Temperature stabilization system with millikelvin gradients for refractometry

Speaker: Patrick Egan, National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg, MD 20899-8211

Authors: Patrick Egan and Jack A. Stone, National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg, MD 20899-8211

Refractometry of air is a central problem for interferometer-based dimensional measurements. Refractometry at the 10^{-9} level is only valid if air temperature gradients are controlled at the millikelvin (mK) level. Very precise tests of second-generation National Institute of Standards and Technology (NIST) refractometers involve comparing two instruments (two optical cavities made from ultralow expansion glass) that are located in nominally the same environment; temperature gradients must be kept below a few millikelvin to achieve satisfactory precision of these tests. In this paper we describe a thermal stabilization scheme that maintains < 1 mK thermal gradients over 100 h in a $0.5\,\mathrm{m} \times 0.15\,\mathrm{m} \times 0.15\,\mathrm{m}$ volume. Our approach uses passive (aluminum envelopes and foam insulation) and active (thermistors, foil heaters, and proportional–integral–derivative control) temperature stabilization. Thermal gradients are sensed with thermocouples and a nanovoltmeter and switch; the reference junctions of the thermocouples being in thermal contact with a thermistor temperature standard. Our < 1 mK gradient performance is limited by the accuracy of the nanovoltmeter and switch.

1 Introduction

High precision measurement of the refractive index of air can be achieved by locking the wavelength of a laser in resonance with an optical cavity [1, 2, 3]. A laser is in resonance with a cavity, as in Figure 1, when an integer number N of half-wavelengths $\lambda/2$ equals the cavity length L. When a laser is locked to the resonance of a dimensionally stable cavity, the wavelength $\lambda = c/(n\nu)$ is fixed and the frequency of the laser ν will track the refractive index n. Locking lasers to cavities and measuring optical frequency can be extended to determine refractive index absolutely. In this approach, a resolution of 10^{-12} in laser frequency is relatively straightforward to achieve, but when it comes to determining refractive index, uncertainties in mechanical and thermal effects can be several orders of magnitude larger.

The refractive index of air has a temperature dependence of about $10^{-6}\,\mathrm{K}^{-1}$ at atmospheric pressure; testing the accuracy of a refractometer to the 10^{-9} level means measuring temperature (and gradients) to the millikelvin (mK) level. A lesser effect is thermal expansion of the cavity: ultralow expansion glass is specified with a coefficient of thermal expansion of $(0 \pm 3 \times 10^{-8})\,\mathrm{K}^{-1}$ from 5 °C to 35 °C, so maintaining the dimensional length of a 0.1 m cavity to the 10^{-9} level means keeping thermal fluctuations below 3 mK. However, we have investigated the thermal expansion effects of our cavities and found the coefficient of thermal expansion to be $3 \times 10^{-9}\,\mathrm{K}^{-1}$ at (20 ± 0.5) °C with a temperature of zero expansion at 19.1 °C, so thermal expansion of the cavity is a much less significant problem than the temperature dependence of refractivity. Another

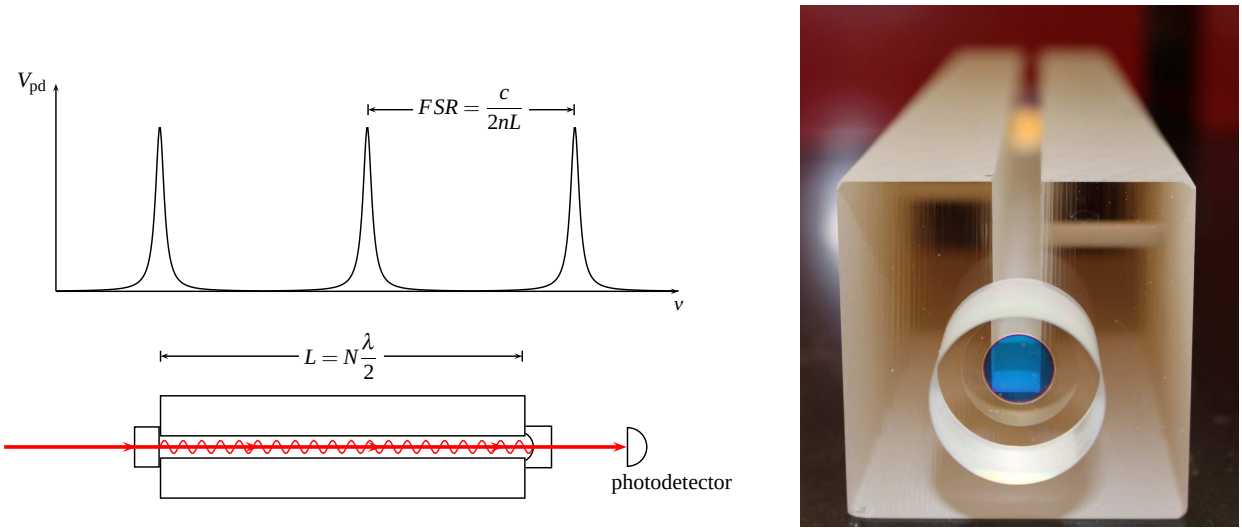

Figure 1: As a laser is swept in frequency, spikes of light intensity are observed at the output of the cavity when the laser achieves resonance. The frequency spacing between these intensity spikes is called the free spectral range.

effect is thermally induced distortion of the cavity mirrors owing to differing material properties between the mirrors and spacer [2, 4]. This effect can also approach the $10^{-8}\,\mathrm{K}^{-1}$ level, but since our second-generation refractometer has a spacer and mirrors both made from ultralow expansion glass, we expect this effect to be much smaller; this is borne out experimentally by observing that a short cavity and a long cavity both have nearly the same thermal behavior, whereas any end-effects should be more pronounced for the short cavity. Lastly, and for completion, thermal effects from the laser beam, such as thermoelastic deflection and laser induced temperature gradients, are negligible ($< 10^{-12}$).

The standard approach to temperature stabilization is to use several stages of passive enclosure (high thermal conductivity metal) in combination with active temperature control at various points and levels of the enclosures. In small volumes ($10^{-4}\,\mathrm{m}^3$) this method has kept thermal gradients to the tens of microkelvin level [5]. Our refractometer employs two optical cavities of lengths 0.15 m and 0.33 m in a vacuum chamber and this requires millikelvin temperature gradients in a volume of $(0.5 \times 0.15 \times 0.15) = 10^{-2}\,\mathrm{m}^3$. The following sections describe how we achieved this.

2 Apparatus

Temperature control was achieved by several stages of enclosure, as shown in Figure 2. Outermost, 45 mm of rigid foam acted as an insulating barrier. Inside the foam box was an actively controlled 12.7 mm thick aluminum envelope. The temperature at eight points around this envelope was sensed with thermistors (98 kΩ at 25 °C). At each of these points the temperature was proportional–integral–derivative (PID) controlled within a millikelvin with foil heaters. Each of the eight heater segments consisted of five 100 mm × 300 mm foil heaters wired in parallel giving a total resistance of 12 Ω, through which the temperature controllers could drive a maximum of 9 V.

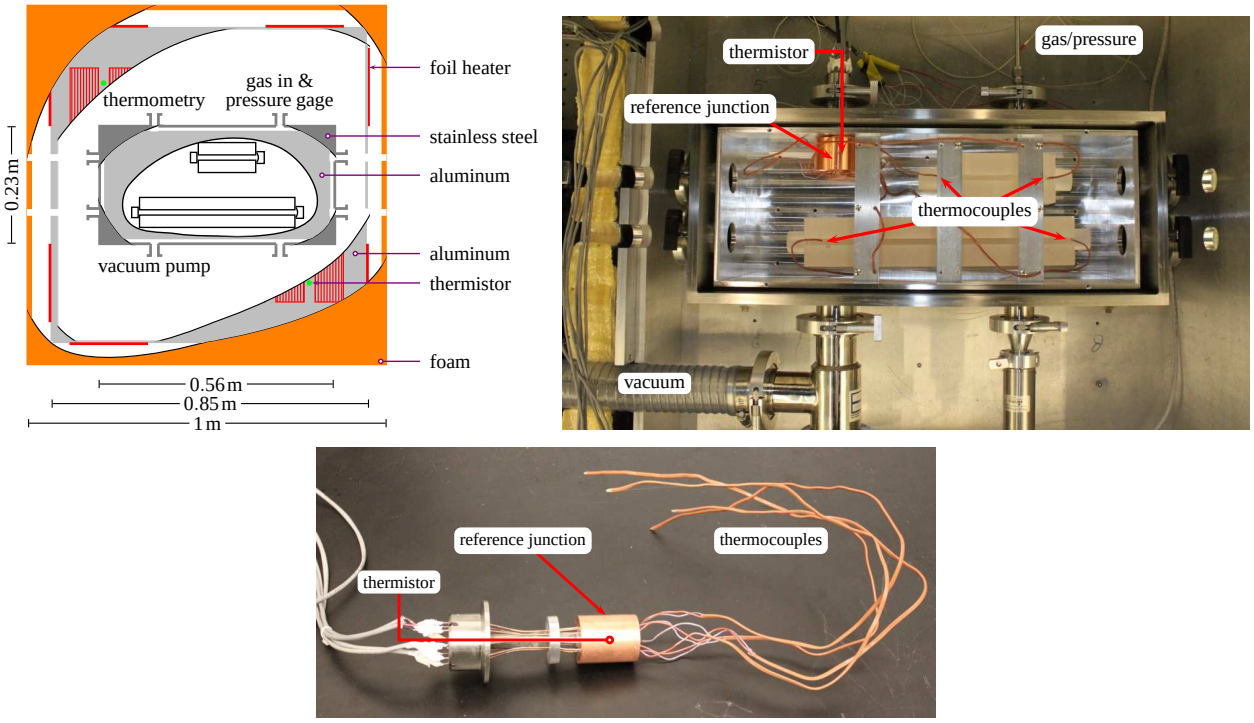

Figure 2: Temperature stabilization and thermometry configuration. Thermistors and foil heaters on the outer aluminum envelope were used to stabilize temperature. Thermocouples inside the vacuum chamber sensed temperature gradients around the optical cavities relative to a thermistor temperature standard.

The surface thermistors were roughly calibrated by adhering them to an isothermal block at 20 °C and reading the resistance of each; in this way the absolute temperature at any two points on the outer aluminum envelope did not differ by more than a few millikelvin.

Inside the temperature controlled aluminum envelope was the 12.7 mm thick stainless steel vacuum chamber. For the purpose of reducing temperature gradients, an aluminum (15 times higher thermal conductivity than stainless steel) vacuum chamber would have been preferable, but a stainless steel chamber is what we currently have. Inside the vacuum chamber was another 6.3 mm thick aluminum envelope, though this envelope had holes drilled in it to feedthrough the thermocouple reference junction and to transmit two laser beams. Each metal box was thermally isolated from the other by nylon and (in vacuum) polyether ether ketone (PEEK) posts. The cavities were suspended at their Airy points by 0.3 mm diameter stainless steel cables hung from aluminum bars running across the inner aluminum box.

We used one thermistor temperature standard to measure absolute temperature and four type-T thermocouples to sense temperature gradients relative to the temperature standard. A cylindrical copper block housed the reference junctions for all thermocouples, and the reference junctions surrounded the temperature standard in the block, as shown in Figure 2. The thermocouples are sheathed in perfluoroalkoxy (PFA) which, like PEEK, is known to have low outgassing effects

[outgassing.nasa.gov]. To ensure temperature uniformity between all reference junctions, a larger (60 mm diameter) cylindrical copper block was slipped over the reference junction block (limited to a 30 mm diameter so as to slide in and out through a KF40 flange) when the thermometry assembly was placed in the vacuum chamber. The thermistor probe slid from outside the chamber into the center of the copper block and, as such, can be easily removed for calibration. A superthermometer read the thermistor and a nanovoltmeter with a solid-state switch system read the thermocouples.

We chose thermocouples instead of thermistors to sense temperature gradients because: (a) a configuration having thermocouples with their reference junctions tied to a temperature standard is particularly well-suited to our application, and (b) we were concerned with resistive heating from thermistors—a thermistor can generate up to 100 μW and, in addition to the likelihood of resistive heating distorting temperature readings at vacuum, having this much heat in proximity to the cavities is undesirable. Thermocouples dissipate negligible power but are much less sensitive to temperature than thermistors, and offsets arising from thermal electromotive forces are a concern. In order to avoid these complications we used copper–copper weld joints where possible when wiring the thermocouples and took care to clean and deoxidize the contacts when crimping to the nanovoltmeter and switch. We note that the thermocouples are inserted into the slot of the cavity spacer; the slot in the cavity can be seen in the end-on photograph of Figure 1. If the walls of the inner aluminum envelope were a couple of millikelvin different than the cavity, the readout temperature of a thermocouple in free space would be influenced by radiation from the aluminum; to have the thermocouple tightly surrounded by the glass cavity is to sense a temperature more representative of the actual conditions in the cavity.

The nonideal features of our system, as referred to occasionally in this text, are the following:

- Thermal shorts: there is a bundle of instrumentation wires fedthrough three 30 mm diameter holes in the outer foam and aluminum envelopes. The bundle includes 16 pairs of metal braided 24-gage wire hooking up foil heaters and thermistors on the outer aluminum box, and 4 pairs of metal braided 24-gage wire reading the thermocouples from a feedthrough flange and a thermistor temperature standard. In addition to this bundle, there is a 9 mm stainless steel gas pipe, an hygrometer, and a 50 mm diameter PVC wire-coiled vacuum pipe. The only thing we might eventually do without is the hygrometer.
- Optical feedthroughs: there are four 30 mm diameter holes through all enclosures (one foam, three metal) for coupling laser to both cavities (four viewport flanges are used at the ends of the chamber).
- Enclosure joints: as a matter of practicality, the aluminum enclosures are made of six and seven pieces, while the foam enclosure is made of two. The chamber has two main pieces (a main body and an o-ring gasket lid) and 8 KF flanges.

3 Performance

3.1 Thermal time constants

The transient response of the system was investigated experimentally and with finite element analysis. For the finite element model shown in Figure 3, we assumed that the enclosure was symmetric and, accordingly, meshed an one-eighth segment employing well-known material properties

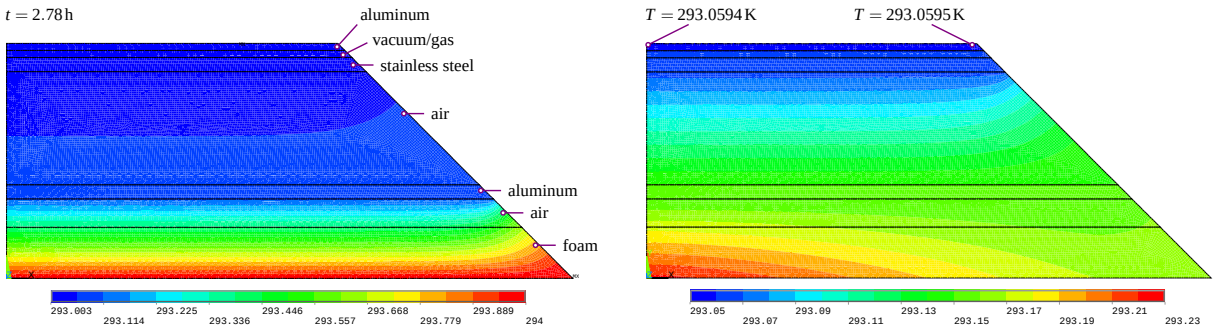

Figure 3: Finite element analysis was used to model (left) the time response of the passive thermal shielding and (right) steady state gradients arising from nonuniform room temperature outside the insulating foam.

for insulating foam, aluminum, stainless steel, and air. The model did not take into account air currents—that is, air is modeled as a solid—but air currents in a tightly enclosed space with small temperature gradients would be small. We calculated the Rayleigh number to be less than 10^5 and assumed that heat transfer was dominated by thermal conduction.

Since our laboratory is temperature controllable to 0.1 K (with a response time of less than 10 min), it was straightforward for us to have the system at a steady state and then increase room temperature by 1.0 K. In this way we observed the step response of the system by monitoring the thermistor temperature standard fedthrough into the vacuum chamber. As shown in Figure 4, we experimentally observed a thermal $1/e$ time constant of 92 h. This agrees quite well with our finite element model which gave a time constant of 116 h. The discrepancy between the experimental and finite element results was attributed to the nonideal features in our experimental system, as itemized in Section 2.

This four day time constant is, on the whole, desirable: it means that small fluctuations in room temperature on the order of several hours, such as those caused by a body working nearby, will have little effect on the temperature distribution around the cavities, and as noted previously, our lab temperature is stable to within 0.1 K. The downside of the long time constant is that temperature jumps on the tens of millikelvin level occurring inside the chamber (such as pumping down to vacuum or venting to air, as will be discussed presently) take a considerable amount of time to settle.

The other time constant we contend with in the case of absolute refractometry is charging or venting the evacuated chamber with a test gas such as air. When the chamber is filled with gas, the gas temperature rises as a consequence of work done by the external pressure resevoir on the gas filling the chamber. The gas temperature increases to $T = \gamma T_0$, where T_0 is the temperature of the gas in its storage cylinder and $\gamma = c_p/c_v$ is the heat capacity ratio defined by the specific heat at constant pressure divided by the specific heat at constant volume for the gas. For air ($\gamma \approx 1.4$) at 293 K, when vented into the evacuated chamber the gas temperature will increase by almost 120 K—this obviously upsets millikelvin temperature stabilization. The small mass of gas $m_{\text{air}} < 10^{-2}$ kg in our chamber means this effect is small, but it is nevertheless a problem at the millikelvin level. For example, the thermal time constant of cooling $\tau = mc_v/hA$ for vented air over an area A with

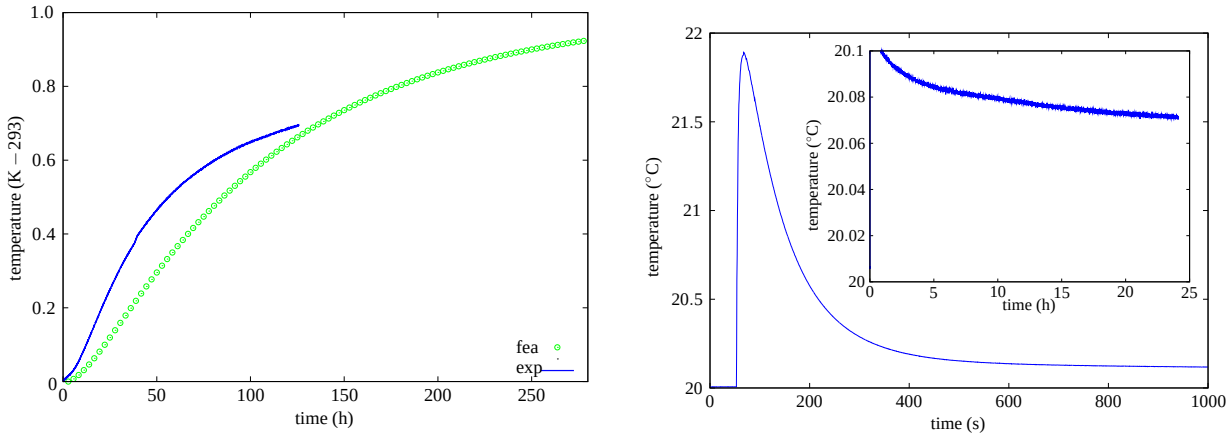

Figure 4: Left: The enclosure time constant was estimated as 92 h experimentally and as 116 h by finite element analysis. Right: Venting from high vacuum up to atmospheric air causes a large temperature jump with a short cooling time constant, but the small increase in equilibrium temperature takes much longer to settle.

heat transfer coefficient h is $\tau \approx 4\,\mathrm{s}$ for our system; so the gas cools rapidly. More importantly, however, since the total internal energy change of the system is zero, the equilibrium temperature can be estimated by solving $\Delta U_{\mathrm{sys}} = \Delta U_{\mathrm{air}} + \Delta U_{\mathrm{chamber}} = 0$ for T_{final}, with $\Delta U = mc(T_{\mathrm{final}} - T_{\mathrm{init}})$ and $T_{\mathrm{init}} = 410\,\mathrm{K}$ for air and $T_{\mathrm{init}} = 293\,\mathrm{K}$ for the chamber. In our system we calculated venting to cause a 58 mK increase in the equilibrium temperature.

The temperature characteristics of venting the chamber from high vacuum to atmospheric air are shown in Figure 4. Since the thermocouple has a response time of about 16 s while the gas cools with a time constant of 4 s, we would expect to only see a $e^{-4} \times 120 = 2.2\,\mathrm{K}$ temperature increase at the thermocouple, and not the full 120 K immediately after venting air. In any event, the cooling air rapidly ($\approx 1000\,\mathrm{s}$) settles to within 100 mK of the initial chamber temperature, but it takes a much longer time for the system to recover from the 80 mK jump in equilibrium temperature, which is consistent with what we calculated. As shown in the inset of Figure 4, one day after venting up to air the system is still more than 60 mK above the 20 °C starting point. Here we are up against the 92 h time constant of the system to pass heat from the warmed cavities to the cooler lab. There are at least three ways to shorten the waiting time for thermal stabilization:

- Increasing the setpoint of the PID temperature controllers by 80 mK about an hour before venting. In this way the hot gas will do the work of the foil heaters. However, for testing our refractometer performance it is expedient to work at the same temperature in vacuum and at atmosphere.
- The addition of more metal placed in the free spaces around the cavities would add thermal mass at our desired chamber temperature and reduce the volume (and hence mass) of gas venting into the cavity. We could implement this solution by adding 30 kg of stainless steel or copper and reduce the change in equilibrium temperature from 58 mK to 17 mK. This, however, is still relatively far from millikelvin stability and we do not have a great deal of

space in the chamber as it is.

- Temporarily remove the insulating box. The effect of removing the insulating box is to reduce the finite element system time constant from 116 h to 63 h, and we would expect a similar 45 % reduction in the 92 h experimental time constant.

In our case, the last approach appears the most convenient way to shorten the time taken for the vented chamber to return to its evacuated temperature.

3.2 Millikelvin gradients

We investigated temperature stability inside the vacuum chamber without any shielding or temperature control and found absolute temperature fluctuations of 60 mK over 20 h[1] 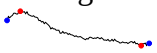 (max: 20.135 °C, min: 20.062 °C) and gradients of up to 200 mK between thermocouples. The fluctuation in absolute temperature is consistent with how well our lab temperature control is expected to perform and we attribute the rather large gradients in the thermocouples to the poor thermal conductivity of stainless steel. With passive stabilization only (that is, one foam and two aluminum enclosures) absolute temperature fluctuations were less than 20 mK over 100 h  (max: 19.813 °C, min: 19.794 °C) and the gradients between thermocouples were less than 2 mK. The interesting point here is that gradients inside the chamber can be kept to the millikelvin level by passive stabilization alone, but to keep absolute temperature stable to any better than the tens of millikelvin requires active temperature stabilization.

The finite element model was used to estimate what effect 0.2 K gradients (characteristic of our lab) on the outside of the insulating foam might have on the inner aluminum shielding. As in Figure 3, a temperature gradient of 0.2 K was applied across the ends of the enclosure segment and the model was solved for steady state. The model predicted gradients of less than 0.1 mK on the inner aluminum envelope. We do not, however, consider this a realistic estimate. For one, the model does not take into account the experimental shortcomings itemized in Section 2, and two, we have used probe thermistors to sense around the outside of the stainless steel chamber and observed gradients of up to 3 mK between various points, whereas the model predicts gradients of only 0.7 mK across the chamber. From this we would argue that wires, pipes, and holes through the enclosures give rise to localized cold or hot spots and that, as opposed to the 0.1 mK gradients around the cavities predicted by finite element, we might expect $0.1 \times 3/0.7 \approx 0.5$ mK gradients.

The PID controllers were capable of stabilizing temperature at the outer aluminum enclosure thermistors to within 1 mK over 100 h 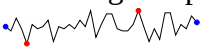 (max: 1.43 mK, min: −1.66 mK, std: 0.81 mK). We tested the ability of the controllers to respond to a 1 K decrease in lab temperature and each controller held its setpoint to 1 mK. Despite this active control, however, the temperature in the vicinity where the wires entered the outer aluminum envelope decreased by 87 mK while the value of the other probe thermistors decreased by about 15 mK of their initial value. Most importantly, the absolute temperature in the chamber decreased by 15 mK over 60 h  (max: 20.018 °C, min: 20.003 °C). In a similar example, we observed a 5 mK increase in absolute temperature during a lab visit when 6 bodies were present in the lab for 1 h. These are extreme cases

[1] The four small infographics that appear inline with the text of this subsection are called sparklines. The duration of the dataset is specified before each sparkline and from this the *x*-scale can be deduced. The maximum and minimum values of the dataset are specified after each sparkline and from this the *y*-scale can be deduced.

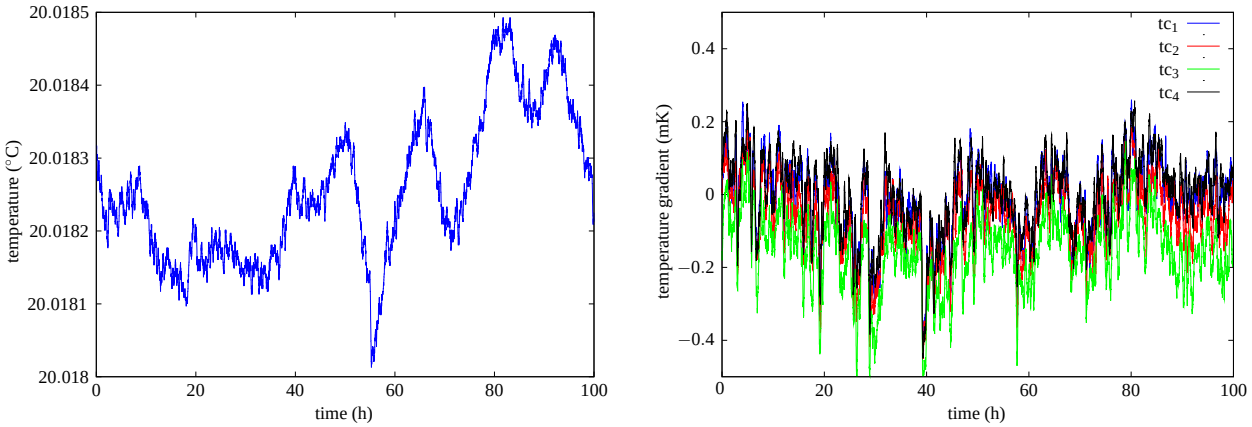

Figure 5: Left: Absolute temperature measured with a thermistor temperature standard remained stable to within 0.5 mK. Right: Gradients between the four thermocouples placed at the ends of the two cavities remained within 1 mK.

(lab temperature is stable to within ± 0.1 K of its setpoint and we almost never have more than two bodies in the lab), but they do suggest conditions when our stabilization system may not perform as desired.

Gradients around the extremeties of the outer aluminum envelope, such as those arising from cold spots, could be reduced using manual feedback from the probe thermistors and adjusting the setpoint of the controllers accordingly. Using this approach, however, we could not reduce gradients on the outer surface of the chamber to any better than 3 mK. Nevertheless, these gradients were stable in time and allowed us to stabilize absolute temperature and gradients between cavity ends to within 1 mK over 100 h as shown in Figure 5. As described in Section 2, the reference junction of each thermocouple surrounded the thermistor temperature standard and so each thermocouple voltage (when scaled by its Seebeck coefficient $\mathrm{d}V/\mathrm{d}T = 40.6\,\mathrm{mV/K}$) represents the difference in temperature between the temperature standard and location of the thermocouple measuring junction (each placed at a cavity end). For the temperature standard the averaging time for the superthermometer was 10 s and a datapoint was taken every 40 s. For the thermocouples the nanovoltmeter made an average of five samples in 3 s and this averaged datapoint was taken every 40 s. For all the data graphed in Figure 5 a 25-sample averaging filter has been applied in order to more clearly see temperature and gradient fluctuations. As well as reducing superthermometer/nanovoltmeter noise, this averaging filter reduces the magnitude of transient gradients occurring on the $25 \times 40 = 1000$ s timescale, but from our experience no significant transients occur that fast. The obvious correlations in the thermocouple data indicate that our gradient measurements are influenced by the nanovoltmeter and switch—it is unlikely that the temperature at each thermocouple fluctuated in the same way, nor do the reference junctions (thermistor data) explain the correlations. The nanovoltmeter and switch are specified to about 1 mK accuracy over 24 hours and this limits how well we can estimate gradients around the cavities. So we would claim that our system achieves < 1 mK gradients, but at the present time we can not reliably measure gradients less than 1 mK.

The sub-millikelvin stability of both the absolute temperature and the temperature gradients

satisfy our refractometer testing requirements, and ultimately our uncertainty in gas and cavity temperatures will be determined by the calibration uncertainty of the thermistor temperature standard (1 mK). We nevertheless note possible room for improvement. The presence of the 3 mK gradients across the outer surface of the chamber suggest that a further stage of active control at the chamber would improve performance. Unfortunately, stainless steel is not the best metal for this, but a sufficiently large foil heater at each corner of the chamber should reduce the 3 mK gradients we currently have, without creating hot or cold spots. We must add, however, that since the setpoint resolution of our controllers is 1 mK and digital, it is unrealistic with current equipment to expect a reduction in chamber gradients to less than 1 mK. Since the nanovoltmeter and switch appear to limit our ability to detect sub-millikelvin gradients, a thermopile may prove useful. We have had a custom thermopile made by wiring 10 type-T thermocouples in series and it awaits deployment. This should give a $\times 10$ boost in sensitivity to temperature gradients between each end of the thermopile. Unfortunately, possible gradients detected with the thermopile will not be referenced to the temperature standard, but it will nevertheless be interesting to investigate sub-millikelvin gradients along a cavity or between both cavity ends.

4 Conclusion

We have designed and tested a temperature stabilization system for refractometry. We have found that a combination of passive and active temperature stabilization can reduce gradients and absolute temperature fluctuations to the millikelvin level in a $0.5\,\mathrm{m} \times 0.15\,\mathrm{m} \times 0.15\,\mathrm{m}$ volume. This satisfies our temperature stability requirements for testing and characterizing our refractometers to the 10^{-9} level. Our performance is currently limited by the accuracy of the nanovoltmeter and switch but we suspect that shortcomings in our stabilization system (instrumentation cable, gas pipes, feedthrough holes, and enclosure joints) would prevent us reducing gradients much below 0.5 mK. Ultimately, however, our uncertainty in temperature measurement will come from the uncertainty of the thermistor temperature standard calibration. Some future ideas we would like to try include sensing temperature gradients with thermopiles and engaging a second stage of active temperature stabilization at the stainless steel vacuum chamber. At the present time, however, our most pressing objective is to measure the refractive index of nitrogen and argon gases at 632.8 nm to the 10^{-9} level and disseminate reference values.

References

[1] M. L. Eickhoff and J. L. Hall, "Real-time precision refractometry: new approaches," *Applied Optics*, vol. 36, no. 6, pp. 1223–1234, 1997. [Online]. Available: http://dx.doi.org/10.1364/AO.36.001223

[2] J. A. Stone and A. Stejskal, "Using helium as a standard of refractive index: correcting errors in a gas refractometer," *Metrologia*, vol. 41, no. 3, pp. 189–197, 2004. [Online]. Available: http://dx.doi.org/10.1088/0026-1394/41/3/012

[3] R. W. Fox, B. R. Washburn, N. R. Newbury, and L. Hollberg, “Wavelength references for interferometry in air,” *Applied Optics*, vol. 44, no. 36, pp. 7793–7801, 2005. [Online]. Available: http://dx.doi.org/10.1364/AO.44.007793

[4] R. W. Fox, “Temperature analysis of low-expansion Fabry-Perot cavities,” *Optics Express*, vol. 17, no. 17, pp. 15 023–15 031, 2009. [Online]. Available: http://dx.doi.org/10.1364/OE.17.015023

[5] R. F. Berg, G. A. Zimmerli, and M. R. Moldover, “Measurement of microkelvin temperature differences in a critical-point thermostat,” *International Journal of Thermophysics*, vol. 19, no. 2, pp. 481–490, mar 1998. [Online]. Available: http://dx.doi.org/10.1023/A:1022521712860