Research Aviation Facility

Algorithm Memorandum:

Subject: Temperature reprocessing

Al Cooper and Dick Friesen

19 May 2014

Background

On RAF aircraft, the primary temperature measurements are produced by platinum-resistance thermometers carried in housings just outside the skin of the aircraft. The housing protects the sensing elements from damage by hydrometeors and other objects while channeling outside air past the sensing wires. The measured variables have names like RTx (formerly TTx) where x describes the measurement with characters like HR1 for the heated sensor #1 mounted on the right side of the fuselage. The change from TTx to RTx recognizes that the actual measurement from these probes is best described as the recovery temperature and this differs from the total temperature if the recovery factor is not unity, which it never is for the probes we use. These measurements of recovery temperature undergo considerable processing in order to produce variables characterizing the actual air temperature, with names like ATx where x has a similar meaning and will be the same as used for RTx for a given sensor. That processing takes into account that the air temperature as sensed has been heated from the temperature in the ambient atmosphere and a correction to the measured temperature must be made to obtain a measurement of the ambient temperature.

The Present Processing Chain:

The processing chain from sensor to archived variable includes the following steps:

1. Determine the resistance-temperature relationship for the sensor. This has historically been done by immersing the sensor in a stirred bath along with a high-quality stem PRT, setting the bath at a series of different temperatures (as indicated by the PRT), and at each temperature measuring the resistance of the sensor. The result is a calibration that consists of a set of corresponding measurements of temperature and resistance, $\{T_i, R_i\}$. This is called the bath calibration.

2. Calibrate the on-board data-acquisition system. The on-board system consists of a special circuit to pass a known current through the sensor and to pre-amplify the resulting voltage in order to measure the resistance of the sensor. That voltage is then digitized by an A-D board and recorded by the data system. For this calibration, to substitute for the need to subject the sensors to various temperatures, the sensor is removed and a resistance box is connected to its socket in the circuit. Calibration consists of setting the resistances from step 1 $\{R_i\}$ into the resistance box and recording the corresponding voltages that are provided by the data system.[1] This calibrates the entire data-acquisition chain including exciter-preamp and A-D conversion. The result is a set of corresponding measurements of resistance and voltage

[1] Because the leads used for this connection have resistance of 0.03Ω, the actual resistances set into the resistance box are 0.03Ω smaller than the resistances from step 1.

$\{R_i, V_i\}$ or, using the correspondence from step 1, a set of corresponding measurements of temperature and voltage $\{T_i, V_i\}$.

3. Fit a quadratic polynomial to these measurements, in the form $T = c_0 + c_1V + c_2V^2$. This then is the calibration used to deduce temperature from recorded voltage for measurements acquired in flight. These fits often have been of marginal quality, with standard errors of >0.1°C and clear evidence in the residuals of a need for higher-order terms in the polynomial.[2] However, this magnitude of error has been deemed acceptable because other sources of error make larger contributions to the net uncertainty. The temperature T determined as above is a measurement of RTx, the total temperature sensed by the sensor. The resulting temperatures are the values archived in final data sets and used for subsequent calculation of the ambient temperature.

4. Convert the recovery temperature to ambient temperature. This step is described in detail in the Technical Note on Processing Algorithms, section 4.4, and won't be repeated here except for some discussion of the "recovery factor" below.

Reasons For Proposing Reprocessing

Over the past year, we have identified some weaknesses in the historical approach, especially as applied to the GV:

1. Weaknesses in the bath calibrations:

The bath calibrations came into question as we compared results from the RAF calibrations to those obtained by ISF, NIST, and DLR. The HARCO and Goodrich (formerly Rosemount) heated sensors are high-quality platinum resistance thermometers (PRTs) that should seldom need recalibration and that should conform to known properties of PRTs even without calibration, unless the probe is damaged or stressed, in which case the element should be discarded rather than recalibrated.

HARCO and Goodrich state that their probe conforms to MIL-P-27723E. [3] The specification states that the temperature-resistance relationship shall be as given by the following equation, called the Callendar - Van Dusen equation,[4] *with the following coefficients:* $\alpha = 0.003925$, $\delta = 1.45$, *and* $\beta = 0.1$ *for T<0 but 0 otherwise:*

$$\frac{R_T}{R_0} = 1 + \alpha\left[T - \delta\left(\frac{T}{100} - 1\right)\left(\frac{T}{100}\right) - \beta\left(\frac{T}{100} - 1\right)\left(\frac{T}{100}\right)^3\right] \tag{1}$$

[2] Fits as obtained using the revised procedures in this memo are much better and do not require higher-order terms.
[3] see http://mil-spec.tpub.com/MIL-P/MIL-P-27723E/MIL-P-27723E00006.htm
[4] Van Dusen, M. S., 1925: *J. Am. Chem. Soc.,* **47,** 326-332

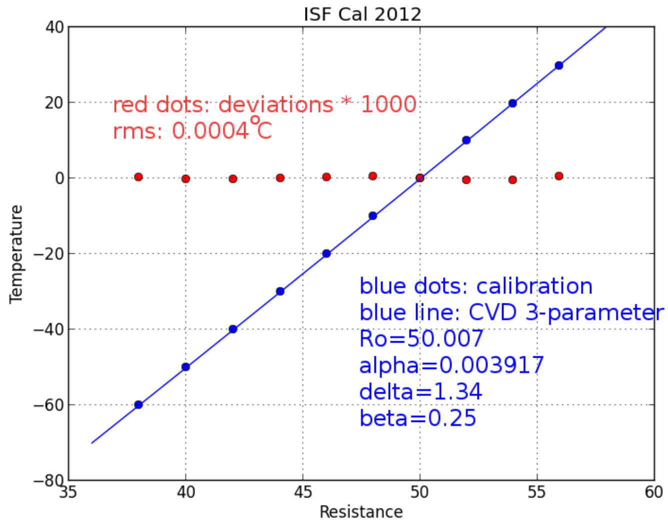

Figure 1: *ISF calibration of sensor 630393, 2012. The blue line is the fit for which all four parameters in the Callendar-Van Dusen equation were allowed to vary. The red dots show the deviations multiplied by 1000; even with this amplification factor, the residuals from the fit are negligible, as is consistent with the very small standard error for the fit (0.0004°C).*

The tolerance allowed by the MIL spec is $\Delta T = \pm(0.25 + 0.005\,|T|)$; i.e., 0.25°C at 0°C increasing to 0.5 at -50°C. This specification applies to the normal flight environment, not just to lab tests. Thus even without calibration the sensor should satisfy (1) to this level of accuracy. Once calibrated, the sensor should be capable of lower uncertainty than this, and unstressed platinum has very stable resistance, so one could argue that calibrations should only be used when it is expected that their uncertainty is less than the tolerance of the MIL SPEC.

As an example of a calibration that shows all indications of providing the highest quality, consider the calibration of sensor S/N 630393 by ISF in 2012 in the ISF stirred bath, with calibration points shown in Fig. 1. For the fit shown in this figure, all the parameters in (1) were allowed to vary (with the restriction that $\beta = 0$ for T>0). The fit is exceptionally good, matching the measured points to a degree that appears higher than expected accuracy of the temperature measurement. This is strong support both for the validity of the calibration and for the representativeness of the Callendar-Van Dusen equation. There are more accurate interpolation formulas in use for precise work, including 9th-order and 12-th order polynomials, but the high accuracy achieved with (1) argues that for our work the added complexity is unnecessary.[5]

The second figure shows the result for a fit in which the parameters δ and β in (1) were fixed at the nominal values and only R_0 and α were allowed to vary. The standard error remains very

[5] This conclusion may need further consideration because the comparison is between two PRTs that can have the same inaccuracies while agreeing with each other to high precision. It will be worth revisiting the ITS-90 temperature scale and associated interpolation formulas at some point, but all indications and some specific guidelines from NIST suggest that within 0.1C there won't be any important differences.

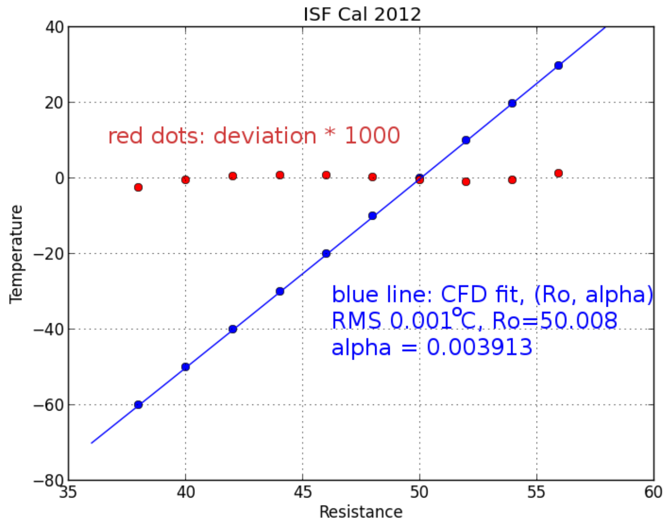

Figure 2: *ISF calibration as shown in Fig. 1 but showing a fit in which the parameters δ and β in (1) were restricted to the nominal values of $\delta = 1.45$ and $\beta = 0.1$ for T<0 and 0 otherwise.*

small, about 0.001°C, so the results are quite insensitive to these small changes in the values of the parameters δ and β. Even if the nominal value of α, 0.003925, is used, the resulting fit to R_0 alone gives a maximum error of only about 0.03°C. Thus this calibration is quite consistent with expected values for a high-quality PRT, and the results can be represented quite well either with the unadjusted Callendar-Van Dusen equation using nominal coefficients or with a fit that slightly improves the representation of the measurements by adjusting the coefficient of thermal resistivity.

This result suggests a way to characterize the various bath calibrations that have been performed in terms of just one parameter, the value of α obtained from a fit of (1) to the measurements with R_0 adjusted to match the particular sensor and δ and β held at their nominal values. Table 1 shows a summary of the coefficients of thermal resistivity obtained from such fits to many of the recent bath calibrations.

The consistency among the ISF 2012, ISF TORERO, and DLR 2011-6 calibrations, and also their consistency with the expected nominal values, strongly suggests that these are reliable calibrations. Conversely, the low values of the RAF calibrations suggests that these are not reliable. A likely explanation is that heat losses through the supporting structure during calibration and inadequate stirring reaching the sensing element is responsible, because the coefficient of thermal resistivity would then be too low as the sensor is actually exposed to a warmer temperature (for the low-T points) than measured by the calibrating PRT immersed alongside the sensor in the bath. It is not clear why the ISF-2011 or NIST-2011 calibrations are not consistent with the top group, but at least in the NIST case there is some suspicion that the immersion of the probes was not adequate. The test set-up they used is designed for stem PRTs, not our sensor configuration.

The unheated Rosemount probes (E102AL) calibrated at DLR appear to be an exception. These calibrations were of five refurbished probes, and all produced values of α near 0.00374. These appear to be valid calibrations of these probes. In contrast, the ISF calibration of S/N 3245 of the same probe gave a value of α very close to the nominal value of 0.003925.

Table 1: Coefficients of thermal resistivity (multiplied by 1000) for some of the bath calibrations. HARCO S/N 630393 unless otherwise specified. The colors separate calibrations meeting the suggested quality check of this section (blue) vs those that should be rejected (red). The black set were calibrations of refurbished probes not build to MIL specs, so there is no reason to reject these on the basis of the value of α.

Bath Calibration	$\alpha\times$**1000**
nominal	3.925
ISF 2012	3.914
DLR 2011-6 S/N 708904 #1	3.916
" " #2	3.916
ISF TORERO 708094 #1	3.917
" " #2	3.912
ISF TORERO 708094 #1 post-cal	3.918
" " #2 post-cal	3.913
ISF Rosemount 2884 TORERO pre-cal	3.919
" " post-cal	3.920
ISF Rosemount 3245 2013-2	3.929
NIST S/N 708904 2011-11	3.813
ISF 2011	3.754
ISF 2011-3 Rosemount Heated #1	3.615
" " #2	3.635
DLR Rosemount E102AL S/N 2603 (unheated)	3.744
" " S/N 2943	3.748
" " S/N 2980	3.741
" " S/N 3109	3.745
" " S/N 3241	3.774
RAF Low-T bath 2011-3	3.665
RAF Low-T bath 2010-6[a]	3.683
RAF Low-T bath 2010-6 RSMT heated #1	3.615
" " #2	3.635
RAF old bath 2010-6	3.715
RAF old bath 2009-03	3.708

[a]used for PREDICT

These invalid calibrations affected all GV projects prior to the TORERO project in 2012 and are the main reason for the reprocessing recommended in this note. C-130 projects are not affected as seriously because the errors arose at low temperature not usually encountered by the C-130, although it may be worth revisiting application of these re-calibrations to the C-130 sometime in the future.

2. Problems with recording calibration data

The procedure described in the section on the processing chain was followed conscientiously and carefully, but because it was assumed that the procedure was complete and valid it is sometimes difficult to determine the values needed to apply corrections. The evidence of the last section suggests that bath calibrations used before 2012 should not be used, so it is worth reprocessing with correct bath calibrations for the temperature sensors. Unfortunately, in many of the calibration records, what was recorded was temperature corresponding to the resistance set in the calibrating resistance box rather than the actual resistance used. This was done because, for a valid bath calibration, this gives the association between temperature and voltage that is the end result of the calibration. However, if the resistances actually used cannot be recovered by linking to some bath calibration, it is not possible to correct for an invalid bath calibration. Both calibration steps are needed, the first to find the association between resistance and temperature (which varies with sensing element) and the second to find the association between resistance and voltage (which is different for each analog card and which sometimes changes with project or time, requiring re-calibration).

This weakness in the records can only be corrected by laborious search through the paper records kept by the technicians, and in some cases it is still difficult to determine which calibration was used. The long-term solution now implemented is the “caledit” program that traces calibrations, links bath with on-board calibrations, and records the results of the calibrations directly in archived data files. This will solve this problem in the future.

3. Temperature-dependence performance of the Analog-to-Digital conversion boards

It was learned in 2010 that the A-D boards used for digitizing signals have a significant temperature dependence, and this is particularly important for boards like those sampling the temperature probes because they are located where they encounter significant temperature changes.[6] This is difficult to remove for early projects because each board has its own calibration and there is enough board-to-board variability to be significant. However, it appears that the variation with temperature is universal among boards, so if it can be determined which boards were used for sampling in each project these corrections can be made. A separate note documents how this correction was

[6] For more information, see this document on the RAF Science wiki.

determined and how it is applied in processing. It is important to note, though, that in old calibrations the temperature of the A-D board was not recorded and some might have been performed at temperatures differing from the desired calibration temperature (+40°C) by perhaps 10°C. This much change could introduce a bias error into the calibration of about 0.3°C, and there is no way to identify if that error is present.

4. Recovery-factor dependence on Mach number

To convert from the measured recovery temperature to the desired ambient temperature, one must use the "recovery factor" which describes by what fraction the air in thermal contact with the sensor is slowed. The recovery factor is defined from

$$\alpha_R = \frac{T_r - T_0}{T_t - T_0} \tag{2}$$

where T_r is the recovery temperature (i.e., the measurement, apart from errors that might arise from self-heating and similar causes of deviation from the accurate value), T_t is the total temperature (if the air were slowed to rest), and T_0 is the true ambient temperature. From this definition,

$$T_r = T_0(1 + \alpha_R \frac{(\gamma - 1)}{2} M^2) \tag{3}$$

where γ is the ratio of specific heats for air (1.4 for dry air, only slightly lower for moist air) and M is the Mach number.

Constant recovery factors were used for processing the measurements from the GV, sometimes varying with project. However, at Mach numbers typical of GV flight, studies of temperature probes documented in Goodrich Technical Report 5755 indicate that the recovery factor can be dependent on Mach number, esp. for heated probes. If α_R is constant, it is straightforward to fit to speed runs, but if it depends on Mach number then the approach is more complicated.

This correction is alternately handled through definition of a "recovery correction," defined as

$$\eta(M) = \frac{T_t - T_r}{T_t} \tag{4}$$

Then the recovery temperature is

$$T_r = T_0 \left[1 + \left(\frac{\gamma - 1}{2}\right) M^2 (1 - \eta) - \eta\right] \tag{5}$$

which also could be used to fit to speed runs to find η, although the fit is more complicated than the usual approach to determine the recovery factor. An advantage of this form is that the recovery

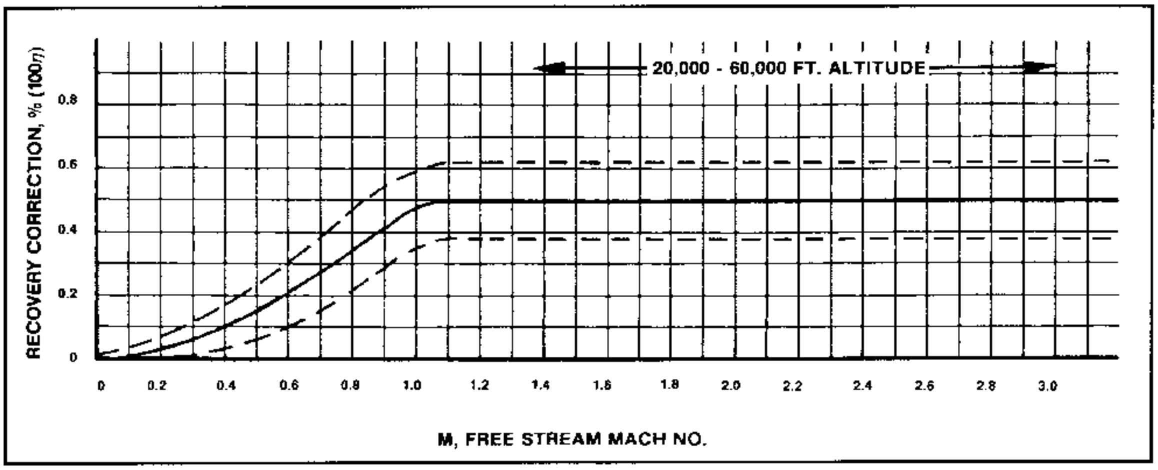

Figure 7: Wind Tunnel Data; Model 101 Recovery Corrections

9

Figure 3: *Recovery correction for the Rosemount 101 temperature sensor. Taken from Goodrich Technical Report 5755, 2003.*

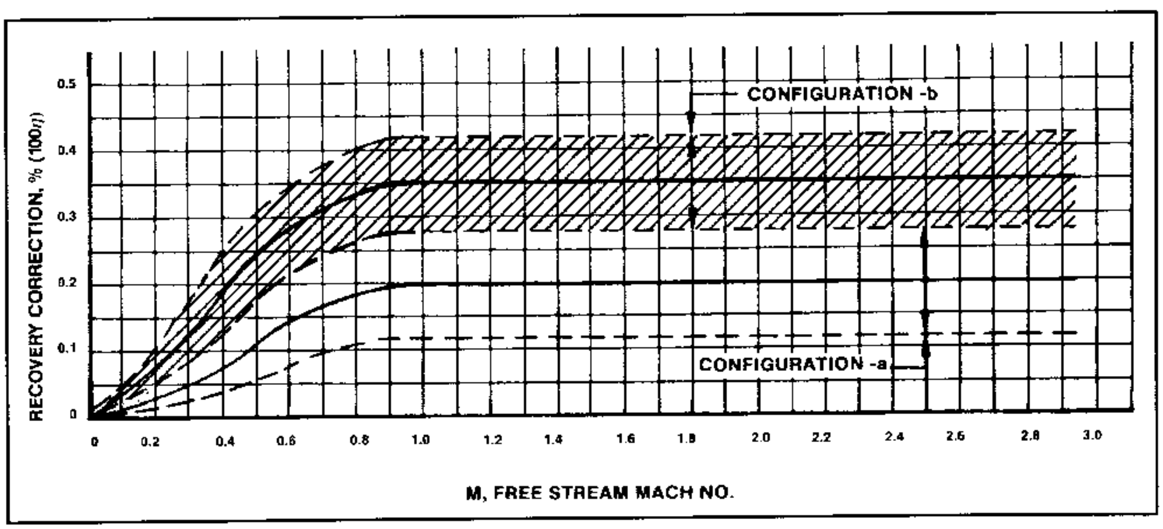

Figure 14: Wind Tunnel Data; Model 102 Recovery Corrections

Figure 4: *Recovery correction for the Rosemount 102 temperature sensor. Taken from Goodrich Technical Report 5755, 2003.*

correction is available, for some sensors, from wind-tunnel tests, e.g., for the Rosemount 101 non-deiced and the Rosemount 102 deiced sensor. Two plots that show such calibration, taken from the Goodrich descriptions (Technical Report 5755, 2003), are shown in Fig. 3 and 4.

These generally show quadratic behavior at low Mach number but level out to a constant correction at high Mach number. The HARCO probe differs slightly in geometry from the Rosemount probe, so use of this information for the HARCO probes may introduce an error, but similar information is not available for the HARCO geometry. It may be that the best approach is to use the Goodrich/Rosemount information as guidance to the functional form expected but fit to flight data to determine or check the details.

The following equations provide conversion between the recovery factor α and the recovery correction η:

$$\alpha_R = 1 - \eta \left(1 + \frac{2}{(\gamma-1)M^2}\right) \tag{6}$$

$$\eta = \frac{(1-\alpha_R)\frac{\gamma-1}{2}M^2}{1+\frac{\gamma-1}{2}M^2} \ . \tag{7}$$

Normally α_R is near unity and η is small (<1%).

The wind-tunnel data shown in Figs. 3 and 4 were converted from $\eta(M)$ to $\alpha_R(M)$ as shown in Figs. 5 and 6. For this purpose, the data plotted in Figs. 3 and 4 were digitized to determine tabular relationships, and from those the recovery factor was determined as a function of Mach number, as shown in figures 5 and 6.

These figures suggest that the unheated probe has, within the accuracy of this determination, a recovery factor that can be taken as constant with Mach number, with the approximate value of 0.97 throughout the normal GV flight range. However, the heated probe (102, configuration) shows a substantial dependence of recovery factor on Mach number, such that neglect of this dependence would introduce an error. As shown in Fig. 7, the dependence is represented well by the following equation, where M is the Mach number:

$$\alpha_R = 0.988 + 0.053\log_{10}M + 0.090(\log_{10}M)^2 + 0.091(\log_{10}M)^3 \tag{8}$$

This is then the formula recommended for the recovery factor for Rosemount temperature sensors. For the HARCO sensors, this could be used tentatively pending further study once reprocessed data are available.

5. Some other minor or limited-time problems, not improved by reprocessing

Limited range of calibrations used in on-board calibrations

Because the bath historically used for calibrations at the RAF covered only a limited temperature range, the on-board calibrations of the A-D cards often spanned only a range corresponding to

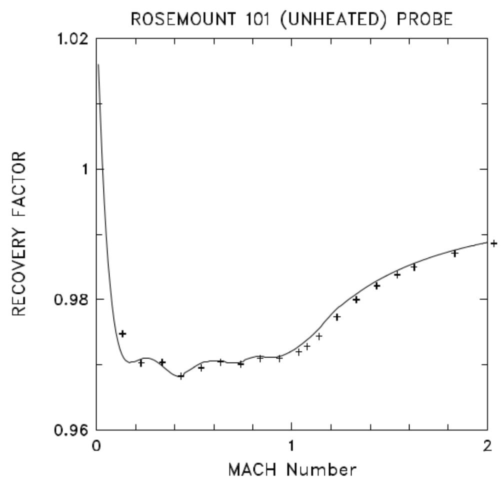

Figure 5: *Recovery factor for the Rosemount 101 (unheated) sensor. The points denoted by the symbol '+' were digitized from Fig. 3, and the curve was obtained by five-point Lagrange interpolation.*

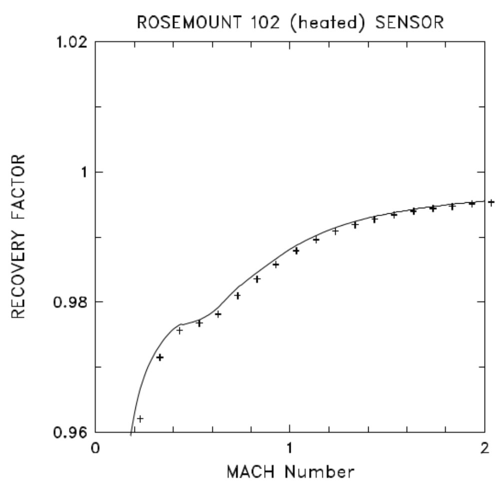

Figure 6: *Recovery factor for the Rosemount 102a (heated) sensor. The points denoted by the symbol '+' were digitized from Fig. 4 and the smooth curve was obtained by five-point Lagrange interpolation.*

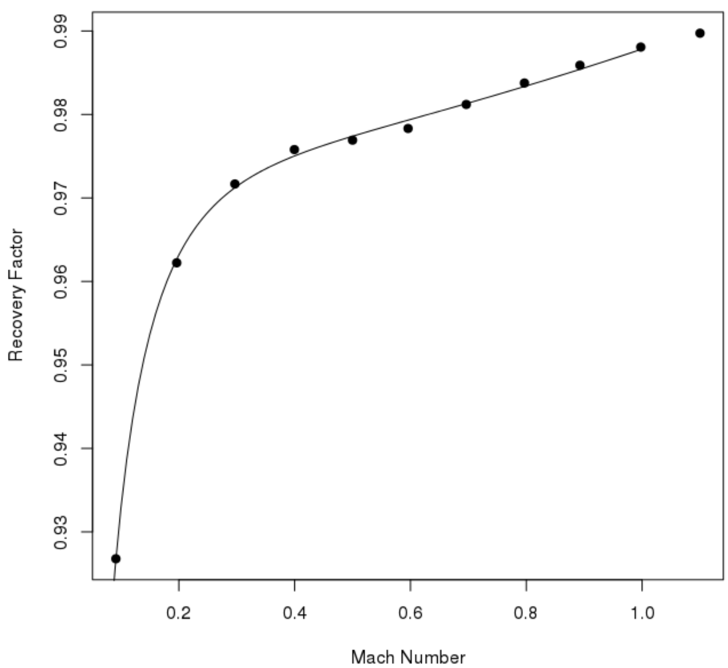

Figure 7: *Fit to recovery factor for the Rosemount 102a (heated) sensor. The dots are selected points read from Fig. 4, and the curve is Eq. (8).*

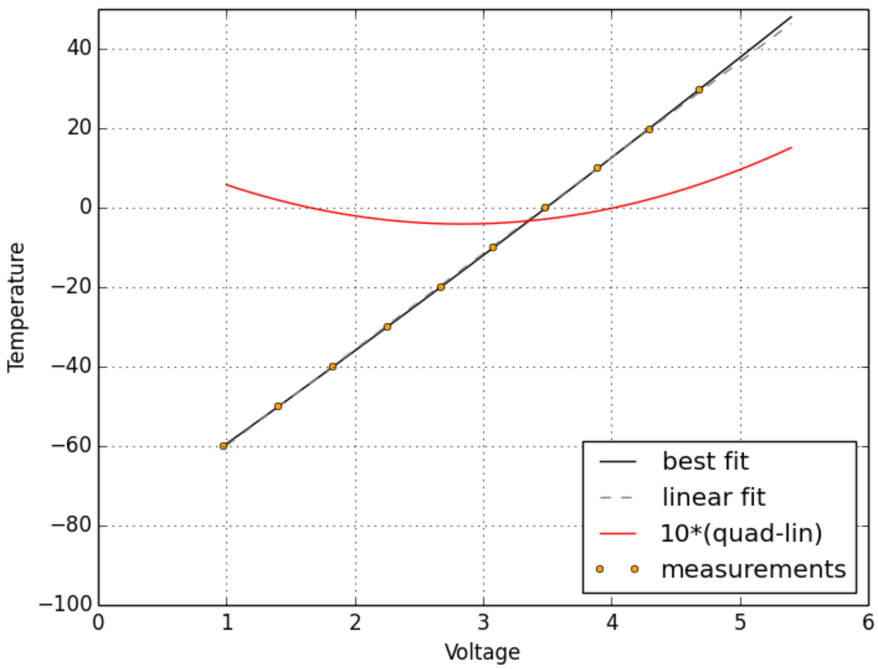

Figure 8: *Linear and quadratic polynomial fits to the calibration measurements for the HARCO S/N 812452A sensor as calibrated after the CONTRAST field experiment. The red line shows the difference, multiplied by 10 to make it easier to read, between values obtained by the quadratic and linear fits.*

temperatures from -40° to $+40^\circ$C. However, the GV often flies where the recovery temperature is -50°C or lower, so the result is that the on-board calibration did not cover the full range of values measured by the GV. The resulting quadratic fits then had to be extrapolated beyond the range of measurement, which is always a bad practice for higher-order polynomial fits.

An example of a good calibration where the calibration covers the normal range of measurements is shown in Fig. 8. The RMS error in the fit is 0.01°C for the quadratic fit and 0.36°C for the linear fit, and the maximum error in the normal range of measurements (-60 to $+20^\circ$C) exceeds 0.5°C, so the quadratic term is needed for adequate representation of the calibration.

Figure 9 shows an example where the calibration range for the A-D card covered only part of the range. The range covered by extrapolation to temperatures below -40°C looks reasonably similar to the case where this temperature range was covered by calibration, although the difference between the quadratic and linear fits is slightly larger in this case. Because different A-D cards can show different quadratic dependence in even valid calibrations, it is difficult to estimate the error that might arise from interpolation, but this example suggests that the error introduced by extrapolation may be smaller than 0.5°C for the low-temperature limit.

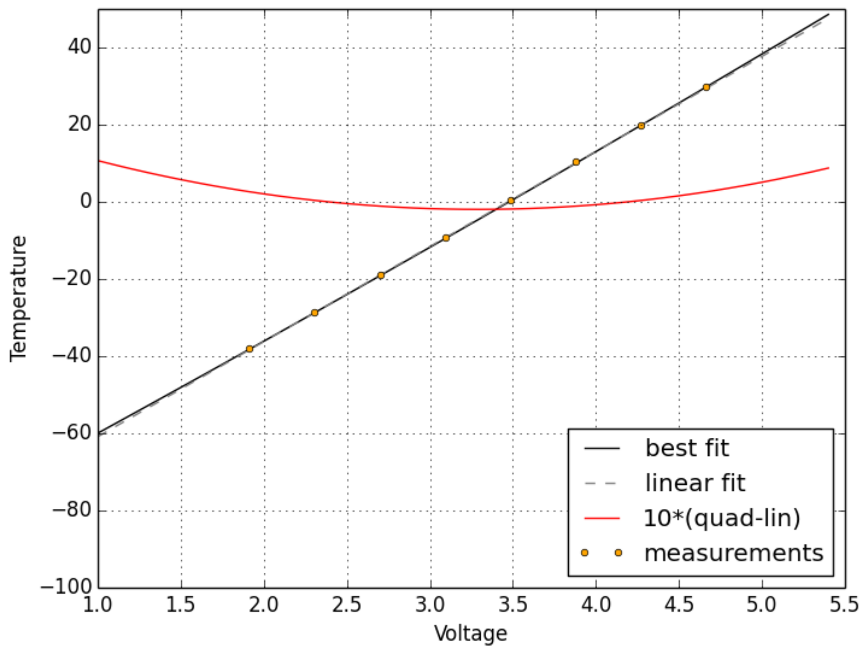

Figure 9: *As in Fig 8 but for a sensor calibrated only over the range of about* -40 *to* $+40^{\circ}C$*.*

Effect of a potentiometer installed on the A-D boards

Prior to the CONTRAST experiment in 2013, a potentiometer was installed on some of the A-D boards including the one used for temperature measurement in order to increase the voltage range. The reason was to attempt to avoid the extremes of the range in normal calibration, but the effect was to include a low-quality component with high potential to drift in the otherwise top-quality A-D cards. Some drift was observed in the CONTRAST project that could have resulted from this component, so the component was removed.

Extreme GV flight conditions

Procedures appropriate for the C-130 and earlier aircraft may not be appropriate for the GV because of its high flight speed (creating 20-25°C of dynamic heating) and the low temperature in the upper parts of the troposphere or the lower stratosphere where the GV can fly. Calibration procedures appropriate for $\pm 30^{\circ}$C can fail when the total temperature is $< -50^{\circ}$C and the ambient temperature is $< -75^{\circ}$C, as is common during GV flight. More attention to calibration and recovery factors is needed for the GV, and there is more potential for airspeed effects to introduce errors.

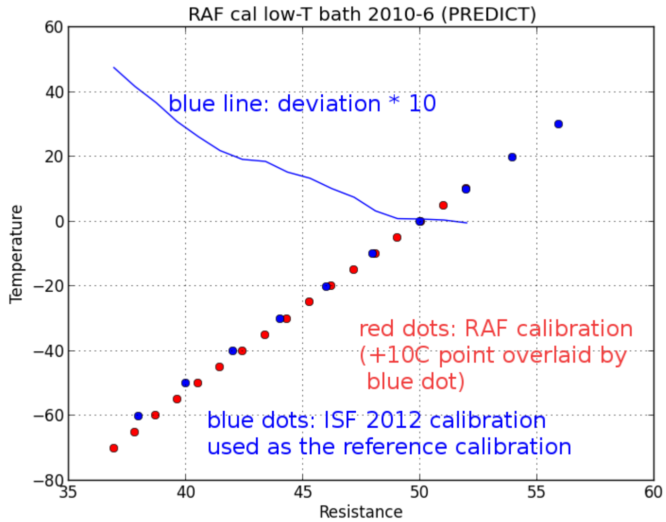

Figure 10: *RAF calibration of 2010-6 in the low-T RAF bath, as used for PREDICT (red dots), in comparison to the ISF-2012 calibration (blue dots). The blue line represents 10 times the error that results if the ISF-2012 calibration is correct and the RAF calibration is used instead; e.g., a typical error at -50°C is about 3°C. The sign is such that the temperature produced by the RAF calibration would be too low by the indicated amount.*

Magnitude of Errors: PREDICT as an Example

In Table_1, the bath calibration that was the basis for the PREDICT on-board calibrations is marked by a footnote. The error introduced by this calibration, if it is assumed that the 2012 ISF calibration is correct, is shown in Fig. 10. The errors are substantial at low T, amounting to about 3°C at typical low GV recovery temperature around -50°C. The errors at higher temperature are dependent on how the fit is propagated through the system via the on-board calibration because there are no measurements above 10°C in the RAF calibration so the on-board fit will have to be extrapolated to higher temperature, with associated extrapolation errors likely because the assumed coefficient of thermal resistivity is too low.

Recommendations: Reprocessing and Future Calibration

Reprocessing of Measurements Made Prior to 2012

Because the errors shown in the preceding section are significant, reprocessing of the data files would be appropriate to provide corrected versions of the data files already released. For this reprocessing, two calibration steps are needed:

Sensor	Calibration	R_0	α
	nominal values	50.0	0.003925
Rosemount 50738A	ISF Mar 2014	49.9960	0.003935
Rosemount 50738B	"	50.0146	0.003936
Rosemount 2984	ISF pre-TORERO	50.5493	0.003919
Rosemount E102AL #3245	ISF Feb 2013	50.0477	0.003929
HARCO 812452A[a]	"	50.0333	0.003913
HARCO 812452B	"	50.0227	0.003915
HARCO 630393A	ISF Mar 2012	50.0081	0.003914
HARCO 630393B	"	50.0053	0.003911
HARCO 708094A[b]	DLR June 2011	49.9873	0.003916
HARCO 708094B	"	49.9871	0.003916

Table 3: *Recommended calibration coefficients for use with the Callendar-Van Dusen equation (1) that relates temperature to resistance of a platinum resistance sensor.*

[a]Almost identical results were obtained Jan 2013
[b]almost identical results from ISF cal pre-TORERO

- Determine an appropriate relationship between resistance and temperature to use for each probe. When possible, this should be based on a recent ISF bath calibration that is checked for consistency with expectations for the Callendar-Van Dusen equation. If this is not available, then nominal values for the coefficients in that equation could be used.
- Determine the relationship between resistance and voltage as obtained in an "on-board" calibration of the appropriate analog-to-digital board and the temperature sensing circuit on the aircraft. This calibration will vary for different projects, and one of a set of available calibrations (usually including pre-project and post-project calibrations) needs to be selected. It is essential that the resistance values used for this calibration be determined; some records only included the temperatures that were thought to correspond to those resistances, and those temperatures are often in error, so retrieval of the right information from the technicians' records was necessary.

Table 3 lists available calibrations that can be used for the first step.

For each project, the procedure for determining the calibration coefficients for use in reprocessing is then as follows:

1. Find the actual resistances R_i used for the on-board calibration (corrected by $+0.03\Omega$ if taken from the screenshots) and the corresponding voltages V_i. Some old records do not include the resistances so it is difficult to recover for those, but recent calibration records include them and they are available for the calibrations relevant to CONTRAST. Dick Friesen

Measurement	**Sensor SN**	c_0	c_1	c_2
TTRL	unheated Rosemount 3245	-83.304	23.222	0.19054
TTHR1	HARCO 630393A	-83.326	23.120	0.24174
TTHR2?[a]	HARCO 630393B	-82.265	22.636	0.32405

Table 4: TREX

[a]labeled in notes as TTHR3 but there is no such probe; check results

compiled a set of calibrations applicable to each GV project, and those can be used to select the calibration for use in reprocessing. Table XX lists the calibrations selected jointly by Dick Friesen, Cory Wolff, and Al Cooper in May 2014.

2. Find the appropriate CVD coefficients (R_0, α) for each probe used in the project. If no valid calibration is available, use the nominal values of 50.0 and 0.003925. These values are also listed in Table XX.

3. Using inversion of the CVD equation, find the set of temperatures $\{T_i\}$ that correspond to $\{R_i\}$ for each calibration. For cases where the correct bath calibration is used, these will be very close to those recorded in the screenshots, but they may differ slightly because they come from the CVD fit instead of from the individual measurements. For cases where the wrong bath calibration was used, it is still possible to determine temperatures corresponding to the resistances actually used. A program that reads sets of calibration coefficients for each project and produces these temperatures is ~cooperw/TCalFit.py and calibration data must be in the subdirectory Calibrations and must follow the model of the project files in that directory.

4. Fit a quadratic polynomial to obtain the calibration in the form $T = f(V)$. This is the corrected on-board calibration. In cases where the original on-board calibration only covered the temperature range from $-40°$–$+40°$C, a linear fit should be obtained also and compared to the quadratic fit to check for possible large errors arising from extrapolation.

A remaining uncertainty is that many of these on-board calibrations were performed before the temperature dependence of the A-D cards was understood. Recent calibrations (since before TORERO) have been corrected for this and the technicians have tried to conduct the calibrations with board temperature near the standard (40°C), but earlier calibrations may have variations resulting from varying temperature at the time of calibration. A 10°C departure from the standard temperature, a reasonable estimate of what might have occurred, could cause a change in the calibration of typically 0.01V, which corresponds to a change in measured temperature of about 0.3°C. The old calibration records did not include the board temperature so this uncertainty remains in the calibrations recommended here.

Tables 4–13 list calibrations that result from this procedure. These are the calibrations recommended for reprocessing. (HIPPO4 and HIPPO5 are still needed here.) The results should be

Measurement	Sensor SN	c_0	c_1	c_2
TTFR	unheated Rosemount 3245	-82.130	22.422	0.30632
TTHR1	HARCO 630393A	-82.959	22.892	0.26392
TTHR2	HARCO 630393B	-82.328	22.675	0.29493

Table 5: PACDEX

Measurement	Sensor SN	c_0	c_1	c_2
TTFR?[a]	unheated Rosemount 3245	-82.198	22.316	0.31273
TTHR1	HARCO 630393A	-82.314	22.430	0.30844
TTHR2	HARCO 630393B	-82.207	22.494	0.30833

Table 6: HEFT-08 / GISMOS

[a] Notes say SN 3245 and TTHL1 but there is no such variable in the netCDF files and SN 3245 is not heated. There is a variable TTHL2 in the netCDF files, called a de-iced measurement, but that variable doesn't occur in the calibration records; the only heated probe is 630393, listed with TTHR1,2 and called "de-iced right". There are variables TTFR and TTHL2, both called de-iced left, There are also variables TTFH1,2, labeled as "de-iced left Rosemount" but the SN-3245 Rosemount is not heated and there is no other Rosemount in the calibration records. It is therefore unclear where this calibration was used, and it should be ignored in reprocessing unless this can be untangled from the project records.

Measurement	Sensor SN	c_0	c_1	c_2
TTFR	unheated Rosemount 3245	-82.108	22.410	0.30155
TTHR1	HARCO 630393A	-83.704	22.428	0.35397
TTHR2	HARCO 630393B	-82.644	21.961	0.43736

Table 7: START08

Measurement	Sensor SN	c_0	c_1	c_2
TTFR[a]	unheated Rosemount 3245	-82.415	22.482	0.29216
TTHR1	HARCO 630393A	-82.441	22.550	0.29465
TTHR2	HARCO 630393B	-82.066	22.467	0.31373

Table 8: HIPPO1

[a] One cal calls this TTHL1 but SN 3245 is an unheated probe; why an H in the name, usually used for heated?

Measurement	Sensor SN	c_0	c_1	c_2
TTHL1	HARCO 708094A	-82.378	22.760	0.30230
TTHL2	HARCO 708094B	-82.245	22.689	0.30490
TTHR1	HARCO 630393A	-82.441	22.734	0.31780
TTHR2	HARCO 630393B	-82.826	22.999	0.28923

Table 9: HIPPO2

Measurement	**Sensor SN**	c_0	c_1	c_2
TTHL1	HARCO 708094A	-82.301	22.731	0.30567
TTHL2	HARCO 708094B	-81.619	22.321	0.36363
TTHR1	HARCO 630393A	-82.068	22.389	0.37123
TTHR2	HARCO 630393B	-82.282	22.761	0.31150

Table 10: HIPPO3

Measurement	**Sensor SN**	c_0	c_1	c_2
TTHL1[a]	HARCO A50738A	-83.803	25.523	-0.30633
TTHL2[b]	HARCO A50739B	-82.009	22.501	0.30480
TTHR1	HARCO 630393A	-82.154	22.501	0.33251
TTHR2	HARCO 630393B	-82.255	22.704	0.30903

Table 11: PREDICT and HEFT10

[a] This calibration has what appears to be an outlier point at the lowest calibration and perhaps should not be used for that reason.

[b] There is another calibration giving -83.974, 25.525, -0.30784. It appears to have an outlier point at the lowest temperature in the calibration.

Measurement	**Sensor SN**	c_0	c_1	c_2
TTHL1	Rosemount A50738A[a]	-83.734	22.175	0.28489
TTHL2	Rosemount A50738B[b]	–	–	–
TTHR1	HARCO 708094A	-75.948	20.355	0.33556
TTHR2	HARCO 708094B	-76.781	20.490	0.34571

Table 12: HIPPO4

[a] limit temperature range; check

[b] cals seem to have problems

Measurement	**Sensor SN**	c_0	c_1	c_2
TTFH1	Rosemount A50738A	-87.788	26.510	-0.31981
TTFH2	Rosemount A50738B	-84.308	23.102	0.30528
TTHR1	HARCO 630393A	-82.990	22.512	0.39721
TTHR2	HARCO 630393B	-83.566	22.704	0.39350

Table 13: HIPPO5

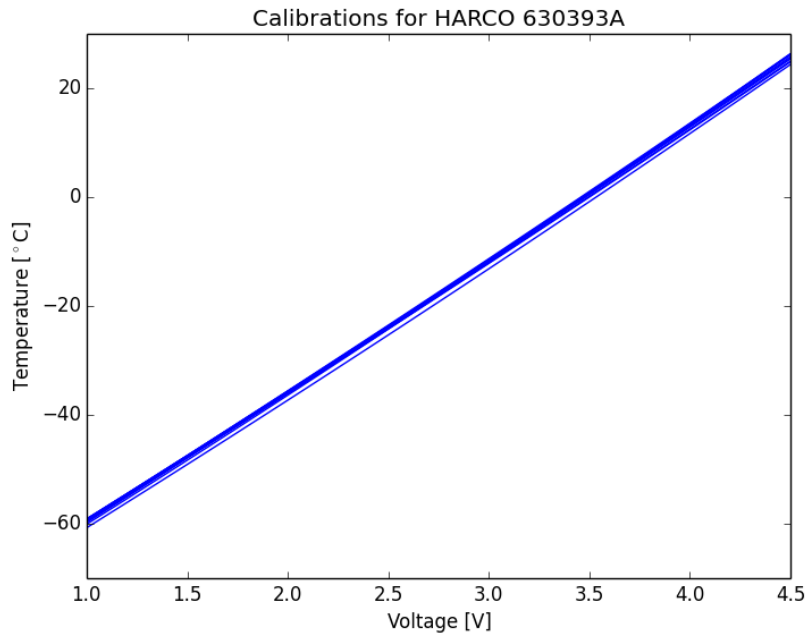

Figure 11: The ten calibrations of HARCO sensor 630393A listed in the tables above. The average difference between all pairs of calibrations over this voltage range was 0.2°C.

reviewed with particular attention to the calibrations that have footnotes in the tables. That review should be based on consistency among the sensors and also reasonable agreement, expected to within about 1°C, with AT_A.

It is difficult to tell from lists of calibration coefficients how different the calibrations are, because the coefficients can be correlated, esp. with changes in c_0 compensated by changes in c_1. Furthermore, the same sensor may be connected to a different A-D card, and the A-D cards vary in calibration, so a constant calibration across all projects would not be expected. Nevertheless, the A-D cards are designed to be the same, so variations may be minor. To check how much variation there is in the coefficients determined as in this note, the 10 calibrations listed for the HARCO 630393A sensor were used to construct a plot of temperature vs. voltage to check for consistency. Figure 11 shows that the calibrations are all quite consistent despite some variation in coefficients. To evaluate the differences, 100 points spanning the plotted voltage range were selected and the calibrations were evaluated in pairs, giving 45 estimates of the RMS difference between pairs of calibrations. The average value of that RMS difference was 0.2°C. The comparable value for the original calibrations was much larger and those calibrations systematically differed from those plotted in Fig. 11, especially for low voltage and low temperature where the differences were typically about 4°C lower in the original calibrations. Lower temperature is expected in the original calibrations if the explanation offered for the source of the problem is correct, because the temperature at calibration was actually higher than the temperature set for the bath. The consistency of the plotted results and the direction of the correction support the approach taken in this memo.

Recommended Procedures for Future Calibrations

1. Perform bath calibrations using ISF facilities only occasionally (perhaps annually or bi-annually) to check for consistency with the bath calibrations listed in Table 3. Change these only on clear evidence of a change in sensor characteristics.

2. For on-board calibrations, I suggest changing the calibration software as follows:

   (a) Incorporate a table like Table 3 into the calibration program.

   (b) Use a serial number, entered by the technicians, to select a set of corresponding temperatures and resistances to use. That set should be constructed from the CVD fit using the coefficients in the table, rather than from the direct measurements entering the calibration, because the fitted representation reduces errors in the individual points. The range used should span the range of temperatures of interest and so provide a calibration of the A-D card that does not require interpolation. Appropriate temperatures at which to calculate these values would be in 10°C increments from -60 to $+30$°C. These then need to be converted to resistance values and displayed for the technicians' use. These resistances can be calculated and stored so that they do not need recalculation each time. The set of temperatures would then be used with the measured voltages to determine the calibration.

3. The auto-cal feature of the A-D cards should be employed frequently to check for drift in the calibrations and to provide a history of A-D performance for each project. Until it is established that these cards are stable enough not to require this, this should be done for each flight.

4. When quality-control review indicates that the set of measured temperatures differ by more than 1°C without extenuating circumstances (like measurement in cloud or ice accumulation), this should trigger an investigation that seeks to learn the cause of the disagreement.

5. Because AT_A has undocumented characteristics and is clearly filtered and delayed, it should not be used as the primary temperature for projects except as a last resort when all other temperatures cannot be used.

—- End of This Memo —-