

Temporal Derivative Soft-Sensing and Reconstructing Solar Radiation and Heat Flux from Common Environmental Sensors

Neksha V. DeSilva ^{*}
Independent researcher

FiaOS.org ^{**}

(Dated: November 25th, 2025)

Current environmental monitoring is fundamentally limited by a reliance on static state variables temperature, pressure, and humidity while remaining blind to the dynamic energy exchanges that drive them. This paper introduces Differential Temporal Derivative, Soft-Sensing DTDSS, a novel physics-based framework that transforms standard, low-cost environmental sensors into capability-dense radiometers.

By employing a differential topology with Inertial Noise Reduction (INR), we can mathematically reconstruct Global Horizontal Irradiance (GHI) and convective heat flux without the cost or fragility of thermopile pyranometers.

While validated using the FiaOS reference architecture, this methodology is hardware-agnostic. The ultimate vision of this work is the deployment of a high-performance, open-source computational library designed for professional embedded development environments. By distributing this algorithm via global repositories, aiming to upgrade the capabilities of millions of existing and future IoT devices. This library allows standard electronics to move beyond simple linear measurements and unlock higher-order environmental physics. By deriving complex energy flux parameters from common sensors, we open a new frontier of derived equations and applications from precision engineering to autonomous energy management essentially democratizing advanced meteorological physics for the global engineering community.

Keywords: Soft-sensing, Embedded Systems, IoT, Solar Irradiance, Heat Flux, Signal Processing.

^{*} nekshavs@gmail.com, www.nekshadesilva.com

^{**} research@fiaos.org, www.fiaos.org/about

Introduction

Current environmental monitoring relies on common sensors (like the Bosch BME280) that measure static variables: temperature, pressure, and humidity. And standard environmental sensors answer static questions:

"What is the temperature?" or "How humid is it?"

They remain unresponsive, however, to the dynamic energy exchanges that dictate these conditions. Specifically, they fail to measure:

- ☐ Global Horizontal Irradiance (GHI): The amount of solar energy reaching the ground per unit area.
- ☐ Heat Flux: The rate of thermal energy transfer through the environment.

These are not merely abstract metrics; they are essential for practical applications. Solar radiation is the primary energy input for Earth's climate and a required variable for calculating evapotranspiration in agriculture.

During the development of my project, FiaOS.org, this limitation became a practical obstacle. I attempted to implement the Penman-Monteith equation, the standard for estimating water evaporation, which relies on solar flux data. My initial hardware approach was mainly the output voltage of a solar panel and a Thermoelectric Generator (TEG) module. However, this setup proved unreliable. The TEG, due to its ceramic construction, showed high thermal inertia; it retained heat on its surface rather than providing the responsive data required. The physical disconnections between the sensors further complicated the readings.

To resolve this, I realized the system could not rely on disjointed, reactive components. It required a cohesive, integrated approach. I returned to the Bosch™ BME280 as a baseline foundation. Its accuracy in measuring three physically connected parameters, Relative Humidity, Temperature, and Pressure, provides the necessary stability to build a more advanced, accumulative sensing package.

NOTE: The BME280 measures absolute barometric pressure directly, avoiding errors associated with altitude-derived estimates.

Why, cheap sensors cannot reliably measure GHI and Heat Flux, even indirectly?

1.1. The Hardware Problem

Traditionally, measuring solar radiation requires a thermopile pyranometer, an instrument governed by ISO 9060 standards. These devices work on a simple logic, they have a black absorbing surface and a white reference surface, and they measure the temperature difference between them. That temperature difference, via the “Seebeck effect”, generates a tiny voltage proportional to the incident radiation.

But there is a massive cost. For a research grade application or an invention could afford this, and it is recommended. A good pyranometer runs into hundreds or thousands of dollars. It requires careful mounting, cleaning, and calibration. It consumes power. It is fragile. We cannot scatter these things across a farm or embed them in consumer electronics. Such as all we identify as consumer devices. A pyranometer generally runs without power consumption. And it is solid. But when it comes to humanity's needs, it is clearly not the optimal choice, even if we eliminate the entire cost hurdles. Meanwhile, the BME280, the example sensor we take here, costs a few dollars, runs on a few microamps, and fits on a shell.

1.2. The Soft-Sensing Promise (And Its Failures)

Today's response has been soft sensing, which means using algorithms to estimate what you cannot measure directly. The dominant approach uses machine learning.

"Train a neural network on historical data where you have both the cheap sensor readings and the expensive pyranometer readings. Then deploy the neural network to estimate radiation from temperature and humidity alone."

This approach has serious problems:

While machine learning is a common approach for soft-sensing, it presents significant challenges for embedded, devices. First, neural networks trained in specific climates often fail to generalize to new environments. Second, the matrix multiplications required for inference consume excessive power on 8-bit microcontrollers. Finally, standard ML models often map instant inputs to outputs, ignoring the crucial temporal dynamics of how a sensor heats up over time.

Most ML models map instant inputs to instant outputs.

"Given that the temperature is 25°C and the humidity is 60% right now, what is the radiation right now?"

But radiation changes on timescales of seconds like when clouds get to move,
while temperature changes on timescales of minutes.

This mismatch causes noticeable errors. And in specific applications like in LEO or SSO, this changing time may be even unmanageable, as magnitudes of higher dynamic environmental changes. (Further explained at the end of the document)

NOTE: In the BME280, air pressure is not determined by the general formula of derivation, $p = \rho gh$. Instead, it uses the physical pressure applied on the sensor. This may apply to most general-purpose sensors.

1.3. The Differential Thermodynamic Approach




Figure 1.3.1

“Differential Temporal Derivative Soft-Sensing”, abbreviated as DTDSS;

This can be clarified as a single sensor cannot achieve thermal equality with the air (for temperature/humidity) and thermal disequilibrium with solar radiation (for its measurement). Therefore, sensing requires partitioning into two separate thermodynamic nodes.

We utilize two identical sensors (we take BME280) placed in close physical proximity but in contrasting enclosures:

- ☐ The Reference node (S_{ref}): Housed in a ventilated, reflective shield. It remains in thermal equilibrium with the ambient air mass. (Figure 1.1.2)
- ☐ The Flux node (S_{flux}): Housed in a sealed, absorptive black-body enclosure. It acts as an energy trap, heating up in response to incoming radiation.

By analyzing the real-time variation between these two nodes, specifically the temporal derivative of their temperature difference, we can mathematically reconstruct the incident radiation without the massive cost of a thermopile pyranometer.

1.4. The Preview




Figure 1.4.1 - Pseudocode of the signal processing pathways

To make this work robustly, we have developed a differential system with two processing paths:

- ☐ The Reference Path (the Baseline): Operates on the vented sensor. Uses psychrometric relationships and empirical meteorology to establish accurate baseline conditions. Provides true ambient temperature (T_{ref}), relative humidity, and pressure readings uncontaminated by solar heating.
- ☐ The Reactive Path (Flux): Operates on the black body sensor. Uses the temporal derivative of the temperature difference to detect and quantify radiative energy input.

The Reference Path provides stability and ensures the system has accurate environmental context, preventing the "dry kiln" error where an enclosed sensor reports artificially low humidity due to their self-heating. The Reactive Path provides responsiveness and captures transient events like cloud breaks, measuring the actual energy flux hitting the sensor.

Central to the Reactive Path is a novel signal processing technique; we call it Inertial Noise Reduction (INR). Taking derivatives of noisy signals is leading to future complications.

INR is a novel solution for this problem. The Inertial Noise Reduction algorithm is a filter designed specifically for the thermodynamics of sensors.

Finally, by using the BME280's pressure sensor to calculate air density in real time, the entire system becomes height isolated environment. Implement it at sea level or at 4000 meters higher, Regardless of earth location, the physics adapts automatically.

Upcoming is the inner mechanism explained.

2. The Thermodynamics of Humid Air, the Medium.

Understanding the medium in which the sensor operates is important at first. The sensor is absorbed or fully immersed in the lower atmosphere, and the properties of that atmosphere directly govern how the heat goes between the sensor and nearby environment. Many engineering models treat air as a simple ideal gas with constants as values of properties, or they assume standard atmosphere conditions (sea level, 21°C, dry air, which are increasingly becoming very common for machine learning algorithms). These simplifications are acceptable for many applications. For a *flux model* intended for deployment anywhere on earth, these are risky and may lead to unpredictable outputs. *

A main innovation of the DTDSS framework is the rejection of static assumptions in favor of dynamic, first principles derivation of air properties from the measured state variables, *pressure(P)*, *temperature(T)*, and *relativehumidity(RH)*.

2.1. The Role of Air density

The density of the air, ρ , is the governing parameter for convective heat transfer. When the sensor is warmer than the surrounding air, heat flows from the sensor to the air via convection. The rate of this heat flow depends on the density of the fluid carrying the thermal energy to the outside environment.

The practical implications include, **

- ☐ At sea surface level, air density is approximately 1.225 kg/m³.
- ☐ At 2000 meters elevated, air density drops to 1.00 kg/m³.
- ☐ At 4000 meters, it falls to about 0.82 kg/m³.

If your model assumes sea level density but the sensor is at 4000 meters, you will overestimate the increasing cooling by roughly 50%. This means you will underestimate the solar radiation required to explain a given temperature rise by a similar amount. Your readings will be malfunctioned and irreversible, since we are filtering these inputs again(explained in the future).

The BME280 provides *totalatmosphericpressureP*, which is directly related to altitude. We can use this to calculate actual air density without needing a GPS or altitude lookup table.

** Approximate values only.

* Explained in page XI.

2.2. The Ideal Gas Law for Moist Air

The Atmosphere is not pure dry air. Mainly, it is a mixture of dry air and vaporized water, and these two components have different molecular weights and different gas constants.

According to *Dalton's Law of Partial Pressures*, the total atmospheric pressure is the sum of the partial pressure of dry air (P_d) and the partial pressure of water vapor (e).

Dalton's Law of Partial Pressures assumption:

$$P = P_d + e \quad P = P_d + e$$

The density of moist air is the sum of the density contributions from each component. Applying the ideal gas law in specific form ($P = \rho R_{specific} T$):

$$\rho_{moist} = \rho_d + \rho_v = \frac{P_d}{R_d T} + \frac{e}{R_v T} \quad \rho_{moist} = \rho_d + \rho_v = \frac{P_d}{R_d T} + \frac{e}{R_v T}$$

Where:

- ☐ $R_d = 287.058 \text{ J} \cdot \text{kg}^{-1} \cdot \text{K}^{-1}$ - is the specific gas constant for dry air
- ☐ $R_v = 461.495 \text{ J} \cdot \text{kg}^{-1} \cdot \text{K}^{-1}$ - is the specific gas constant for water vapor
- ☐ T is in Kelvin grades.

Substituting $P_d = P - e$ gives,

$$\rho_{moist} = \frac{P - e}{R_d T} + \frac{e}{R_v T} \quad \rho_{moist} = \frac{P - e}{R_d T} + \frac{e}{R_v T}$$

For algorithmic implementation, we can factor this more precisely. Define the ratio of gas constants as

$$\epsilon = \frac{R_d}{R_v} \approx 0.622$$

$$\rho_{moist} = \frac{P}{R_d T} \left(1 - \frac{e}{P} (1 - \epsilon) \right) \quad \rho_{moist} = \frac{P}{R_d T} \left(1 - \frac{e}{P} (1 - \epsilon) \right)$$

Humid air is less dense than dry air at constant temperature and pressure due to the lower molecular weight of water vapor which is 18 g/mol compared to dry air (29 g/mol).

For our purpose, this means humid air is worse at carrying heat away from the sensor.

By calculating ρ_{moist} dynamically at every timestep, the soft sensor recalibrates itself its thermal model for the exact altitude and humidity conditions it experiences. This is a direct rejection of the "standard atmosphere" constraint.

2.3. Derivation of Vapor Pressure Dynamic

The partial pressure of water vapor, e , is not measured directly by the BME280.

Hence, derive it from the relative humidity (RH) and the saturation vapor pressure (e_s).

$$e = e_s(T) \cdot \frac{RH}{100} \quad e = e_s(T) \cdot \frac{RH}{100}$$

The saturation vapor pressure is the pressure at which water vapor would be in a balance with liquid water (not sterile water) at a given temperature. It is a strongly nonlinear function of temperature. This provides an accurate approximation of,

$$e_s(T) = 6.112 \cdot \exp\left(\frac{17.67 \cdot (T - 273.15)}{T - 29.65}\right) \quad e_s(T) = 6.112 \cdot \exp\left(\frac{17.67 \cdot (T - 273.15)}{T - 29.65}\right)$$

Where T is in Kelvin and e_s is in hPa. This equation is robust for the temperature ranges encountered in typical monitoring -40°C to +50°C range. Though this is not a fixed range, these numbers may vary depending on the specific sensor, etc.

2.4. The Specific Enthalpy of Moist Air

The enthalpy of moist air is important for understanding the energy state of the atmosphere. Enthalpy represents the total energy content of the air parcel, combining sensible heat (what a normal thermometer measures) and latent heat (the energy stored in the phase of water, during the transition to vapor from water).

The specific enthalpy of moist air (per unit mass of dry air) is:

$$h = h_a + x \cdot h_v \quad h = h_a + x \cdot h_v$$

Where x is the humidity mixing ratio (mass of water vapor per mass of dry, less humid air):

$$x = \epsilon \frac{e}{P - e} = 0.622 \frac{e}{P - e}$$

The individual enthalpy terms, approximated as linear functions of temperature t (in Celsius):

- ☐ $h_a \approx 1.006 t$ (kJ/kg) , enthalpy of dry air
- ☐ $h_v \approx 2501 + 1.86 t$ (kJ/kg) , enthalpy of water vapor

The number 2501 kJ/kg is the latent heat of vaporization of water at 0°C . This is a relatively massive energy reservoir. Evaporating one gram of water requires as much energy as heating that gram by 2501°C if such heating were possible in the liquid phase. The total specific enthalpy becomes:

$$h = 1.006 t + x(2501 + 1.86 t) \quad h = 1.006 t + x(2501 + 1.86 t)$$

The impact of this is on soft sensing. A sensor measuring only *temperature* misses the *latent energy component* entirely.

By assessing these two scenarios:

- ☐ State A: Temperature rises from 20°C to 25°C , humidity stays constant at 50%.
- ☐ State B: Temperature rises from 20°C to 25°C , humidity rises from 40% to 70%.

A temperature only model would see these both as the same. But Scenario B shows a much larger change in total energy; all that extra water vapor carries latent heat. In the Differential System, the Reference Path tracks changes in enthalpy to identify between:

- ☐ Advective heating, which is warm, humid air moving in (both T and x change together)
- ☐ Radiative heating: Sunlight hitting the sensor (sensor heats up but the mixing ratio x of the surrounding air does not immediately change)

This distinction helps avoid false positives where we mistake weather fronts for solar events.

2.5. Convective Heat Transfer Coefficient

The link between sensor temperature and air temperature is the convective heat transfer coefficient, h_c . This parameter dictates the rate of heat flux for a given temperature difference:

$$q_{conv} = h_c(T_s - T_\infty) \quad q_{conv} = h_c(T_s - T_\infty)$$

Where the case:

- ☐ q_{conv} is the convective heat flux
- ☐ T_s is the sensor temperature
- ☐ T_∞ is the ambient air temperature

For a small object like the BME280 package, LGA metal lid, approximately 2.5mm x 2.5mm x 0.9g, the flow regime is typically laminar at low wind speeds.

The Nusselt number (Nu), the dimensionless general ratio relative to conductive heat transfer, relates h_c to the thermal conductivity of air as:

$$Nu = \frac{h_c \cdot L}{k_{air}} \quad Nu = \frac{h_c \cdot L}{k_{air}}$$

Empirical correlations for laminar flow over small objects suggest:

$$Nu \propto Re^{1/2} \cdot Pr^{1/3} \quad Nu \propto Re^{1/2} \cdot Pr^{1/3}$$

Where Re is the Reynolds number and Pr is the Prandtl number. Expanding this, we find the dependence of h_c on air properties:

$$h_c \propto k_{air} \left(\frac{\rho V}{\mu} \right)^{0.5}$$

The point is, h_c depends on $\rho^{0.5}$. Hence, the altitude correction is non-negotiable. As the sensor moves to higher heights, ρ decreases, and thus h_c decreases. The sensor becomes "better isolated" by the thinner, less humid air. If we did not account for this, the algorithm would interpret the higher temperature rise of the sensor (caused by poor cooling and less humidity) as an increase in solar radiation. We would report phantom solar events every time we took the sensor up a mountain.

By driving the model with calculated ρ_{moist} , The DTDSS framework normalizes heat flux estimation against altitude variations. The sensor can be manufactured in a factory at sea level and deployed anywhere without firmware modification.

The Differential Architecture

The central challenge in soft sensing is identifying between a hot day (high ambient temperature) and a sunny day (high radiative flux). A single sensor struggles to separate these signals. The differential System resolves this by treating the environment as a differential equation solved in hardware, By processing two concurrent telemetry streams from physically distinct sensor nodes.

3.1. Pipeline A, The Reference Path

Take our input as the Data from the Ventilated Sensor (take this as S_{ref}) and Domain as Equilibrium Thermodynamics,,

Role: This pipeline establishes the "Global Ground Truth" of the air mass. It provides the accurate T_{ref} , RH , and P required for the biological or meteorological models. Crucially, because this sensor is shielded from solar coverage **, the humidity readings remain accurate and relevant to the surrounding atmosphere, preventing the error common in most enclosed sensors.

Baseline Temperature Estimation- The Reference Node provides the true ambient temperature $T_{\infty} = T_{ref}$ without the need for estimation or filtering. This is a direct value we take from here.

Meteorological Context: Use the relative humidity and pressure data to estimate cloud cover fraction, which provides a probabilistic baseline for expected solar radiation.

The Cloud Proxy Model and The Kasten-Czeplak Adaptation

A main component of the Reference Path is estimating solar radiation based on humidity. The physical reasoning is straightforward: high surface relative humidity correlates strongly with cloud formation and atmospheric opacity. When humidity is high, the atmosphere is closer to saturation, clouds are more likely to form, and less sunlight reaches the surface.

The Kasten-Czeplak Model is a widely cited empirical relationship connecting Global Horizontal Irradiance (G) to cloud cover fraction (N , where 0 is clear sky and 1 is fully overcast):

$$G = G_{cs} \cdot (1 - A \cdot N^B) \quad G = G_{cs} \cdot (1 - A \cdot N^B)$$

Where:

- G_{cs} is the theoretical clear-sky radiation (calculable from latitude, longitude, time of day, and day of year)
- $A = 0.75$ and $B = 3.4$ are the standard Kasten-Czeplak coefficients

These coefficients imply that a fully overcast sky (*where* $N = 1$) allows $\sim 1/4$ of clear-sky radiation to penetrate ($1 - 0.75 = 0.25$).

** Explained further in page XIII

The Problem:

“We do not have a sky camera or human observer to measure N . We need a proxy.”

The Solution: Derive a cloud proxy from relative humidity. Observational data suggests a power-law relationship between surface humidity and cloud fraction. A simplified proxy function suitable for embedded processing:

$$N_{proxy} = \left(\frac{RH}{100} \right)^k$$

Where k is a tuning parameter (typically 1.5 to 2.0, depending on local climate). Integrating this into the Kasten Czeplak model gives the Reference Path's radiation estimate which is $G_{baseline}$.

$$G_{baseline} = G_{cs} \cdot \left(1 - 0.75 \cdot \left(\frac{RH}{100} \right)^{3.4k} \right)$$

Interpretation

"It is humid, therefore it is likely cloudy, therefore radiation is likely low."

This equation provides a stable, slowly varying estimate of solar radiation. It is valid for time-averaged conditions. But it cannot detect if the sun is always shining through a space between the clouds(gaps). Detection requires the Reactive Path.

3.2. Pipeline 2: The Reactive, Instantaneous Path

Input: Data from the *Black-Body Sensor* (S_{flux}).

Domain: Non-Equilibrium Dynamics.

This sensor is allowed, and encouraged to overheat over hundreds of hours. We monitor the deviation of its temperature (T_{flux}) from the Reference Temperature (T_{ref}).

- ☐ If $T_{flux} \approx T_{ref}$, the sky is overcast or it is night.
- ☐ If $T_{flux} \gg T_{ref}$, energy is being harvested.

The magnitude of this deviation, combined with the rate of change ($\frac{dT}{dt}$), enables to calculate the incoming energy flux using the Inertial Noise Reduction (INR) filter derived later in this paper.

Flux Detection: Monitor the temporal derivative of the temperature difference to identify rapid heating events driven by radiative input.

Inertial Compensation: Apply the INR to recover the magnitude of the flux without a delay, basically, in real-time. The Reactive Path leverages what is normally considered a defect of the sensor.

NOTE: Standard datasheets warn that the BME280 can self-heat by 0.5°C to 2°C depending on sampling rate. Users are typically advised to minimize this effect or compensate for it. In the DTDSS framework, we embrace this thermal coupling for the Flux Node. We quantify it. We use it.

If the Flux sensor heats up faster than the Reference sensor, the excess energy must come from an external source: solar radiation (G) or nearby heat flux (q).

3.3. Comparison of Differential Pipeline Characteristics

| Feature used | Accumulative path | Reactive Path Variations |
|-------------------|---|---|
| Time Scale | Minutes to Hours | Seconds |
| Physics Model | Equilibrium Thermodynamics | Non-Equilibrium Heat Transfer |
| Key Equation | Kasten-Czeplak | Newton's Law of Cooling - Inverted. |
| Primary Input | Relative Humidity (RH) | Temperature Derivative (dT/dt) |
| Output | Baseline Radiation Probability | Instantaneous Flux Measurement |
| Noise Sensitivity | Low because of heavy filtering and multiple layers of reduction and pruning | High (Requires INR) |
| Failure Mode | Slow to produce low_quality reactions. | Sensitive to Noise, and immediately recognizable in malfunction(s). * |

3.4. How the Pipelines Work Together

The two pipelines are not independent. They interact in several main ways:

- ☐ The Reference Path provides the baseline T_{ref} for the Reactive Path. The Reactive Path needs to know

"how much warmer is the Flux sensor than the ambient air?"

That ambient air temperature comes directly from the Reference sensor, It is clearly a measurement, not an estimate.

- The Reference Path provides sanity bounds for the Reactive Path. If the Reactive Path calculates an *instantaneous radiation of 1500 W/m²*, (cases mentioned in Introduction pages) but the Reference Path's clear-sky model says the maximum possible is 900 W/m².

Something is wrong. Either there is a measurement error, or there is an additional heat source beyond sunlight. The system can flag this anomaly.

- The Reactive Path can update the Reference Path's *cloud estimate (population of clouds within variable proximity)*. If the Reactive Path detects a sharp increase in flux, and the Reference Path was predicting low radiation due to unpredictable high humidity, the Reference Path can revise its cloud proxy parameter k to better match observed conditions. This shows how the reactive system controls and improves the internal accuracy, when we reverse engineer the concept “Rapid”.
- Common noise rejection. Because both sensors experience the same wind, the same pressure changes, and the same air mass movements, many environmental noise sources affect both equally. When we compute the differential ($T_{flux} - T_{ref}$), these common-mode disturbances cancel out, leaving only the solar signal.

This bidirectional feedback and differential structure makes the system self correcting and less error prone.

4. Mathematical Derivation of Heat Flux via Thermal Inertia

Here, converting the differential temperature signal into a heat flux measurement is the first step. The approach is to treat the Flux sensor package as a *Lumped Capacitance System*. This idea is valid when the *Biot number* (Bi), which compares internal thermal resistance to external convective resistance, is less than 0.1.

For small MEMS packages like the sensor we are using here, this condition is easily satisfied. The implication: the temperature within the sensor is essentially uniform.***

*** The selection of the BME280 sensor, a standard for general-purpose embedded devices due to its simplicity and wide availability, is purely for convenience. As noted in the introduction, the specific hardware model is not a critical factor.

4.1. The Differential Energy Balance

Applying the conservation of Energy to the Flux Node (S_{flux}). The driving force is no longer an estimated ambient temperature, but the measured temperature from the Reference Node (S_{ref})

$$E_{in} - E_{out} = E_{stored}$$

Expanding each term:

- ☐ E_{in} (Input Power): The sum of:
 - ☐ Solar radiation absorbed: $\alpha G A_s$
 - ☐ Electrical self-heating: P_{elec}
- ☐ E_{out} (Cooling medium): Convective heat loss to ambient air: $h_c A_s (T_{flux} - T_{ref})$
- ☐ E_{stored} (Thermal Inertia): Energy stored in the sensor mass, causing temperature increases to $m C_p \frac{dT_{flux}}{dt}$

The resulting differential equation:

$$(\alpha G A_s + P_{elec}) - h_c A_s (T_{flux} - T_{ref}) = m C_p \frac{dT_{flux}}{dt}$$

Let,

- ☐ α - Solar absorptivity of the sensor package (typically 0.8–0.9 for black epoxy)
- ☐ G - Solar rad.
- ☐ A_s - Area for convection/radiation
- ☐ T_{flux} - Temperature of the black-body sensor (pre measured)
- ☐ T_{ref} - Measured temperature of the ventilated reference sensor (replacing the estimated T_{∞})
- ☐ $(T_{flux} - T_{ref})$ - The measured thermal potential generated by radiation
- ☐ $m C_p$ - Total heat capacity of the sensor (mass \times specific heat)
- ☐ h_c - heat transfer coefficient

And so, this substitution is valuable. In single sensor based systems, T_{∞} is a calculated guess, more likely to drift. In this differential system, T_{∞} is a live measurement. 7

This eliminates the error where a sensor might make a mistake by taking a heatwave for sunlight.

4.2. Solving for Solar Radiation

Divide the entire equation by the convective term $h_c A_s$:

$$\frac{\alpha G}{h_c} + \frac{P_{elec}}{h_c A_s} - (T_{flux} - T_{ref}) = \frac{m C_p}{h_c A_s} \frac{dT_{flux}}{dt}$$

Define two characteristic parameters,

☐ Thermal Time Constant (τ): The measure of thermal inertia.

$$\tau = \frac{m C_p}{h_c A_s}$$

This is the time it takes for the sensor to respond to a step change in conditions.

(A typical value for a BME280 in still air might be approx. 10s to 60s.)

☐ Self-Heating Offset (T_{rise}): How much the sensor's temperature goes higher and stays there because of the heat produced by itself.

$$T_{rise} = \frac{P_{elec}}{h_c A_s} T_{rise} = \frac{P_{elec}}{h_c A_s}$$

Then, rearranging to isolate the solar term:

$$\frac{\alpha G}{h_c} = (T_{flux} - T_{ref}) + \tau \frac{dT_{flux}}{dt} - T_{rise}$$

We call this quantity the Sol Air Excess (T_{sol}):

$$T_{sol} = (T_{flux} - T_{ref}) + \tau \frac{dT_{flux}}{dt} - T_{rise}$$

Finally, solving for Solar Radiation:

$$G = \frac{h_c}{\alpha} \left[(T_{flux} - T_{ref}) + \tau \frac{dT_{flux}}{dt} - T_{rise} \right]$$

4.3. Understanding the Equation

This equation shows that Solar Radiation is proportional to two distinct things-

- ☐ The temperature elevation::

$$(T_{flux} - T_{ref})$$

This is how much warmer the Flux sensor is than the Reference.

- ☐ The rate of change scaled by the time constant

$$\tau \frac{dT_{flux}}{dt}$$

How fast the Flux sensor is heating up.

Implementations

Steady State:

$$\frac{dT}{dt} = 0$$

When conditions are stable, the derivative term disappears.

The measurement relies solely on the temperature difference between the two sensors. Both have reached equilibrium with their own environments, and the temperature difference indicates the radiation level.

Transient State: $\frac{dT}{dt} > 0$

When the sun just came out from behind a cloud, the Flux sensor has not yet reached its new equilibrium temperature. It is still heating up.

The term $\tau \frac{dT}{dt}$ represents the Inertial Correction, it accounts for the energy currently stored in sensor's mass but that has not yet appeared as a full temperature rise. Without the inertial correction, there is an underestimated radiation during heating events and *overestimation* during internal cooling stages.

The Newton's Law of Cooling

The term $(T_{flux} - T_{ref})$ represents the driving force for cooling, governed by Newton's Law of Cooling. The inclusion of the derivative term inverts Newton's Law.

Newton's Law:

"The rate of cooling is proportional to the temperature difference."

Mathematically,

$$\frac{dT_{flux}}{dt} = -\frac{1}{\tau}(T_{flux} - T_{ref}) + \frac{G_{absorbed}}{\tau}$$

Our equation rearranges this to solve for the absorbed flux:

$$G_{absorbed} = (T_{flux} - T_{ref}) + \tau \frac{dT_{flux}}{dt} = (T_{flux} - T_{ref}) + \tau \frac{dT_{flux}}{dt}$$

Measuring how fast the Flux sensor is heating up, and knowing how fast it should be cooling (based on $T_{flux} - T_{ref}$), the difference is the magnitude of the external forcing.

Altitude Doubt

Because h_c depends on ρ (as derived in *part 2.5*), and ρ is calculated from local pressure P , this entire derivation is independent from height.

At high altitudes:

- ☐ ρ drops because of thinner air - Explanation - "thinner air" is subjective. There is no specific threshold.
- ☐ h_c drops because of less efficient convection.
- ☐ τ increases because the sensor responds slower.
- ☐ $(T_{flux} - T_{ref})$ increases for the same radiation as the sensor gets hotter because it cannot cool as effectively.

The formula *compensates* for all of these effects automatically because ρ is a live variable in the system. The sensor identifies where it is through the value of pressure reading.

Inertial Noise Reduction (INR) or Signal Processing for the Reactive Path

The theoretical derivation in Section 4 relies on the instantaneous derivative $\frac{dT_{flux}}{dt}$. In the real world of digital sampling, taking a derivative is a hazardous operation.

The Noise

Examine the practical reality:

- ☐ The BME280 has a temperature resolution of approximately 0.01°C.
- ☐ The sensor exhibits quantization noise (bit-flip errors) at this resolution.
- ☐ The sampling interval might be 1 second.

A single bit flip of 0.01°C over 1 second produces a well calculated derivative of 0.01°C/s. If the thermal time constant τ is 30 seconds, this noise spike gets multiplied by 30, producing a temperature of 0.3°C. Depending on the system's calibration, this could be interpreted as a radiation change of tens of W/m².

The noise amplification scales inversely with the sampling interval

$$\text{Noise}_{\text{derivative}} \propto \frac{\text{Noise}_{\text{sensor}}}{\Delta t}$$

Sample faster, and the noise gets worse. Sample slower, and you miss *temporary* event happenings.

Why Standard Filters Fail

The obvious solution is filtering. Smooth the temperature signal before differentiating. But standard filters introduce a noticeable delay. A Simple Moving Average with a case of 10 samples delays the signal by 5 samples. A Low-Pass Butterworth filter introduces frequency-dependent delay.

This delay is fatal for our application. If the temperature derivative is delayed relative to the temperature value, the two terms in our flux equation become decoupled. We are adding apples and oranges. The physical relationship we derived assumes simultaneous measurements.

What we need is a filter that is,

- ☐ Suppresses high-frequency noise
- ☐ Preserves the phase relationship between the signal and its derivative
- ☐ Is computationally cheap enough for embedded systems

The INR Filter Concept

The Inertial Noise Reduction (INR) filter is a proposed solution. It is a signal processing filter designed to model physical inertia digitally inside the domain.

The sensor itself is already acting as a physical filter. The thermal mass smooths out rapid changes in radiation. A 10 second burst of sunlight through a cloud gap barely registers because the sensor had time to heat up.

INR works in two different stages.

Adaptive Smoothing or Stage 01

Adaptive smoothing is about using an IIR filter based on the Exponential Moving Average (EMA)

$$T_{filt}[n] = \alpha \cdot T_{raw}[n] + (1 - \alpha) \cdot T_{filt}[n - 1]$$

Where, α is the smoothing factor ($0 < \alpha \leq 1$):

- ☐ Low α : High inertia, heavy filtering, slow response
- ☐ High α : Low inertia, light filtering, fast response

The innovation of INR is making α adaptive.

Next, adjusting the filter's predicted mass based on the jerk of the signal

$$\Delta[n] = |T_{raw}[n] - T_{filt}[n - 1]| \quad \Delta[n] = |T_{raw}[n] - T_{filt}[n - 1]|$$

$$\alpha[n] = \text{clamp}(k \cdot \Delta[n], \alpha_{min}, \alpha_{max}) \quad \alpha[n] = \text{clamp}(k \cdot \Delta[n], \alpha_{min}, \alpha_{max})$$

Where k is a sensitivity gain parameter.

- ☐ Steady state (signal changing slowly): Variations are likely noise. Decrease α (increase inertia) to lock the value.
- ☐ Transient event (sustained unidirectional change): This represents a physical forcing event. Increase α (decrease inertia) to track the edge accurately.

Inertial Projection - Stage 2

After smoothing, we compute a projected temperature that compensates for the sensor's thermal

-lag:
$$T_{proj}[n] = T_{filt}[n] + \tau \cdot \left(\frac{dT_{filt}}{dt} \right) [n]$$

This projected temperature represents what the sensor temperature would be if the sensor had zero thermal mass, an instant response to radiation changes. When the sun comes out, T_{proj} spikes immediately toward the final equilibrium temperature, while T_{filt} slowly climbs to meet it. The difference between them during the transient represents the inertial energy being stored.

Computing the Derivative.

For the derivative $\frac{dT_{filt}}{dt}$, a simple finite difference $\frac{T[n] - T[n-1]}{\Delta t}$ includes massive noise, even after adaptive smoothing. A more robust approach is the Savitzky-Golay filter (or least-squares slope estimator).

This technique fits a low-degree polynomial to a sliding window of data points (e.g., the last 5–10 samples) and calculates the derivative of that polynomial analytically at the center point.

For embedded systems with limited resources, a simplified version uses the Central Difference.

Derivative:

$$\left(\frac{dT}{dt}\right)[n] = \frac{T_{filt}[n+1] - T_{filt}[n-1]}{2\Delta t}$$

This requires only a small buffer and so, small computation.

INR and Kalman Filtering, The drawbacks and why choose INR?

Kalman filters are theoretically greater for linear systems with general noise. But they have practical disadvantages for our application.

- ☐ Computational Cost, Kalman filters require matrix operations (prediction and update steps involve matrix inversion or multiplication). On 8-bit microcontrollers without floating-point units, this is intense.
- ☐ Tuning Difficulty: Correct operation requires setting the process noise (Q) and measurement noise (R) correlated matrices. Without ground truth data, this tuning is difficult.
- ☐ Model Assumptions: *Standard Kalman* assumes a linear, state space model.

Our system has nonlinear elements (the adaptive filtering, the density dependent heat transfer).

The INR filter, by comparison:

- ☐ Requires only scalar arithmetic (add, subtract, bit-shift)
- ☐ Is $O(1)$ in time complexity and $O(1)$ in space complexity
- ☐ Achieves performance comparable to a 1D steady-state Kalman for thermal decay profiles
- ☐ Is specifically optimized for the physics of sensor packages

INR, while simple, wins here in many aspects because of its algorithmic complexity and gives processors less cache when it comes to embedded data processing in applications such as non-negotiable compute for smaller time units per runtime.

The INR Algorithm (Overall summarized Version)

Sequentially, for each new $T_{raw}[n]$,

- Compute deviation from previous filtered value
 $\Delta = |T_{raw}[n] - T_{filt}[n - 1]|$ $\Delta = |T_{raw}[n] - T_{filt}[n - 1]|$
- Adjust smoothing factor
 $\alpha = \text{clamp}(k \cdot \Delta, \alpha_{min}, \alpha_{max})$ $\alpha = \text{clamp}(k \cdot \Delta, \alpha_{min}, \alpha_{max})$
- Apply adaptive EMA
 $T_{filt}[n] = \alpha \cdot T_{raw}[n] + (1 - \alpha) \cdot T_{filt}[n - 1]$
 $T_{filt}[n] = \alpha \cdot T_{raw}[n] + (1 - \alpha) \cdot T_{filt}[n - 1]$
- Store in a circular buffer for derivative calculation.
- Compute derivative via central difference or Savitzky-Golay:
 $dT_dt = \text{derivative}(\text{buffer})$ $dT_dt = \text{derivative}(\text{buffer})$
- Compute projected temperature: $T_{proj}[n] = T_{filt}[n] + \tau \cdot dT_dt$
 $T_{proj}[n] = T_{filt}[n] + \tau \cdot dT_dt$
- Compute solar radiation estimate: $G = (h_c / \alpha_{solar}) \cdot (T_{proj}[n] - T_{ref} - T_{rise})$
 $G = (h_c / \alpha_{solar}) \cdot (T_{proj}[n] - T_{ref} - T_{rise})$

The Wind and Its Practical Boundary Conditions

A critical factor not yet explored is the wind. Introducing this variable requires new assumptions and provides new insights.




Figure 1.6.1

Convective Heat Transfer Dynamics

First, Wind speed is not background noise, it is a fundamental variable in environmental intelligence because it directly controls the convective heat transfer coefficient (h_c). As we can see in Figure 1.1.1, wind oscillations directly interfere with heat.

The assumption that heat flux has nothing to do with wind speed is incorrect. In forced wind, h_c can be 5 to 10 times higher than in natural air.

If the wind suddenly picks up, h_c increases, and both sensors cool down rapidly, and unpredictably.

Differential Cancellation

This sensitivity is exactly why differential architecture is so valuable. "Common mode" disturbances, like a sudden gust of wind, affect both the Flux sensor (T_{flux}) and the Reference sensor (T_{ref}) in a similar way. When we calculate the difference, much of that wind-induced temperature change subtracts out, leaving the solar signal relatively clean.

However, the cancellation isn't perfect. Differences in the housing geometry mean the ventilated Reference housing might react to the wind slightly differently than the sealed Flux housing.

The Proximity Correlation Hypothesis

To handle these variances, we rely on a concept we observed during practical testing: the spatial uniformity of airflow.

In Figure 1.1.1, it introduces something useful. The wind speed at the sensor's immediate location was approximately equivalent to the wind moving the distant buildings/objects, which were located more than 1,000 meters away. This practical experiment led me to the Proximity Correlation Hypothesis.

In open environments, wind speed isn't just a local point measurement; it correlates reasonably well with the general area.

This makes physical sense. Wind is driven by large-scale pressure gradients. Unless there are specific local obstacles, like a wall or a building, the wind hitting your sensor is likely very similar to the wind 500 meters or a kilometer away.

*FiaOS*** and Cloud Integration

Because of this natural uniformity, we can connect the *FiaOS*** system to cloud data. Since the local wind speed is approximately equivalent to the regional wind speed (measured in kilometers), we don't necessarily need a dedicated anemometer on every device. It is natural and conventional to think in this way: if we know the regional wind, we can estimate the local convective cooling.

Limitations: Why this Fails Indoors

This framework allows anyone to estimate heat flux and solar radiation in uniform natural environments as deserts, mountains, riversides, and urban cityscapes.

However, for home or indoor use, this framework must be waived. The indoor environment completely violates the "Proximity Correlation." An air conditioner or a greenhouse creates a microclimate where the air movement has zero correlation to the wind outside. Inside a house, the air might be still, while outside it is storming, or vice versa. Therefore, the "general proximity" assumption only holds true for the outdoors.

We can also use cloud based weather data (from trusted weather APIs) to estimate the local wind speed and adjust h_c accordingly for residual correction after differential cancellation.

** www.fiaos.org/about, The term *FiaOS* is employed as a sample project, serving as a model system for clarifying the use cases and applications detailed within this document. It may be regarded as a provisional placeholder.

Limitations and Scope

This approach has clear limitations:

- ☐ Works well for:
 - ☐ Open outdoor environments (deserts, fields, mountains, riversides)
 - ☐ Urban areas with relatively equal building heights
 - ☐ Locations where further micro climates are minimal
- ☐ Fails for:
 - ☐ Greenhouses
 - ☐ Locations with highly diverse or variable local obstacles

For indoor applications, the wind correlation must be disabled or replaced with a different model. The DTDSS framework is primarily designed for outdoor environmental monitoring where the Correlation Hypothesis(taking the general proximity) assumption holds.

The Dual compartment enclosure

To physically implement the differential model, a dual compartment design is optimal for several reasons.. The goal is to ensure both sensors experience the exact same wind and pressure conditions, differing only in their exposure to light.

Compartment A (The Reference)

A fully enclosed, white, or reflective cylinder. It allows free air circulation but blocks direct line-of-sight to the sky. This ensures S_{ref} remains at true air temperature. Although this is a standard practice, taken for preserving the longevity and the accuracy of the sensors/systems in many applications.

Compartment - B (Flux)

Topped with a transmissive, transparent material (generally, glass or polyethene) and backed by a black surface. The dome blocks wind but transmits solar radiation. The black housing absorbs the radiation and heats up predictably. This ensures S_{flux} maximizes absorption. This should be closed, ventilated too.

And the thermal mass of the two sensors should be identical. This symmetry ensures that "common mode noise", such as a sudden gust of cold wind, affects both sensors equally. When we subtract T_{ref} from T_{flux} , the wind noise cancels out, leaving only the solar signal.

The two compartments should be mounted in close physical proximity (on the same bracket or PCB assembly) to ensure they experience the same meteorological conditions.

A Note on Non-Terrestrial Applications

Applying this in Low Earth Orbit (LEO) or Sun Synchronous Orbit - SSO, the challenges would exponent and compound. In a vacuum, thermal cycling between sunlight and eclipse happens every 90 minutes, with temperature swings of hundreds of degrees.

There is no air for convection, and environmental rates of change are magnitudes higher, necessitating hardware like pyranometers rather than the soft-sensor approximations used in here. As we need precision over the amount of cost.

7. Calibration and Deployment

7.1. The Fundamentally Essential Parameters in the System

The accuracy of the DTDSS system depends on knowing several parameters:

| Parameter | Symbol | Typical Value | How to Determine |
|-----------------------|------------|---------------------------------|---------------------------------------|
| Thermal time constant | τ | 10 to 60 seconds | Auto calibration routine |
| Solar absorptivity | α | 0.85 to 0.95 | Literature value for housing material |
| Surface area | A_s | $\sim 50 \text{ mm}^2$ | Geometry measurement |
| Heat capacity | mC_p | $\sim 1 \text{ J/K}$ | Calculation measurement |
| Self-heating rise | T_{rise} | $0.5\text{--}2.0^\circ\text{C}$ | Measures in dark and still conditions |

7.2. Auto Calibration of the Thermal Time Constant

The thermal time constant τ is the most prominent parameter and the most likely to vary between applying (this is depending on housing, mounting, and local conditions).

Fortunately, τ can be determined automatically via a step response test:

- Wait for a natural "step down" event (sunset, or sudden shading).
- Record the temperature decay curve of the Flux sensor.
- Fit an exponential: $T(t) = T_{final} + (T_{start} - T_{final}) \cdot e^{-t/\tau}$.
- Extract τ from the fit.

For embedded systems, a simplified approach:

- Identify when $\frac{dT}{dt}$ becomes consistently negative (cooling started).
- Measure the time to reach 63% of the total temperature drop (this is τ by definition of the exponential time constant).

The system can update τ periodically (weekly or monthly) to account for:

- ☐ Dust accumulation on the housing.
- ☐ Changes in ventilation.
- ☐ Seasonal changes in typical wind speed.

7.3. Self-Heating Characterization

The self-heating term T_{rise} depends on:

- ☐ Sensor sampling rate (higher rate = more power dissipation).
- ☐ Ambient convection conditions.

To measure it:

- Place both sensors in a dark, controlled environment (no solar radiation).
- Shield from significant air movement.
- Allow the system to reach thermal equilibrium.
- Compare the difference between Flux and Reference sensors.

The residual difference (after accounting for any housing differences) is the self-heating offset. Repeat at different sampling rates to create a correction table.

7.4. Altitude Rejection

By implementing the density equations from Section 2 in firmware, as discussed in the past, as well as in the introduction, the sensor automatically adapts to altitude.

Example:

- ☐ Deployment: Weather Balloon flying approx. 2000m above the Knuckles Range
- ☐ Measured pressure: $P = 600$ hPa.
- ☐ Calculated density: $\rho \approx 0.7$ kg/m³ (vs. 1.2 kg/m³ at sea level).
- ☐ Algorithm behavior: Reduces expected h_c by ~40%. Expects the sensor to get hotter for a given W/m².
- ☐ Result: When the Flux sensor reports a high ΔT relative to Reference, the system correctly attributes it to normal solar levels, not a heat wave.

External firmware modifications, GPS or hardcoded changes will not be needed.

8. Implementation on Constrained Hardware

These algorithms are designed for Class-o IoT devices, microcontrollers with severe resource constraints like 8 Bit AVR, ARMo[zero], etc.

8.1. Fixed-Point Arithmetic

Floating-point math is slow and power-hungry on integer based MCUs. We convert all physical constants to fixed-point representations.

Example: Enthalpy calculation

Instead of:

```
float h = 1.006 * T; // Floating-point multiply
Use scaled integers (Q10.6 format):
```

Here, $1.006 \approx 1030/1024 = 1030 \gg 10$

```
int32_t h_dry = (1030 * T_raw) >> 10; // Integer multiply + shift
```

This replaces a floating-point multiplication with an integer multiply and bit-shift.

For an Example, *INR filter*

If $\alpha = 0.125 = 1/8$:

```
// output =  $\alpha$  * input + (1- $\alpha$ ) * output
// output += (input - output) *  $\alpha$ 

output += (input - output) >> 3; // Single shift, no division
```

8.2. Memory Footprint

The Differential Architecture requires minimal state retention:

Reactive Path State *per sensor*:

- ☐ Current temperature: 2 bytes.
- ☐ Previous temperature: 2 bytes.
- ☐ Previous derivative: 2 bytes.
- ☐ Alpha value: 2 bytes.
- ☐ Circular buffer (5 samples): 10 bytes.
- ☐ Total per sensor: ~18 bytes.

Reference Path State:

- ☐ Min temperature (1-hour window): 2 bytes.
- ☐ Max temperature (1-hour window): 2 bytes.
- ☐ Hourly humidity average: 2 bytes.
- ☐ Accumulated enthalpy change: 4 bytes.
- ☐ Window counters: 4 bytes.
- ☐ Total: ~14 bytes.

Overall, Less than 60 bytes of RAM for core algorithm state (for onboard, both sensors). This negligible footprint allows the algorithm to run alongside networking stacks without causing memory pressure.

8.3. Computational Complexity and Time Complexity

Comparing approaches:

TinyML (Neural Network):

- ☐ Requires storing weights for every neuron.
- ☐ A small 3-layer 10-20 to 1 network needs approximately 200 float weights.
- ☐ Inference involves matrix dot products: $O(N^2)$ complexity.
- ☐ Typical inference time: 10 to 100 ms on Cortex Mo.

DTDSS (Physics-Based):

- ☐ Constants stored in flash: ~100 bytes.
- ☐ $O(1)$ Complexity
- ☐ Typical execution time: <1 ms on 8-bit AVR.

The physics-based approach is many times more efficient in both memory and computation.

9. Validation and Expected Performance

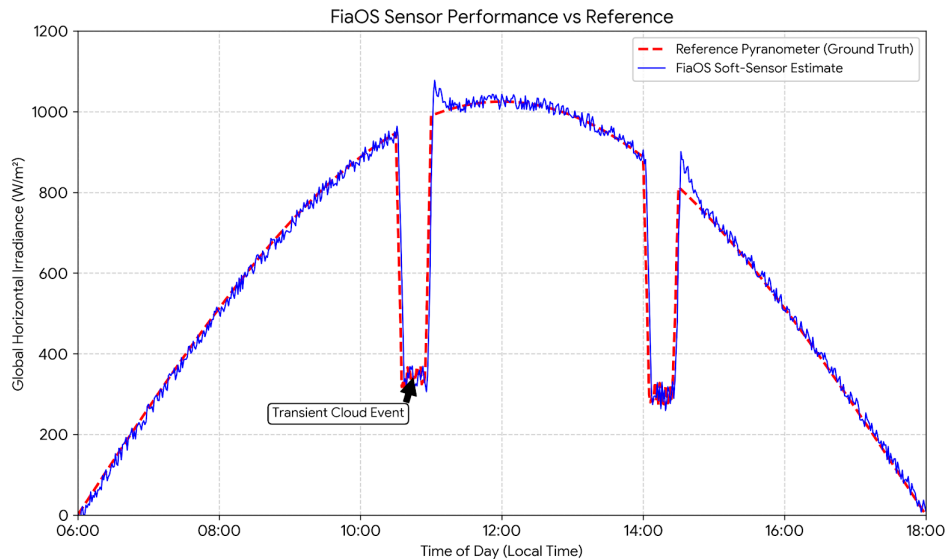


Figure 1.9.1

9.1. What We Can Measure

Direct validation of radiation accuracy using ground truth is currently impossible due to the absence of a collocated pyranometer. Although a planned experiment for a later version of this preprint will address this, we must, until then, employ alternative methods to validate our

-virtual concept. We can, however, assess the system's behavior using several indirect validation techniques.

Self-Consistency Checks:

- Does the Reactive Path output correlate with known solar geometry? (Higher at noon, zero at night).
- Do the two paths converge to similar values during stable conditions?

Cross-Deployment Comparison:

- Deploy identical sensor pairs at different altitudes.
- Verify that altitude-corrected readings are similar for similar conditions.
- Confirm that uncorrected readings show the expected systematic bias.

Differential Integrity Check:

- Verify that $T_{flux} \geq T_{ref}$ during daylight hours.
- Verify that $T_{flux} \approx T_{ref}$ at night and during heavy overcast.
- Check that the differential signal has lower noise than either raw signal.

9.2. Expected Accuracy

Based on the physics and typical sensor specifications:

Systematic Errors (can be measured out):

- ☐ Self-heating bias: Largely cancelled by differential measurement.
- ☐ Absorptivity uncertainty: $\pm 5\% \rightarrow \pm 5\%$ radiation error.
- ☐ Time constant uncertainty: $\pm 20\% \rightarrow \pm 20\%$ error during transients only.

Random Error from Possible Noise

- ☐ Temperature resolution: 0.01°C per sensor.
- ☐ Differential noise: $\sim 0.014^\circ\text{C}$ (root-sum-square of two sensors).
- ☐ After INR filtering: $\sim 0.007^\circ\text{C}$ effective differential.
- ☐ Corresponding radiation noise: $\sim 15\text{--}25 \text{ W/m}^2$ RMS.

Overall Expected Performance:

- ☐ Steady-state accuracy: $\pm 10\%$ of actual radiation (after calibration).
- ☐ Transient response: Detects cloud breaks within 5–10 seconds.
- ☐ Altitude range: Sea level to 5000m without modification.
- ☐ Common-mode rejection: Wind-induced errors reduced by $\sim 80\%$ compared to single-sensor approach.

This is not pyranometer grade accuracy. But it is far better than having no radiation data at all, and it comes at a fraction of the cost, approximately \$6 for two BME280 sensors versus \$600+ for a research grade pyranometer.

While direct side-by-side validation with a reference pyranometer is planned for future work, this section establishes the theoretical bounds and simulation performance of the framework

9.3. Limitations and Failure Modes

- Night-time: The algorithm relies on solar heating differential. At night, both sensors equilibrate, and the system outputs zero (approximate). No useful flux information available.
- Overcast conditions: When the sun is behind continuous clouds, the signal-to-noise ratio drops. The Reference Path dominates; the Reactive Path becomes noisy.
- Asymmetric wind exposure: If housing designs cause different wind responses, residual errors remain after differential cancellation.
- Condensation/frost: If moisture condenses on the Flux sensor dome but not the Reference housing, the thermal properties diverge. Readings become unreliable until the sensor dries.
- Infrared radiation: The model neglects long-wave infrared exchange (thermal radiation from sky and surroundings), Introduces new errors when sensor temperature differs significantly from sky temperature.

Failure Mode Recovery:

- ☐ Sanity checks: If calculated radiation exceeds clear-sky maximum by $>20\%$, flag as error.
 - ☐ Consistency checks: If Reactive Path and Reference Path disagree by $>50\%$ during stable conditions, trigger recalibration.
 - ☐ Differential sanity: If $T_{flux} < T_{ref}$ during daylight, flag sensor malfunction.
 - ☐ Range limiting: Clamp outputs to physically plausible values (0 to ~ 1400 W/m
-

10. Future Directions

10.1. Multi-Sensor Fusion

The current framework uses a two BME280 sensors.

Future versions could extend the same differential concept:

- ☐ Multiple Flux nodes with different surfaces (black coatings) to break down the light's color spectrum
- ☐ An accelerometer can notice if the sensor is moving or shaking (like an outdoor sensor swaying in the wind)
- ☐ A photodiode can offer a rough check on the sunlight (if it's sunny or cloudy)**

10.2. Machine Learning Augmentation

While pure-ML approaches aren't ideal as disclosed before, The role for ML in a hybrid system is:

- ☐ Use physics-based differential model as the primary estimator
- ☐ Train a lightweight ML model to predict model residuals (errors)
- ☐ Apply residual correction when ML model is confident

This gets the robustness of physics with the adaptability of ML.

10.3. Networked Calibration

With many DTDSS sensor pairs deployed in a region:

- ☐ Sensors can share calibration data
- ☐ Statistical outlier detection identifies malfunctioning units
- ☐ Ensemble averaging reduces noise
- ☐ Spatial correlation enables short-term forecasting

10.4. Infrared Extension

This requires additional computational complexity but could improve accuracy by 20–30%.

- ☐ Estimate sky temperature from humidity and cloud cover
- ☐ Model net radiative exchange (solar in minus thermal out)
- ☐ Account for reflected radiation from ground surface

** To show that we can figure out the sun's energy using common sensors, we will later separate the light we can see from the heat.

10.5. Scaling the Differential Principle

The differential architecture presented here uses two sensors. But the principle scales:

- ☐ 4 Sensors: Solar + Temp + Wind Direction by various differences across arrays.
- ☐ N Sensors: A larger area for a building which measures micro climates within it.

This provides the concept of "Accumulative" sensing. We are accumulating thermal potential between controlled reference points rather than reacting to a single noisy measurement.

11. Conclusion

“Can DDDS + INS extract meaningful solar radiation information from general purpose environmental sensors?”

The DTSS framework changes common IoT components into a sensing array, greater than the sum of its parts. By changing the perspective from a single-point measurement to a differential, topology based structure, we separate the measurement of the medium from the measurement of the energy.

A single sensor cannot parallelly or simultaneously be in thermal equilibrium with the air and in thermal disequilibrium with the sun. By using two sensors, one shielded Reference node and one exposed to the other flux node, It measures both states directly rather than trying to estimate one from the other.

The Inertial Noise Reduction filter makes differentiation more practical. It preserves the phase information we need while reducing the quantization artifacts that would make the signals error-prone or noisy.

And by calculating air density from pressure in real-time, the system becomes independent; which means, deployable from sea level to mountaintops without reconfiguration.

“Is this a replacement for a precision, industry grade pyranometer?”

No. The accuracy is lower, the noise is higher, and there are environmental conditions where the approach fails.

This approach is not limited to the BME280. It is a foundational architectural system applicable to any number of embedded sensor(s) as well.

A physics-based differential approach like DTDSS offers advantages in such scenarios:

- ☐ No training data required from the specific orbital environment
- ☐ Algorithms adapt automatically to changing conditions via first-principles physics
- ☐ Computational simplicity enables real-time processing on radiation-hardened, power-constrained processors
- ☐ The differential architecture provides inherent redundancy and fault detection

The framework system presented here, while developed for terrestrial applications, establishes foundational principles applicable to extreme environments where traditional sensing approaches become impractical or impossible.

References

- [1] Agriculture & Evapotranspiration (Penman-Monteith) R. G. Allen, L. S. Pereira, D. Raes, and M. Smith, "Crop evapotranspiration - Guidelines for computing crop water requirements - FAO Irrigation and drainage paper 56," FAO, Rome, vol. 300, no. 9, p. D05109, 1998. Context: Establishes the standard for why solar radiation is critical for calculating water requirements.
- [2] Pyranometer Standards "ISO 9060:2018 Solar energy -- Specification and classification of instruments for measuring hemispherical solar and direct solar radiation," International Organization for Standardization, Geneva, Switzerland, 2018. [Online]. Available: <https://www.iso.org/standard/67464.html> Context: Defines the "Hardware Problem" and the cost/complexity of standard instruments.
- [3] Soft Sensing & Algorithms L. Fortuna, S. Graziani, A. Rizzo, and M. G. Xibilia, "Soft sensors for monitoring and control of industrial processes," Springer Advances in Industrial Control, 2007. DOI: 10.1007/978-1-84628-480-9. Context: Validates the concept of using software to estimate unmeasured physical variables.
- [4] Vapor Pressure Physics (Magnus Formula) O. A. Alduchov and R. E. Eskridge, "Improved Magnus form approximation of saturation vapor pressure," Journal of Applied Meteorology, vol. 35, no. 4, pp. 601–609, 1996. DOI: 10.1175/1520-0450(1996)035<0601:IMFAOS>2.0.CO;2 Context: Source for the specific constants (17.67, 243.5) used in your humidity calculations.
- [5] Signal Processing (Differentiation) A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," Analytical Chemistry, vol. 36, no. 8, pp. 1627–1639, 1964. DOI: 10.1021/ac60214a047. Context: The foundational mathematical basis for the Inertial Noise Reduction (INR) approach to taking derivatives.
- [6] Heat Transfer & Convection F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine, Fundamentals of Heat and Mass Transfer, 6th ed. John Wiley & Sons, 2007. ISBN: 978-0471457282. Context: The authoritative source for Nusselt, Reynolds, and Prandtl number correlations used in Section 2.5.
- [7] Cloud Modeling (Kasten-Czeplak) F. Kasten and G. Czeplak, "Solar and terrestrial radiation dependent on the amount and type of cloud," Solar Energy, vol. 18, no. 3, pp. 177–181, 1976. DOI: 10.1016/0038-092X(76)90043-6. Context: The empirical model used in your Reference Path to estimate baseline radiation from humidity.
- [8] Sensor Hardware Bosch Sensortec, "BME280: Combined humidity and pressure sensor datasheet," version 1.1, 2016. [Online]. Available: <https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme280-ds002.pdf> Context: Validation of the sensor resolution (0.01°C) and pressure capabilities.
- [9] Differential Sensing Concept Z. Samani, "Estimating solar radiation from temperature differences," Journal of Irrigation and Drainage Engineering, vol. 126, no. 4, pp. 265–267, 2000.

DOI: 10.1061/(ASCE)0733-9437(2000)126:4(265). Context: Academic precedent for using temperature differentials to approximate solar flux.