

Using an Integrated Heater to Recover RH Accuracy from High-Humidity Exposure



Alex Thompson, Sakeenah Khan

ABSTRACT

Polymer-based relative humidity sensors can exhibit a recoverable shift in RH error after prolonged exposure to high-humidity conditions. This shift in RH error from the time-zero RH is called RH drift. In systems that cannot rely on passive recovery at normal room conditions or offline oven baking, an integrated heater provides a practical in-system method to accelerate moisture removal from the sensing element and reduce RH offset. This application note addresses two design cases: one-shot recovery after an isolated high-humidity event and periodic mitigation during sustained high-humidity operation. On the 62-mil rigid FR4 HDC3020 board evaluated here, with the thermal pad soldered, one second of full-power heating every minute at 3.3V delivered the best balance among the tested continuous schedules, holding final RH drift near +0.18 %RH. Heater effectiveness improves when PCB thermal mass is reduced, thermal resistance to the environment is increased, and heater power is increased. However, those gains in heater performance must be balanced against power consumption.

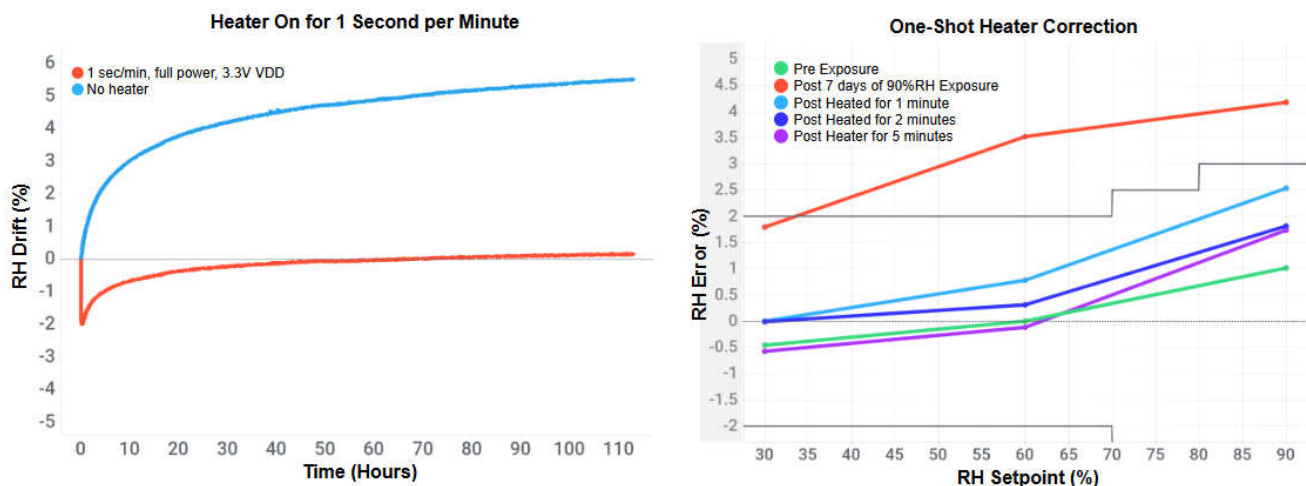


Figure 1-1. Continuous and One-Shot RH Drift Mitigation

1 Introduction

Figure 1-1 illustrates the working principle for capacitive RH sensors. Capacitive RH sensors use a polymer whose dielectric properties change as the polymer absorbs and releases water vapor. When the RH sensor is exposed to high-humidity conditions (>80%RH), the sensing polymer can become saturated with water vapor. The polymer cannot rely on diffusion to evacuate the excess water vapor when the ambient RH remains very high because rate at which water vapor is absorbed into the polymer exceeds the diffusion rate (the rate which the water vapor molecules exit the polymer) at room temperature. Hence, the RH drift will become increasingly positively shifted as the duration of high-humidity exposure grows.




Figure 1-1. How Capacitive RH Sensors Use a Polymer to Capture Water Vapor

RH drift is not the same as RH accuracy. For example, if the minimum and maximum limits for RH accuracy at a given humidity level is $\pm 2\%$ RH, that is not factored into RH drift. If an RH sensor reads 81%RH in an 80%RH environment, RH accuracy has been satisfied. When the high humidity starts to add excess water vapor into the sensor and the RH error begins to increase from the 81%RH number to 83%RH, the RH sensor has experienced a +2%RH drift.

Figure 1-2 illustrates a simplified example of RH drift. In both the blue and red curves, exaggerated "up and down" curves of RH Hysteresis are shown, to make clear that RH drift is a separate phenomenon from RH Hysteresis. The red curve set is the same RH sensor as the blue curve set, but after experiencing RH drift from a high humidity exposure. The left arrow shows how RH sensor shifts up in RH from that environmental condition. Then the right arrow pointing down demonstrates that with high temperature exposure, either through a bake or using the integrated heater inside of TI's RH sensors, the RH drift can be mitigated and removed. The integrated heater provides users with a way to bake out excess water vapor during application use.

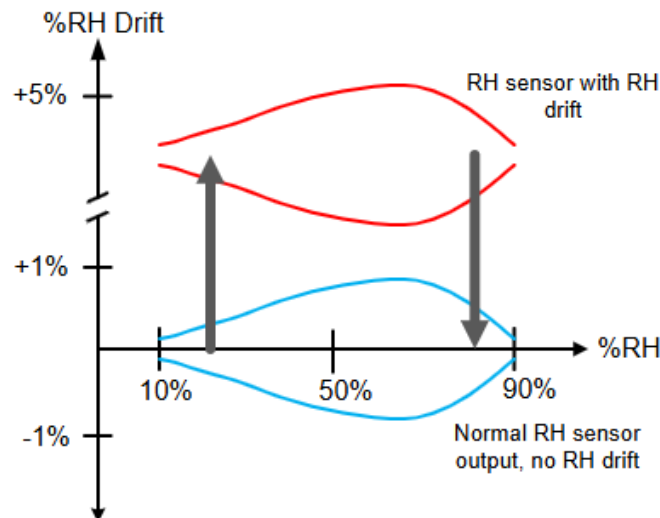



Figure 1-2. RH Drift Example

2 Integrated Heater Operation and Design Levers

An integrated heater is present on each humidity sensor in TI's portfolio. The HDC1x, HDC2x, and HDC302x are digital RH sensors with heaters that can be turned on programmatically, and the HDC3120 is an analog RH sensor whose heater is enabled/disabled with an external pin. The HDC302x has multiple available heater power levels that are configurable by the user. The heater is a resistive element within the die that draws a relatively large amount of current through the V_{DD} line to generate heat. This raises the junction temperature of the RH sensor and heats the sensing polymer so that trapped excess water vapor evaporates back into the air. Temperature and RH can continue to be read during heater operation, so users can observe how hot the sensor gets. The HDC302x heater configuration changes the resistance of the heater, so that different heater power levels can be achieved.

Heater On With Thermal Pad Not Soldered



Heater On With Thermal Pad Soldered




Figure 2-1. How Soldering the Thermal Pad Affects the Heater

Three variables dominate heater effectiveness in practice: thermal path, heater energy, and activation schedule. The thermal path determines how much of the electrical heater power becomes useful junction temperature rise instead of escaping into the PCB. This can be understood through both the thermal resistance to ambient and thermal mass of the PCB.

Table 2-1 demonstrates the difference in $R_{\theta JA}$, when the thermal pad is and isn't soldered. This table is taken from the Thermal Characteristics section of the [HDC3120 Datasheet](#), which has the same package and heater as the HDC302x, just set to a fixed heater power level of 62mW at 3.3V V_{DD} . When the thermal pad is not soldered, the $R_{\theta JA}$ nearly doubles. However, designing for a higher $R_{\theta JA}$ will also slightly increase the temperature and RH response time of the sensor. This is because now there is a higher resistance to ambient temperature and RH changes. Hence, the sensor is slower to respond to those environmental changes. If fast

response time is critical to the application, the drawbacks of not soldering the thermal pad can outweigh the benefits for some users.

Table 2-1. HDC3120 Thermal Information

THERMAL METRIC		HDC3120 DEF (WSON) 8 PINS				UNIT
		THERMAL PAD SOLDERED		THERMAL PAD UNSOLDERED		
		HEATER OFF	HEATER ON	HEATER OFF	HEATER ON	
$R_{\theta JA}$	Junction-to-ambient thermal resistance	89.3	95.0	170.5	176.3	°C/W
$R_{\theta JC(top)}$	Junction-to-case (top) thermal resistance	58.0	62.0	82.9	122.7	°C/W
$R_{\theta JB}$	Junction-to-board thermal resistance	58.0	85.0	117.7	86.5	°C/W
Ψ_{JT}	Junction-to-top characterization parameter	12.4	15.7	21.5	25.2	°C/W
Ψ_{JB}	Junction-to-board characterization parameter	57.6	61.9	117.5	121.5	°C/W
$R_{\theta JC(bot)}$	Junction-to-case (bottom) thermal resistance	37.9	42.0	–	42.0	°C/W
M_T	Thermal Mass	5.7	5.7	5.7	5.7	mJ/°C

To improve the heater temperature rise, users should target increasing $R_{\theta JA}$, the thermal resistance to ambient. When thermal resistance to ambient is greater, a greater portion of the heater power will remain in the sensor since it is harder for the heat to dissipate into the ambient air.

TI recommends minimizing the thermal mass of the PCB. A smaller PCB thermal mass will result in a larger temperature rise when activating the heater because the thermal mass acts as a reservoir for heat. A larger thermal mass will siphon off more of the heat into the PCB. Use flex PCBs or thin PCBs (<32 mil) for RH sensors to maximize ambient sensing capability and the temperature rise from the heater. However, FR4 62 mil PCBs are much more common and cheaper to build, and many users solder the thermal pad for structural stability or to optimize the temperature and RH response time. The temperature rise from the heater will increase much more due to reductions in thermal mass compared to increases in thermal resistance to ambient, so users should prioritize minimizing the PCB thermal mass.

Heater energy is the second major lever. While HDC1x, HDC2x, and HDC3120 do not have configurable heaters, the HDC302x series of RH sensors do have the option to set a desired heater power level. The default power level is relatively low if no configuration step is taken, but the *full-power* setting minimizes the internal heater resistance to maximize the heat generated. For HDC302x at 3.3V, the full-power heater setting is typically 249mW and up to 368mW maximum. Users should verify the supply and V_{DD}/GND traces can handle up to 100 mA when the heater is used. Other heater power settings in the HDC302x will reduce the final achieved heater temperature, but that can be acceptable depending on the RH drift tolerance the user defines for themselves. Furthermore, as voltage supply increases, the heater power increases as well. In that case, a lower heater power setting can be more appropriate to achieve equivalent mitigation to 3.3V full heater power for the same design.

3 Power Consumption

Using the heater also requires additional power consumption compared to normal operation. The purpose of this section is to help users understand what that average power consumption is. For HDC302x, 3.3V V_{DD} and full heater power, the typical heater power is 249mW, and can be as much as 368mW maximum, The current draw on V_{DD} can be over 100mA with those settings. This means that even if operated only briefly, the heater causes the average current to increase significantly compared to normal operations for the HDC302x. To calculate the average IQ, users can use the following equation:

$$I_{Q_{AVG}}(mA) = \left(\frac{HeaterTime_{ON}}{TotalTime}\right) * I_{Q_{HEATER}}(mA) + \left(\frac{HeaterTime_{OFF}}{TotalTime}\right) * I_{Q_{NOHEATER}}(mA) \quad (1)$$

For example, running the heater for 1 second per minute means that the heater is active for 1 second, then for 59 seconds the heater is not active. Even when the RH sensor is cooling down for 1-5 seconds after the heater is activated, the elevated current is no longer being drawn by the heater ceases. The heater-off IQ is the normal average current drawn by the HDC3020, which for a continuous conversion rate of once per second is about 2µA. The heater on current can be as high as 100mA at 3.3V supply. The resulting average current is as follows:

$$I_{Q_{AVG}}(mA) = \left(\frac{1second}{60seconds}\right) * 100mA + \left(\frac{59seconds}{60seconds}\right) * 0.002mA = 1.67mA \quad (2)$$

This is shown in [Figure 3-1](#) to help visualize the power consumption and how the heater is run in short bursts:

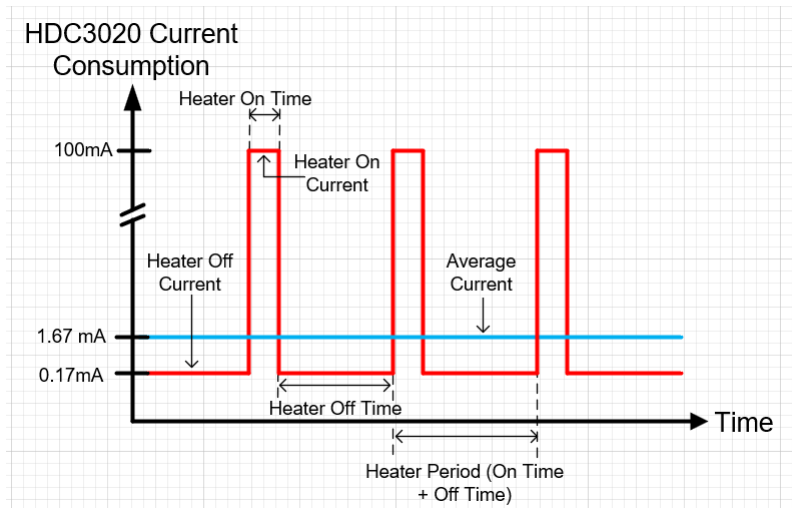


Figure 3-1. Heater Average Power Consumption Example

Users must understand the average power consumption required to run the heater and the system power budget that can be allocated to the RH sensor. The heater power setting and the voltage supply can be lowered to reduce the power consumption, but both reduces the heater's effectiveness.

To address heater activation schedule, this application note seeks to provide a starting point for users to mitigate RH errors caused by high-humidity exposure by seeking to answer two questions: how long the heater must be activated for and how often should the heater be run to achieve optimal RH drift mitigation. The findings in this application note saves users time by helping to eliminate different testing scenarios in their system while also providing proven recommendations to mitigate RH drift in common scenarios.

4 Test Conditions

This application note will demonstrate the effectiveness of using the heater to mitigate RH drift experimentally. Two scenarios were tested using the heater: a one-shot heater application that seeks to answer how much heating is needed after a discrete high-humidity event, and a continuous correction that seeks to answer what periodic heating schedule contains RH drift while the sensor remains in a high-humidity environment.

One-shot heater correction must be used in situations where the RH sensor is exposed to high RH for periods of time but is not a normal occurrence. If the high-humidity condition is not a normal occurrence, there is no assumption of a periodic heating routine to continuously mitigate RH drift. The high-humidity will cause some amount of RH drift, and that can be corrected by using the heater after the high-humidity condition has subsided. RH drift can resolve on its own slowly over time, so using the heater will accelerate the diffusion of water vapor out of the sensor back into the air to improve the RH accuracy as quickly as possible.

Continuous heater correction must be used when high-humidity conditions are frequent occurrences or even constant. The heater will not just be run as needed like in one-shot heater correction but needs to be run on a consistent time interval to continually mitigate the RH drift as it is occurring in real time. Continuous heater correction does not mean that the heater is on continuously, but rather the high humidity exposure is continuous.

The test board shown in [Figure 4-1](#) is a PCB with a single HDC3020, is composed of FR4 core, and is 62 mil thick. This PCB also has the thermal pad soldered. This PCB is a non-optimal design according to TI's recommendations, but this was a deliberate choice. This PCB is a more common design that better reflects what most users can encounter in system testing and hence the results are likely to be broadly applicable to most users.




Figure 4-1. HDC3020 Individual Test Board

Most tests were run at 3.3V, the nominal V_{DD} supply for HDC3020. Lowering the voltage supply will reduce the heater power, and increasing the voltage supply increases the heater power. 25°C was chosen as the temperature for all testing and stressing, since as a common *room* temperature the findings can be most broadly applicable. The ambient temperature affects the heater's performance as well. For a given design, the heater achieves the same final steady-state temperature regardless of the ambient temperature. This means that lower ambient temperatures results in a larger net temperature rise, and higher ambient temperatures results in a smaller net temperature rise. Users can use this 25°C data as a baseline for their expectations, then adjust their experiments accordingly to match the application temperature.

5 One Shot Heater Correction

One-shot heater recovery is the appropriate mode when the sensor experiences an occasional high-humidity event rather than a regularly recurring one. The experiment in Table 1 emulates that use case: measure initial RH error, soak the devices at 25°C and 90%RH for seven days, measure the shifted RH error, then apply one full-power heater pulse of varying duration and measure RH error again across 30%RH, 60%RH, and 90%RH. The results captured in Figure 5-1 are not captured over time, but over RH setpoint. Because the high-humidity exposure is temporary, the results in Figure 5-1 must be treated as a recovery experiment, not merely a sequence diagram.

Table 5-1. One-Shot RH Drift Mitigation Procedure

Step	Description	Supply Voltage (V _{DD})	Heater Power	Temperature	%RH
1	Initial RH error test	3.3V	n/a	25°C	30%, 60%, 90%
2	90%RH soak for 7 days	n/a	n/a	25°C	90%
3	Post exposure RH error test	3.3V	n/a	25°C	30%, 60%, 90%
4	Split parts into 3 groups	n/a	n/a	25°C	n/a
5	Apply heater for 1 minute to group 1, 2 minutes to group 2, 5 minutes to group 3	3.3V	Full power, approximately 250mW	25°C	50%
6	Post heater RH error test	3.3V	N/A	25°C	30%, 60%, 90%

90%RH was chosen because it is a known high-humidity condition that induces RH drift in a matter of hours without risk of condensation occurring during the experiment. This was done for seven days since that is long enough to induce enough RH drift to push the HDC3020 out of the RH accuracy specification. The HDC3020 has configurable heater power levels, so the full power setting was chosen to achieve the largest temperature rise possible. A larger temperature rise increases the diffusion rate of the water vapor molecules, hence providing better RH drift mitigation. One minute is usually sufficient to achieve a steady-state temperature rise. To explore the question of how long the heater must be run at steady state, some devices were run for longer to see how much benefit was gained.

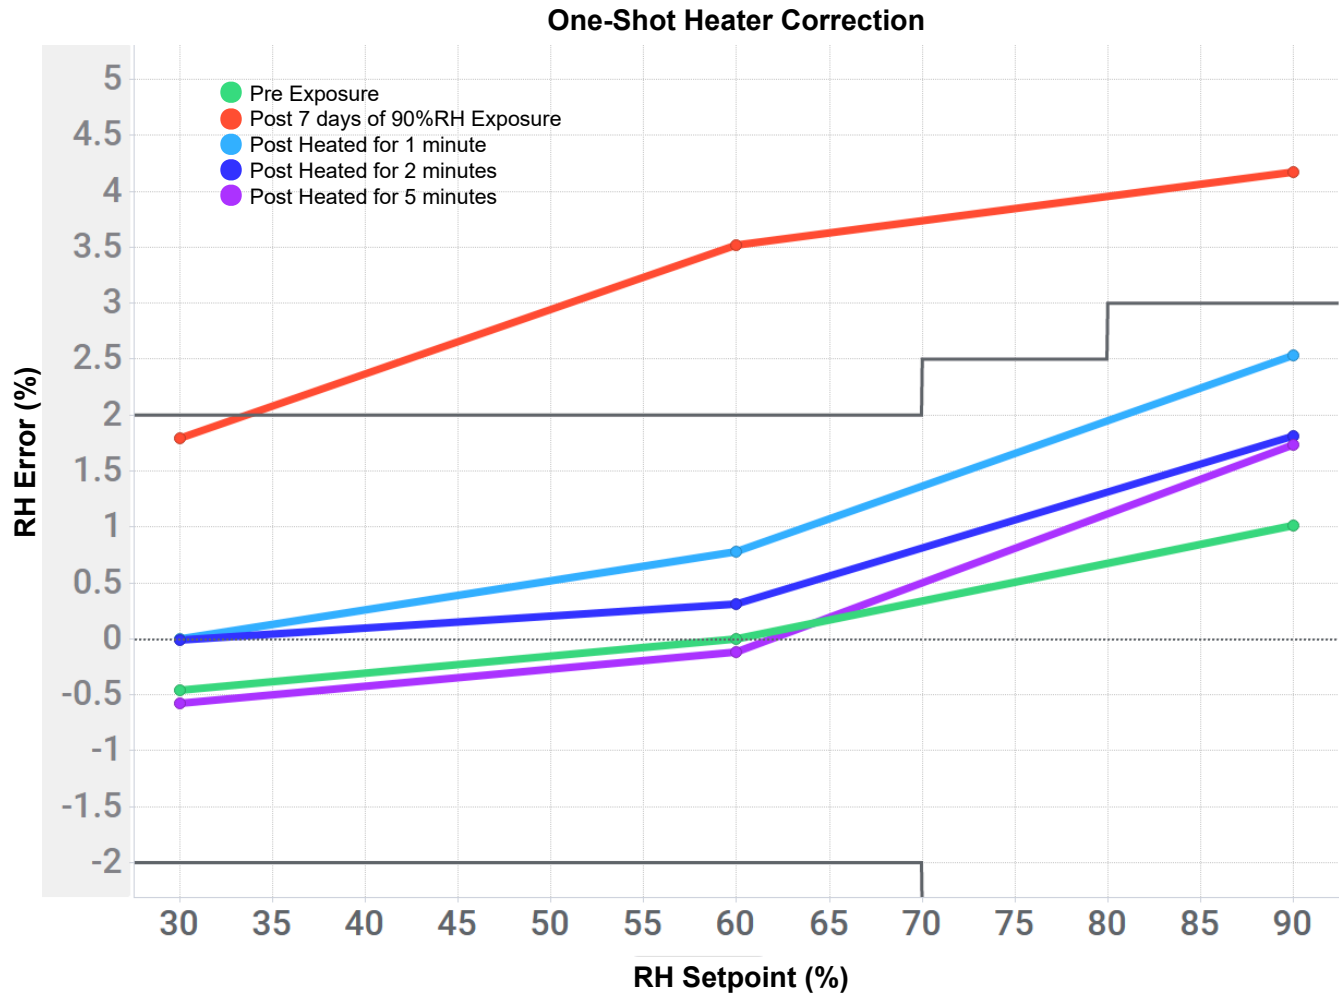


Figure 5-1. One-Shot Heater Correction: RH Error vs RH Setpoint

In [Figure 5-1](#), RH error is plotted against RH setpoint. This is not showing how the RH is changing from the initial result over time but instead showing how the RH error has changed across RH when the different steps are applied. Each line represents an average of the results for 17 HDC3020s. The RH error was measured at 25°C across multiple RH levels. The green line shows the initial time-zero RH error. The red line shows the RH error after 7 days of exposure at 25°C/90%RH, with no heating done to remove the RH drift. The blue & purple lines show the RH error after running the heater for variable durations.

One minute of full-power heating already pulled the averaged HDC3020 much closer to the original performance, but two minutes delivered a materially better result at 60%RH and 90%RH. Five minutes only marginally improved the 90%RH result further. The strongest engineering takeaway is therefore simple: on this board, 2 minutes at full heater power is the best starting point for one-shot recovery, while 5 minutes is better treated as an optional high-RH refinement when power and cycle time are less constrained. If judging solely at ambient RH levels between 30%RH and 60%RH, 1 minute of heating offers almost the same level of correction as the longer heating times while saving power consumption by running the heater for at least half as long as the other options.

6 Continuous Heater Correction

Periodic heater correction is the appropriate mode when high humidity is expected regularly or continuously. HDC3020 devices were operated in groups of five parts under 25°C / 90%RH conditions, with three heated devices and two unheated controls. [Figure 6-1](#), [Figure 6-2](#) and [Figure 6-3](#) show the average results for the two groups of HDC3020s. These graphs show the RH drift over time to show how the RH drift is mitigated in the heated devices as the high-humidity exposure continues. Below are the different tested scenarios:

Table 6-1. Continuous High-Humidity Exposure Heater Correction Attempts

Scenario	Supply Voltage (V _{DD})	Heater Power Setting	Ambient Temp	Heater Current	Heater Power	Heater Time On	Heater Time Off	Heater Period	Final temperature due to heater activation	Test duration
1	3.3V	Full power	25°C	70mA	231mW	1 minute	59 minutes	1 hour	51.7°C	72 hours
2	3.3V	Full power	25°C	68mA	225mW	5 minutes	55 minutes	1 hour	55.7°C	96 hours
3	3.3V	Full power	25°C	70mA	231mW	15 seconds	59 minutes, 45 seconds	1 hour	43.7°C	72 hours
4	3.3V	Default Power	25°C	19mA	62mW	1 second	59 seconds	1 minute	7°C	90 hours
5	1.8V	Full power	25°C	40mA	72mW	1 second	59 seconds	1 minute	8.8°C	46 hours
6	3.3V	Full power	25°C	80mA	264mW	1 second	59 seconds	1 minute	27.8°C	113 hours
7	3.3V	Full power	25°C	80mA	264mW	1 second	29 seconds	30 seconds	27.8°C	67 hours

Default heater power is the HDC3020 heater power setting if no configuration step is taken. The results in this experiment are specific to the PCB shown in [Figure 4-1](#).

For all continuous heater correction scenarios, the test and measurements always followed this sequence:

1. Measure and record HDC3020 & calibrated reference temperature and %RH
2. Activate heater at desired power level for specified amount of time
3. Turn off the heater.
4. Wait until the "heater off time" has elapsed.
5. Return to step 1
 - a. This loop continues until the total test duration has been completed.

When the heater is deactivated, it will begin to cool down. The time for the sensor temperature to match the environment again is variable, depending on each individual sensor and the final temperature the heater was able to achieve (as discussed above, the temperature rise of the heater can be determined by voltage supply, heater power setting, and PCB design). This sequence ensures the HDC3020 is cooling off for the duration of the *heater off time* before the next measurement occurs. If the measurement is taken before the HDC3020 has fully cooled down (returned to ambient temperature), the RH results is artificially low, and the effectiveness of the RH drift mitigation can be overestimated.

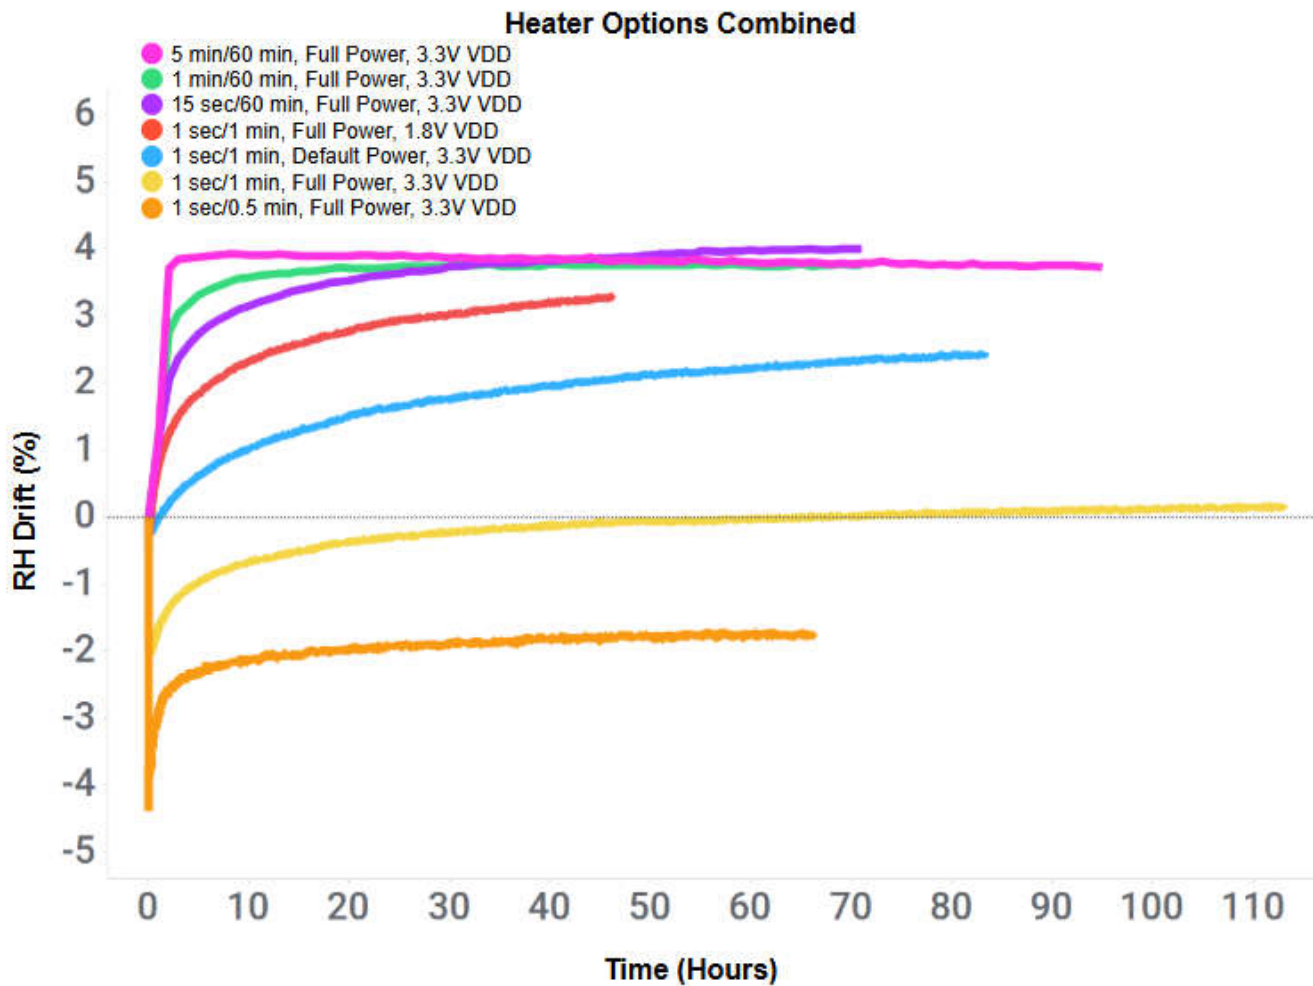


Figure 6-1. Continuous Heater Correction All Scenarios Combined

In [Figure 6-1](#), each continuous heater method is displayed on one graph. The red, blue, yellow, and orange curves are different configurations of running the heater for one second per minute, and all of those appear to be more effective than running the heater for longer durations every hour. The ideal option for this PCB design is to activate the heater for one second every minute at 3.3V V_{DD} with full heater power enabled to minimize RH drift.

[Figure 6-2](#) and [Figure 6-3](#) show the individual results of all seven continuous heater correction scenarios, with the heated devices plotted against the unheated controls. Each graph below shows the shift in RH error from the initial result, with the heated and control devices averaged for simplicity. Heated parts are colored red, controls with no heating applied are colored blue.

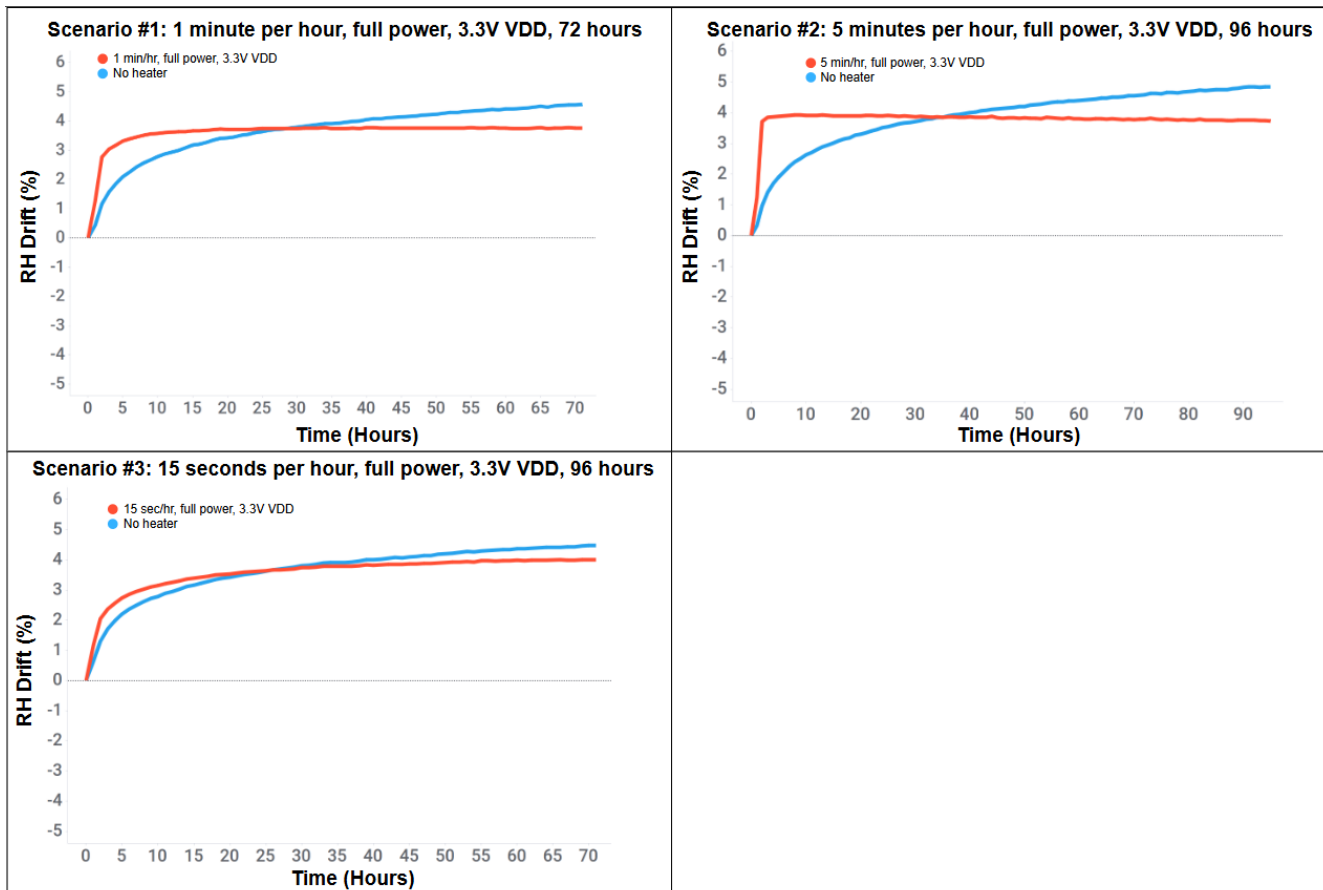


Figure 6-2. Continuous Heater Correction Scenarios 1, 2 and 3

Scenario 1 shows the heated DUTs still experience significant RH drift, but they plateau between 3-4%RH error. The controls continue to drift under high-humidity exposure. The heater aids in reducing the amount of RH drift, but the remaining RH drift is only about 1% lower than the unheated parts, which is very unlikely to be satisfactory. Scenario 2 shows the heated DUTs still experience large RH drift, but they also plateau to about ~4%RH error. The heated devices also show a slow decline in RH drift as time goes on, but that decline will take place over a long and undetermined amount of time. Scenario 3 shows that 15 seconds per hour provides less RH drift mitigation than Scenario 1, which is expected. In Scenario 3 the heated devices do not demonstrate any plateauing of RH error; hence this option must be considered ineffective and is not recommended.

In Scenarios 1-3 at the very beginning of each run the RH drift of the heated parts is accelerated compared to the control units. This is most pronounced with Scenario 2, where the heater is on for 5 minutes per hour. This is occurring because of diffusion of water vapor into the sensing polymer. When the heater is initially activated, it removes water vapor molecules from the sensing polymer which dries out the sensor as intended. However, the drier the polymer is, the faster the diffusion of water vapor molecules into the sensor is. Hence the RH drift accelerates beyond what rate it would normally change at. After enough saturation a balance is achieved where the RH drift is no longer accelerating but is just replenishing the water vapor lost to heat. In the case of Scenario 2, the polymer is losing more water vapor to heat than it can replenish in 55 minutes, so the RH drift is decreasing. While the initial RH drift acceleration appears worse than doing nothing, the goal with continuous heater correction is long term performance, not how the sensor behaves in the first few hours. This behavior can be entirely avoided when running the heater in shorter one second bursts.

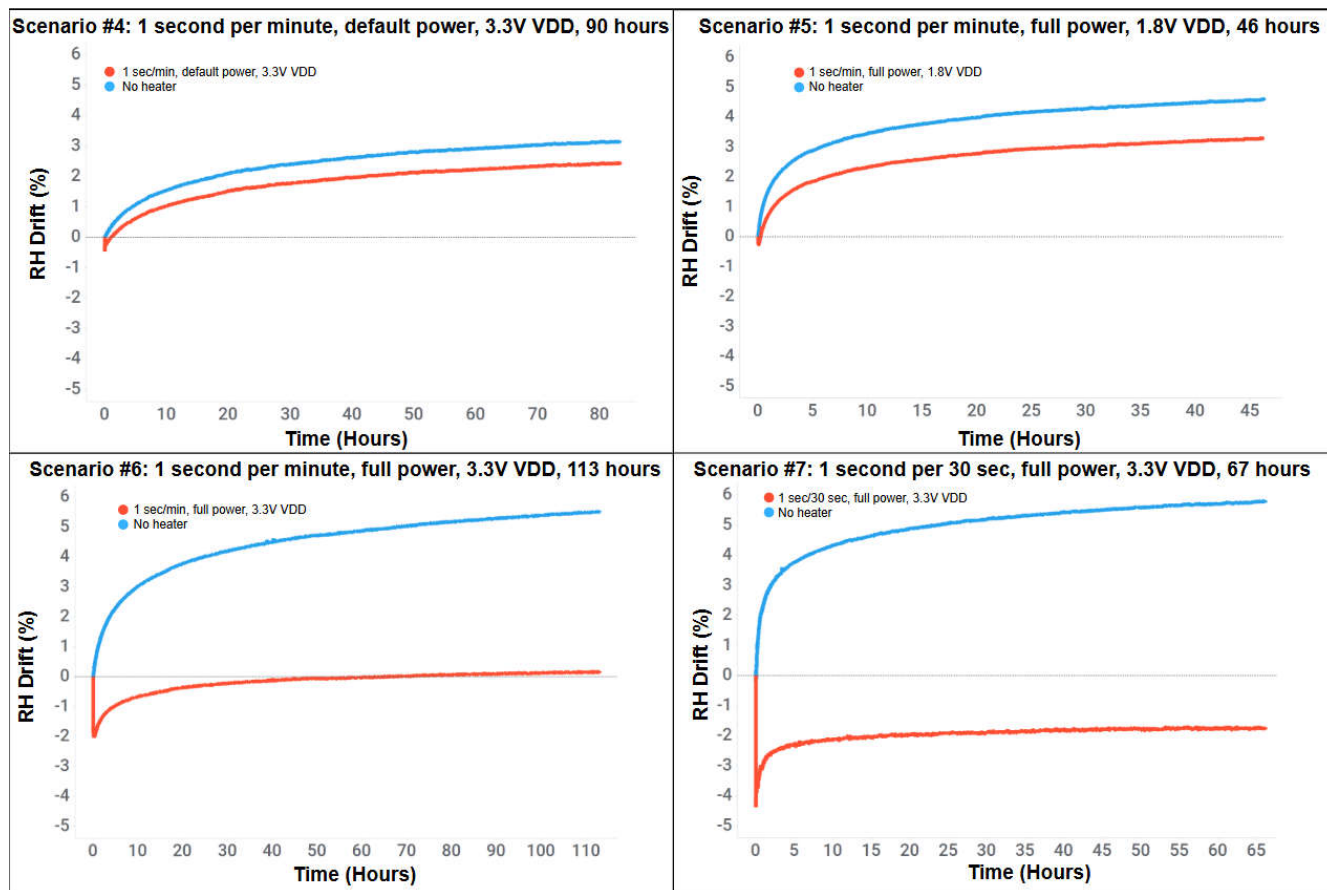


Figure 6-3. Continuous Heater Correction Scenarios 4, 5, 6 and 7

Scenario 4 shows that the heater reduces the overall level of RH error, both at the beginning of the test and after several days. After nearly 4 days, the RH error has been contained to about +2.5%RH, but it is still rising at the end of the test. Using the default heater power will save on power consumption but comes at the cost of less RH error mitigation. Activating the heater more often than once per minute may improve RH drift mitigation at the default heater power level. Scenario 5 shows slightly worse performance than the default heater power setting at 3.3V. At 1.8V the heater power is much lower so the heater impact on RH error is worse. If operating at 1.8V, consider activating the heater more often than once per minute to improve the RH drift mitigation. The test was stopped at 46 hours to save time since the results were not demonstrating adequate correction.

Scenario 6 shows that the RH drift is kept much closer to 0% with this setting. RH drift remains near 0% after nearly five days at 90%RH, and the rate of increase is nearly flat. In Scenario 7, the RH error of the heated parts drops quickly then stays at a consistent RH level with a -2%RH drift. While this setting may cause a negative RH drift, the effects of positive RH drift effects excess water vapor absorption are very well mitigated. However, compared to the results in Scenario 6, current consumption has effectively doubled because the heater pulse is occurring twice as often. For this PCB, using 1 second per minute at full power provides the optimal RH drift mitigation.

As seen in Scenarios 4-7, there is an initial drop in RH drift in the first few minutes of activating the continuous heater, with Scenario 7 being the most significant. This occurs due to the initial heater pulse removing the water vapor in the sensor before the high humidity can start shifting the RH readings. This drop is only temporary, vanishing after a few minutes of heater use, users must again focus on the long-term correction benefits and not the transient results at the beginning of the test.

7 Summary

Table 7-1. One-Shot Heater Correction Results Summary

%RH Setpoint	Initial %RH Error (%RH)	Post 90%RH Exposure %RH Error (%RH)	Post 1 minute of full power heating %RH Error (%RH)	Post 2 minutes of full power heating %RH Error (%RH)	Post 5 minutes of full power heating %RH Error (%RH)
30%RH	-0.46	1.79	-0.01	-0.01	-0.58
60%RH	0	3.52	0.77	0.31	-0.12
90%RH	1.01	4.16	2.53	1.81	1.73

Table 7-2. Continuous Heater Correction Results Summary

Scenario	Continuous Heater Timing	Final RH Drift (%RH)
Scenario 1	1 minute per hour, full power, 3.3V V_{DD}	3.75
Scenario 2	5 minutes per hour, full power, 3.3V V_{DD}	3.73
Scenario 3	15 seconds per hour, full power, 3.3V V_{DD}	4.00
Scenario 4	1 second per minute, default power, 3.3V V_{DD}	2.42
Scenario 5	1 second per minute, full power, 1.8V V_{DD}	3.27
Scenario 6	1 second per minute, full power, 3.3V V_{DD}	0.18
Scenario 7	1 second per 30 seconds, full power, 3.3V V_{DD}	-1.77

There are many variables that will dictate the effectiveness of using the heater to mitigate RH error from high-humidity conditions. Testing by users is necessary to determine the right heater setting and schedule for their system. To optimize the heater performance, users need to focus on the variables that are within their control to best mitigate RH drift. The two main areas where users have control are PCB design and the heating schedule.

The PCB design comes down to two factors, thermal resistance to ambient ($R_{\theta JA}$) and thermal mass (M_T). The thermal mass of the PCB will also determine the final rise time and can have a very large impact on the heater performance. PCBs with smaller thermal mass will have a reduced heatsinking effect when the heater is active, resulting in much more heat being retained in the sensor. This difference in heat can be seen in Table 6, where the Delta T ($^{\circ}C$) column shows that regardless of V_{DD} , a flex PCB has a much larger temperature rise than the rigid FR4 PCB shown in [Figure 4-1](#).

The thermal resistance to ambient will determine how well the heat from the heater is retained in the sensor. A larger $R_{\theta JA}$ results in more heat retained, which means a better temperature rise. A lower $R_{\theta JA}$ means more heat dissipates into the air, so less is retained in the sensor resulting in a worse temperature rise. Soldering the thermal pad decreases the $R_{\theta JA}$ since it provides a low thermal resistance path for energy to escape the sensor. Therefore, to achieve a higher final temperature with the heater it is recommended to leave the thermal pad unsoldered. Practically speaking, if the temperature rise from ambient is larger, the heater will likely not need to be activated as often during continuous correction. In contrast, if the temperature rise is small, the heater will be ineffective at RH drift mitigation. This is evidenced by Scenario 4, where the default heater power level cannot provide sufficient heating to remove the water vapor as fast as it is entering the sensor.

Table 5 illustrates that the calculated thermal resistance for a flex PCB is much higher than that of the rigid FR4 test PCBs used in all the heater correction data collected for this document. Coupled with a much lower thermal mass, the flex PCB results in a much larger temperature rise when using the heater. Table 5 also shows that even for the rigid FR4 PCB, increasing the supply voltage from 3.3V to 5V also significantly increases the final heater temperature. For the purposes of this document, all RH sensors on the PCBs from [Table 7-3](#) have soldered thermal pads.

Table 7-3. Heater Rise and Thermal Resistance Calculations for Rigid & Flex PCB

		V _{DD} (V)	I Avg (mA)	P (mW)	Delta T (°C)	R _{θJA} (°C/W)
Heater Off	Flex	5	0.17	0.87	0.61	700.17
		3.3	0.16	0.54	0.38	694.40
	Rigid FR4 PCB	5	0.17	0.87	0.22	253.03
		3.3	0.16	0.54	0.12	229.49
Heater On	Flex	5	23.5	117.50	78.89	671.44
		3.3	17.5	57.75	39.42	682.74
	Rigid FR4 PCB	5	26.2	131	32.39	247.21
		3.3	18.7	61.71	14.53	235.58

Lower heater power levels, whether through not using the full heater power configuration or using a lower voltage like 1.8V, reduce the ability of the heater to mitigate RH error increases over time. The experimental results using FR-4 PCBs showed that neither the 1.8V full heater power nor the 3.3V default heater power provided sufficient heating to prevent the RH drift from growing during the test. In situations where users are locked into using lower heater power levels due to current consumption constraints or lower V_{DD} levels, further optimizations with PCB layout and assembly can be done to increase the efficiency of the heater. For example, using a flex PCB or even thinner FR4 PCB (31 mils for instance) will reduce the thermal mass of the PCB, allowing more heat to be retained within the RH sensor and hence resulting in a larger temperature rise.

To successfully mitigate RH drift, these PCB design elements need to be deliberate choices to balance heater capability with the broader system constraints instead of arbitrary decisions or afterthoughts. Once a PCB design has been settled on, optimizing for the right heater schedule may require iterative testing to use the heater as often as necessary, while not overusing to minimize power consumption if that is a concern. Users should have some understanding of the end application temperature and humidity conditions that the sensor is deployed in. These conditions should be replicated as closely as possible in a laboratory environment, ideally with a temperature and humidity chamber.

If recurring condensation, wash exposure, or water ingress risk is part of the use case, consider protected-cover device variants and mechanical shielding as part of the design rather than relying on heater action alone.

For continuous high-humidity correction, performing frequent short bursts of heat more often proved to be much more effective than longer durations of heater activation with longer time intervals. This demonstrates that users do not need to allow the heater to reach a steady-state elevated temperature to be effective. Running the heater in one-second pulses every minute proved an effective strategy to mitigate RH drift, and running the heater for 1 second per minute (at 3.3V and full heater power) was the most effective case, and where TI recommends users begin when debugging how to mitigate continuous RH drift with the heater. This gave the best balance of RH drift mitigation while efficiently utilizing the heater and avoiding excess current consumption.

For one-shot heater correction, where RH drift happens due to an isolated high-humidity period, a longer heating duration of 2 or 5 minutes at full heater power proved better than no correction or just 1 minute. 5 minutes of heating was a marginal improvement compared to 2 minutes, but 2 minutes of heating was sufficient to return the sensor to specification at 50%RH. Finding the right amount of time to heat will again require testing due to system differences between users, but achieving a steady state heating temperature for several minutes is most useful for one-shot RH drift recovery.

7.1 Step-By-Step Walkthrough

The following guide provides instructions on how users can apply the findings and measurement techniques of this application note into their design. The RH drift numbers and current consumption levels are estimated for the purposes of these examples to demonstrate the iterative analysis approach; testing is required to determine the true values in a user's design and environmental conditions. TI recommends that RH sensors that have not been drifted before for each test. Using the same parts again through the testing iterations may have compounding RH drift effects that makes it difficult to determine how successful a given heater correction scheme is.

For a first example, consider a design that is using the HDC3020. The HDC3020 is deployed in an environment of 25°C and 90%RH. Because of the high RH, some RH drift is presumed to occur during environmental

exposure, but the exact amount is unknown. In addition, a limited power budget dictates an average current draw no greater than 1mA.

Step 1: Define requirements

- Acceptable RH drift: At this time, the actual RH drift the HDC3020s experience in this environment is unknown, however a pass/fail criterion must be defined. In this case, RH drift must be no greater than +2%RH.
- Power Budget: Average current draw must be no greater than 1mA.

Step 2: PCB Considerations

- Ideally the PCB design will be optimized for heater performance, as detailed in both the [Section 2](#) and [Section 7](#). However, often users are locked into at least some aspects of the PCB design for reasons that extend beyond RH drift mitigation optimization. In this example, the system design must utilize a 62-mil thick PCB made of FR4. However, optimizations can still be made.
 - **Minimizing thermal mass:** If the PCB structure could be modified, making the PCB thinner would lower the thermal mass. Since it cannot be made thinner or on a flex PCB, creating a tab cutout where the sensor is placed will reduce the amount of PCB material surrounding the sensor, and slightly reduce the thermal mass, hence reducing the ability of the heater's energy to be siphoned away from the sensor.
 - **Maximizing thermal resistance to ambient:** Since the thermal mass is difficult to modify if the PCB structure cannot change, maximizing the thermal resistance to ambient will help the heater temperature rise. Choosing to not solder the thermal pad in this instance will allow more heater power to be retained in the sensor, improving the RH drift correction performance.
 - **Note:** Minimizing thermal mass and maximizing thermal resistance to ambient is not a "one or the other" case. Both must be done when possible to get the best heater performance possible.

Step 3: Choose Initial Approach: Continuous Heater Correction

- In this case, the HDC3020 is being exposed to the high humidity conditions constantly, so finding a way to use the heater periodically during the high humidity exposure is appropriate. A one-shot heater approach would be appropriate if the high humidity exposure was not a common occurrence so a single correction would be sufficient.

Step 4: Initial Continuous Correction Settings & Testing

- Based on the findings of this application note, use these starting conditions:
 - **Voltage & Heater Power:** This app note recommends users start with a 3.3V voltage supply and using the full heater power setting. However, the average current consumption of those settings together, 1.67mA, will exceed the established current limit of 1mA (see [Section 3](#) for details on how this current is arrived at). Hence, either voltage supply or heater power level must be lowered to meet the current limit. TI recommends modifying the heater power first. This is because the heater power level is set programmatically and is relatively easy to change between testing iterations. Changing voltage supply is another option as well. If the first test does not demonstrate sufficient RH drift mitigation and the heater is already at full power setting, increasing the voltage supply will increase the heater power.
 - To lower the current consumed, the heater power will be lowered to half power setting. Half power setting is chosen in this example because the power level is detailed in the datasheet making the heater current consumption for a given supply voltage easy to calculate. Voltage supply is left at 3.3V for now, to have a better comparison to the results in this application note.
 - **Schedule (Starting Point):** A 1-second heater pulse is initially set to occur every minute.
 - At this heater schedule with 3.3V supply and half heater power, the average current consumption is about 0.69mA.
- **Setup:** The testing environment must be able to replicate the anticipated environmental conditions that will cause RH drift or be done in the environment if available. TI recommends performing these tests in a temperature and humidity chamber. The RH sensors need to be evaluated against an ISO 17025 calibrated humidity reference.
- **Monitoring:** Test the RH sensors for 72 hours. 72 hours provides enough time for most of the RH drift to be induced while not being excessively long. RH readings must be collected once a minute, according to the

methodology described in [Section 6](#) of this application note. The drift test must also include controls which receive no heater correction to understand what the RH drift will be from this environmental condition.

- **Results after 72 hours:** the control sensors have drifted by +5%RH, and the heated devices exhibit a positive RH drift of about +2.5%.

Step 5: Iteration & Optimization

- The average current consumption requirement has been fulfilled, but the RH drift remains too high.
- Based on these results, either the heater power must be increased, or the heater schedule must be modified to have the heater run more often.
 - When multiple variables could be changed to impact the RH drift outcome, start with the easiest variable to change. In this case, changing the heater schedule to have it run slightly more often is easier. Increasing the heater power level to another setting between full and half power will have an unknown current consumption that will need to be characterized to make the determination if its appropriate for the power budget in this system. It is easier to leave the heater power at half power and have the heater pulse for slightly more often to evaporate the excess water vapor more frequently.
 - The calculation shown in Equation 1 is done again, this time with a shorter heater off time. Using 1 second pulses every 30 seconds gives an average current draw of 1.38mA, which is still too high. Increasing the heater off time to 45 seconds gives an average current draw of 0.92mA. This is within the average current draw limit.
- The heater is left at half power but now it is activated every 45 seconds instead of every 60 seconds. A second 72-hour test is conducted.
 - This test has improved RH drift mitigation, now limiting the RH drift to +1.5%RH after 72 hours. The average current consumption is 0.92mA. The RH drift has improved further, and the current consumption remains within the limit.

Based on the RH drift testing and the system limitations, using the continuous heater correction method with a 1 second heater pulse every 45 seconds at half heater power provided the best balance between RH drift mitigation and power consumption in this scenario.

Some applications require temperature and humidity environments that are more extreme than what has been presented in this application note so far. When high temperature is combined with high humidity, RH drift will become worse because the amount of water vapor in the air is so much higher than at 25°C/90%RH, so a much higher rate of diffusion is necessary to keep that water vapor from accumulating in the sensor. In these conditions though, the heater effectiveness is reduced because the high ambient temperature will result in a smaller temperature rise. In this scenario, system constraints such as voltage supply limits or current consumption limits often need to change to find a heater power level and schedule that can produce useful RH drift mitigation:

- **Power Consumption:** Increasing the limit of how much power can be consumed so that the heater can be run more often, for longer, and/or at a higher power level.
- **Voltage:** 3.3V supply is too low to generate the temperature rise necessary for the heater to have enough of a temperature rise to diffuse the excess water vapor at a meaningful level. The supply voltage must be increased to 5V or even 5.5V to increase the power of the heater.
- **Schedule:** The heater could be run every 30 seconds instead of every minute, or for a slightly longer pulse duration.
- **PCB:** The PCB may need to be changed as well to optimize the heater performance. Start by using a new sensor on the PCB, but without the thermal pad soldered. If that is insufficient, a thinner PCB or a flex PCB may be needed to get a better temperature rise out of the heater.
- **RH Drift requirements:** Depending on which of these requirements can and can't be changed, the amount of tolerable RH drift may need to be adjusted in the system if the heater proves insufficient to mitigate RH drift.

Continue to iterate over what requirements can be changed, starting with the easiest to adjust, and evaluate the RH drift of the heated devices to see what the best RH drift mitigation achievable is. Some environmental conditions are more challenging like shown in the above example, but regular heater usage can still help improve the RH drift in many environments.\

8 References

1. Datasheets
 - a. [HDC1080](#)
 - b. [HDC2080](#)
 - c. [HDC3020](#)
 - d. [HDC3120](#)
2. Silicon User Guides
 - a. [HDC2x Silicon Users Guide](#)
 - b. [HDC3x Silicon Users Guide](#)
3. Application Notes
 - a. [Humidity Sensor Comparison Guide](#)
 - b. [TI Humidity Sensors: Programming & Integration Guide](#)
 - c. [How to Debug RH Accuracy Issues in RH Sensors](#)

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you fully indemnify TI and its representatives against any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to [TI's Terms of Sale](#), [TI's General Quality Guidelines](#), or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products. Unless TI explicitly designates a product as custom or customer-specified, TI products are standard, catalog, general purpose devices.

TI objects to and rejects any additional or different terms you may propose.

Copyright © 2026, Texas Instruments Incorporated

Last updated 10/2025