

Quantifying Cryocooler Contamination

P. S. Spoor

Chart Industries
Troy, NY 12180

ABSTRACT

Contamination of the cryocooler working gas (usually helium) with other species results in performance degradation if those species freeze at or below the cooler's intended process temperature. Even relatively small amounts of contaminant can result in harmful blockage of the fine fluid passages in the cold zones^{1,2}. Fortunately, avoiding contamination is fairly straightforward: use clean parts and low outgassing materials, and vacuum-bake (or hot dwell-and-purge¹) the assemblies to remove the remaining volatiles. Verifying cleanliness, however, may require many months of run-time to establish temperature stability. For low-cost mass production and rapid R&D cycles, it is very desirable to have a method for actually quantifying the level of contamination in a cryocooler against a pass/fail specification at the time of manufacture. This paper proposes a very simplified model of outgassing that leads to a quantitative estimate of available contaminants in milligrams, based on the accumulated vapor-pressure and the RGA measurements on evacuated units. This estimate can be compared to a figure-of-merit derived from contamination studies in the literature, hence yielding a pass-fail test for cleanliness.

INTRODUCTION

The issue of contamination of working gas in cryocoolers is widely recognized^{1,2}, as are many of the best practices for minimizing it. The impact of specific levels of contamination has been studied, as well¹. The measurement of contamination level in cryocoolers is not as well understood, however. In the cited studies, the contamination is deliberately introduced, so its level is at least approximately known before the start of the test. On the other hand, contamination coming from within the cooler itself is challenging to measure, because it is stored not in the working gas but on surfaces, in non-metallic or porous solids, and in any trapped volumes. Analyzing a sample of the helium from a cryocooler does not tell us the actual contaminant level of interest, because:

1. the cryocooler acts as its own cryopump, so as long as the cooler is running, it continually freezes out any contaminants in the working gas.
2. when the cryocooler is not running, the partial vapor pressure of any contaminant is limited by the balance of release vs. reabsorption.

What we really want to know is the quantity of contamination that could potentially outgas over a long period of time, e.g., how much would outgas if we could continually keep the vapor pressure low.

THEORY

To begin, we start by making some simplifying assumptions. By far the most common and stubborn contaminant is water; occasionally there are other volatiles such as methanol from cur-

ing epoxies, but we can take the analysis pretty far by assuming water as the only contaminant of interest (for now). There may be several simultaneous outgassing processes, each with its own characteristics; nonetheless, it is generally true that the release of contaminants slows down as the amount available is reduced. If we make the very simplifying assumption that the overall outgas rate is proportional to the amount of (unfrozen) contaminant in a cooler, we can write:

$$m'(t) = m_0(1 - e^{-t/\tau}) \tag{1}$$

where

$m(t)$ = the mass of contaminant frozen in the cold zones at time t ;

m_0 = the total mass of contaminant in the cryocooler available for outgassing at time $t = 0$; and

τ = a time constant, characteristic of a given size and type of cryocooler.

In other words, the outgassing process is assumed to obey a relaxing exponential, with an asymptote at m_0 , and where $m(t)$ reaches 63% of m_0 at time $t = \tau$. The buildup of pressure in an initially evacuated vessel due to outgassing is fundamentally similar, except that the pressure saturates not because the available contaminant has been exhausted, but because the release and reabsorption of contaminant are in equilibrium. Nonetheless, there is an important connection between them: if (vapor) pressure is uniform in a cryocooler during an outgas test (a big if, we should note) then the initial rate of pressure buildup roughly corresponds to the rate of frozen mass accumulation that the cryocooler would experience if running. In other words, if the cryocooler is evacuated, then isolated from vacuum, the pressure will rise according to

$$p(t) = p(\infty)(1 - e^{-bt}) \tag{2}$$

where the relationship between m_0 from Eq. (1) and $p(t)$ is

$$m_0 \simeq \tau \left. \frac{dp}{dt} \right|_{t=0} \cdot V \frac{M}{RT_H} \tag{3}$$

Here V is the cryocooler internal gas volume, R is the universal gas constant, T_H is the temperature of the outgas test (which is chosen to equal the expected, or, to be conservative, the hottest expected operating temperature of the warm parts of the cryocooler), and M is the molecular weight of the dominant vapor (presumed to be water in most of our analysis).

Finding m_0 and τ

Without knowing τ , we can't get m_0 directly from Eq. (3). However, τ can be deduced from a sequence of two outgas tests, where the cryocooler is vacuum baked at the same temperature T_H for a time Δt in-between the two tests. In this case, the two outgas tests are only run long enough to establish an initial slope, for a much shorter time than Δt . Using the familiar dot notation for time-derivative, and the subscript i for "initial", we can say that

$$\frac{\dot{p}_{i,2}}{\dot{p}_{i,1}} \simeq e^{-\Delta t/\tau} \tag{4}$$

In principle, the initial rate of increase of vapor pressure and the ultimate saturation pressure are both proportional to the source mass, so either one can be used to find the parameters of interest. This redundancy permits independent calculations for the same value, which can be compared for consistency. Disagreement between the calculations can lead to insights and ways to improve the method beyond the very simple version presented here (for instance, how to handle multiple outgassing processes with different time constants).

For example, if there exists a bakeout temperature $T_B > T_H$ that can accelerate the bakeout process (in other words, if the cryocooler can be baked out at a higher temperature than its expected operating temperature) then m_0 can be deduced from a sequence of two outgas tests at T_H bracketing a dwell (where the vapor pressure is allowed to build) at T_B . Knowing the volume, the temperature,

and the gas species reveals how much mass Δm was removed between outgas tests; the higher bakeout temperature ensures that Δm is much greater than what is removed during the outgas tests at T_H (recall that vapor pressure tends to rise exponentially with temperature, so even a modestly higher bakeout temperature will have a big effect). Supposing we use the saturation pressures in the outgas tests before and after the dwell; then

$$\frac{p(\infty)_2}{p(\infty)_1} \simeq \frac{m_0 - \Delta m}{m_0} \tag{5}$$

In a similar fashion, if a vacuum bake (or other process, such as a hot pressurized dwell-and-purge²) is inserted between two outgas tests, and run for a controlled time, the reduction in outgas saturation pressure or initial outgas rate can be used to predict the time required to reduce the contamination level below a given threshold.

Threshold Criteria for Contamination

Equations (1) through (5), combined with the proper series of outgas tests, can determine the amount of contaminant in a given cooler and how quickly it can be eliminated. It remains to establish how much contamination can be tolerated. For that we turn to the literature, and subsequently, to our own experiments running coolers with contamination that has been quantified. There are numerous studies published on the effects on contamination on cryocoolers; one of the most thorough is a 1998 study¹ by Hall, Ross *et al*, detailing the effects of known amounts of deliberately introduced contamination on a pulse-tube cryocooler’s ultimate (no-load) temperature. This study concluded that amounts of water vapor (the most prevalent and most stubborn contaminant) below about 50 ppm would keep the temperature from rising more than 1K. Our own experiments, on the other hand, seem to indicate that the threshold is closer to 150 ppm. One intuitively appealing way to reconcile these results is to state the contamination level a different way. Because the assumed mechanism by which contamination harms performance is by freezing and blocking flow passages, we can argue that the important statistic is not ppm level in the working fluid (or theoretical ppm level that might exist if all the available contaminant could be suspended as vapor in the working fluid at one time) but the actual mass per unit area of heat exchanger or regenerator area.

On that basis, our results and those of Hall & Ross are fairly consistent, and suggest that a contaminant load of ~0.01 mg per mm² of regenerator area will cause a ~1K rise in no-load temperature (H&R observe how astounding it is that such a small amount, covering the RGR cross-section with a layer a only few molecules deep, could have any measurable effect). Hence:

$$T_c(t) \sim T_{c,i} + \frac{m'(t)}{A_{RG} \cdot 0.01 \frac{\text{mg}}{\text{mm}^2}} \tag{6}$$

If Equation (1) is used for $m'(t)$, then the empirical knowledge from the literature is combined with the simple theory expressed by Equations (1) through (5) to yield a prediction that can be tested against actual cryocooler behavior.

RESULTS

The method outlined here was used to verify that the contamination load in a particular cryocooler design was more than sufficient to explain its warming trend, and to identify which components were the major sources of contaminants. Not surprisingly, the single biggest source turned out to be the stator VPI compound, which had been recommended to us by the manufacturer as “low outgassing” but was actually incurable by any bakeout process we could discover.

We decided to set our warming target to $\Delta T_c \leq 1\text{K}$ (i.e. the cryocooler should not warm up more than 1K due to contamination); in this particular cryocooler design, with a capacity of ~30W at 77K, the regenerator area is ~1460 mm², so Eq. (6) can be rewritten as

$$\Delta T_c(t) \sim 0.07 \frac{\text{K}}{\text{mg}} m'(t) \tag{7}$$

Eqn. (7) suggests that for this cryocooler, to keep $\Delta T_c \leq 1\text{K}$, m_0 must be ≤ 15 mg. We performed series of outgas tests on all the cryocooler components, to obtain a ‘contamination inventory’.

Figure 1 shows the results for a pair of motor stators; a sequence of two outgas tests were done at 60°C (the expected max temperature of operation) with a ‘dwell’ at 90°C in-between (the solid lines are exponential fits to the data; note here that our DAQ system was still getting debugged and there are frequent dropouts).

From the volume of the vacuum vessel used in these tests, the species of contaminant removed (identified by RGA) and the temperature, we can tell how much mass was released during the dwell (and subsequently pumped out before the second 60 °C outgas test). Because the vapor pressure was so much higher during the 90 °C dwell (~9000 mTorr. maximum) we can neglect the small amount evolved during the 60 °C tests. The modest change in slope between the two tests tells us that only a small fraction of the total contaminant load was removed. Overall, the tests suggest that the stators start out with over 90 mg of contaminant, and that it would take over 20 cycles of hot dwell and pump-out to reduce the level down to single digits (or equivalently, over a month of constant vacuum bakeout at 90 °C). The stators were found to be the biggest single source of contaminant, and the cleanliness of the units was vastly improved by changing VPI compounds.

This is not *prima facie* a surprising result, that the stator potting compound would be a significant source of contamination for a cryocooler with motors housed inside the pressure vessel; but quantifying the contamination and the minimum processing time to remove it helped to justify the cost of discarding existing inventory and buying new parts.

Contaminated vs. clean cryocooler tests

Our outgas tests measured the effects of various bakeouts and dwells, and hastened refinements to our manufacturing and processing, resulting in coolers with an estimated contamination load at or below 15 mg. Not only was the stator VPI compound changed, but we also began baking out coldheads separately from compressors (they can stand a higher temperature), among other process improvements. The original, heavily contaminated prototypes provided a good means of contrasting the ‘before’ and ‘after’ cryocooler temperature stability. In particular, we could use Eq. (7) to predict the temperature rise of each cryocooler, and compare it to actual measurements as a means of evaluating the validity of the overall approach to understanding contamination.

Figure 2 shows the comparison for the first week or so running a cryocooler with a total estimated contaminant load of 160 mg. Two things are apparent from this graph; one, the agreement appears excellent, and two, the warmup rate at this level of contamination is easily detectable within a few days, since the fastest warmup occurs in the beginning (if we believe the functional form of Eq (1) and its cousins).

However, tests on ‘clean’ cryocoolers provide a more nuanced view. One challenge that became apparent to us in the process of setting up a fleet of cryocoolers to run long-term is that the running conditions have to be tightly controlled in order to verify that a change of less than 1K in coldtip

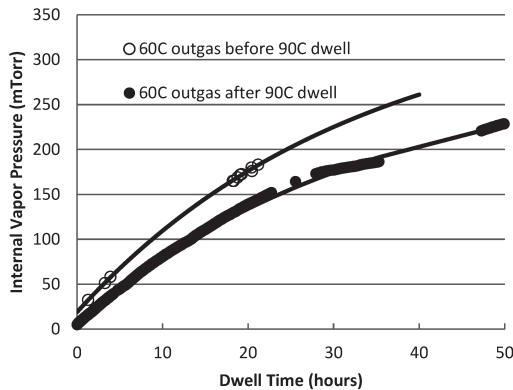


Figure 1. Vapor pressure versus time for two 60°C outgas tests, one before and one after a “dwell” at 90°C, of two linear motor stators in a 3.5-liter vacuum vessel.

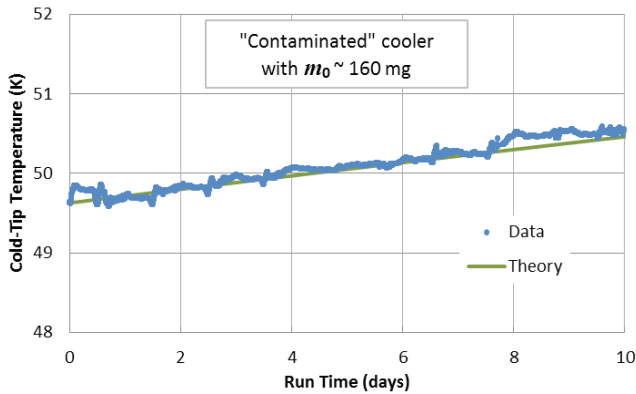


Figure 2. Contaminated cooler, for its first few days of running, compared to theory.

temperature is indeed due to contamination. The cryocooler of Figure 2 was water-cooled, with a dedicated controlled-temperature bath; but for a fleet of many coolers this would not be practical (and the fundamental design for these assumes air-cooling). Changes in room temperature cause changes in internal charge pressure, which result in power consumption fluctuations as well as variations in reject temperature. Mapping the performance of each cooler over various conditions enables some correction for these variations, but these corrections are imprecise. Furthermore, there are many variables that can affect the no-load temperature of a cryocooler, especially during its early burn-in period, which are hard to control or correct for (slight wear of clearance seals that have incidental contact, for instance).

Finally, as the contamination load is reduced to fairly low levels, we find that the estimates based on initial outgas rate and the estimates based on saturation vapor pressure don't agree particularly well. This is due to a variety of factors; for instance, the vapors outgassed in one of these coolers include quite a bit of hydrogen from the stainless-steel. We could try to de-gas the SS, but the hydrogen will not affect the performance of these coolers. It does, however, mean that the evolution of the 'light' components has to be somehow separated from the 'heavy' components, using RGA spectra, etc., and this makes resolving the initial slope of the heavy-component vapor-pressure curve particularly challenging.

Figure 3 shows a comparison of theory and experiment for a relatively 'clean' cooler, over its first few days of running. Here, the contamination load is given as a range of values, since the two methods of measuring contamination do not agree especially well. This cooler does not have a particularly steady temperature, due to fluctuating conditions; it is difficult to draw a conclusion about

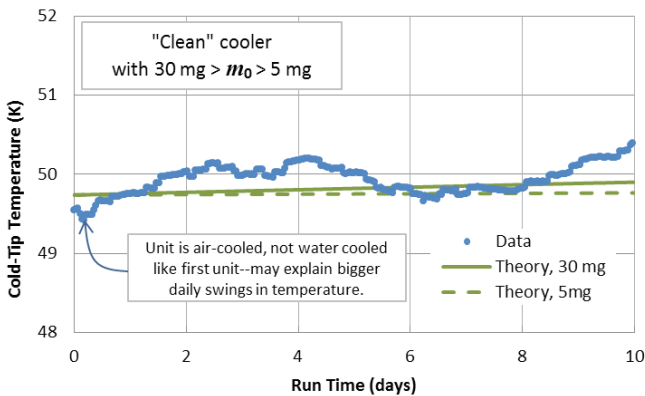


Figure 3. "Clean" cooler over 10 days.

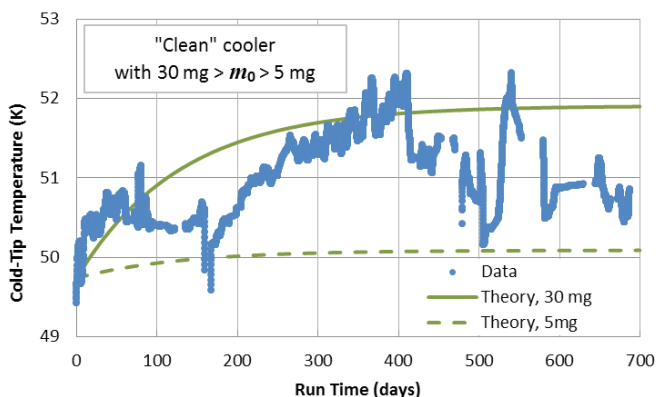


Figure 4. “Clean” cooler over two years.

the comparison of theory to experiment in this case (indeed, the amount of expected temperature rise could hardly be resolved in this length of time).

As it happens, this paper is written roughly two years after the bulk of the work described here was accomplished; and so we have a much longer record for this and other cryocoolers in this study. Figure 4 shows the cooler of Figure 3 but after almost 700 days of running. During that time there were observed some rather significant fluctuations in cold-tip temperature, not all of them understood. One likely cause of some (perhaps all) of these fluctuations is variation in vacuum level in the vessel enclosing the coldhead. In our run-in racks, several coolers’ vacuum vessels are served by a common vacuum pump; whenever a cooler is added or removed from the rack, the individual coolers’ vessels must be valved off from the common manifold. Occasionally one would be left valved off for some time, resulting in steady degradation of the vacuum surrounding the coldhead. The data over time clearly show these trends, yet quantifying the effects of vacuum loss on coldhead performance have proved challenging, and the inexpensive (Pirani) gauges we typically use have shown inconsistencies as well, so the data cannot be rigorously corrected for changes in vacuum level.

The overall trend for two years, however, seems to land squarely between the contamination load estimates. In a sense, this justifies our reliance on the contamination measurement as a leading indicator of cryocooler health, before we had enough long-term data to smooth out the short-term changes.

Further refinements to this method (in analysis and measurement techniques) may enable confident predictions regarding cryocoolers or components with very low levels of contamination. Its most obvious utility at this point is in estimating performance penalties or processing times in systems with much higher contamination levels.

CONCLUSIONS

Based on the results in the previous section, it would appear that the method we have outlined for quantifying cryocooler contamination is reasonably accurate and can reliably predict the long-term warming trend of a cryocooler. However, this appearance must be at least somewhat coincidental. There are a myriad of simplifying approximations that have been made in order to arrive at a method that is amenable to simple measurements and rules of thumb. Some of these are discussed in the beginning of the “Theory” section, such as the assumption that there is one dominant outgassing process, with one overall time constant. Others involve how the tests are performed.

For one, we use Pirani gauges for all the vapor-pressure measurements described here; but these are not especially accurate above 1000 mTorr and are especially inaccurate for water vapor, the main species of interest. For another, we find that the evolution of vapor pressure in a device or vacuum vessel containing a component often does not conform to the simple relaxing-exponential form of Equation (2); sometimes there are very clearly two relaxing exponentials (in other words,

two superimposed exponentials provide a much better fit to the data), and sometimes there is a linear component that must be subtracted out, as described in “Results”. This may be due to a small leak, or in the case of a cryocooler or compressor, it may be hydrogen outgassing, which comes to dominate the pressure trace (strictly speaking, the hydrogen outgassing rate is slowing down with time, but the curvature is not always detectable on the time scale of our experiments). An exponential component is not always visible against this background. For another, as a device or component becomes clean, it may well be cleaner than the vacuum lines used to connect it with a gauge or RGA, especially if these are the same lines used for bakeout. Notwithstanding attempts to maintain the vacuum lines themselves at an elevated temperature, it is often the case that water-vapor outgassing from a fully processed cryocooler can’t be detected above the background in the vacuum lines.

Rather than a definitive study, this work should be viewed as an introduction to what is possible, a frame on which refinements can be added. Good baseline measurements with capacitive manometers in place of Pirani gauges, and a more sophisticated view of outgassing mechanisms will perhaps take this method to a point where it is more generally useful.

ACKNOWLEDGMENT

This work benefited greatly from discussions with colleagues knowledgeable about vacuum systems and cryocoolers. Thanks in particular are due to Keith Gustafson, of Chart Industries, and Willy Gully (formerly of Ball Aerospace, now retired).

REFERENCES

1. Hall, J. L. and Ross, R. G. Jr. “Gas Contamination Effects on Pulse Tube Performance,” *Cryocoolers 10*, Springer US (1999), pp. 343-350.
2. Castles, S., Price, J. D., Glaister, D. S., Gully, W., Reilly, J., Nast, T., Kotsubo, V., “Space Cryocooler Contamination Lessons Learned and Recommended Control Procedures,” *Cryocoolers 11*, Springer US (2001), pp. 649-657.