°C target.

Practical Viability at Scale

The system is net energy-negative at single-cell scale is expected rather than disqualifying. This is because the fan's power draw is essentially fixed while the recoverable power scales with the area being cooled. A deployed 300 W poly-Si module recovering the same 3.9% at high temperature would regain approximately 11.7 W against a fan load on the order of 1 to 2 W. Thereby, implying a plausible net gain exceeding 9 W per module. This extrapolation carries significant caveats. It assumes that the percentage recovery observed on a single cell transfers to full-module scale, that fan duty cycle under real diurnal heating resembles the continuous-heating case tested here, and that a small number of fans can serve an entire module's rear surface. None of these assumptions is verified by the present data. What the present data does establish is the correct framing for Aim 4: net energy viability is a question of scale. The single-cell laboratory answer is negative, and the module-scale answer is plausible but unproven.

Limitations

A key limitation of this study concerns replication. Although each condition was run three times, complete serial logs were retained for only one trial per condition. Thus, all quantitative results rest on single continuous recordings. Consecutive readings within a trial are serially correlated, which means the large reading counts ($n = 122$ to 208) overstate the effective quantity of independent evidence, a form of pseudoreplication. The statistical approach mitigates this where possible: band-level power readings showed negligible autocorrelation. Therefore, it is treated as approximately independent, whilst voltage results are explicitly downgraded to descriptive status. Nonetheless, this study demonstrates its effects within single trials of a single cell, and cross-trial replication with fully retained logs is required before the effect sizes reported here can be generalised.

Relatedly, as shown in Table 1, the first logged open-circuit voltages differed across conditions: 20.54 V for the no-heat-sink trial, 19.86 V for the heat-sink fan-off trial, and a transient-depressed 16.17 V at the start of the fan-on trial. Part of this gap is explained by differing starting temperatures, since logging began at 38.9°C for the no-heat-sink trial but at approximately 45.8°C and 45.4°C for the two heat-sink trials. Adjusting the fan-off baseline to 38.9°C by using the measured coefficient of -0.0322 V/°C accounts for roughly 0.22 V of the 0.68 V difference. The remaining 0.46 V is not explained by temperature and is consistent with cumulative thermal history affecting the cell across the fixed trial sequence, with small changes in electrical contact or thermal-interface conditions introduced when the heat sink was bonded, or with a combination of both.

As the same cell was tested sequentially in a fixed order, drift of this kind cannot be separated from condition effects within the present design. This affects absolute cross-condition comparisons, such as those in Table 1, more than it affects the fan-on against fan-off comparison in Table 2. This compares two adjacently run heat-sink trials beginning from nearly identical temperatures and matched within

temperature bands. Future work should record a standardised room-temperature baseline before every trial and randomise or counterbalance condition order where the apparatus permits.

A further limitation is the use of a heat gun as the thermal stress source rather than a spectrally calibrated lamp at $1,000 \text{ W/m}^2$. While the heat gun allows precise, repeatable thermal load application and eliminates irradiance variation, it applies heat primarily to the rear surface of the cell rather than replicating the photon-induced heating that occurs in real solar operation. In a real solar cell exposed to sunlight, heat does not only come from the outside, it is generated internally. This is because electrons get excited and recombine inside the cell during normal operation. The heat gun that is used in this study only replicates the external heating. It cannot reproduce that internal electrical heating, since that only happens when the cell is actually being illuminated. Due to this, the temperature coefficient and power values measured here can describe how the cell responds to outside heat alone, and not the combined effect of heat and light working together. This is what real solar cells experience. This means that the numbers in this study are not perfectly comparable to the official temperature coefficients that are listed on manufacturer datasheets, since those are measured by heating the cell with actual simulated sunlight (AM1.5G), and not an external heat source.

Trial duration is a related constraint. Beyond the replication issue described above, each logged trial lasted under four minutes. To properly measure cumulative energy gains over a real day and check whether the power saved by cooling actually outweighs the power that the fan itself uses, would need to be tested over several hours under a simulated full-day heat and light profile.

A final limitation concerns a water-based cooling setup that was attempted but did not produce usable results. The plan was to dunk only the 3D-printed heat sink in a water bath, whilst keeping the cell, sensors and electronics dry. This is because water conducts heat away faster than air and should cool the cell more than the fan does. The heat sink could not be sealed off from the rest of the setup and the water leaked into the electronics, breaking the sensor readings and the Arduino's logging before any usable data could be collected. Consequently, the water setup was dropped in favour of the fan-based system that was used in the main results. This means the study cannot directly compare air cooling to water cooling for this specific heat sink design, even though other studies, such as those of Zondag (2008) and Zanlorenzi et al. (2018), have reported that water cooling achieves larger temperature reductions than forced air.

Future versions of this setup should use a properly sealed enclosure that keeps the water bath fully separate from all electronics, such as a watertight thermal interface or an external water channel running outside the electrical components. This is to test whether water cooling's extra benefit can be captured safely.

Implications and Future Directions

The findings of this study matter for real-world solar deployment in places that are extremely sunny and hot. If the 3.9% power gain observed at $70\text{--}80^\circ\text{C}$ holds up over full days of operation, a temperature range routinely reached by deployed arrays in India, Pakistan and the Gulf, then the energy recovered

across gigawatt-scale installations would be meaningful (Dubey et al., 2013). The 3D-printed PETG design used here can keep cost and complexity low enough that it could realistically be built into affordable solar panel designs. This is without needing the specialised supply chains that proper heat-exchange hardware usually requires.

Future work should tackle three gaps that this study points to. First, swapping the simple ON/OFF fan control for PWM speed control that is driven by a PID loop. This would reduce the temperature overshoot that has been observed in the research. Thereby, allowing researchers to properly test whether temperature can be held steady rather than just tracked as it rises. Second, testing under a calibrated solar simulator producing actual AM1.5G-equivalent light rather than a heat gun would produce temperature coefficients that can be directly compared to published numbers and manufacturer datasheets. Third, measuring fill factor, which requires full I-V curve sweeps at fixed temperature points rather than continuous logging, would capture the kind of curve-shape degradation that this study's voltage and current logging alone cannot pick up.

CONCLUSION

This study tested whether a 3D-printed, microcontroller-controlled cooling system could meaningfully improve the electrical performance of a polycrystalline silicon solar cell under thermal stress, whilst keeping irradiance constant throughout. Three conditions were compared: no cooling, a passive 3D-printed heat sink alone, and the same heat sink with an Arduino-controlled fan engaging automatically above 45°C.

Summary of Findings

The evidence has demonstrated that active cooling does improve performance. However, the improvement is modest and depends heavily on temperature. At 70–80°C, the hottest band tested, the active cooling condition delivered a statistically significant power gain of roughly 3.9% over the passive heat sink alone (35.84 mW vs 34.48 mW; $p = 0.034$), with a directionally consistent voltage advantage (19.33 V vs 19.07 V). At 60–70°C, there was no statistically detectable difference that was found between active and passive cooling ($p = 0.56$ for power). Below that, in the 50–60°C band, the comparison is not meaningful, as those readings capture the seconds immediately after fan switch-on rather than a settled state.

Therefore, this pattern supports the study's hypothesis only partially. Active cooling does recover some of the efficiency lost to heat, and the size of that recovery grows as the cell gets hotter. This is exactly as the underlying physics would predict. Hotter cells have more heat to remove, so a fixed-output fan has more to work with. However, the recovery achieved here is small in absolute terms, and the fan never actually brought the cell's temperature back down below the 45°C threshold once the heat gun was running continuously. The system reduced the rate at which the cell overheated more than it reversed the

overheating itself.

The bigger picture is that the gap between a solar panel's rated performance and what it actually delivers in hot climates is not fixed: it can be narrowed, even with cheap and accessible parts. A \$20 fan, a 3D-printed heat sink and a basic microcontroller produced a measurable recovery in power output. That recovery came at a parasitic cost far exceeding it: at single-cell scale, the net energy balance of the tested apparatus is negative. However, the fixed nature of the fan's draw against the scalable nature of the recovery means the same architecture plausibly turns net-positive at module scale. This is a claim that this study framed but does not prove. It matters most in places, such as South Asia, the Middle East and other high-irradiance regions, where polycrystalline silicon dominates installed capacity specifically because it is cheap, and where the same cost pressure that makes poly-Si attractive also makes expensive cooling infrastructure unlikely to get built. This study has pinpointed that a meaningful first step toward closing that gap does not require expensive equipment.

Implications

For manufacturers and system integrators, the results indicate that low-cost active cooling is a viable design direction worth taking further. For national energy planners, particularly in countries with large poly-Si fleets and high ambient temperatures, this points to a source of recoverable energy yield that current performance models likely underestimate. This is because most yield forecasts are built on rated capacity rather than real-world thermal losses.

Limitations and Future Work

The principal constraints on the findings are the use of a heat gun rather than a calibrated solar simulator, the reliance on a single retained trial per condition, and trial durations under four minutes. Future work should address these by testing under AM1.5G-equivalent illumination, retaining logs from all repeat trials, running multi-hour trials to capture full-day energy yield, replacing ON/OFF control with proportional PWM control, and rebuilding the water-cooling variant with proper electrical isolation. Additionally, fill factor measurement, which the continuous-logging setup could not capture, remains a clear next step.

Even a small, low-cost intervention measurably changed how this cell handled heat. On its own, that result is reason enough to keep building on this approach rather than treating thermal losses in polycrystalline silicon as simply the cost of doing business in hot climates.

REFERENCES

1. Amalu, E. H., & Fabunmi, O. A. (2022). Thermal control of crystalline silicon photovoltaic (c-Si PV) module using Docosane phase change material (PCM) for improved performance. *Solar Energy*, 234, 203–221.
2. Baghel, N. S., & Chander, N. (2022). Performance comparison of mono and polycrystalline

silicon solar photovoltaic modules under tropical wet and dry climatic conditions in east-central India.

Clean Energy, 6(1), 165–177. <https://doi.org/10.1093/ce/zkac001>

3. Bahaidarah, H., Subhan, A., Gandhidasan, P., & Rehman, S. (2013). Performance evaluation of a PV (photovoltaic) module by back surface water cooling for hot climatic conditions. *Energy*, 59, 445–453.
4. Chaich, Z. (2025). Experimental evaluation of water cooling effects on photovoltaic module performance in a hot climate. *International Journal of Energy and Water Resources*. <https://doi.org/10.1007/s42108-025-00346-y>
5. Chandel, S. S., & Agarwal, T. (2017). Review of cooling techniques using phase change materials for enhancing efficiency of photovoltaic power systems. *Renewable and Sustainable Energy Reviews*, 73, 1342–1351.
6. Chandrasekar, M., & Senthilkumar, T. (2015). Experimental demonstration of enhanced solar energy utilization in flat PV (photovoltaic) modules cooled by heat spreaders in conjunction with cotton wick structures. *Energy*, 90, 1401–1410.
7. Chapin, D. M., Fuller, C. S., & Pearson, G. L. (1954). A new silicon p-n junction photocell for converting solar radiation into electrical power. *Journal of Applied Physics*, 25(5), 676–677. <https://doi.org/10.1063/1.1721711>
8. Dubey, S., Sarvaiya, J. N., & Seshadri, B. (2013). Temperature dependent photovoltaic (PV) efficiency and its effect on PV production in the world: A review. *Energy Procedia*, 33, 311–321.
9. Duffie, J. A., & Beckman, W. A. (2013). *Solar engineering of thermal processes* (4th ed.). Wiley.
10. Goswami, D. Y. (2023). *Principles of solar engineering* (4th ed.). CRC Press.
11. Green, M. A. (1982). *Solar cells: Operating principles, technology, and system applications*. Prentice-Hall.
12. Green, M. A., Dunlop, E. D., Hohl-Ebinger, J., Yoshita, M., Kopidakis, N., & Hao, X. (2021). Solar cell efficiency tables (version 57). *Progress in Photovoltaics: Research and Applications*, 29(1), 3–15.
13. Hamed, A. H., Sharkawy, A., Hamdan, I., & Maghrabie, H. M. (2025). Experimental study on performance enhancement of smart PV water cooling using a built ON–OFF temperature control application in a digital embedded system. *Environmental Progress & Sustainable Energy*, 45(3). <https://doi.org/10.1002/ep.70162>

14. Huang, M. J. (2011). The effect of using two PCMs on the thermal regulation performance of BIPV systems. *Solar Energy Materials and Solar Cells*, 95(3), 957–963.
15. Hudişteanu, V.-S., Cherecheş, N.-C., Ţurcanu, F.-E., Hudişteanu, I., & Romila, C. (2024). Impact of temperature on the efficiency of monocrystalline and polycrystalline photovoltaic panels: A comprehensive experimental analysis for sustainable energy solutions. *Sustainability*, 16(23), 10566. <https://doi.org/10.3390/su162310566>
16. International Energy Agency. (2024). *World energy investment 2024*. IEA. <https://www.iea.org/reports/world-energy-investment-2024>
17. Jordan, D. C., & Kurtz, S. R. (2013). Photovoltaic degradation rates: An analytical review. *Progress in Photovoltaics: Research and Applications*, 21(1), 12–29. <https://doi.org/10.1002/pip.1182>
18. Kavlak, G., McNerney, J., & Trancik, J. E. (2018). Evaluating the causes of cost reduction in photovoltaic modules. *Energy Policy*, 123, 700–710. <https://doi.org/10.1016/j.enpol.2018.08.015>
19. Köntges, M., Kurtz, S., Packard, C., Jahn, U., Berger, K. A., Kato, K., Friesen, T., Liu, H., & Van Iseghem, M. (2014). *Review of failures of photovoltaic modules* (IEA-PVPS Task 13, Report T13-01:2014). International Energy Agency.
20. Kusuma, V. A., Aprillia, H., Suprpto, S. S., Ramadhani, M. N., Firdaus, A. A., & Putra, D. F. U. (2023). Analysis of the effect of a microcontroller-based solar panel cooling system on temperature and power output. *International Journal of Applied Power Engineering*, 12(2), 119. <https://doi.org/10.11591/ijape.v12.i2.pp119-125>
21. Maxim Integrated. (2019). *DS18B20 programmable resolution 1-Wire digital thermometer* [Datasheet].
22. Moharram, K. A., Abd-Elhady, M. S., Kandil, H. A., & El-Sherif, H. (2013). Enhancing the performance of photovoltaic panels by water cooling. *Ain Shams Engineering Journal*, 4(4), 869–877.
23. Niknam, S. A., Mortazavi, M., & Li, D. (2021). Additively manufactured heat exchangers: A review on opportunities and challenges. *The International Journal of Advanced Manufacturing Technology*, 112, 601–618. <https://doi.org/10.1007/s00170-020-06372-w>
24. Pavgi, A., Oh, J., & TamizhMani, G. (2021). Thermally conductive backsheets (TCB) of PV modules: Positive impacts on performance, lifetime and LCOE. *Energies*, 14(5), 1252. <https://doi.org/10.3390/en14051252>
25. Raju, K., Priyadarshini, M. S., Asrani, A., & Bajaj, M. (2024). Optimizing solar panel cooling: A

smart approach using Arduino and IoT integration. *E3S Web of Conferences*, 591, 07003.

<https://doi.org/10.1051/e3sconf/202459107003>

26. REN21. (2025). *Renewables 2025 global status report*. REN21 Secretariat.
27. Sarhaddi, F., Farahat, S., Ajam, H., Behzadmehr, A., & Adeli, M. M. (2010). An improved thermal and electrical model for a solar photovoltaic thermal (PV/T) air collector. *Applied Energy*, 87(7), 2328–2339.
28. Skoplaki, E., & Palyvos, J. A. (2009). On the temperature dependence of photovoltaic module electrical performance: A review of efficiency/power correlations. *Solar Energy*, 83(5), 614–624.
29. Texas Instruments. (2015). *INA219 zero-drift, bidirectional current/power monitor with I2C interface* [Datasheet].
30. Wedepohl, K. H. (1995). The composition of the continental crust. *Geochimica et Cosmochimica Acta*, 59(7), 1217–1232. [https://doi.org/10.1016/0016-7037\(95\)00038-2](https://doi.org/10.1016/0016-7037(95)00038-2)
31. Zanlorenzi, G., Szejka, A. L., & Canciglieri, O. (2018). Hybrid photovoltaic module for efficiency improvement through an automatic water cooling system: A prototype case study. *Journal of Cleaner Production*, 196, 535–546.
32. Zhao, D., Aili, A., Zhai, Y., Xu, S., Tan, G., Yin, X., & Yang, R. (2019). Radiative sky cooling: Fundamental principles, materials, and applications. *Applied Physics Reviews*, 6(2), 021306.

About the Author

Naavya Garg is a Grade 12 student at Aditya Birla World Academy in Mumbai, India. She is interested in renewable energy and applied engineering, particularly in how low-cost, accessible hardware can address performance problems in solar technology across hot climates. She plans to study engineering at university, with a focus on sustainable energy systems.