https://doi.org/10.62836/iaet.v5i1.0010

Global Science Publishing

IAET

Innovations in Applied Engineering and Technology

https://ojs.sgsci.org/journals/iaet

Article

Design and Simulation of Combined Coriolis and Thermal Flow Sensors

Jiatong Zhang

Monta Vista High School, Cupertino, CA 95014, USA

Abstract: Liquid flow measurement is important in automated systems that need to deliver repeatable amounts of liquid. A dispenser may need to handle both rapid filling and slow delivery, making measurement over a wide flow range important. This study combines Coriolis and thermal flow estimates to extend that range and improve precision where both sensors are useful. A common signal model links vibration timing difference and temperature rise to volume flow, while accounting for heating and tube pressure loss. With reference water, 10 W of effective heating and 15-reading averages, the predicted range meeting a 10% relative-standard-deviation target is 12.91 – 92.30 mL/s for Coriolis and 0.240 – 23.96 mL/s for thermal sensing. Together, the sensors cover a predicted measurement range of 0.240 – 92.30 mL/s (a maximum-to-minimum flow ratio of approximately 385:1). Simulations then evaluated the rules for selecting or combining the sensor estimates at 75 selected flow rates within this predicted range. All four rules, including switching, gave relative root mean squared error (RMSE) below 10% at these points. At 20 mL/s, uncertainty weighting reduces RMSE by 16.7% relative to the better individual sensor. Blending and joint likelihood give similar results. Tests with five sets of liquid properties show how specific heat capacity, density and viscosity change the measurement range and calibration requirements. Longer averaging reduces random scatter, but an incorrect calibration can leave repeated estimates centred on the wrong flow.

Keywords: Coriolis flow sensing; thermal flow sensing; sensor fusion; uncertainty; calibration; simulation

1. Introduction and Literature Review

1.1. Need for Flow Measurement

Flow measurement is important when a system needs to dispense a repeatable amount of liquid. For a small dispenser, the measured flow can be integrated over time to estimate delivered volume, so errors in the flow estimate can affect the final amount. Both thermal and Coriolis meters are already used for liquid-flow measurement, but they infer flow from different physical signals [1,2].

The useful range of a proposed meter depends on more than the sensing principle. Tube diameter and flow rate affect pressure loss, while heating, signal strength, and noise affect the uncertainty of the flow estimate. These limits need to be considered together before choosing a physical design.

1.2. Existing Thermal and Coriolis Sensors

A Coriolis meter vibrates a tube and uses the flow-induced phase difference between vibration signals to

Received: 14 October 2026; Accepted: 29 October 2026.

Corresponding: Jiatong Zhang (20090323az@gmail.com)

determine mass flow. A thermal meter uses heat transfer and the temperature response, or the power required to maintain a temperature difference, to infer flow [1, 2]. Thermal methods depend on the liquid's thermal properties and the sensor arrangement. Coriolis sensing measures mass flow directly, but conversion to volumetric flow requires density.

The two principles have also been combined experimentally. Lötters et al. integrated thermal and micro-Coriolis sensors on a chip and provided overlapping ranges that allow the thermal sensor to be calibrated against the Coriolis sensor. Their thermal channel measured a local temperature response in a microchannel [3].

Here, the thermal model uses the temperature increase between the inlet and outlet of a heated section of a millimetre-scale water-flow path. This arrangement allows the heating requirement and the pressure requirement to be examined together.

1.3. Motivation for a Combined Sensor

The motivation for combining the sensors is that their uncertainties change differently with flow. In the model considered here, thermal sensing is more useful at lower flow, while Coriolis sensing becomes more useful as its timing signal grows. Using the more suitable channel at each flow could extend the measurement range beyond that of either sensor alone. Where both channels are useful, combining their readings could also reduce uncertainty.

Uncertainty weighting gives more influence to the flow estimate with the smaller expected spread. It uses inverse-variance weights: each contribution is proportional to one divided by the variance of that estimate. When estimates are unbiased and their errors are independent, these weights minimise the variance of a linear combination whose weights add to one [4]. Appendix B gives the derivation used here.

The physical design determines whether that favourable region can be used. The liquid must flow through the tube within the pressure budget, and the heater must produce a measurable temperature rise. Calibration adds a further requirement: a thermal reading can be repeatable but consistently shifted away from the true flow. Such a shift can become important after averaging has reduced the random fluctuations.

1.4. Research Question and Objectives

The research question is: how far can Coriolis and thermal sensing extend each other's measurement range, what additional precision can combining their readings provide, and when is recalibration needed to retain these benefits?

The study first establishes how the two signals and their uncertainties change with flow, and identifies where each sensor meets a chosen uncertainty target within the heating and pressure limits. Switching, blending, uncertainty weighting and joint likelihood are then compared with two sensor baselines using the same simulated observations. The analysis tests both the combined measurement range and the additional precision obtained by combining readings. Changes in signal noise and liquid properties show how those results depend on the intended operating conditions. Finally, tests with an unchanged calibration and different averaging lengths identify when changing the liquid or reducing random noise requires a more accurate conversion from signal to flow.

2. Basic Models of the Sensors

2.1. Coriolis Sensing Principle and Model

The Coriolis flow meter consists of a single tube or parallel tubes that carry the liquid, a drive actuator, position sensors, and a transmitter. A small actuator maintains the vibration of the tube or tubes. When fluid flows through the vibrating tube, Coriolis forces cause a small twist in the tube. This creates a phase shift between the inlet and outlet sensors located on either side of the curve. This shift can be used to calculate the mass flow rate, with a larger shift indicating more flow [5].

The timing difference Δt is the separation between corresponding points in the two vibration signals, as illustrated in Figure 1. Flow through the vibrating tube produces this difference. It is not the time taken by liquid to travel between the sensors. A simplified relation used during the design calculations connects this timing

difference to mass flow and tube stiffness [6]:

$$\Delta t = \frac{8 r_g{}^2 \dot{m}}{K}$$

Here, r_g is half the spacing between the tube legs, $\dot{m}$ is mass flow in kg/s, and K is torsional stiffness in N m/rad. The timing difference in this equation is in seconds. Greater leg spacing increases the response, whereas a stiffer tube reduces it. This relation explains why mechanical design influences sensitivity. For the numerical sensor comparisons, a calibration coefficient X for a fixed device expresses that response directly in the selected signal and flow units:

$$\Delta t = XQ$$

Here, Q is in mL/s, the timing difference Δt is in μs, and X is in μs/(mL/s). The selected coefficient is 0.10 μs/(mL/s), so a flow of 20 mL/s produces a 2.0 μs timing difference. Calibration would determine this coefficient by comparing timing measurements with independently known flows. This reference-water coefficient remains fixed for water comparisons. Section 2.5 explains how it changes with density when the liquid changes. The illustrative geometry calculation and the selected calibration are listed separately in Appendix A.

2.2. Thermal Sensing Principle and Model

The thermal sensing model uses a heater located upstream and a temperature sensor located downstream. As water flows through the tube, it carries thermal energy from the heater towards the sensor. The temperature rise relative to the incoming water can be used to estimate flow [7]. With a fixed amount of heater power transferred to the water, a higher flow rate would produce a smaller temperature rise because the heat is spread across a greater volume of water, and vice versa.

The sensor comparisons use the temperature increase from the inlet to the outlet once flow and heating are steady. At each location, bulk temperature means a mean across the tube cross-section weighted by the flow through each part of it. Using this temperature difference in a heat balance connects heater power, liquid properties and measured temperature directly to flow. Appendix F.2 describes the empirical uncertainty model used during the initial design exploration and the low- and high-flow effects it was intended to represent. For steady flow with constant properties, the balance is [7]:

$$P_{\text{eff}} = \dot{m} c_p \Delta T = \rho c_p Q_{\text{SI}} \Delta T$$

Here, P_{eff} is the effective heater power transferred to the water in W, $\dot{m}$ is the mass flow in kg/s, ρ is density in kg/m^3, c_p is specific heat capacity in J/(kg K), and ΔT is the outlet-minus-inlet bulk temperature rise in K. The symbol Q_{SI} denotes the same volume flow as Q, expressed in m^3/s instead of mL/s. Since 1 mL equals 10^{-6} m^3, rearranging the heat balance for Q in mL/s gives:

$$\Delta T = \frac{A}{Q},\ A = \frac{P_{\text{eff}}}{\rho c_p \times 10^{-6}}$$

The coefficient A therefore has units K mL/s. For water at the reference conditions, 5 W gives A = 1.198 K mL/s, while 10 W gives A = 2.396 K mL/s. At 20 mL/s, the corresponding temperature rises are 0.0599 K and 0.1198 K. Doubling the effective heating power doubles the temperature signal at the same flow.

The selected temperature-rise range is 0.10–10 K. Its lower bound sets a minimum signal for thermal use, while its upper bound limits heating of the liquid. For the reference-water design with a 5 mm internal diameter and 0.50 m tube length, 5 W of effective heating places this range between approximately 0.120 and 11.98 mL/s. At 10 W, the corresponding range is 0.240–23.96 mL/s. These values use the power actually transferred to the liquid. Tube geometry sets the pressure limit and influences heat transfer in a physical device. The heat balance calculates the temperature rise from the effective heating power and the flow.

2.3. Signal Noise and Averaging

A measurement window contains N independent readings taken while the flow is steady. Their mean timing difference is denoted by y_C, and their mean temperature rise by y_T. The subscripts C and T identify the Coriolis and thermal channels. The simulation represents each window by:

$$y_C = XQ + \varepsilon_t, y_T = \frac{A}{Q} + \varepsilon_T$$

The random errors ε_t and ε_T have zero mean and are independent Gaussian variables. If individual readings have standard deviations σ_t and σ_T, the standard deviations of the averaged timing and temperature readings are $\frac{\sigma_t}{\sqrt{N}}$ and $\frac{\sigma_T}{\sqrt{N}}$, respectively [8]. Each simulated window is represented by one timing value and one temperature value generated with these reduced standard deviations.

The reference single-reading values are σ_t = 0.50 μs and σ_T = 0.020 K. The latter describes noise in the measured temperature difference. It is distinct from the 0.10 K minimum accepted temperature rise: one describes random scatter, while the other defines when the thermal signal is used. Averaging 15 independent readings reduces the signal standard deviations to approximately 0.129 μs and 0.00516 K. The water comparisons use N = 1, 15, and 150. A fixed calibration error is represented separately as a change in the signal conversion. It remains fixed across the readings in a test condition, so averaging reduces random scatter while leaving that systematic shift. Section 5.2 tests an extra source of random temperature error, and Sections 5.4 and 5.5 examine calibration shifts.

2.4. Converting the Signals into Flow and Uncertainty

Each channel is converted into a flow estimate by reversing its calibration relation:

$$Q_C = \frac{y_C}{X}, Q_T = \frac{A}{y_T}$$

Both estimates are now in mL/s and can be compared or combined. A timing error is divided by X, so the Coriolis flow standard deviation is constant at fixed averaging. Inverting the temperature signal is nonlinear: the same temperature error causes a larger flow error when the temperature rise is small. First-order uncertainty propagation gives [8]:

$$s_C = \frac{\sigma_t}{X\sqrt{N}}, s_T \approx \frac{\sigma_T Q^2}{A\sqrt{N}}$$

The symbols s_C and s_T describe standard deviations of flow estimates in mL/s, whereas σ_t and σ_T describe noise in the original signals. The thermal expression is a local approximation around a specified flow. It shows why thermal uncertainty grows approximately as the square of flow. Relative standard deviation is obtained by dividing either flow standard deviation by Q and multiplying by 100.

The uncertainty curves meet at a flow Q_x:

$$Q_x = \sqrt{\frac{A\sigma_t}{X\sigma_T}}$$

At this crossover, neither sensor has a precision advantage over the other, making it a useful place to investigate combination. Averaging both channels by the same factor changes the size of their uncertainties but leaves the crossover unchanged. The switching threshold is selected separately using validation measurements.

2.5. How Liquid Properties Enter the Model

Changing the liquid changes both the thermal response and the conversion from mass flow to volume flow. For a fixed Coriolis mass-flow sensitivity, the volume-flow coefficient scales with density. The thermal coefficient depends on density and specific heat capacity:

$$X_l = X_w \frac{\rho_l}{\rho_w}, A_l = \frac{P_{eff}}{\rho_l c_{p,l} \times 10^{-6}}$$

The subscript l identifies a liquid and w identifies reference water. At the same volume flow, a denser liquid carries more mass per second and therefore produces a larger timing signal in this fixed-sensitivity model. Density multiplied by specific heat capacity is the volumetric heat capacity: the heat needed to raise a unit volume of liquid by one kelvin. A smaller value gives a larger temperature rise for the same power and volume flow. Dynamic viscosity has another role: it changes Reynolds number, friction and the pressure required to drive that flow. These effects must be evaluated together before deciding whether a liquid offers a wider measurement range.

The liquid comparison uses five sets of properties for water, milk, ethanol, glycerin and mineral oil, listed in Appendix A. Each set is held constant during a test. Thermal conductivity is used in the early exploration of heat conduction described in Appendix F.2. The steady heat balance uses density and specific heat capacity, while the pressure-loss calculation uses density and viscosity.

3. Design of the Combined Sensor

3.1. Proposed Physical Arrangement

The proposed design places Coriolis and thermal sensing sections along the same liquid flow path (Figure 1). The Coriolis section uses an actuator to vibrate the tube and two vibration sensors to measure the flow-dependent timing difference. Downstream, a heater adds heat to the liquid and a temperature sensor measures the response 20 mm away. An upstream temperature sensor provides the incoming temperature needed to calculate the temperature rise. A signal-processing unit converts the signals to flow estimates and combines the readings.

The inlet and outlet sensors are assumed to read the flow-weighted mean temperatures defined in Section 2.2. Both sensing sections are evaluated at the same steady flow. The 20 mm separation specifies the schematic layout, while the steady heat balance uses the effective heating power and the inlet-to-outlet temperature difference.

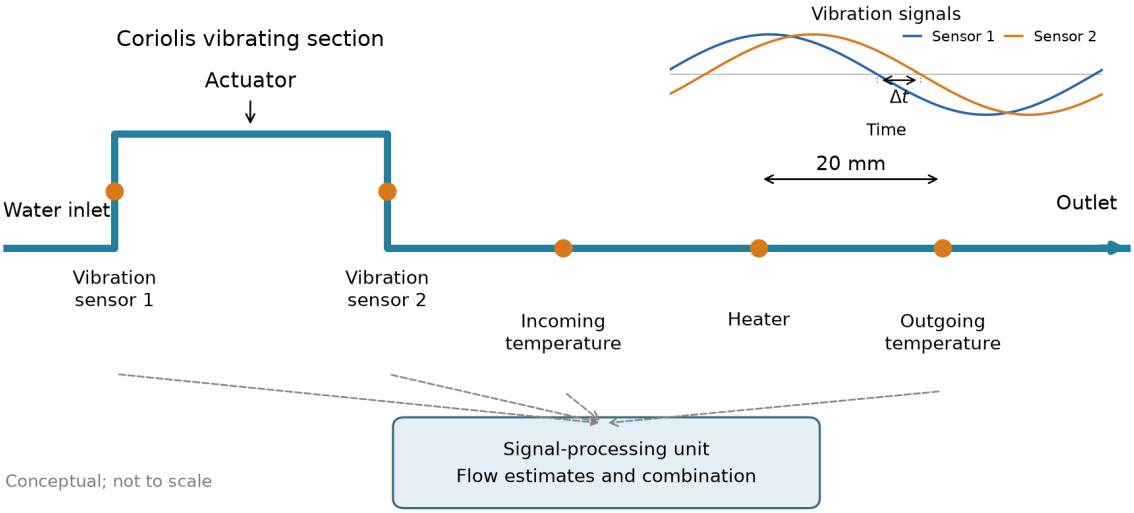


Figure 1. Conceptual combined flow meter, not to scale. The blue path indicates liquid flow, and dashed arrows carry readings to the signal-processing unit. The Coriolis section measures tube vibration, and the inset illustrates the timing difference Δt between its two sensing positions. The thermal section measures the outgoing temperature relative to the incoming temperature after heat is added.

3.2. Pressure and Velocity Requirements

The simulation uses a circular tube with an inner diameter of 5 mm and length of 0.50 m. Pressure drop is calculated using the Darcy–Weisbach equation, with additional losses included for bends, the inlet, and outlet [9]:

$$\Delta P = (f\frac{L}{d} + K_{\text{loss}})\frac{\rho v^2}{2}$$

Here, L is tube length, d is internal diameter, v is mean water speed, f is the Darcy friction factor, and K_{loss} is the combined minor-loss coefficient. The velocity follows from flow divided by the tube cross-sectional area. Reynolds number (Re) is used to identify the flow regime and calculate friction. The model uses laminar friction below Re = 2300 and the Haaland approximation above Re = 4000 [10]. Between these bounds, the simulation gradually blends the two equations. Appendix A gives the equations and parameter values.

The maximum allowed pressure drop is 0.5 bar, and the maximum mean speed is 10 m/s. The allowable upper flow is whichever of these two requirements is reached first. Diameters of 2 and 10 mm are also evaluated with the same length, roughness, and minor-loss coefficient. This comparison isolates hydraulic effects while the sensor coefficient X remains fixed.

3.3. Sensor Fusion

The two channels produce flow estimates with different uncertainties. The following rules decide which estimate to use, how to combine them when both are informative, and when to fall back to Coriolis because the temperature signal is outside its accepted range.

3.3.1. Switching

Switching selects the thermal reading at or below a chosen flow threshold and the Coriolis reading above it. This method is simple, but an unsuitable threshold can switch to the less reliable sensor. Here, the flow estimated by the thermal sensor, Q_T, is used to decide which side of the threshold the measured flow lies on. A separate validation dataset is used to choose the threshold. The selected rule is then evaluated on the test dataset.

Every method that uses the thermal channel first checks whether the observed rise y_T is between 0.10 and 10 K. If it falls outside this range, the method uses the Coriolis estimate. This common fallback keeps weak or excessive temperature signals from controlling the output. At flows below the lower thermal-flow limit, maintaining the temperature-rise limit would require reducing or switching off the heater. For reference water, estimates are restricted to $0.01 – 92.30$ mL/s, using the nearest endpoint when an estimate lies outside this interval. When calibration is updated for a different liquid, the upper output bound is set by that liquid's pressure-loss and velocity limits.

3.3.2. Blending and Uncertainty Weighting

Blending avoids an abrupt switch by gradually increasing the Coriolis contribution as flow rises and decreasing the thermal contribution. A smooth S-shaped curve, called a logistic transition, determines the thermal weight from the measured flow Q_T. Its midpoint and width are selected on the validation data. A wider transition uses contributions from both sensors over a broader range.

Uncertainty weighting instead gives more influence to the sensor with the smaller estimated uncertainty at each flow. Both sensors contribute most meaningfully where their uncertainties are similar. When one has much larger estimated random uncertainty, the combination relies mainly on the other. With w_T denoting the thermal weight, the combined estimate is:

$$Q_F = (1 - w_T)Q_C + w_T Q_T$$

For independent unbiased estimates with fixed weights calculated from known variances, the inverse-variance weight and the resulting variance are:

$$w_T = \frac{s_C^{\,2}}{s_C^{\,2} + s_T^{\,2}}, s_F^{\,2} = \frac{s_C^{\,2} s_T^{\,2}}{s_C^{\,2} + s_T^{\,2}}$$

Here, s_F denotes the standard deviation of the combined flow estimate. The weight lies between zero and one. Equal uncertainties give each sensor half the weight. If the thermal standard deviation is twice the Coriolis standard deviation, the thermal weight is one fifth. In the implemented uncertainty-weighting rule, the thermal standard deviation s_T is calculated from the flow estimated by the thermal sensor, Q_T. A noisy temperature reading therefore changes both the estimated flow and the weight assigned to it. Repeated simulations evaluate the resulting flow error. Appendix B provides the calculation and the logistic blending equation.

3.3.3. Joint Likelihood Estimation from Both Signals

The previous methods first turn each signal into a flow estimate and then choose or combine the estimates. Joint likelihood provides another comparison: for each candidate flow q, it predicts the timing and temperature signals and compares them with the observed pair. The chosen flow minimises:

$$J(q) = \frac{(y_C - Xq)^2}{\sigma_t^{\,2}/N} + \frac{(y_T - \frac{A}{q})^2}{\sigma_T^{\,2}/N}$$

Each squared discrepancy is divided by the variance of the corresponding window-mean signal. A discrepancy in a less noisy signal therefore receives a larger penalty. Minimising this expression gives the flow most compatible with both Gaussian signal models [11]. The search uses the same flow bounds and thermal

fallback as the other methods. This comparison tests whether treating the nonlinear temperature conversion directly offers a useful improvement over the simpler combination rules.

3.4. Measurement Range and Precision Gain

Two questions are evaluated separately. First, over which flows is the measurement accurate enough? For the design calculation, a channel qualifies when its predicted standard deviation is no more than 10% of the true flow and its physical operating conditions are met. The combined measurement range includes every flow at which at least one channel qualifies. The implemented rules are then tested using relative RMSE, which is RMSE divided by the true flow. This must remain below 10%, with all simulated windows included whether the thermal reading is used or rejected. Second, how much additional precision can be obtained by combining the readings?

Combining two independent readings can reduce random uncertainty, but the size of the reduction depends on their relative precision. We define the predicted fractional gain g as:

$$g = 1 - \frac{s_F}{\min(s_C, s_T)}$$

For example, $g = 0.10$ means a 10% lower flow standard deviation than the better individual sensor at that same flow. The largest predicted gain is 29.29%, reached when both sensor standard deviations are equal. When one sensor is much more precise, adding the other brings little extra information.

A precision-gain interval contains the flows where the predicted gain reaches a chosen target and the temperature, pressure and velocity conditions are all satisfied. The expected temperature rise A/Q must be between 0.10 and 10 K. The main gain target is a 10% reduction in standard deviation relative to the better sensor at the same flow. This differs from the 10% relative-error target used to define the measurement range. Appendix B also reports gain targets of 5% and 20% and derives how the intervals change with heating power.

4. Baseline Simulation and Results

4.1. Baseline Simulation Setup

The baseline uses water at approximately 20 °C in a tube with a 5 mm internal diameter and a length of 0.50 m. The reference thermal model uses 5 W transferred to the water, and a 10 W case examines the effect of stronger heating. Geometry, liquid properties, and model coefficients remain fixed during each comparison. Appendix A lists the inputs and compares the reference-water properties with NIST data [12].

The test uses 66 flow points between 0.2 and 90 mL/s. At each point, 2000 independent measurement windows are generated for each of five random seeds, giving 10,000 windows in total. The six methods receive the same signal pairs. They comprise two baselines, Coriolis alone and thermal with Coriolis fallback, and four combination rules: switching, blending, uncertainty weighting and joint likelihood. The second baseline uses the thermal estimate whenever the temperature reading is accepted and otherwise uses Coriolis. It is labelled “Thermal + fallback” throughout the figures and tables. The true flow determines the expected signals and is used to calculate errors, while the estimation rules receive only the observed signals and fixed model parameters. The test sweep includes temperature signals outside the selected range to examine fallback behaviour. The conditions allowing continuous heating are determined separately using the temperature and pressure limits in Section 3.4.

Switching and blending are tuned using 41 logarithmically spaced validation flows from 0.23 to 87 mL/s, with 3000 windows per flow and a separate random seed. Their settings minimise average squared flow error over those validation points. This gives greater representation to low-flow intervals than a linearly spaced grid. The settings are then held fixed for the test. Appendix C records the grids, seeds, and selected parameters.

4.2. What the Error Measures Mean

For a fixed test flow, the mean estimate shows where the measurements are centred. Bias is that mean minus the true flow. Standard deviation (SD in equations and compact table headings) measures the spread around the mean. Root mean squared error (RMSE) measures the combined effect of bias and scatter relative to the true flow:

$$\mathrm{RMSE} = \sqrt{\frac{\Sigma(\hat{Q}_i - Q)^2}{M}}$$

Here, M is the number of repeated measurement windows, $\hat{Q}_i$ is the estimate in window i, and Q is the true flow. Each window averages N individual readings. With the standard deviation calculated using divisor M, the error separates exactly into:

$$\mathrm{RMSE}^2 = \mathrm{Bias}^2 + \mathrm{SD}^2$$

A narrow set of readings can still have a large RMSE if it is centred on the wrong value. When bias is small, RMSE and standard deviation are nearly equal. The results report mean, bias, and standard deviation alongside RMSE so that a consistent offset can be distinguished from random scatter. Repeatedly generating noisy signals and calculating the resulting estimates is a Monte Carlo evaluation [13]. It includes the nonlinear conversion from temperature to flow and the decision to accept or reject each thermal reading, whereas the standard-deviation formula in Section 2.4 approximates the response near a specified flow.

4.3. Tube Geometry and Individual Sensor Performance

A larger tube reduces pressure drop at the same flow. The baseline 5 mm tube reaches the assumed 0.5 bar pressure budget near 92.3 mL/s. The 2 mm tube reaches the pressure budget at approximately 9.43 mL/s, compared with 451.72 mL/s for the 10 mm tube. For all three diameters, the pressure budget is reached before the velocity limit. Figure 2 brings together the diameter comparison and the individual sensor uncertainties calculated from the common signal model.

The Coriolis timing signal becomes larger relative to its fixed signal noise as flow increases, so relative flow uncertainty falls. The thermal signal decreases with flow, causing relative uncertainty to rise. Their uncertainty curves cross at 17.30 mL/s for 5 W and 24.47 mL/s for 10 W. The selected thermal upper flows are 11.98 and 23.96 mL/s. The black curve in Figure 2b shows the wider coverage obtained by using thermal sensing at low flow and Coriolis at high flow.

4.4. Extending the Measurement Range

At the reference noise levels and with 15-reading averages, Coriolis meets the predicted 10% relative-standard-deviation target from 12.91 mL/s up to the pressure limit of 92.30 mL/s. At 10 W, thermal sensing meets the same target throughout its accepted temperature range, from 0.240 to 23.96 mL/s. Their measurement ranges overlap, giving a continuous combined measurement range of 0.240 – 92.30 mL/s (a maximum-to-minimum flow ratio of approximately 385:1, calculated as 92.30/0.240). Table 1 compares this with the 5 W case.

At 5 W, the two qualifying measurement ranges are separated by a small gap between 11.98 and 12.91 mL/s. These predictions use the expected signals. The implemented rules make their decisions using noisy observations, so their actual errors are checked in a separate simulation.

An additional measurement-range check evaluates 25,000 windows at each flow, extending the grid towards the temperature and pressure boundaries without changing the algorithm settings chosen for the baseline comparison. At 10 W, all four combination rules remain below 10% relative RMSE at all 75 tested flows from 0.241 to 92.2983 mL/s. The largest relative RMSE is 5.63% for switching and 5.04% for uncertainty weighting. Switching is therefore sufficient to obtain the wide range, while combining readings can improve precision where the thermal signal is accepted. At 5 W, a denser check between the two predicted measurement ranges gives peak relative RMSEs close to 10%, leaving much less margin. Appendix C.5 gives the boundary checks and sampling details.

4.5. Precision Gain from Combining the Sensors

At 5 W, the predicted gain grows as flow approaches the thermal upper limit. Its highest value within the selected thermal range is 9.82%, just below the chosen 10% target. At 10 W, the interval with at least 10% lower standard deviation than the better sensor is 17.03–23.96 mL/s. Both endpoints lie well below the 5 mm tube's pressure limit.

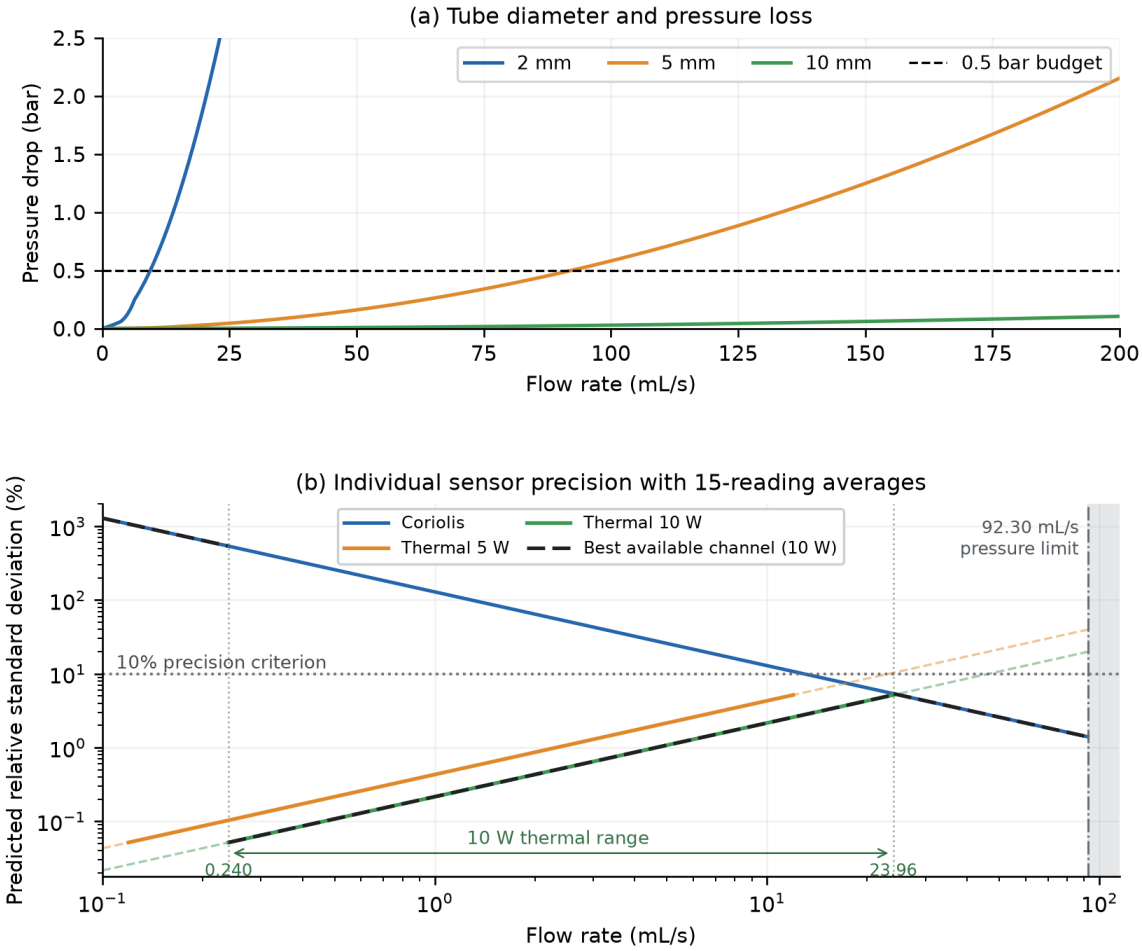


Figure 2. Geometry and individual sensor comparisons. (**a**) Pressure drop for 2, 5, and 10 mm tubes, with water, 0.50 m length, 1.5 μm roughness, and minor-loss coefficient 2.0. The horizontal line marks the 0.5 bar budget. The 10 mm crossing is beyond this plotted range. (**b**) Predicted relative flow standard deviation with 15-reading averages. Thermal curves are solid within the selected temperature range and dashed outside it. The black line follows the smaller standard deviation among the available channels at 10 W. It follows Coriolis when the temperature rise is outside the accepted range. The horizontal 10% line defines the measurement-range target, and grey shading marks flows above the 5 mm tube's 92.30 mL/s pressure limit.

Table 1. Predicted measurement ranges meeting a 10% relative-standard-deviation target for reference water, the 5 mm tube, and 15-reading averages. All ranges are in mL/s.

Channel/Measurement Range	5 W	10 W
Coriolis	12.91–92.30	12.91–92.30
Thermal	0.120–11.98	0.240–23.96
Either available channel	0.120–11.98 and 12.91–92.30	0.240–92.30

The gain changes continuously with flow and power. The 10% reduction is the selected comparison target, and Appendix B shows how choosing 5% or 20% changes the precision-gain interval.

The same analysis connects heating to tube diameter. For the reference noise levels, obtaining a 10% predicted reduction in standard deviation at any flow requires a flow of at least 12.11 mL/s and effective heating power of at least 5.05 W. The 2 mm tube reaches its pressure budget at 9.43 mL/s, below that minimum flow. The 5 mm and 10 mm tubes can reach flows within the precision-gain interval. Appendix B derives the combinations of power and flow that meet the gain target and the physical limits.

4.6. Comparing the Combination Rules

Figure 3 compares test RMSE across flow rates for 15-reading windows. At low flows within the thermal range, the thermal estimate provides most of the useful information. At high flows, the thermal-assisted methods increasingly use the Coriolis fallback. The enlarged 10 W view focuses on the intervening region, where combining accepted readings can improve the estimate.

The 11.98 and 23.96 mL/s markers are upper flows for the expected thermal signal. The rule instead checks each noisy temperature observation, so the transition to fallback is gradual. Curves continue beyond these markers because Coriolis still supplies an output: 40 mL/s is below the tube's 92.30 mL/s pressure limit, but above the selected thermal range at either power. These tails contribute to the wider measurement range through Coriolis fallback. Figure 4 identifies the smaller interval where using both accepted readings gives the specified extra precision.

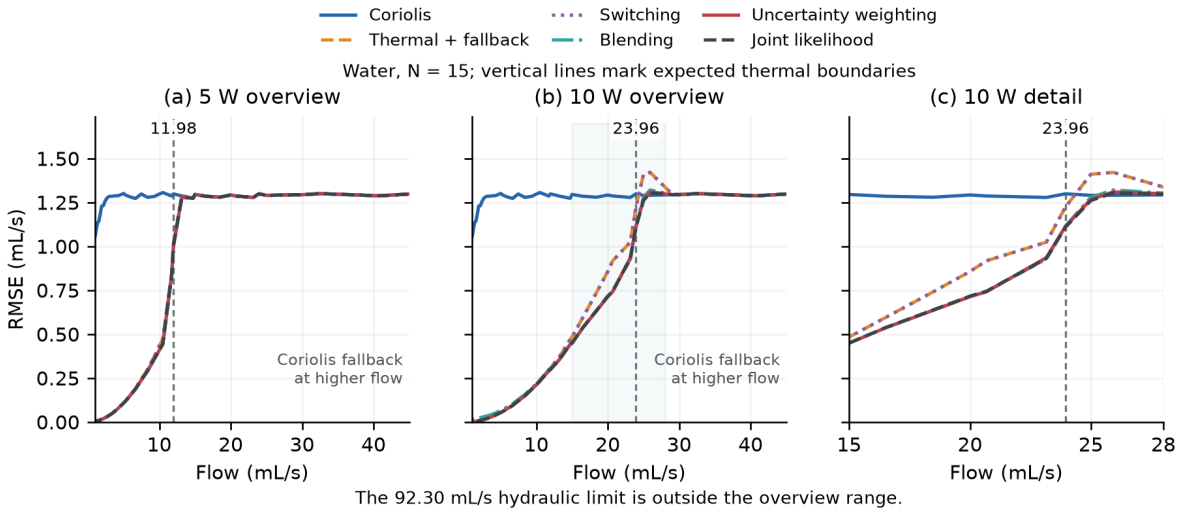


Figure 3. Test RMSE with 15-reading averages. The 5 W and 10 W overviews include flows where the thermal reading is rejected and Coriolis supplies the output. The enlarged 10 W view shows the region where combining accepted readings improves precision. Vertical boundaries mark expected temperature rises of 0.10 K, not pressure cutoffs. Each point pools 10,000 windows from five seeds.

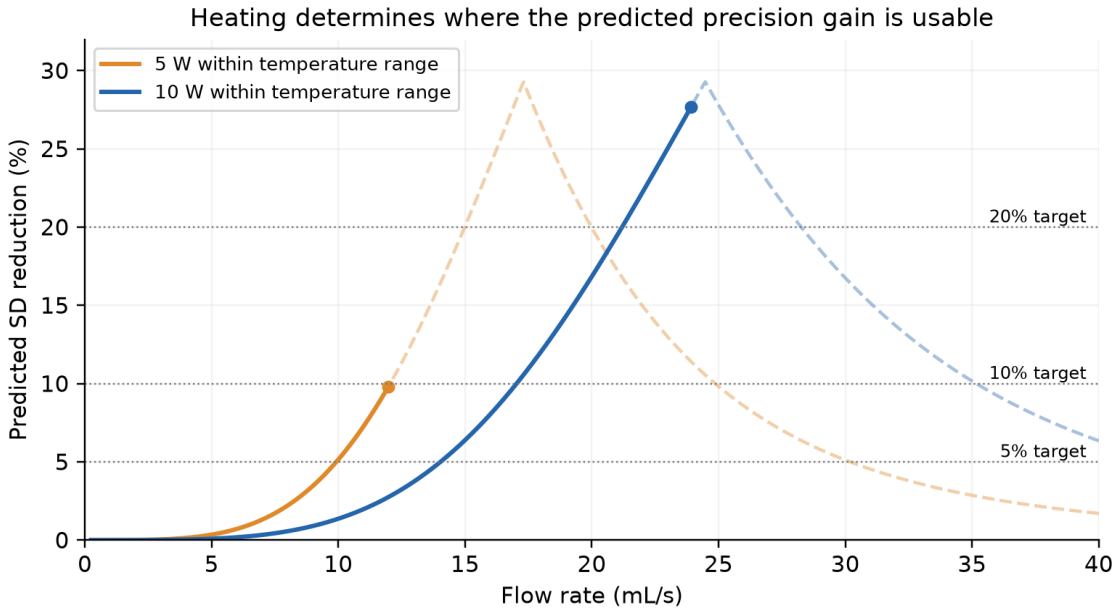


Figure 4. Predicted reduction in flow standard deviation relative to the better individual sensor. Solid portions satisfy the selected temperature, pressure and velocity limits. Dashed portions show the predicted reduction outside the temperature range. The 5%, 10%, and 20% reference lines show how the precision-gain interval depends on the design target. The largest reduction within the allowed 5 W thermal range is 9.82%. Increasing power keeps the temperature signal usable at higher flows, where both sensors can contribute.

At 20 mL/s, 10 W, and $N = 15$, all tested thermal windows pass the acceptance rule. This flow lies inside the predicted precision-gain interval and provides a representative comparison. Table 2 shows the complete error breakdown. The mean estimates are close to the true flow, and their biases are small relative to their standard deviations. Thermal sensing has a lower RMSE than Coriolis, while blending, uncertainty weighting, and joint likelihood reduce the remaining error further.

Table 2. Flow estimates at a true flow of 20 mL/s, 10 W, and 15-reading averages.

Method	Mean	Bias	SD	RMSE
Coriolis	19.994	−0.006	1.293	1.293
Thermal + fallback	20.042	+0.042	0.857	0.858
Switching	20.042	+0.042	0.857	0.858
Blending	20.006	+0.006	0.715	0.715
Uncertainty weighting	19.997	-0.003	0.714	0.714
Joint likelihood	20.020	+0.020	0.717	0.718

All numerical columns are in mL/s. Each method uses 10,000 windows.

Uncertainty weighting gives an RMSE of 0.714 mL/s, compared with 0.858 mL/s for Thermal + fallback. Across the five runs with different random seeds, each comparing the methods on the same observations, the mean reduction is 16.7%. The 95% confidence interval for this mean improvement is approximately 15.8 – 17.7%, calculated from the variation across the five simulation runs. Blending and joint likelihood give RMSEs of 0.715 and 0.718 mL/s. Their small separation indicates similar performance in this example, so algorithm complexity can be considered alongside error reduction.

Switching and Thermal + fallback coincide in these 15-reading comparisons because the selected switching rule uses every accepted thermal estimate and returns to Coriolis when the temperature reading is rejected. This explains the overlapping curves and the identical rows in Table 2. Appendix C.1 gives the selected thresholds, while Appendix C.3 explains how thermal acceptance changes when only one reading is used.

5. Sensitivity to Measurement Conditions

5.1. Timing Noise

The reference timing noise is a selected model input. Figure 5 shows how changing it moves the uncertainty crossover before considering the physical operating limits. To assess its effect on the findings, the single-reading timing standard deviation is varied from 0.25 to 0.50 and 1.00 μs, while temperature noise remains at 0.020 K. For each noise level, switching and blending are tuned on validation data and then tested with the same flow grid and repeated-seed procedure. The algorithms use the noise level assigned to that case.

Table 3 shows how timing precision changes the flow range in which combining the sensors reduces predicted standard deviation by at least 10% relative to the better individual sensor. With a timing standard deviation of 0.25 μs, 5 W already provides a range meeting this target. At 1.00 μs, even 10 W reaches a maximum reduction of only 9.82% within the selected temperature limits. The required minimum effective powers for the three timing-noise levels are approximately 2.53, 5.05, and 10.11 W, respectively, provided that the tube can carry the corresponding flows within the pressure and velocity limits.

The simulations with repeated noisy measurements support the same dependence. At 20 mL/s, 10 W, and $N = 15$, uncertainty weighting reduces RMSE relative to the better single-sensor baseline by approximately 19.7%, 16.7%, and 5.2% for timing standard deviations of 0.25, 0.50, and 1.00 μs. The better baseline is Coriolis in the first case and Thermal + fallback in the other two. Blending and joint likelihood remain close to uncertainty weighting across these cases. Appendix C gives results for all six methods.

Reducing timing noise lets Coriolis meet the relative-error target at lower flows and moves the predicted precision-gain interval towards lower flows.

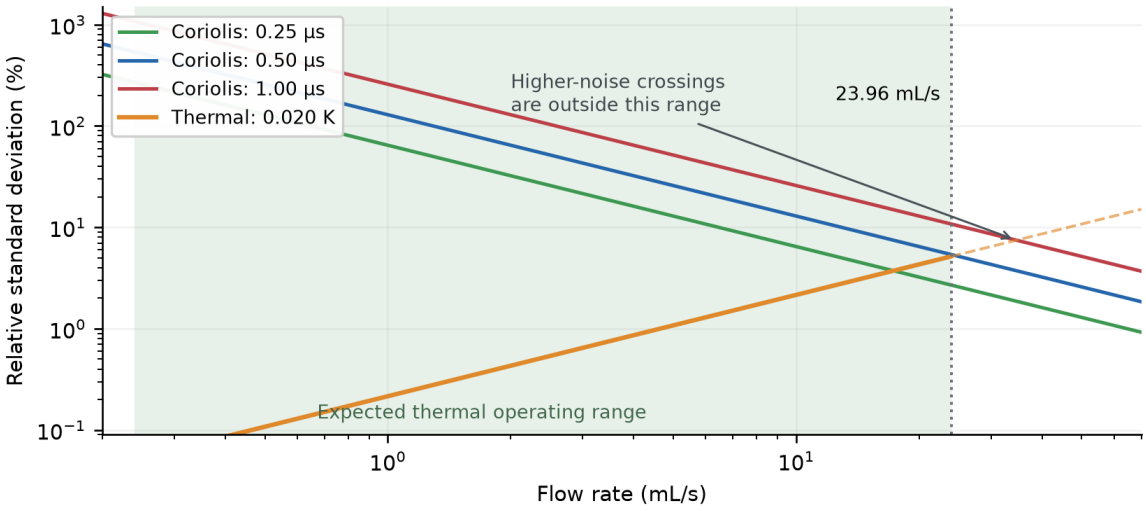


Figure 5. Predicted relative flow standard deviations at three single-reading timing-noise levels, with 10 W and 15-reading averages. The thermal curve is shared because temperature noise remains at 0.020 K. Shading marks the selected thermal range. At crossings above its 23.96 mL/s upper boundary, the expected temperature rise is below the 0.10 K minimum, so these crossings lie outside the predicted thermal operating range at this heating power.

Table 3. Timing-noise dependence of the predicted ⩾10% standard-deviation reduction interval relative to the better individual sensor, for a 5 mm tube.

Timing SD(µs)	5 W Interval(mL/s)	10 W Interval(mL/s)
0.25	8.52–11.98	12.04–23.96
0.50	Empty	17.03–23.96
1.00	Empty	Empty

5.2. Sensitivity to an Additional Random Error

The previous comparison assumes a fixed standard deviation in each temperature-difference reading. To test how strongly the results depend on that choice, an additional independent random error is introduced with standard deviation equal to 2.5% or 10% of the expected temperature rise. The added error has zero mean and is averaged once with the other reading noise. All algorithms continue to use the original calibration coefficients and noise estimates. Switching and blending also retain their original threshold and transition settings. This tests their behaviour when the measurements contain extra noise that the algorithms have not been told to expect. Figure 6 shows the resulting RMSE across flow rates.

At low flow, a temperature error proportional to the expected rise produces an approximately constant relative flow error. It therefore changes the shape of the uncertainty curve rather than simply shifting the original curve vertically. At the 20 mL/s, 10 W comparison point with 15-reading averages, the RMSE of uncertainty weighting rises from 0.717 mL/s without the added component to 0.722 and 0.795 mL/s for the 2.5% and 10% components. The RMSE of Thermal + fallback rises from 0.871 to 0.879 and 1.009 mL/s, while Coriolis remains at 1.300 mL/s. Fusion continues to reduce error in this example, but the absolute accuracy worsens. A larger percentage improvement over a degraded baseline would therefore not mean that the measurement itself had improved.

This additional-noise comparison uses a separate set of test flows and random samples, specified in Appendix F. At each test flow, all methods receive the same pairs of timing and temperature observations.

5.3. Different Liquids with Correct Calibration

The same tube and heater are next tested with water, milk, ethanol, glycerin and mineral oil. For each liquid, density and specific heat capacity determine the signal conversions, while density and viscosity determine pressure loss. Geometry and single-reading noise stay unchanged. In this comparison the liquid is known: each

estimator uses its liquid-specific calibration, and switching or blending is selected on separate validation data for that liquid. Appendix A.3 lists the selected model inputs. Published sources document composition-dependent milk properties [14], heat-capacity data for ethanol and glycerin [15 – 17], and grade-dependent mineral-oil properties [18].

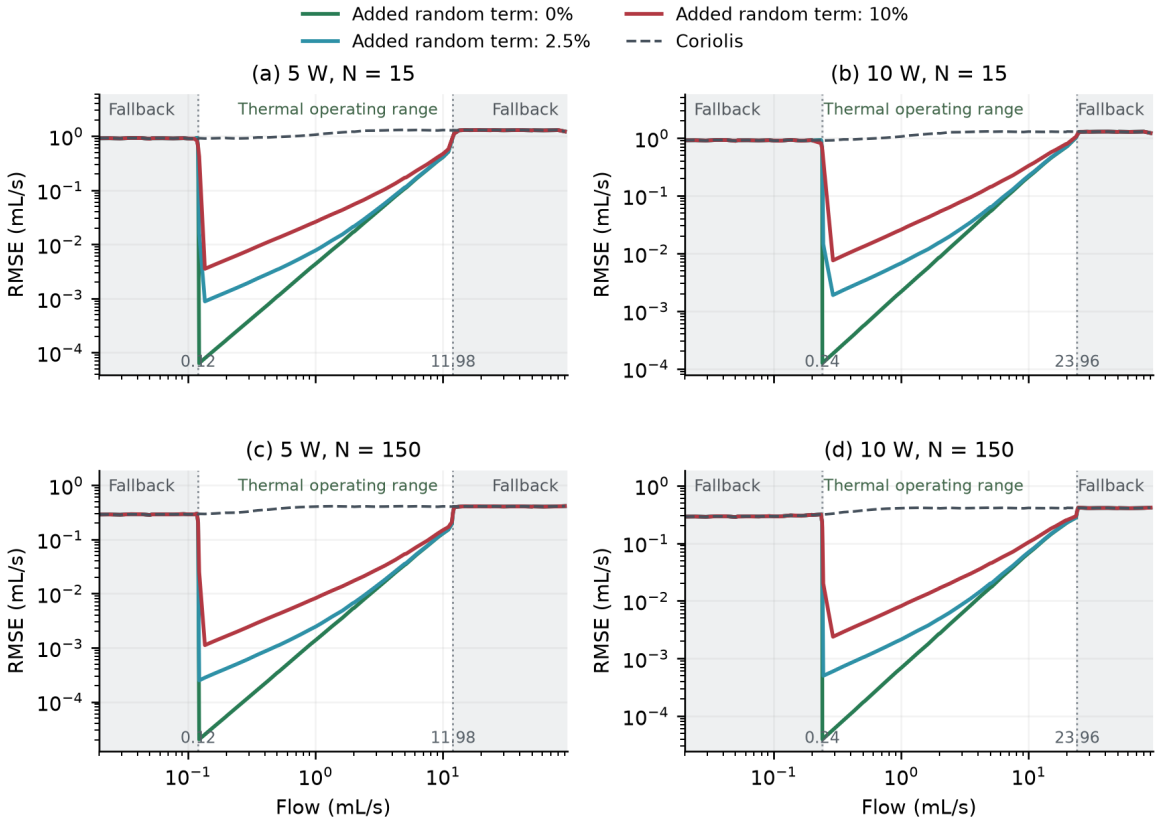


Figure 6. RMSE of uncertainty weighting with additional independent temperature-signal noise of 0%, 2.5% or 10% of the expected rise per reading. The 0% case still includes the original 0.020 K temperature noise and 0.50 µs timing noise. The dashed curve is Coriolis. When the observed temperature rise is outside 0.10–10 K, the algorithm rejects the thermal reading and uses Coriolis alone. Outputs from these windows remain in the RMSE calculation.

For ethanol, the product of density and specific heat capacity is about 46% of that of water. The same 5 W therefore produces a larger temperature rise and extends the upper flow for thermal sensing from 11.98 to 25.97 mL/s (Figure 7). After also updating the Coriolis volume-flow sensitivity, combination is predicted to reduce standard deviation by at least 10% relative to the better individual sensor from 19.95 to 25.97 mL/s. Thus the change in liquid extends thermal coverage and allows both sensors to contribute meaningfully at the lower heating power.

Glycerin is instead limited by pressure loss. Its selected viscosity is roughly 1500 times that of water. The 0.5 bar pressure budget limits flow to 1.03 mL/s. Considering sensor uncertainties alone, a predicted standard-deviation reduction of at least 10% would begin at 12.57 mL/s with 5 W, far above that pressure-limited flow. Within the flows that the tube can carry and that meet the temperature-rise limits, thermal sensing is much more precise and the Coriolis timing signal contributes very little additional information. Increasing heating to 10 W cannot overcome the pressure limit.

For mineral oil, the upper flow at which the thermal sensor can be used is limited by heating at 5 W and by pressure loss at 10 W. At 10 W, its expected temperature rise remains acceptable up to 58.82 mL/s, but the pressure limit is reached at 53.59 mL/s. The tube therefore sets the upper flow for thermal sensing in this case. Milk remains closer to water: the flow interval giving at least a 10% predicted standard-deviation reduction from combination is narrow at 5 W and wider at 10 W.

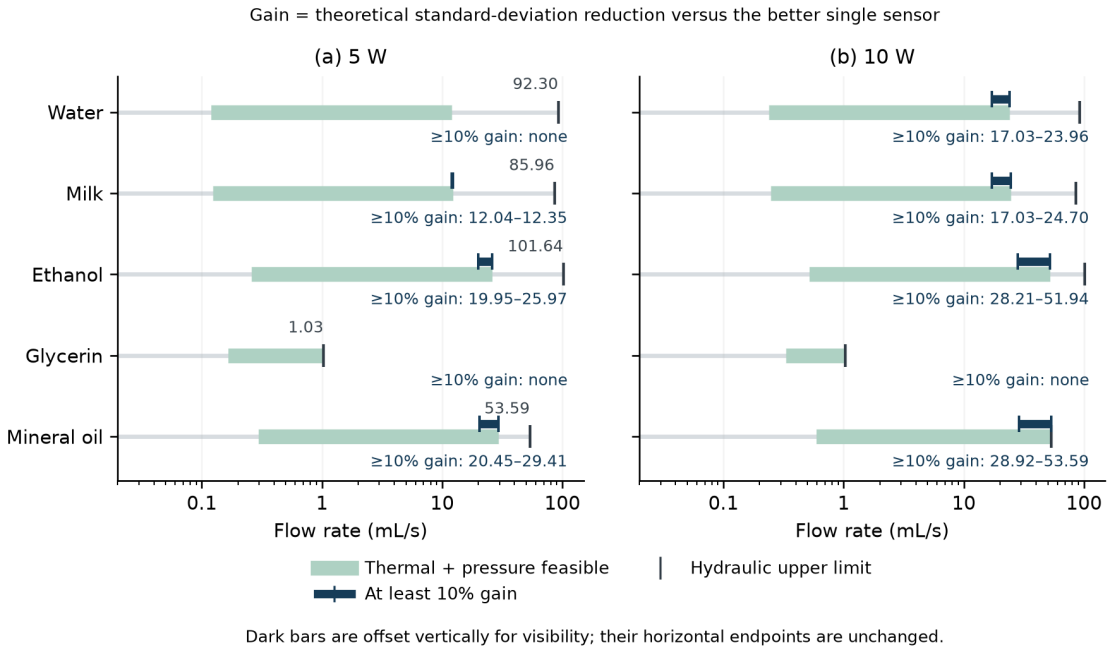


Figure 7. Liquid-dependent operating ranges for the fixed 5 mm tube. Pale-green bars show flows that satisfy both the thermal temperature-rise limits and the tube's pressure and velocity limits. The upper flows allowed by the tube are marked separately. Offset dark segments and their numeric labels show where combining the sensors reduces predicted flow standard deviation by at least 10% relative to the better individual sensor. At 5 W, milk meets this target only from 12.04 to 12.35 mL/s, while water and glycerin have no qualifying interval. Segment widths follow the logarithmic flow axis.

For ethanol at 20 mL/s, 5 W and 15-reading averages, Thermal + fallback gives an RMSE of 0.801 mL/s and uncertainty weighting gives 0.715 mL/s. At 10 W these become 0.399 and 0.387 mL/s. Stronger heating improves absolute accuracy but leaves less extra information for Coriolis to contribute at this flow.

Appendix F.3 shows the relative RMSE of uncertainty weighting across each liquid's attainable flow range and compares all six methods. Where liquids share the same test flow, they are compared at the same heating power and number of averaged readings. For each liquid, all methods receive the same pairs of timing and temperature observations.

5.4. Changing Liquid and Retaining an Earlier Calibration

The preceding results use the correct conversion coefficients for each liquid. Changing liquid without updating those coefficients can instead shift the mean readings. Two cases separate the effects. Coriolis-only recalibration updates the Coriolis density conversion and the estimator's upper flow bound for the new liquid's pressure limit, while retaining the water thermal coefficient. Unchanged water calibration retains both water signal coefficients, the water output bounds and the switching and blending settings previously selected for water. In both cases, the simulated signals and allowable test flows follow the actual liquid's properties and pressure limit, even when the estimator retains the water output bounds.

For example, at a true ethanol flow of 20 mL/s and 10 W, unchanged water calibration gives a noise-free Coriolis output of approximately 15.81 mL/s, while the thermal output is 9.22 mL/s. The Coriolis shift follows the ratio of ethanol density to water density. The thermal shift instead follows the product of density and specific heat capacity for ethanol divided by the same product for water. Figure 8 shows how these incorrect conversions affect the combined estimate.

With 15-reading averages and correct calibration, uncertainty weighting gives a mean of 20.01 mL/s and an RMSE of 0.387 mL/s. Unchanged water calibration gives a mean of 9.36 mL/s and an RMSE of 10.64 mL/s. Coriolis-only recalibration leaves the combined result near 9.36 mL/s: the algorithm still treats the incorrect thermal estimate as highly precise and gives it most of the weight.

With unchanged water calibration, increasing averaging to 150 readings reduces the standard deviation of

the combined estimate from 0.192 to 0.061 mL/s, but its mean remains near 9.35 mL/s. The repeated readings become more consistent around the wrong value. Appendix D gives the conversion errors expected without random noise, and Appendix F.3 records representative outputs for all five liquids.

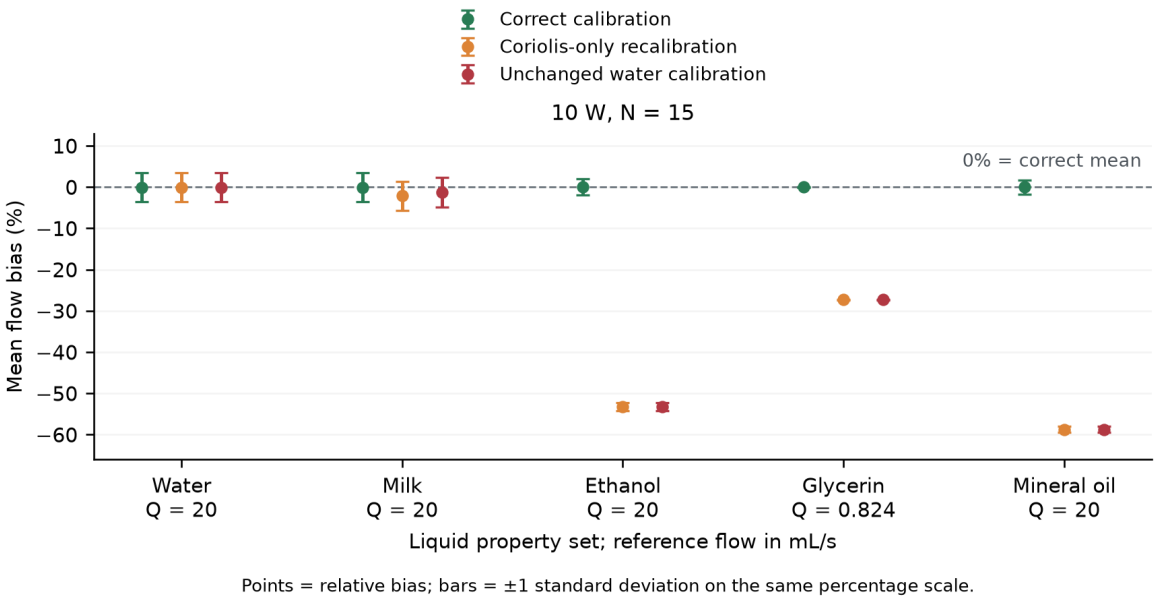


Figure 8. Uncertainty weighting after changing liquid at 10 W with 15-reading averages. Points show signed relative bias, and error bars show ±1 standard deviation on the same percentage scale. Zero means that the mean estimate equals the true flow. The three cases are Correct calibration, Coriolis-only recalibration and Unchanged water calibration, as defined in Section 5.4. Glycerin uses 0.824 mL/s, while the other liquids use 20 mL/s. Appendix F gives full settings and numerical results.

Recalibration is needed when the old conversion no longer meets the required accuracy. At the representative flows in Table 4, unchanged water calibration gives a relative RMSE of 3.84% for milk, within a 10% target. The ethanol, glycerin and mineral-oil cases give 53.22%, 27.26% and 58.74%, respectively, so they require an updated conversion to meet that target. Correct calibration brings all five cases below 4%.

Table 4. Calibration decisions at 10 W with 15-reading averages using uncertainty weighting. Relative RMSE is 100 × RMSE/true flow.

Liquid	True Flow (mL/s)	Unchanged Water CalibrationRelative RMSE (%)	Meets 10% Target?
Water	20.000	3.59	Yes
Milk	20.000	3.84	Yes
Ethanol	20.000	53.22	No
Glycerin	0.824	27.26	No
Mineral oil	20.000	58.74	No

At fixed effective power, the thermal conversion depends on density times specific heat capacity, while the Coriolis volume conversion depends on density. A viscosity change alters the attainable flow through pressure loss. Updating the calibration and checking the pressure limit therefore address different consequences of changing liquid. The next test examines a smaller change in thermal response to find when recalibration is needed to retain the precision benefit of fusion.

5.5. Why Averaging Changes the Calibration Requirement

This test returns to reference water and varies the thermal response while keeping the Coriolis calibration

correct. At fixed heating power and volumetric flow, a smaller product of density and specific heat capacity gives a larger ideal temperature rise, as expressed by the coefficient A. Changes in effective heating power, density, specific heat capacity, or the gain of the temperature readout can therefore alter the conversion from temperature to flow. Their combined effect is represented as a fractional change δ in the thermal response:

$$A_{\text{true}} = (1 + \delta)A_{\text{cal}}$$

The true thermal signal is generated using A_{true}, while the estimator continues using the calibration value A_{cal}. The test changes the thermal response alone and keeps the Coriolis calibration fixed, isolating the effect of error in the thermal conversion. A positive mismatch means that the observed temperature rise is larger than the calibration predicts at the same true flow. Dividing the unchanged coefficient by this larger rise makes the thermal estimate too low. In the absence of random noise, its flow bias is:

$$b_T = \frac{-\delta Q}{1 + \delta}$$

For example, a +5% response mismatch at a true flow of 20 mL/s gives a thermal estimate of 20/1.05 = 19.05 mL/s, a bias of approximately −0.95 mL/s. The mismatch is held fixed throughout each test condition, while random noise is generated independently for each window. This separates a systematic calibration shift from random measurement scatter [19].

Averaging reduces random fluctuations, while a fixed calibration shift remains. Figure 9 shows how this changes both the spread and the total error. With correct calibration at this flow, increasing N from 15 to 150 readings reduces RMSE for all six methods. With a +5% thermal response mismatch at 20 mL/s and 10 W, the RMSE of uncertainty weighting improves from approximately 0.985 to 0.731 mL/s when N increases from 15 to 150. Over the same change, Coriolis RMSE improves from approximately 1.293 to 0.409 mL/s. The combined result becomes more accurate in absolute terms but loses its advantage over Coriolis because the latter improves more.

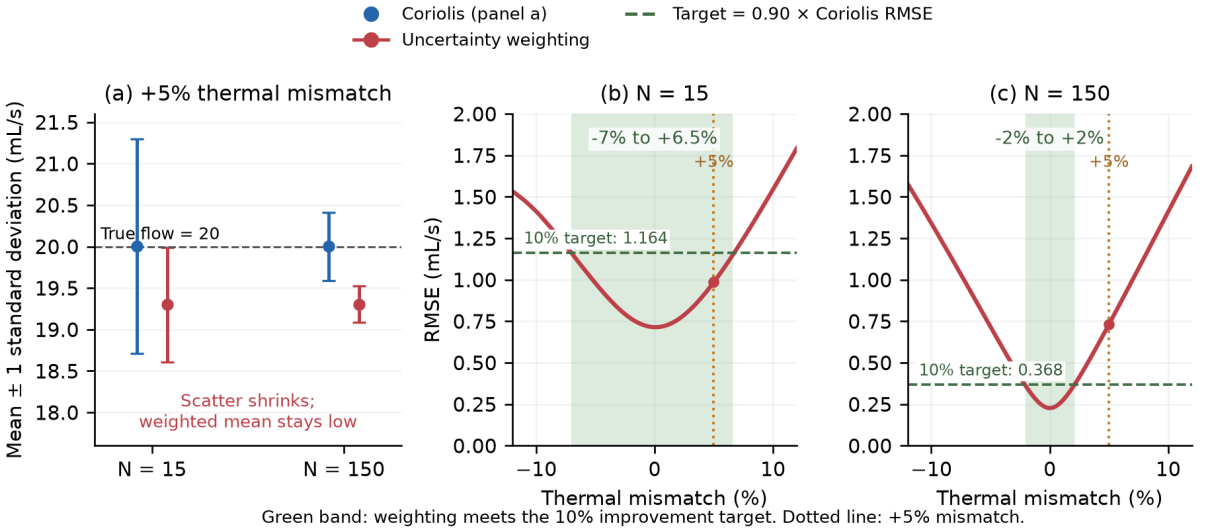


Figure 9. Averaging and thermal response mismatch at 20 mL/s and 10 W, pooling 25, 000 windows per condition from Appendix C.2. (**a**) At +5% mismatch, uncertainty weighting remains centred below the true flow even as its error bars shrink. (**b**, **c**) The 15- and 150-reading panels compare its RMSE with 90% of Coriolis RMSE, the target for a 10% improvement over Coriolis. Shading marks the mismatch values that meet this target. The +5% markers show where the example in panel (**a**) lies on these curves. Figure A1 gives the corresponding joint likelihood results.

The calibration tolerance is found by starting at zero response mismatch and moving in either direction until the combined estimate no longer retains an RMSE at least 10% below that of Coriolis. Coriolis is the reference here because it is unaffected by the imposed thermal calibration change: this asks when adding the thermal channel remains worthwhile. At 20 mL/s and 10 W, uncertainty weighting meets this target from −7.00% to +6.50% mismatch for N = 15. Joint likelihood meets it from −6.50% to +6.75%. With N = 150, both ranges narrow to approximately −2.00% to +2.00%. Mismatch is tested in steps of 0.25 percentage points. Appendix C

reports how these endpoints vary across repeated simulations with different random seeds.

With longer averaging, random scatter becomes smaller, so a fixed calibration error accounts for more of the total error. Keeping a 10% RMSE advantage over Coriolis therefore requires the thermal response coefficient to remain closer to its calibrated value. Appendix B derives how bias and random variance contribute to this condition. Appendix D examines what happens when sensor uncertainty is estimated incorrectly or when part of the heating power is lost to the surroundings.

A corresponding test that changes only Coriolis sensitivity also finds a narrower calibration tolerance with longer averaging. When both channels have a flow bias in the same direction, their combination remains biased in that direction, while opposite biases partially cancel but do not generally disappear. Appendix D. 5 reports these comparisons.

6. Discussion

This project began with the need to measure both slow and rapid liquid delivery using one flow path. The combined design uses thermal sensing for low-flow resolution and Coriolis sensing for higher flows. The results show how far that combination extends the measurement range, what extra precision is available when both sensors can be used, and how a change of liquid or calibration affects both.

6.1. What the Results Suggest

The main benefit is the wider measurement range shown in Figure 2. For the 10 W reference-water case, thermal sensing covers the low flows at which Coriolis has large relative uncertainty, while Coriolis continues above the thermal heating limit. The implemented switching and combination rules meet the 10% relative RMSE target at every tested flow from 0.241 to 92.30 mL/s. Within this range, the overlap provides a second benefit: uncertainty weighting reduces the remaining random error when both readings carry useful information.

Switching is easier to implement but depends strongly on the selected threshold. Blending makes the transition smoother, although smoothness alone does not ensure lower uncertainty. With the blending settings selected on separate validation data, blending and uncertainty weighting perform similarly to joint likelihood in the representative 20 mL/s case. This gives a practical reason to evaluate a simple rule before adopting a more computationally involved estimator.

6.2. Useful Operating Range and Design Choices

For the reference-water design, the upper combined flow is set by the pressure budget. The thermal lower flow is set by the maximum permitted temperature rise, and its upper flow by the minimum accepted temperature signal. These boundaries explain why the 10 W design can span a broad range even though its substantial precision-gain interval is much narrower.

At 5 W, the flow where the two sensors have equal predicted uncertainty lies above the selected upper thermal-flow limit, so most of the potentially larger fusion gain falls outside the accepted thermal range. Increasing effective heating to 10 W moves the usable thermal range into a region where both sensors contribute. Increasing the tube diameter addresses a different restriction: it permits a larger flow at the same pressure budget. The equations in Appendix B give the flows at which combination meets the chosen standard-deviation reduction target. These flows must also satisfy the tube's pressure and velocity limits.

At a single-reading timing standard deviation of 0.25 μs, a 2 mm tube with 5 W has a flow interval of 8.52–9.43 mL/s in which combination reduces predicted standard deviation by at least 10% relative to the better individual sensor. Increasing power to 10 W moves the flows meeting that reduction target above the maximum flow allowed by the tube's pressure limit. Stronger heating improves the thermal signal, but a stronger thermal sensor also leaves less extra information for Coriolis to add at the same flow. The best choice therefore depends on whether the aim is lower total measurement error or a larger improvement from adding a second sensor.

Once the two sensors provide a useful improvement together, calibration determines whether that improvement can be maintained. Tests of thermal response and Coriolis sensitivity both show that, as random

noise is reduced, smaller calibration errors can erase the advantage over the correctly calibrated sensor alone. If independent readings are collected at 150 Hz, 15 and 150 readings represent 0.1 and 1.0 s, respectively. A longer averaging interval also delays an updated estimate and includes readings from a longer period over which the flow may change. Choosing averaging therefore requires both an accuracy target for steady flow and a requirement for how quickly changes must be detected.

For ethanol, the lower product of density and specific heat capacity produces a stronger temperature signal, allowing both sensors to contribute meaningfully at 5 W. For glycerin, a favourable temperature signal does not help fusion when high viscosity prevents access to the required flows. For mineral oil, increasing power changes what limits the upper flow for thermal sensing: heating is the limit at 5 W, while pressure loss is the limit at 10 W. These examples show why thermal response, the conversion from Coriolis mass flow to volume flow, and pressure loss must be checked for the chosen liquid. Recalibration and increased averaging address different problems: recalibration corrects the signal-to-flow conversion for the liquid, while averaging reduces independent random scatter around the resulting mean.

6.3. Limitations

The numerical results use assumed signal sensitivities and noise levels, constant liquid-property sets, and a steady bulk heat balance. The retained liquid-property sets do not specify composition, purity or grade, so their correspondence to a particular sample depends on these factors and temperature. The model requires representative inlet and outlet temperatures; local wall temperatures, axial conduction, mixing, thermal dynamics, and mechanical coupling require a spatial or experimental model. Wall heat transfer and axial conduction can influence the temperature response [20]. Tube-diameter comparisons hold timing sensitivity fixed, whereas a redesigned vibrating tube would require recalibration. Correlated noise, time-dependent drift, bubbles, and flow transients can change the benefit of averaging and the performance of the acceptance rule. Physical performance remains to be evaluated with a prototype.

6.4. Next Steps

A prototype should be calibrated by collecting liquid over a measured time interval, weighing it, and using the measured density to convert mass flow to volume flow. Tests should begin near the flow where the two sensors have similar uncertainties, then introduce sudden flow changes, temperature changes, vibration, and controlled bubbles. The present results suggest specific conditions to test: water at 10 W over the 17.03–23.96 mL/s interval, the liquid-dependent limits in Figure 7, and thermal-response changes of a few percent with 15- and 150-reading averages. These measurements would test both the predicted reduction in random uncertainty and how closely calibration must be maintained to preserve it.

7. Conclusions

This study explores a combined thermal and Coriolis flow meter through simulation. At 10 W and with 15-reading averages, the predicted channel ranges together cover approximately 0.240–92.30 mL/s for reference water. Simulated switching and combination rules meet a 10% relative RMSE target at all 75 tested flows from 0.241 to 92.30 mL/s. Within that wider range, 17.03–23.96 mL/s provides at least a 10% predicted reduction in standard deviation relative to the better individual sensor. The tube-diameter and timing-noise analyses show how pressure loss and readout precision change the flows at which this reduction can be achieved.

Repeated tests at 20 mL/s and 10 W show similar RMSE reductions for blending with settings selected on validation data, uncertainty weighting, and joint likelihood. For uncertainty weighting at this operating point, the thermal response can differ from its calibration by −7.0% to +6.5% with 15-reading averages while retaining a 10% RMSE advantage over Coriolis. With 150-reading averages, that interval narrows to approximately −2.0% to +2.0%. The five-liquid comparison shows that both sensor conversions and pressure loss must be evaluated for the chosen liquid. Ethanol permits a useful contribution from both sensors at 5 W, whereas glycerin reaches its pressure limit before that contribution becomes substantial. Tests of changing liquid without recalibration

further show that averaging can make a seriously biased combined estimate more repeatable without making it accurate. Selecting heating power and averaging therefore requires checking the attainable flow range, the improvement over the better individual sensor, and how accurately each sensor remains calibrated.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The reported simulation results are presented in the figures and tables in this article and its appendices. Appendices C, E and F describe the simulation settings and implementation.

Conflicts of Interest

The author declares no conflict of interest.

Appendix A. Tube Geometry and Model Inputs

Appendix A.1. Geometry and Hydraulic Calculation

The tube-diameter comparison evaluates internal diameters of 2, 5, and 10 mm. Tube length remains fixed at 0.50 m, with the same water properties, surface roughness, and minor-loss coefficient. Each diameter is evaluated against a pressure-drop budget of 0.5 bar and a maximum mean water speed of 10 m/s. Increasing the diameter reduces water speed at a given volumetric flow rate and lowers the calculated pressure drop (Table A1).

Table A1. Calculated hydraulic limits for the three tube diameters.

Internal Diameter(mm)	Flow at 0.5 Bar(mL/s)	Flow at 10 m/s(mL/s)
2	9.43	31.42
5	92.30	196.35
10	451.72	785.40

The velocity and Reynolds number are calculated as follows, using Q_{SI} in m^3/s:

$$v = \frac{Q_{SI}}{\pi d^2/4}, \mathrm{Re} = \frac{\rho v d}{\mu}$$

The Darcy friction factor is 64/Re for laminar flow. For Re ⩾ 4000 the Haaland approximation is used:

$$\frac{1}{\sqrt{f}} = -1.8\log_{10}[(\frac{\varepsilon}{3.7d})^{1.11} + \frac{6.9}{\mathrm{Re}}]$$

The roughness ε is 1.5 μm. Between Re = 2300 and 4000, the laminar and Haaland values are linearly blended with a weight rising from zero to one. For the 5 mm tube and reference water properties, these boundaries correspond to approximately 9.07 and 15.77 mL/s. The pressure-limit crossing is located numerically and compared with the velocity-limit flow.

The worksheet relation in Section 2.1 provides a separate illustrative mechanical calculation. With half-spacing r_g = 0.030 m, torsional stiffness K = 10 N m/rad and density 1000 kg/m^3, it predicts approximately 7.2 μs at 10 mL/s, or a volume-flow slope of 0.72 μs/(mL/s). The signal comparisons use the independently selected reference slope of 0.10 μs/(mL/s). The symbol r_g for half-spacing is distinct from the uncertainty ratio r in Appendix B.

Appendix A.2. Parameters and Assumptions

Table A2 distinguishes water properties, design choices and assumed signal characteristics. The water properties are held constant at values appropriate to near-room-temperature conditions. The density, specific heat capacity and viscosity agree with NIST water data at 293.15 K and 0.101325 MPa to within 0.05% [12]. Geometry, heating power, noise levels and acceptance limits are selected inputs to the design comparison. Effective heating power is the power transferred to the flowing water.

Table A2. Fixed model inputs and comparison cases.

Quantity	Value	Basis
Water temperature	20 °C	Reference condition
Density ρ	998.2 kg/m^3	Reference water input
Specific heat capacity c_p	4182 J/(kg K)	Reference water input
Dynamic viscosity μ	0.001002 Pa s	Reference water input
Internal diameter d	5 mm; comparisons 2 and 10 mm	Design choice
Tube length L	0.50 m	Design choice
Roughness ε	1.5 μm	Hydraulic assumption
Minor-loss coefficient K_{loss}	2.0	Hydraulic assumption
Pressure and speed limits	0.5 bar; 10 m/s	Design requirements
Timing coefficient X	0.10 μs/(mL/s)	Assumed calibration
Timing noise σ_t	0.50 μs; sensitivity 0.25 and 1.00	SD of one timing reading
Temperature noise σ_T	0.020 K	SD of one temperature-difference reading
Effective power P_{eff}	5 and 10 W	Heating cases
Temperature-rise limits	0.10–10 K	Signal and heating requirements
Independent readings N	1, 15, 150	Averaging cases
Flow output bounds	0.01–92.2983 mL/s	Positive range and pressure limit

Appendix A.3. Liquid Property Sets

Table A3 lists the density, viscosity, specific heat capacity and thermal conductivity used for five liquids. These simulation inputs are taken from the project's graphical user interface (GUI) at a nominal temperature of 20 °C. The water values are the same as those used in the baseline calculation.

Measurements of milk show that its specific heat capacity, thermal conductivity and density depend on composition and temperature [14]. NIST provides heat-capacity data at specified temperatures for ethanol and glycerin [15–17], while mineral-oil properties depend on the oil grade [18].

Table A3. Liquid properties used in the simulations at a nominal temperature of 20 °C.

Liquid	Density(kg/m^3)	Viscosity(Pa s)	Specific heat capacity(J/(kg K))	Thermal conductivity(W/(m K))
Water	998.2	0.001002	4182	0.598
Milk	1030.0	0.0021	3930	0.53
Ethanol	789.0	0.0012	2440	0.171
Glycerin	1260.0	1.49	2410	0.285
Mineral oil	850.0	0.025	2000	0.145

Thermal conductivity is used to calculate heat conduction in the GUI's exploratory model. The steady bulk-temperature model uses density and specific heat capacity to relate heat input to temperature rise.

For each liquid, the calculation updates the timing coefficient, thermal coefficient, Reynolds number and pressure limit. Both heating cases use the same selected effective powers of 5 and 10 W across the property sets.

Appendix A.4. Symbols Used in the Analysis

The channel labels C, T, and F denote Coriolis, thermal, and fused estimates. The symbols σ_t and σ_T describe the standard deviations of individual timing and temperature-difference readings, in μs and K respectively. The symbols s_C, s_T, and s_F describe the standard deviations of the resulting flow estimates in mL/s. N is the number of independent readings averaged within one window, whereas M is the number of windows used to evaluate an estimator. Q denotes true flow in a simulation, Q_C and Q_T are the Coriolis and thermal flow estimates, and q is a candidate flow in the likelihood calculation. The thermal coefficient A is distinct from tube cross-sectional area, which is written explicitly as $\pi d^2/4$.

Appendix B. Combination Equations and Design Criteria

Appendix B.1. Switching and Blending

Let B_q restrict a flow estimate to the interval from 0.01 mL/s to the hydraulic upper flow, 92.2983 mL/s for reference water in the 5 mm tube. An estimate outside this interval is replaced by the nearest bound. The Coriolis output is obtained by applying B_q to y_C/X. An accepted thermal observation first gives $Q_T = A/y_T$, before any output bounds are applied. The thermal output is then restricted with B_q. The thermal estimate before restriction is used to choose the switching branch and calculate the blending and uncertainty weights.

Switching uses the bounded thermal output when Q_T is at or below the selected threshold, and the bounded Coriolis output above it. The logistic thermal weight is:

$$w_T = \frac{1}{1 + \exp[(Q_T - Q_0)/b]}$$

Here, Q_0 is the midpoint and b is the width parameter, both in mL/s. At the midpoint, each sensor contributes 50%. A larger width parameter gives a more gradual transition. The thermal weight multiplies the thermal output after its bounds are applied, while one minus that weight multiplies the bounded Coriolis output. If the observed temperature rise is outside 0.10–10 K, the rule uses the Coriolis output alone instead of combining the two readings.

Appendix B.2. Inverse Variance Derivation

For a weighted combination of independent, unbiased sensor estimates, with a fixed thermal weight w_T, the variance is:

$$s_F^2 = (1 - w_T)^2 s_C^2 + w_T^2 s_T^2$$

Differentiating with respect to the weight and setting the result to zero gives $w_T = s_C^2/(s_C^2 + s_T^2)$. Substituting it into the variance expression gives $s_F^2 = s_C^2 s_T^2/(s_C^2 + s_T^2)$. This derivation explains why the less variable estimate receives the greater weight. To implement the rule, the observed thermal flow estimate Q_T is substituted into the formula for the thermal flow standard deviation, using the specified signal-noise levels and averaging settings. The weight therefore changes from one measurement window to the next. The design intervals below use the theoretical expression with a fixed weight at each flow. Simulated RMSE evaluates the implemented rule, including the variation in weights caused by noisy observations.

Appendix B.3. Gain Interval and Required Heating

Define r as the ratio s_T/s_C. Let g be the desired fractional reduction in standard deviation compared with the more precise individual sensor. For a target greater than zero and no greater than $1 - 1/\sqrt{2}$ (approximately 0.2929), this reduction is achieved when:

$$a \le r \le \frac{1}{a},\ a = \sqrt{(1-g)^{-2} - 1}$$

The dimensionless quantity a specifies how similar the two sensor uncertainties must be. Since $r = (Q/Q_x)^2$, the flows meeting the target reduction before applying the physical limits are:

$$Q_x\sqrt{a} \le Q \le \frac{Q_x}{\sqrt{a}}$$

The precision-gain interval is the part of this range that also satisfies the temperature condition $A/\theta_{max} \leqslant Q \leqslant A/\theta_{min}$ and the hydraulic flow limits. It therefore contains flows at which the target reduction in standard deviation and the physical requirements hold together. Here, θ_{min} = 0.10 K and θ_{max} = 10 K are the selected temperature-rise limits. At a 10% target, a is approximately 0.4843. A larger target requires the two uncertainties to be closer and therefore narrows the interval.

The same condition can be rearranged to find the allowed heating at a specified flow. Let $c = \sigma_t/X$, the single-reading Coriolis flow standard deviation. The required interval for A is:

$$\max(\theta_{min}Q, \frac{a\sigma_T Q^2}{c}) \le A \le \min(\theta_{max}Q, \frac{\sigma_T Q^2}{ac})$$

Multiplying these bounds by $\rho c_p \times 10^{-6}$ converts them to effective power in W. Both requirements must hold at once, so the larger lower bound and the smaller upper bound set the allowed heating. For a nonempty power interval, the lower and upper flow conditions are:

$$\frac{ac\theta_{min}}{\sigma_T} \le Q \le \frac{c\theta_{max}}{a\sigma_T}$$

Intersecting that range with the hydraulic limits determines whether any heating choice can meet the gain target. The minimum possible effective power, when the hydraulic limits permit the associated flow, is:

$$P_{critical} = \rho c_p \times 10^{-6} \times \frac{ac\theta_{min}{}^2}{\sigma_T}$$

At the reference noise levels and a 10% target, the minimum flow is 12.1081 mL/s and the minimum power is 5.0545 W. A power exceeding this minimum must still be paired with a flow satisfying the full interval. Table A4 shows the effects of varying the selected gain target.

Table A4. Predicted flow intervals meeting the target reduction in standard deviation and the physical limits, at the reference noise levels and 5 mm diameter.

Target Reduction in SD(%)	Power(W)	Precision-Gain Interval(mL/s)	Minimum Flow to Meet Target at Any Power (mL/s)
5	5	9.92–11.98	8.22
5	10	14.03–23.96	8.22
10	5	Empty	12.11
10	10	17.03–23.96	12.11
20	5	Empty	18.75
20	10	21.19–23.96	18.75

The last column gives the lowest flow at which any heating power can meet the target and the physical limits. This calculation allows power to vary rather than restricting it to the listed 5 and 10 W cases.

Appendix B.4. Calibration Bias and Averaging

With a fixed thermal response mismatch, the noise-free thermal estimate is $Q/(1 + \delta)$, so its bias is $-\delta Q/(1 + \delta)$. At a fixed flow, consider combining an unbiased Coriolis estimate with a thermal estimate having fixed weight w_T and bias b_T. If the two random errors are independent, the combined mean squared error is:

$$MSE_F = (1 - w_T)^2 s_C{}^2 + w_T{}^2 s_T{}^2 + w_T{}^2 b_T{}^2$$

The first two terms describe random variance, while the final term is the squared bias of the combined estimate. Averaging more independent readings reduces the random terms but leaves the fixed bias contribution unchanged. With the nominal inverse-variance weight and random variances, fusion has lower MSE than unbiased Coriolis when $b_T^2 < s_C^2 + s_T^2$. Meeting a specified improvement target, rather than merely outperforming Coriolis, requires a stricter limit on bias. The reported tolerance to response mismatch is calculated directly from simulated RMSE. This includes weights calculated from noisy readings, the nonlinear

conversion from temperature to flow, and the rule for accepting or rejecting thermal observations.

Appendix B.5. Joint Likelihood Calculation

The objective in Section 3.3.3 is minimised over the bounded flow interval. Let $c_{obs} = y_C/X$ be the Coriolis flow estimate before output bounds are applied, and let $s_{signal} = \frac{\sigma_T}{\sqrt{N}}$ be the standard deviation of the mean temperature-difference reading within a window, in K. The Coriolis flow standard deviation s_C is defined in Section 2.4. Stationary points satisfy:

$$q^4 - c_{obs} q^3 + \beta q - \gamma = 0$$

$$\beta = A y_T (\frac{s_C}{s_{signal}})^2,\ \gamma = A^2 (\frac{s_C}{s_{signal}})^2$$

The coefficients β and γ collect the measured temperature and the relative scales of the timing and temperature noise.

Although a fourth-order polynomial has a closed-form solution, its expressions are lengthy. A numerical procedure makes it straightforward to compare the admissible candidates within the flow bounds. The solver partitions the positive flow interval at the turning points of this polynomial, locates each candidate local minimum by bisection, and compares its objective value with those at both flow bounds. The chosen candidate has the smallest objective. Rejected thermal observations use the same bounded Coriolis output as the other thermal-assisted methods.

Appendix C. Evaluation Settings and Additional Results

Appendix C.1. Sampling and Parameter Selection

The main test grid consists of 55 logarithmically spaced flows from 0.2 to 90 mL/s plus 1, 5, 10, 11.98, 15, 20, 23.95, 25, 30, 50, and 80 mL/s, giving 66 unique points. Each flow has 2000 windows for each of seeds 104729, 130363, 155921, 180503, and 205019. The 5 and 10 W cases each use N = 1, 15, and 150. Mean squared errors and means are pooled across the equal-sized seed runs before calculating pooled RMSE and standard deviation.

Validation uses seed 271828 and 3000 windows at each of 41 logarithmically spaced flows from 0.23 to 87 mL/s. The switching threshold and blending midpoint are each chosen from 60 logarithmically spaced flow values, from 0.3 mL/s to the smaller of four times Q_x and the hydraulic upper limit. Blending widths are 0.03, 0.06, 0.12, 0.25, and 0.50 times Q_x. The parameters are selected separately for each method to minimise the unweighted mean squared error over this validation grid. Table A5 records the reference-noise settings.

Table A5. Selected switching and blending parameters at the reference noise levels.

Power (W)	N	Switch Threshold	Blend Midpoint	Blend Width
5.0	1	12.001	13.160	2.077
5.0	15	12.001	14.432	2.077
5.0	150	12.001	14.432	2.077
10.0	1	21.510	21.510	2.937
10.0	15	26.120	26.120	6.118
10.0	150	23.703	26.120	6.118

Thresholds and widths are in mL/s.

Appendix C.2. Repeated Seed Comparisons and Calibration Endpoints

The reported 16.7% improvement is the mean of five results, one for each random seed. For each seed, the percentage improvement is 100 times one minus the ratio of uncertainty-weighting RMSE to the RMSE of Thermal + fallback, using the same simulated observations for both methods. The interval 15.8–17.7% uses the standard error of these five improvements and a Student t multiplier with four degrees of freedom. It quantifies

variation between the random simulation runs.

The thermal-response mismatch test uses Q = 20 mL/s, 10 W, N = 15 or 150, and 97 equally spaced response mismatches from −12% to +12%. Five seeds (86028121, 104395303, 122949829, 141650939, and 160481183) provide 5000 windows per seed and mismatch. The algorithm settings selected with correct calibration remain fixed as the actual thermal response is changed. The reported tolerance is the continuous interval containing zero mismatch throughout which RMSE is no more than 0.9 times Coriolis RMSE. Table A6 gives the interval endpoints after pooling all five runs, together with the range of endpoints obtained from the individual runs.

Appendix C.3. Thermal Acceptance and Interpretation

The thermal-acceptance fraction is the proportion of measurement windows whose observed temperature rise is between 0.10 and 10 K. A theoretical temperature boundary uses the expected signal A/Q, whereas the acceptance rule acts on the noisy observation y_T. Near a boundary, some thermal observations are accepted and others are rejected, in which case the rule uses Coriolis alone. At 20 mL/s and 10 W the acceptance fraction rises from 83.82% for N = 1 to 100% in the tested N = 15 and 150 windows. At sufficiently high flow, the methods that use thermal observations give nearly the same result as Coriolis because nearly all thermal observations are rejected.

Table A6. Response-mismatch ranges retaining at least a 10% RMSE improvement over Coriolis at 20 mL/s and 10 W.

N	Method	Pooled Interval	Lower Endpoint Range	Upper Endpoint Range
15	Uncertainty weighting	−7.00 to +6.50	−7.00 to −6.75	+6.50 to +6.75
15	Joint likelihood	−6.50 to +6.75	−6.50 to −6.25	+6.75 to +7.00
150	Uncertainty weighting	−2.00 to +2.00	−2.00 to −2.00	+2.00 to +2.00
150	Joint likelihood	−2.00 to +2.00	−2.00 to −2.00	+2.00 to +2.00

Entries are percentages.

For 15-reading windows, validation selects switching thresholds of 12.00 mL/s at 5 W and 26.12 mL/s at 10 W. These are above the largest flows inferred from an accepted temperature reading, 11.98 and 23.96 mL/s. Switching therefore uses every accepted thermal estimate and returns to Coriolis when the temperature reading is rejected, giving the same output as Thermal + fallback.

With one reading per window at 20 mL/s and 10 W, joint likelihood has about 1.6% higher RMSE than Thermal + fallback. Near the minimum accepted temperature rise, the nonlinear conversion to flow and the decision to accept or reject a reading affect the comparison. Averaging reduces that noise, making the standard-deviation approximation in Section 2.4 more informative.

Appendix C.4. Timing Noise Sensitivity Protocol

The timing-noise comparisons set the standard deviation of an individual timing reading to 0.25, 0.50, or 1.00 μs. They use the same physical parameters, validation grid, rule for choosing candidate thresholds and widths, and five test seeds as the baseline comparison. At each noise level, switching and blending are tuned using validation observations with that noise level and correct calibration. Their selected settings are then held fixed while calibration tolerances are checked at 20 mL/s and 10 W for N = 15 and 150. Table A7 gives RMSE at 20 mL/s and 10 W for 15-reading windows across the three timing-noise levels.

Appendix C.5. Verification of the Combined Measurement Range

The measurement-range check uses reference water, 5 and 10 W, and 15 readings per window. It retains the baseline switching and blending settings in Table A5 without further tuning. Each flow pools 25,000 windows: 5000 from each of seeds 86028121, 104395303, 122949829, 141650939 and 160481183. A logarithmically spaced grid is supplemented with selected test flows across the measurement range and closely spaced points at

the predicted temperature and pressure boundaries. The 10 W test contains 84 flow values, including points used to inspect behaviour near those boundaries. Of these, 75 lie in the reported interval 0.241 – 92.2983 mL/s. All windows contribute to RMSE, whether the thermal observation is accepted or rejected.

Table A7. RMSE in mL/s at 20 mL/s, 10 W, and 15-reading windows across timing-noise cases.

Method	0.25 µs	0.50 µs	1.00 µs
Coriolis	0.647	1.293	2.586
Thermal + fallback	0.858	0.858	0.858
Switching	0.648	0.858	0.858
Blending	0.519	0.715	0.813
Uncertainty weighting	0.519	0.714	0.813
Joint likelihood	0.519	0.718	0.816

Blending, uncertainty weighting and joint likelihood remain close in RMSE across the timing-noise cases.

At the exact predicted thermal lower limit, 0.239551 mL/s at 10 W, approximately half the noisy temperature observations exceed 10 K and are rejected. Using Coriolis alone in those windows gives a large relative error, so the range verified by the implemented rules is reported from the tested 0.241 mL/s point. At the upper pressure limit, capping estimates at the output bound reduces their error, but the Coriolis relative standard deviation before this cap is already only 1.40%. The upper end of the measurement range is therefore set by pressure loss. The maximum pooled relative RMSE over the reported 10 W grid is 5.63% for switching, 5.54% for blending, 5.04% for uncertainty weighting and 5.02% for joint likelihood. The maxima in each individual seed run are also below 10%.

At 5 W the transition is close to the 10% error target, so a refinement uses 41 flows from 12.60 to 13.00 mL/s and 50,000 windows per seed, giving 250,000 windows at each flow. The maximum pooled relative RMSE is 9.99% for switching and 9.88% for uncertainty weighting. One switching seed reaches 10.02%. This transition has little margin, whereas the 10 W case remains comfortably below the target.

Appendix C.6. Definitions of Mean, Bias and Standard Deviation

For the M estimates obtained at a fixed true flow Q, the mean and bias are:

$$\overline{\hat{Q}} = \frac{1}{M}\Sigma\hat{Q}_i, \text{Bias} = \overline{\hat{Q}} - Q$$

$$\text{SD} = \sqrt{\frac{\Sigma(\hat{Q}_i - \overline{\hat{Q}})^2}{M}}$$

The divisor is M, matching the empirical RMSE definition in Section 4.2 and making $\text{RMSE}^2 = \text{Bias}^2 + \text{SD}^2$ exact for the evaluated windows.

Appendix D. Additional Model Sensitivity

Appendix D.1. Incorrectly Estimated Temperature Noise

The temperature-noise standard deviation used to generate signals is fixed at 0.020 K, while the value supplied to uncertainty weighting and joint likelihood is multiplied by 0.5 or 2. Underestimating the noise gives the thermal channel too much influence, whereas overestimating it reduces that influence. This test changes what the algorithms assume about the noise without changing the noise in the generated observations. In the timing-noise comparison, by contrast, the algorithms are supplied with the noise level actually used to generate their data. The selected flows are 5, 10, 15, 20, and 23 mL/s at 10 W, for N = 1, 15, and 150. Table A8 gives the results at 20 mL/s and 10 W for 15-reading windows.

Table A8. Effect of temperature-noise assumptions at 20 mL/s, 10 W, and 15-reading windows.

Assumed Noise Multiplier	Method	Bias	SD	RMSE
0.5	Uncertainty weighting	+0.016	0.786	0.786
0.5	Joint likelihood	+0.020	0.787	0.787
1	Uncertainty weighting	-0.005	0.714	0.714
1	Joint likelihood	+0.020	0.718	0.718
2	Uncertainty weighting	-0.011	0.887	0.887
2	Joint likelihood	+0.042	0.888	0.889

Bias, standard deviation, and RMSE are in mL/s.

Appendix D.2. Heat Loss

The heat-loss test represents power lost to the surroundings as U times the outlet-minus-inlet temperature rise. The selected coefficient U, in W/K, combines heat loss through different paths into one parameter. The surrounding reference temperature is set equal to the inlet temperature. Here P is the supplied heating power before this loss, and the true temperature rise becomes:

$$P = (\rho c_p Q_{SI} + U)\Delta T,\ \Delta T = \frac{P}{\rho c_p Q_{SI} + U}$$

The estimator retains its loss-free calibration, with $A = P/(\rho c_p \times 10^{-6})$. Its noise-free thermal flow estimate is then higher than the true flow by $U/(\rho c_p \times 10^{-6})$. The selected coefficients 0.02 and 0.10 W/K correspond to thermal-flow offsets of approximately 0.00479 and 0.02396 mL/s. These small offsets give a concrete mechanism through which a change in the thermal environment can shift the reading. Local calorimetric studies also identify wall heat transfer and axial conduction as influences on temperature response [20].

Appendix D.3. Connection to the Liquid Comparison

Section 5.3 changes density, specific heat capacity and viscosity together using the liquid properties in Table A3. Section 5.4 then examines what happens when the actual liquid changes but some or all of the estimator's calibration coefficients are still those for water. In the Unchanged water calibration case, effective power is unchanged and both channels use the water coefficients. Before temperature acceptance checks or output bounds are applied, their noise-free flow estimates are:

$$\hat{Q}_C = \frac{\rho_l}{\rho_w} Q, \hat{Q}_T = \frac{\rho_l c_{p,l}}{\rho_w c_{p,w}} Q$$

These relations predict whether each calibration error makes the flow estimate too high or too low, providing independent checks on the simulation. Updating the density used in the Coriolis conversion corrects that channel's flow estimate. Correcting the thermal estimate also requires the specific heat capacity, because its response depends on volumetric heat capacity, the product of density and specific heat capacity. The liquid-change tests therefore examine calibration changes associated with material properties. The proportional-noise test instead adds random temperature error, while the thermal-response mismatch test changes the actual temperature response by a fixed fraction.

Appendix D.4. Additional Proportional Noise Protocol

Each temperature-difference reading receives an extra independent Gaussian error whose standard deviation is α times the expected temperature rise, with α = 0, 0.025 or 0.10. After averaging a window of N readings, this extra error has standard deviation $\alpha(A/Q)/\sqrt{N}$. The simulation draws it independently of the fixed temperature noise and adds the two errors, so their variances add. The generated observations contain the extra noise, while the algorithms retain their baseline noise assumptions and switching or blending parameters. The 5 and 10 W cases use 15 and 150 readings, the five liquid-study test seeds and 2000 windows per flow and seed. When the extra noise is zero, the outputs are checked against the corresponding baseline water results.

The added component acts on the temperature signal alone, keeping Coriolis unchanged, and has the same averaging behaviour as the other independent reading noise.

Appendix D.5. Calibration Errors in Either Channel

The Coriolis-response mismatch test changes the true Coriolis sensitivity to $(1 + \delta_C)X$ while keeping the estimator's coefficient X fixed. The thermal response remains correctly calibrated. It uses the same 20 mL/s flow, 10 W power, 15- and 150-reading windows, seeds, 5000 windows per seed and 0.25-percentage-point spacing between mismatch values as Section C.2. The same random draws are reused at each mismatch value, and the combination settings selected with correct calibration remain fixed. All six estimation methods are evaluated. Figure A1 shows uncertainty weighting and joint likelihood for the tests in which either the thermal response or the Coriolis response is changed.

When only the Coriolis response is changed, the comparison baseline is Thermal + fallback: the thermal estimate is used whenever accepted, with Coriolis used otherwise. Every tested window in this sweep accepts the thermal reading, so the baseline is unaffected by the Coriolis mismatch. Uncertainty weighting retains a 10% RMSE advantage from −6.00% to +4.00% sensitivity mismatch at $N = 15$ and from −1.50% to +1.25% at $N = 150$. Each panel asks how much error can be introduced into one channel before combining it with the correctly calibrated channel loses the target benefit. The interval widths depend on the error of the reference method in each panel, as well as on the changed channel. They therefore compare the benefit of combination against different baselines, rather than directly ranking the calibration tolerance of the two sensors.

Tests with bias in both channels set their noise-free outputs to either 19 or 21 mL/s at a true flow of 20 mL/s. When both channels read too high or both read too low, uncertainty weighting retains a bias close to ±1 mL/s as averaging increases from 15 to 150 readings. When one reads too high and the other too low, the errors partly cancel. They do not cancel completely because the two weights are unequal and are recalculated from each noisy observation.

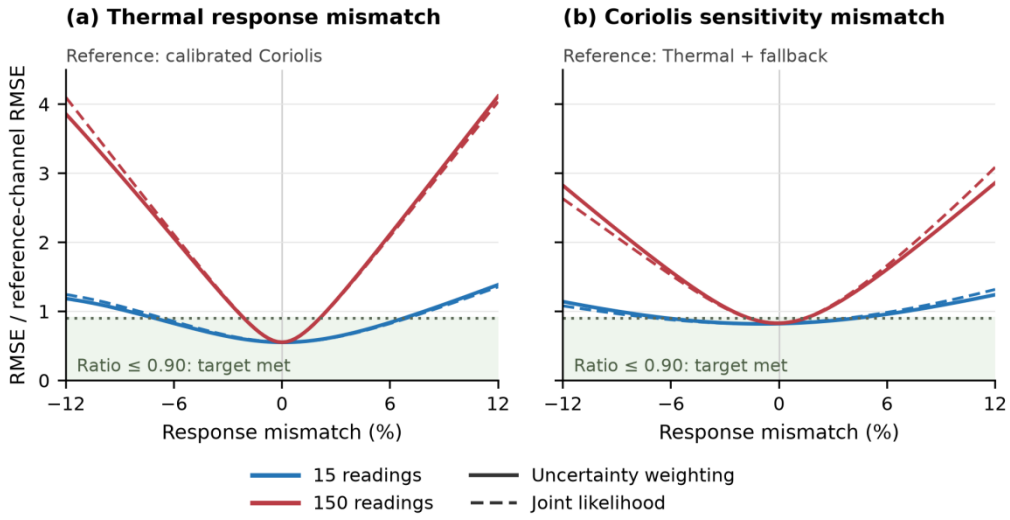


Figure A1. Effect of changing either sensor's calibration response at 20 mL/s and 10 W. In the thermal-response mismatch test, each combination method's RMSE is divided by Coriolis RMSE. In the Coriolis-response mismatch test, it is divided by the RMSE of Thermal + fallback, which uses Coriolis only if a thermal observation is rejected. A ratio at or below 0.90 meets the 10% improvement target. Solid and dashed curves show uncertainty weighting and joint likelihood, respectively. The reference channel is correctly calibrated in each panel, and all windows contribute to the error calculation.

Appendix E. Code Implementation

The simulations were implemented in Python 3.12.14. The baseline and sensitivity calculations used NumPy 2.5.4. SciPy 1.18.1 supplied the Student t quantile for the confidence interval on the improvement calculated

from paired random-seed runs. Figures were produced with Matplotlib 3.11.2. The measurement-range tests in Appendix C.5 used NumPy 2.3.5 with the same model and fixed algorithm settings.

NumPy arrays hold the mean timing difference and mean temperature rise for each simulated measurement window. Seeded random draws generate independent Gaussian errors for these window means. Their standard deviations are obtained by dividing the single-reading noise levels once by the square root of the number of readings in a window. Every estimator receives the same pair of observations. The code applies the common rule for accepting thermal observations and the specified output bounds, computes each flow estimate, and pools the means and squared errors across seeds before reporting bias, standard deviation and RMSE. Switching and blending use parameters selected on separate validation observations. Joint likelihood uses the bounded numerical search described in Appendix B.5.

Implementation checks cover unit conversion, heat balance, hydraulic limits, averaging, and consistent application of the thermal-acceptance rule and output bounds across methods. Independent evaluations of the joint-likelihood objective check that the numerical solver selects its minimum. The measurement-range tests in Section C.5 also check how often thermal observations are accepted and how often estimates are capped at the output bounds near the temperature and pressure limits. The calculated errors are checked against the identity $RMSE^2 = Bias^2 + SD^2$. Appendices C, D and F specify the tested flow values, random seeds and calibration cases.

Appendix F. Parameter Exploration and Liquid Comparisons

Appendix F.1. Selecting the Transition Parameters

The interactive project explored where to switch between sensors and how gradually to blend their estimates. The liquid-property and additional proportional-noise comparisons select these settings using separate validation observations. Figure A2 shows the parameter selection for water using the protocol in Appendix F.3, which starts at a lower validation flow and considers lower candidate thresholds than the baseline protocol in Appendix C. The selected settings give the lowest mean squared error on the validation observations and are then used on independent test data. The colour scale shows the square root of this error in mL/s, so its minimum occurs at the same parameter values. The baseline water comparison retains the separate settings in Table A5.

Appendix F.2. Interactive Exploration

The project GUI provides controls for flow, liquid properties, noise, averaging and combination strategy. It

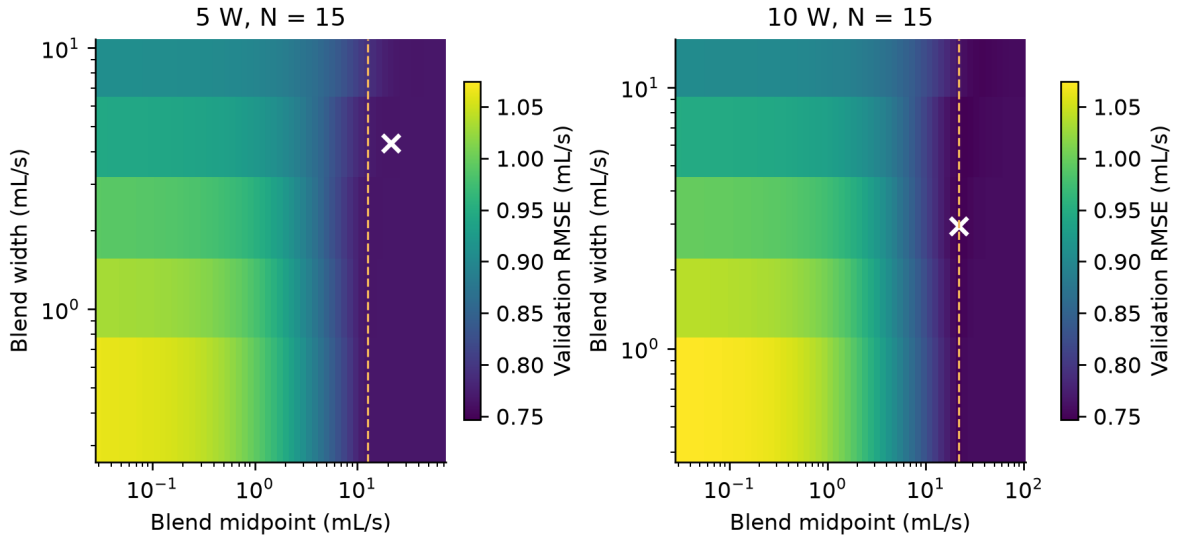


Figure A2. Validation root mean squared error for candidate blending midpoints and widths, using water and 15 readings per window under the liquid-property and additional proportional-noise protocol. Each panel uses the same observations for every candidate. The white cross marks the selected blending parameters, while the dashed gold line marks the independently selected switching threshold. Both are selected before evaluation on test observations. The baseline water comparison uses the separate settings in Table A5.

also displays hydraulic and thermal quantities, making it useful for exploring how a parameter change affects several parts of the design. Figure A3 shows the interface used for this exploratory stage.

The initial exploration considered why a thermal measurement could become less reliable at both ends of the flow range. At very slow flow, conduction can carry heat between sensing locations before the liquid does. At high flow, a fixed heater produces a smaller temperature rise, and mixing can broaden a travelling temperature pulse. An empirical relative-uncertainty curve represented these trends using a base term, a low-flow term and a high-flow term:

$$u_T = 0.025 + 0.05\frac{Q_{low}}{Q} + 0.08\left(\frac{Q}{Q_{high}}\right)^{1.5}$$

In this expression, u_T is a fractional standard deviation, Q_{low} = 0.01 mL/s, and Q_{high} = 9.07 mL/s. The numerical coefficients set exploratory error scales. The low-flow term decreases as flow increases, while the high-flow term grows, raising predicted relative uncertainty at either end of the range.

Appendix F.3. Full Liquid Comparisons

The liquid-property comparison uses 45 logarithmically spaced test flows between 0.02 mL/s and the smaller of 200 mL/s and 98% of the actual liquid's hydraulic limit. It also tests 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10, 15, 20, 25, 30, 50 and 80 mL/s whenever those values lie within this interval. Further test points are placed at each temperature boundary and at the flow where the two predicted standard deviations are equal, together with points 2% above and below each of these values. Only points within the test interval are retained. The grid also includes the representative flow used in Table A9: the smaller of 20 mL/s and 80% of the hydraulic limit. The powers are 5 and 10 W and the averaging counts are 15 and 150. Five seeds (104729, 130363, 155921, 180503 and 205019) provide 2000 windows per flow and seed. Every method and calibration case receives the same observations within a given condition.



Figure A3. Interactive simulator used in the Week 7 project, shown in its heat-pulse exploration view. Its controls allow the user to change liquid properties, noise and combination rules. The steady bulk-temperature model described in Section 2 is used for the quantitative comparisons in this paper.

Table A9. Flow estimates and errors from uncertainty weighting at 10 W with 15 readings per window.

Liquid	True Flow	Calibration	Mean	Bias	SD	RMSE
Water	20.000	Correct calibration	19.992	−0.008	0.717	0.717
Water	20.000	Coriolis-only recalibration	19.992	−0.008	0.717	0.717
Water	20.000	Unchanged water calibration	19.992	−0.008	0.717	0.717
Milk	20.000	Correct calibration	19.993	−0.007	0.696	0.696
Milk	20.000	Coriolis-only recalibration	19.569	−0.431	0.699	0.821
Milk	20.000	Unchanged water calibration	19.743	−0.257	0.723	0.768
Ethanol	20.000	Correct calibration	20.007	+0.007	0.387	0.387
Ethanol	20.000	Coriolis-only recalibration	9.362	−10.638	0.193	10.640
Ethanol	20.000	Unchanged water calibration	9.358	−10.642	0.192	10.644
Glycerin	0.824	Correct calibration	0.824	−0.000	0.001	0.001
Glycerin	0.824	Coriolis-only recalibration	0.599	−0.224	0.001	0.224
Glycerin	0.824	Unchanged water calibration	0.599	−0.224	0.001	0.224
Mineral oil	20.000	Correct calibration	20.001	+0.001	0.338	0.338
Mineral oil	20.000	Coriolis-only recalibration	8.251	−11.749	0.148	11.750
Mineral oil	20.000	Unchanged water calibration	8.254	−11.746	0.148	11.747

Correct calibration uses the liquid-specific coefficients for both sensors and the liquid-specific hydraulic bounds. Coriolis-only recalibration updates the Coriolis coefficient and hydraulic bounds but retains the water thermal coefficient. Unchanged water calibration retains the water coefficients for both sensors and the water hydraulic bounds. All flow and error columns are in mL/s. The glycerin comparison uses 80% of its hydraulic limit instead of 20 mL/s.

Validation uses 41 logarithmically spaced flows from 0.023 mL/s to the smaller of 197 mL/s and 97% of the hydraulic limit assumed by the estimator, with seed 271828 and 3000 windows per flow. The switching threshold and blending midpoint are each chosen from 60 candidate flow values, running from 0.03 mL/s to four times the crossover flow predicted from that calibration, where the two sensor standard deviations are equal. These candidate flow values can lie above the physically attainable flow range, allowing a rule to favour the thermal channel throughout that range when it is consistently more precise. Blending widths are 0.03, 0.06, 0.12, 0.25 and 0.50 times the same predicted crossover flow. Each rule's parameters minimise mean squared flow error, with every validation flow receiving equal weight. In the Correct calibration case, switching and blending are tuned separately for each liquid. The Coriolis-only recalibration and Unchanged water calibration cases retain the settings selected for water. This liquid-specific validation and test grid is evaluated separately from the baseline water comparison in Appendix C.

The theoretical precision-gain intervals are calculated from the expected temperature rise and each liquid's actual hydraulic limit. The implemented rules instead decide whether to accept each observed temperature rise and apply the output bounds specified by their calibration. In the Unchanged water calibration case, the output bounds remain those for water even when the actual liquid has a different hydraulic limit. The resulting estimates are then compared with the true simulated flow to calculate their errors. Figure A4 shows the relative RMSE of uncertainty weighting across the attainable flow range of each liquid.

Figure A5 compares all six estimation methods across the five liquids at 5 W using 15 readings per window and correct calibration.

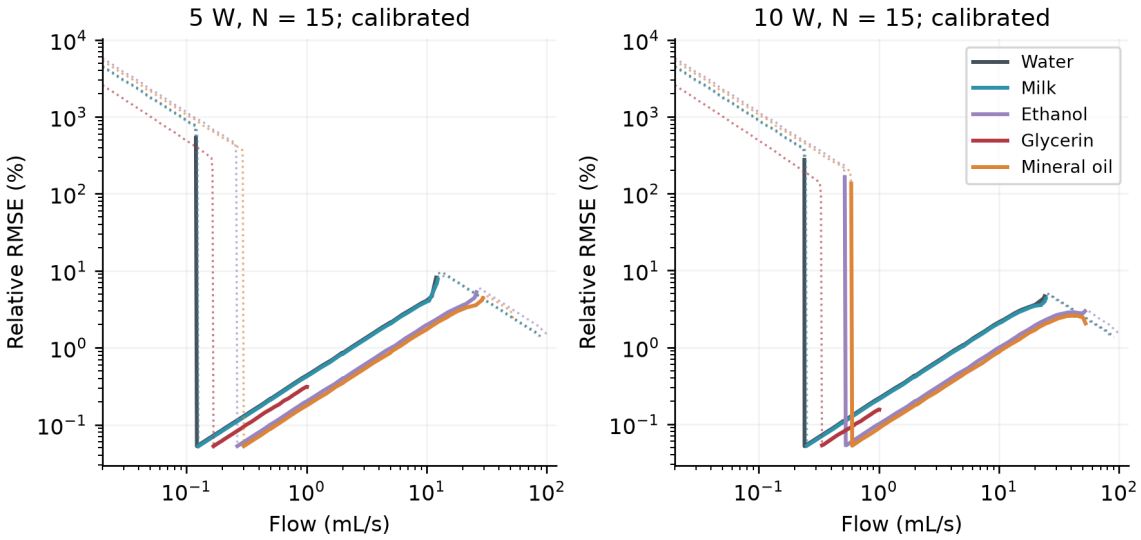


Figure A4. Relative RMSE of uncertainty weighting for the five liquids under Correct calibration, at 5 and 10 W with 15 readings per window. Relative RMSE is 100 times RMSE divided by true flow. Each curve stays within its liquid's hydraulic range. Faint dotted portions identify flows where the expected temperature rise is outside the acceptance limits.

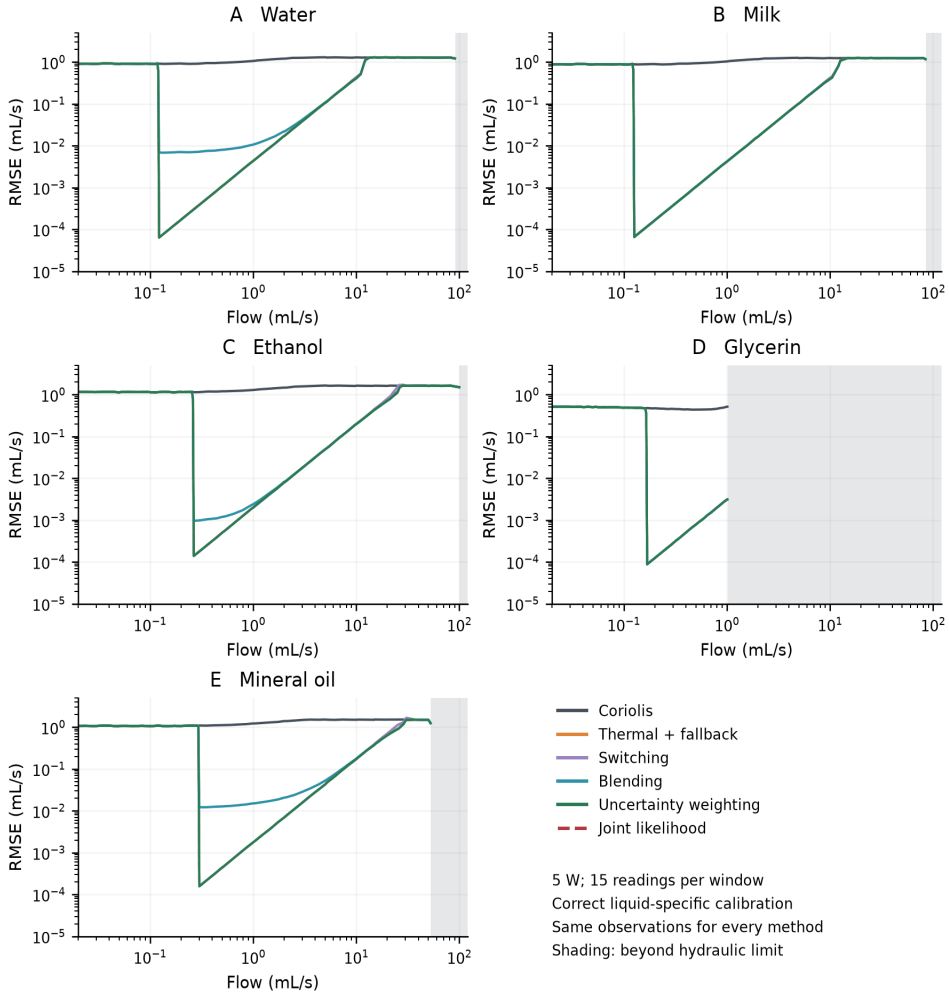


Figure A5. Comparison of six estimation methods for the five liquids under Correct calibration, at 5 W with 15 readings per window. All methods use the same simulated observations. Those using the thermal channel apply the same rule for accepting or rejecting its readings. Grey regions mark flows above each liquid's hydraulic limit, and the curves stop within that limit.

References

1 Bronkhorst. Coriolis Flow Meter Working Principle. Available online: https://www. bronkhorst. com/knowledge-base/coriolis-flow-meter-working-principle/ (accessed on 23October 2026).

2 Bronkhorst. How Does a Thermal Mass Flow Meter Work? Available online: https://www.bronkhorst.com/knowledge-base/thermal-mass-flow-measurement/ (accessed on 23October 2026).

3 Lötters JC, Lammerink TSJ, Groenesteijn J, *et al.* Integrated Thermal and Microcoriolis Flow Sensing System with a Dynamic Flow Range of More Than Five Decades. *Micromachines* 2012; **3**(**1**): 194–203. https://doi.org/10.3390/mi3010194.

4 Lyons L, Gibaut D, Clifford P. How to Combine Correlated Estimates of a Single Physical Quantity. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1988; **270**(**1**): 110–117. https://doi.org/10.1016/0168-9002(88)90018-6.

5 Wang T, Baker R. Coriolis Flowmeters: A Review of Developments over the Past 20 Years, and an Assessment of the State of the Art and Likely Future Directions. *Flow Measurement and Instrumentation* 2014; **40**: 99–123. https://doi.org/10.1016/j.flowmeasinst.2014.08.015.

6 Kolhe V, Pawar S, Chaudhari V, *et al.* Parameter Optimization of Coriolis Mass Flow Meter in Laminar Flow Regime Using Doe-Taguchi Method. *Journal of Thermal Engineering* 2023; **9**(**4**): 1026–1040. https://doi.org/10.18186/thermal.1335677.

7 *U. S. Department of Energy. DOE Fundamentals Handbook: Thermodynamics, Heat Transfer, and Fluid Flow*; Volume 2; DOE-HDBK-1012/2-92; U. S. Department of Energy: Washington, DC, USA, 1992. Available online: https://www. energy. gov/sites/default/files/2026-04/DOE-HDBK-1012-92_VOL2. pdf (accessed on 20 November 2026).

8 Joint Committee for Guides in Metrology (JCGM). *Evaluation of Measurement Data—Guide to the Expression of Uncertainty in Measurement*; JCGM 100:2008; JCGM: Sèvres, France, 2008. https://doi.org/10.59161/JCGM100-2008E.

9 *U.S. Department of Energy. DOE Fundamentals Handbook*: *Thermodynamics, Heat Transfer, and Fluid Flow*; Volume 3; DOE-HDBK-1012/3-92; U.S. Department of Energy: Washington, DC, USA, 1992. Available online: https://www.energy.gov/documents/doe-hdbk-1012-92vol3 (accessed on 20 November 2026).

10 Haaland SE. Simple and Explicit Formulas for the Friction Factor in Turbulent Pipe Flow. *Journal of Fluids Engineering* 1983; **105**(**1**): 89–90. https://doi.org/10.1115/1.3240948.

11 Cowan G. Statistics (Chapter 40). Available online: https://pdg.lbl.gov/2024/reviews/rpp2024-rev-statistics.pdf (accessed on 20 November 2026).

12 Lemmon EW, Bell IH, Huber ML, et al. Thermophysical Properties of Fluid Systems. In *NIST Chemistry WebBook*; NIST Standard Reference Database Number 69; National Institute of Standards and Technology: Gaithersburg, MD, USA. Available online: https://webbook. nist. gov/chemistry/fluid/ (accessed on 20 November 2026).

13 Joint Committee for Guides in Metrology (JCGM). *Evaluation of Measurement Data—Supplement 1 to the Guide to the Expression of Uncertainty in Measurement—Propagation of Distributions Using a Monte Carlo Method*; JCGM 101:2008; JCGM: Sèvres, France, 2008. https://doi.org/10.59161/JCGM101-2008.

14 Minim LA, Coimbra JSR, Minim VPR, et al. Influence of Temperature and Water and Fat Contents on the Thermophysical Properties of Milk. *Journal of Chemical & Engineering Data* 2002; **47**(**6**): 1488–1491. https://doi.org/10.1021/je025546a.

15 Domalski ES, Hearing ED. Condensed Phase Heat Capacity Data. In *NIST Chemistry WebBook*; NIST Standard Reference Database Number 69; National Institute of Standards and Technology: Gaithersburg, MD, USA. Available online: https://doi.org/10.18434/T4D303 (accessed on 20 November 2026).

16 NIST Chemistry WebBook. Ethanol. Available online: https://webbook. nist. gov/cgi/cbook. cgi? ID=C64175&Mask=2 (accessed on 25 October 2026).

17 NIST Chemistry WebBook. Glycerin (Glycerol). Available online: https://webbook.nist.gov/cgi/cbook.cgi?ID=C56815&Mask=2 (accessed on 25 October 2026).

18 ExxonMobil. White Oils Products. Available online: https://www. exxonmobil. com/en/whiteoil/products (accessed on25 October 2026).

19 Taylor BN, Kuyatt CE. *Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results*; NIST Technical Note 1297; National Institute of Standards and Technology: Gaithersburg, MD, USA, 1994. https://doi.org/10.6028/nist.tn.1297.

20 Gelderblom H, Van Der Horst A, Haartsen JR, *et al.* Analytical and experimental characterization of a miniature calorimetric sensor in a pulsatile flow. *Journal of Fluid Mechanics* 2011; **666**: 428–444. https://doi.org/10.1017/s0022112010004234.

© The Author(s) 2026. Published by Global Science Publishing (*GSP*).

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.