Leaks: The Good, the Bad and the Ugly

What leaks are, how to classify and size them, and how to make useful leaks.

Steve Hansen

I. INTRODUCTION

When a vacuum chamber is evacuated with a pumping system, the rate of pressure decline will slow and eventually, for all practical purposes, cease. The minimum pressure that the system reaches is called the base pressure.

The job of the vacuum pump is to remove gas molecules from the system. In theory, a high vacuum pump should be able to remove each and every molecule that wanders into its inlet. In practice, the system itself is continuously contributing a seemingly infinite number of molecules and the pump has to contend with this load.

What this means is that a 1 liter chamber has more gas than the 1 liter's worth. There is gas that has adhered (adsorbed) to the walls of the chamber. There is gas that is within some of the materials that make up the system and that will, at reduced pressures, evolve into the chamber. Some materials will turn to gas (vaporize) at low pressures. Finally, some gas will enter the system through holes, cracks and other gaps in the system's walls.

In some cases gases are introduced intentionally into a vacuum system. This is the case with sputtering systems, ion sources and chemical vapor deposition systems, to name but a few examples.

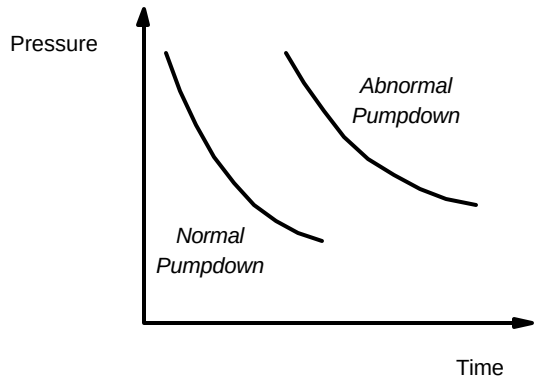

Figure 1 - Normal & Abnormal Pumpdowns

Each of these gas sources is a leak. The bad ones are termed real leaks. In these the offending gas is transmitted into the system through an actual channel from the outside world. The ugly ones are a result of gas sources within the system. These may result from poor materials choices or contamination of the vacuum surfaces by, for example, finger prints. These internal sources of gas are termed virtual leaks. The intentional leaks, the good ones, are the ones that are used to introduce process gases into the system.

Good, bad or ugly, the common denominator is that leaks represent the ingress of gas molecules, at some given rate, into the vacuum system.

In this article we'll discuss how to differentiate between real and virtual leaks, how leaks are sized, and how to make and size intentional, predictable leaks.

II. IDENTIFYING LEAKS

The performance of a vacuum system is dependent upon a whole host of parameters: the type and size of the pump, the size of the chamber, the length and directness of the lines between the pump and the chamber, and so forth.

If you have a system that you run on a regular basis, you will get an understanding of how quickly it pumps down and what the base pressure is. If you make a change to the system (perhaps it's just a matter of having opened the system or perhaps you put some new fixtures in or appendages on the system) you might notice that it doesn't behave as well as it did. What you might observe is a slower pumpdown and a poor base pressure. The differences are illustrated in Figure 1.

A lesson to be learned here is that it is useful and frequently important to keep notes of your system's performance. If you don't have a baseline, you won't know when something is amiss.

So, at this point we know something is wrong with the system. However, the pumpdown profile won't reveal what the problem is. It could be contaminated pump oil or it could be a partially closed valve. Or, it could be a leak.

Here is where an additional system feature is valuable. This feature is an isolation valve that is located at or near the pump inlet. If the system is pumped to its base pressure and the isolation valve is closed, the

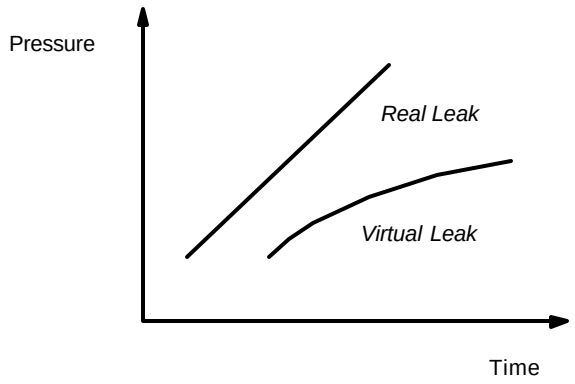

Figure 2 - Differentiating Leaks

pressure in the system above the valve will rise at some rate. (Obviously this requires that the gauge be located above the valve where it should be anyway.) The rate and form of this pressure rise profile is indicative of the size and type of leak.

Here is another place where a baseline is important: every system will exhibit some pressure rise. The baseline will let you know when it is abnormally high. If the pressure rate-of-rise is not abnormal, then you will have to look elsewhere for the problem. This could be a pump problem or a leak below the isolation valve.

If the pressure is rising abnormally, use your watch and record the pressure at a succession of equal time intervals. Try to let the pressure rise a couple of orders of magnitude and then plot the results on a linear scale.

Referring to Figure 2, a real leak will yield a linear pressure rate-of rise curve. The slope of the curve is a function of the leak rate and the volume of the system: the larger the leak, the steeper the slope; the larger the volume, the shallower the slope. We will quantify this later in the article.

The leak can be through an actual hole or channel, say a crack in a weld, a badly seated gasket or some other aperture. Or the leak can be a result of the gas from the outside room permeating through some component of the system. Elastomers, as are used in O-ring seals, are permeable to gases. There's no actual hole but gas molecules can work their way through the bulk of the material. Different elastomers have varying permeabilities: silicone is really bad with regard to this while Viton is relatively non-permeable. Permeation is selective as well: helium will work its way through a gasket much more easily than nitrogen.

Minimizing permeation leaks is a matter of selecting the correct elastomers (or other organics) and then minimizing the exposed areas of those permeable materials.

Permeation is not restricted to organics. At very low pressures, getting into the ultrahigh vacuum realm, permeation through glass and metals can start to become a problem.

Getting back to the rate-of-rise test, if the slope gets shallower as time goes on, you probably have a virtual leak. For example, gas trapped in a threaded hole under a bolt will leak out slowly causing the pressure to rise. Eventually the gas will all leak out (since this is a finite source) and the leak will appear to go away. This sort of situation can be verified by repumping the system and

Material	**10^{-6}**	**10^{-2}**	**1**	**10**
Water	-110	-60	-15	15
Mercury	-40	45	120	190
Octoil	35	125	200	235
DC-704	50	145	210	250
Cadmium	125	270	380	-----
Zinc	175	350	-----	-----
Lead	425	710	-----	-----
Silver	680	1050	-----	-----
Tin	820	1200	----	-----
Copper	850	1250	-----	-----
Carbon	1900	2400	-----	-----
Tungsten	2400	3250	-----	-----

Table 1 - Vaporization Temperatures for Selected Materials at Various Pressures. Numbers across the top are pressures in Torr. Values in the columns below the pressures are the vaporization temperatures in °C. These values have been derived from vapor pressure curves from several sources and are rounded off.

then looking again at the pressure rate-of-rise. Since the gas has already been dissipated, the virtual leak will no longer be present. It is for this reason that you can buy special bolts and screws for vacuum use that have holes drilled through them.

Moisture and surface contamination (finger prints, etc.) will have a similar effect. Eventually these will cease to be gas sources if the system is pumped for a long enough period of time. (Depending upon the level of vacuum required, it may be a *very* long time.)

You have to be cautious of the vapor pressures of the various materials that make up the system components as inappropriate choices can result in virtual leaks. Organics within the chamber will be sources of virtual leaks. If you are trying to get vacuum levels below 10^{-6} Torr, you have to be concerned about metals and metal alloys that have high vapor pressure constituents.

Table 1 lists a number of representative materials that might be found in vacuum systems with their vaporization temperatures at several selected pressures. It can be seen that water is a problem at almost any pressure: at 10^{-2} Torr, water will vaporize at any temperature over -60 °C.

Brass, an alloy of copper and zinc, is usually an acceptable vacuum material for medium vacuum work. However, zinc has a fairly high vapor pressure. Thus, if a brass component is used at 10^{-6} Torr in an application where the brass will be heated to a temperature of over 165 °C, the zinc will freely vaporize. A similar situation will be seen with cadmium plated hardware.

Copper is a pretty safe material given its low vapor pressure. Aluminum, stainless steel and other "high vacuum" materials similarly have very low vapor pressures.

If your virtual leak source is due to a basic materials incompatibility, then no amount of pumping is going to rid the system of that problem.

To conclude this section, what you'll probably find in your rate-of-rise test is a mixture of low-level real and virtual leaks that are harmless to your application (and which would probably take forever to fix anyway), and occasional nasty leaks that have to be fixed. These nasty things will also be combinations of real and virtual leaks. Now we'll look at how to size a leak.

III. SIZING LEAKS

As stated before, leaks represent molecules entering the system. Therefore, the proper way to specify a leak is in language that relates to how many molecules per unit time are being admitted. However, talking about molecules per second is a bit inconvenient as no other common vacuum measurement uses molecules as a unit. To get around this we have to go to the gas laws and Avogadro's number. The relationship that everyone learns in high school chemistry goes as follows:

> A container of 22.4 liters volume at 0 °C and one standard atmosphere pressure (760 Torr) will contain Avogadro's number of molecules, i.e. about 6.02×10^{23} molecules. This number of molecules is termed a mole and will have a mass in grams equivalent to the atomic mass of the particular gas as measured in atomic mass units (AMU).

For example, a 22.4 liter vessel of nitrogen (atomic mass 28) at standard conditions (0 °C and 760 Torr) will contain Avogadro's number of molecules and the gas will have a mass of 28 grams. The same container of helium will contain 4 grams of that gas.

If the temperature and pressure in the vessel are not standard, then the ideal gas law is invoked to adjust for the deviation from standard conditions. If the gas is at a higher pressure (but at standard temperature), there will be a proportionally larger number of molecules. If the gas is at a higher temperature (but at standard pressure), then there will be fewer molecules. (One caveat: make sure that you make temperature adjustments with Kelvin (absolute) units.)

The bottom line here is that if volume, temperature and pressure are specified, it is then possible to determine how many molecules are in the volume. This gets us to terms that we can use in vacuum practice as we always talk in terms of volumes, pressures and temperatures.

For vacuum purposes, the above relationship is normalized to more convenient volume units, i.e. liters or cubic centimeters. What we end with are *standard condition* volumes.

Using liters, the standard condition volume term is the Torr-liter. This represents the molecules contained in a one liter volume at a pressure of 1 Torr. Not stated but understood is a standard temperature of 0 °C. Adjusting from the relationship of 22.4 liters of gas at 760 Torr and 0 °C equals 6.02×10^{23} molecules, one Torr-liter would then contain about 3.5×10^{19} molecules.

Using cubic centimeters, the standard condition volume term is the std cc (scc). This represents the number of molecules contained in a 1 cc volume at a pressure of 760 Torr. A standard temperature of 1 °C is also understood. Going back to the pressure/temperature/volume relationship for a mole of gas, a std. cc contains about 2.7×10^{19} molecules.

In vacuum practice there is a term called *throughput* which is usually abbreviated as *Q*. This was briefly discussed in the very first issue of tBJ. Throughput is simply the number of standard condition volume units

that flow in a given period of time. Thus we have the usual terms of Torr-liters/second (T-l/sec) and standard cubic centimeters per second (Std. cc/sec) or minute (sccm).

The convention for which set of units you pick is usually based on the application. Q within the vacuum system is usually specified in terms of T-l/sec since pressures are usually measured in Torr and volume in liters. Leaks are usually quantified in terms of Std. cc/sec. Q from intentional leaks (i.e. flow control devices) is usually specified as Std. cc/minute (sccm) or, for high flows, Std. liters/minute (slm). Some handy conversions are:

1 Torr-liter = 1.32 Std. cubic centimeter (scc)
1 Torr-liter/sec = 79 Std. cubic centimeters/min (sccm)

Another term you will run into for leak rates is Atmosphere cc/sec (or Atm. cc/sec). This is about equivalent to sccm except that the pressure reference is the prevailing ambient pressure (usually near 760 Torr unless you live in Denver or Albuquerque).

Since all of these terms relate to molecular quantity, thence to the actual mass of the gas, standard condition volumes are also called *mass quantities* and standard condition volumetric flow is also called *mass flow*.

IV. RELATIONSHIP BETWEEN Q, PRESSURE and SPEED

Again going back to Volume 1, Number 1, there is a simple way of tying some basic vacuum units together in a useful way. In the term T-l/sec we have pressure (Torr) times speed (liters/sec). Thus, mass flow is equal to speed (e.g. pumping speed or line conductance) times pressure:

$$Q = P \times S$$

Please refer back to that issue for a few examples of how this is applied within vacuum systems. Suffice it to say that this is like Ohm's law:

$$I = \frac{E}{R}$$

where I (current), which represents electrons per unit time, is analogous to Q (molecules per unit time), E (voltage) is analogous to pressure, and $1/R$ (reciprocal resistance or conductance) is analogous to speed.

V. DETERMINING THE SIZE OF A LEAK

With that bit of grounding behind us, it is fairly straightforward to quantify the size of a leak in units such as sccm by using the pressure rate-of-rise measurement discussed in Section II. As noted:

- A pressure rise over time indicates a leak.
- A linear rate-of-rise indicates a constant leak rate, typical of real leaks.
- The slope of the curve is a function of both the size of the leak and the volume of the system.

Knowing that the Q of the leak is related to pressure, volume and time, we can determine the size of the leak from the rate-of-rise curve (pressure change with time) assuming we know the volume of the system. Of course, since we are measuring the mass flow, Q, as standard condition volumetric flow, we also have to adjust for these standard conditions. The relationship that does this is:

$$Q(sccm) = 79\left[\frac{273}{273+T}\right]\left[V\frac{\Delta P}{t}\right]$$

The guts of this equation is in the right hand set of brackets where ΔP is the change in pressure in Torr, t is the time, in seconds, over which the change in pressure occurred, and V is the volume of the chamber in liters. The units here are Torr $\times$ liters $\div$ seconds. The left hand brackets give us the temperature adjustment factor where 273 is 0 °C in Kelvin units and T is the temperature of the gas in the chamber, also in K. Finally, the 79 is the conversion from T-l/sec to sccm.

The biggest uncertainty here is the volume of the system. With a rule and a bit of time, a reasonable volume determination can be made.

VI. THERMAL MASS FLOW CONTROLLERS

Now we'll start to talk about the good leaks, the intentional sort that are used to admit gases into vacuum systems. The gases may be used to control pressure or they may be used to produce a plasma or promote a chemical reaction. There are various forms of leaks that are suitable for doing this. Needle valves are one common form. The main problem with needle valves or other simple low conductance leaks is that it is not easy to determine with any sort of precision the Q of the leak at any particular setting.

Precision leaks are a big deal in industry. Semiconductor vacuum process equipment, as one example, would not produce anything but junk were it not for the accurate and repeatable leaks that are used to

admit the reaction gases to the systems. These are a far cry from the common needle valve.

The most common type of device for this purpose is the thermal mass flow controller (T-MFC or simply MFC). As the name implies, these rather clever instruments work by measuring the rate at which the molecules flowing through a sensor tube transport heat. Figure 3 shows a typical 3 inch footprint commercial MFC with a schematic diagram of the internal workings.

There are 4 major components in an MFC: the mass flow sensor, the flow bypass, a control valve and the electronics.

The mass flow sensor consists of a short length (on the order of an inch) of small-bore capillary tubing. The small diameter of the tube ensures that the gas traverses the tube in viscous, laminar flow. This simply means that the molecules flow parallel to the axis of the tube without turbulence. This makes the device repeatable.

Placed around the tube are a pair of windings that act as heaters and, by monitoring the resistances of the windings, the temperature of the tube at those locations.

Each of the windings is attached to a bridge circuit that maintains each of the windings at a specific temperature (i.e. each winding is held at a constant resistance). When the gas enters the sensor tube, the first heater warms the gas to some temperature. The molecules of the gas transport the heat down the tube.

At any given gas flow, each of the heaters will have some amount of power delivered to it to maintain the desired temperature. If the gas flow rises, the first heater will require more power to maintain the constant temperature. From this, more heat will be transported to the second heater (more molecules per unit time) and its power will be decreased. The bridge circuit looks at how much power is required by the heaters and provides an output that, after linearization and scaling, is directly proportional to the gas' mass flow.

The skinny tube can only transport so much gas, usually just a few sccm. To handle larger mass flows a bypass element is placed in parallel to the flow sensor tube. The bypass is just another laminar flow path and it is sized so that there is a constant splitting ratio between the two paths. The bypass serves exactly the same function as a shunt does in a ampere meter. By using bypasses of various sizes, MFCs can be produced that have full scale flows from one to thousands of sccm.

What I've described thus far is the metering size of the device. To actually control the mass flow a control valve is introduced along with the appropriate closed-loop circuitry. To make the device work, you

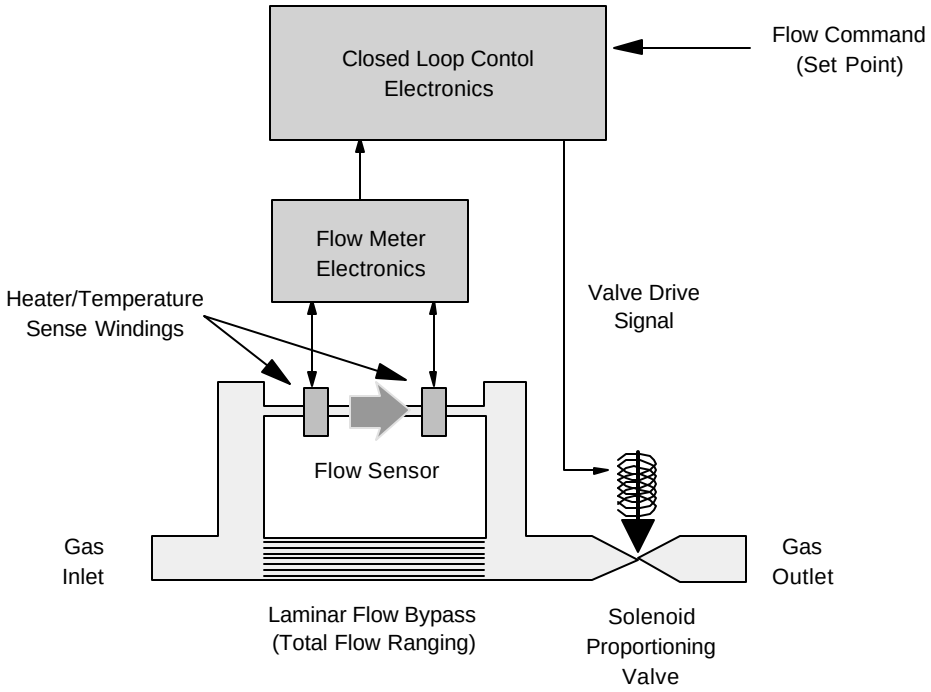

Figure 3 - Thermal Mass Flow Controller Schematic

place it in the gas line and apply a control voltage (usually a 0-5 volt signal) to get the desired flow output.

Since MFCs work on the basis of the thermal properties of the gas, they are sensitive to the gas type. When you buy one you have to specify not only the sccm range but also the gas type. Manufacturers provide correction factors for hundreds of different gases. Each correction factor is a constant that is referenced to nitrogen = 1.

Thermal MFCs are great instruments but their prices ($500 range for bargain basement instruments to $1000 for a good general purpose MFC) are prohibitive for most amateur uses. Next we'll discuss a simple approach.

VII. SIMPLE CAPILLARY LEAKS

A piece of capillary tubing can be used as a reasonably good and predictable leak. First, the diameter and aspect ratio of the tube (small and long with respect to diameter) assures that the gas is in laminar flow for virtually its entire length and, by knowing the dimensions and inlet and outlet pressures, it is possible to calculate the Q of the tube with reasonable precision.

The equation that is used to calculate Q for long, round tubes was developed independently by Hagen and Poiseuille and is therefore called the Hagen-Poiseuille equation. It is:

$$Q = \left(\frac{\pi d^4}{128 \eta l} \right) \left(\frac{P_1 + P_2}{2} \right) (P_1 - P_2)$$

where d is the diameter of the tube, η is the viscosity of the gas, l is the length of the tube, P_1 is the inlet pressure and P_2 is the outlet pressure.

When using this (or any other) equation, you have to get the units right. Most of the recent references that I have seen state the Q in Pascal-meters/second (one Pa is equivalent to 7.5 milliTorr). Given that, d and l are in meters, η is in Pa-sec, and P_1 and P_2 are in Pa. I've listed the viscosities of several gases in Table 2 in units of Pa-sec. along with the equivalents in Torr-min. This will permit you to get Q in sccm using units of cm and Torr.

As an example, a 36 inch length of 0.007 inch inside diameter capillary tubing, with vacuum at one end and 760 Torr at the other, will have a calculated Q of about 4.5 sccm with nitrogen at room temperature.

VIII. MAKING A SIMPLE CAPILLARY LEAK

Figure 4 shows a simple implementation of a controllable capillary leak. The leak itself if a length of stainless steel hypodermic tubing. I got a selection of tubes in 2 and 3 foot lengths from Small Parts, Inc., 13980 N.W. 58th Court, P.O. Box 4650, Miami Lakes, FL 33014, (800)-220-4242.

After selecting the correct size and length of tubing for my application (more on this in the next section), I sealed each end of the tube in a brass hose barb. While doing this I placed a piece of 1/8" id neoprene automotive hose over the capillary. This was mainly for mechanical protection. The downstream end of the capillary was screwed into a brass KF16 flange. The upstream end was fitted to an inline sintered metal filter (Speedaire or equivalent, available from hardware stores), a pressure gauge, and an inlet connector. Gas is supplied from a regulated source.

I only needed one gauge as my application, which will be discussed in the next section, has the downstream end of the tube at vacuum. Therefore P_2 is essentially zero.

By varying the inlet pressure over a moderate range, the Q of the leak can be changed quite substantially.

Table 2 - Viscosity Data for Selected Gases

Gas	Viscosity ($x10^{-4}$ Pa-sec)	Viscosity ($x10^{-8}$ T-min)
Argon	0.229	0.286
Butane	0.076	0.095
Deuterium	0.126	0.158
Helium	0.199	0.249
Hydrogen	0.090	0.112
Nitrogen	0.180	0.225
Oxygen	0.207	0.259
Propane	0.082	0.102

Source: Perry's Chemical Engineers' Handbook, Robert H. Perry and Don W. Green, 6th Ed., McGraw-Hill Book Co., 1984. Data valid at 300K.

Doubling the inlet pressure from 0 psig (nominally 760 Torr at sea level) to 15 psig (nominally 1520 Torr) will result in a 4x variation in mass flow.

One problem with using a conventional Bourdon-type gauge is that it will not compensate for atmospheric pressure variations: it's referenced to the prevailing atmospheric pressure. If your needs don't include high precision you can probably live with this. If not there are a couple of alternatives. One would be to have a barometer nearby. You could then develop a correction table for variations in atmospheric pressure. More convenient would be to use a simple differential pressure gauge, such as a Dwyer Magnehelic® gauge with one port connected to the vacuum chamber, the other to the inlet end of the capillary. As long as the vacuum is on the order of a few Torr or less, the gauge will accurately (±2% of full scale per Dwyer's specification sheet) monitor the absolute upstream pressure. Magnehelics with ranges up to 30 psi are available.

IX. AN APPLICATION

My specific application is shown in Figure 5. This shows the inlet manifold that I have added to my 2 inch CVC diffusion pump. The arrangement was designed to permit me to introduce a variety of gases into the system, monitor the inlet pressure of the diffusion pump (to ensure that it reaches a good base pressure without any gas being introduced and that the inlet doesn't rise to too high a pressure with gas flowing) and control the pressure at the chamber. I also added a simple vent valve.

The manifold is a stackup of several standard vacuum fittings. Starting at the pump, I made an adapter for the CVC's inlet flange from a 6 inch to 2-3/4 inch "zero-length" CF adapter that I had obtained as surplus. This large flange also is used to mount the entire assembly through the top of my vacuum bench.

Above the adapter is a 2-3/4 CF to KF40 adapter. Using all KF components above this part makes it easy to rearrange the reset of the manifold. At the pressures I'm working at (10^{-5} Torr, give or take) the elastomeric seals don't pose a problem.

The high vacuum gauge I use is an old CVC Philips-type cold cathode gauge. It has a 1/2 inch tubulation which I soft silver soldered into a KF40 to 1/4 NPT (female) adapter tee (MKS Instruments, Inc. part number 100319605). I drilled out the threaded hold to accept the gauge connection.

Above this is a KF40 butterfly valve. This is a surplus unit that I got with a defective internal O-ring seal. The valve serves as a throttle and does not need to fully seal anyway. Next is another KF40 to 1/4 NPT adapter tee. To the 1/4 inch port I have attached a Speedaire filter with a flare adapter, flare cap and O-ring. This is the vent. Also attached is the KF16 flange that goes to the capillary leak. Operation is as follows:

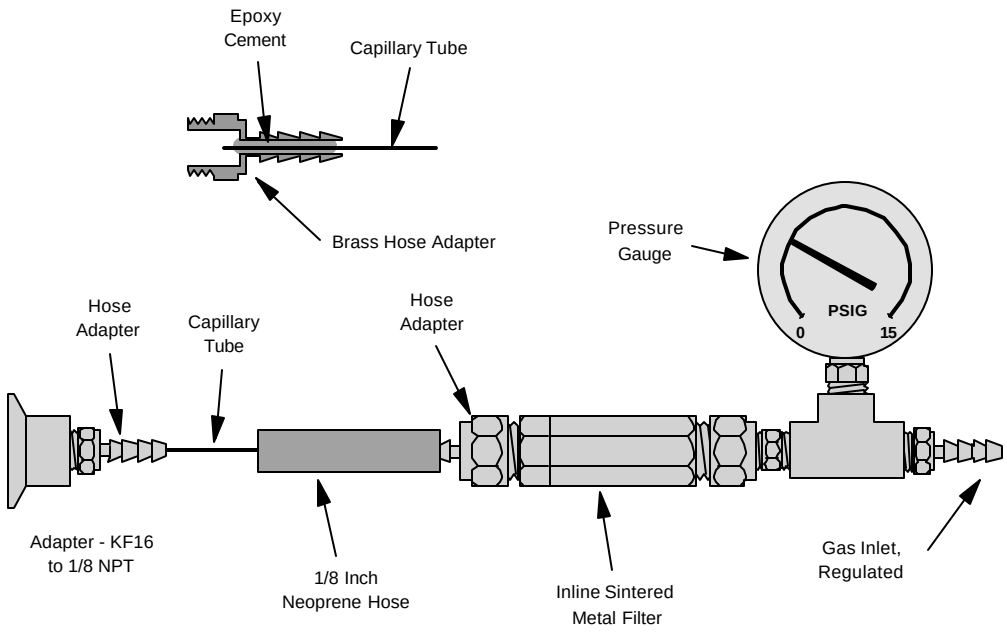

Figure 4 - Capillary Leak

- With gas closed off to the leak (via a valve, not shown, at the upstream end of the capillary), the system is pumped to base pressure. The butterfly valve is open.
- When the base pressure has been reached, the gas flow is started by opening the valve upstream of the capillary. Pressure at the inlet of the capillary is adjusted to just above atmospheric pressure (low flow). Chamber pressure (gauge not shown) and pump inlet pressure (ion gauge) are monitored.
- If things have been sized correctly, the ion gauge should indicate a low enough pressure for proper operation of the pump and the chamber pressure should be lower than desired.
- The butterfly valve is partially closed to raise the chamber pressure. If this is not adequate, the gas pressure is increased in small increments to admit more gas into the manifold.

While the above could be accomplished with a needle valve, the capillary ensures that the flow rate will never be so high as to cause a large gas burst to enter the system. I hate cleaning up after a diffusion pump dump.

Looking at some calculations, the pump has a rated speed of 105 liters/sec at the pump's inlet flange. With 4.5 sccm of nitrogen flowing, the inlet pressure ($P = Q/S$) should be about 0.5 milliTorr. To achieve a 10 milliTorr chamber pressure at the 4.5 sccm flow rate, the manifold conductance will have to be reduced to about 6 liters/sec. The stack of KF40 manifold fittings probably gets the conductance down to around 10 liters/sec by themselves. The butterfly valve does the rest.

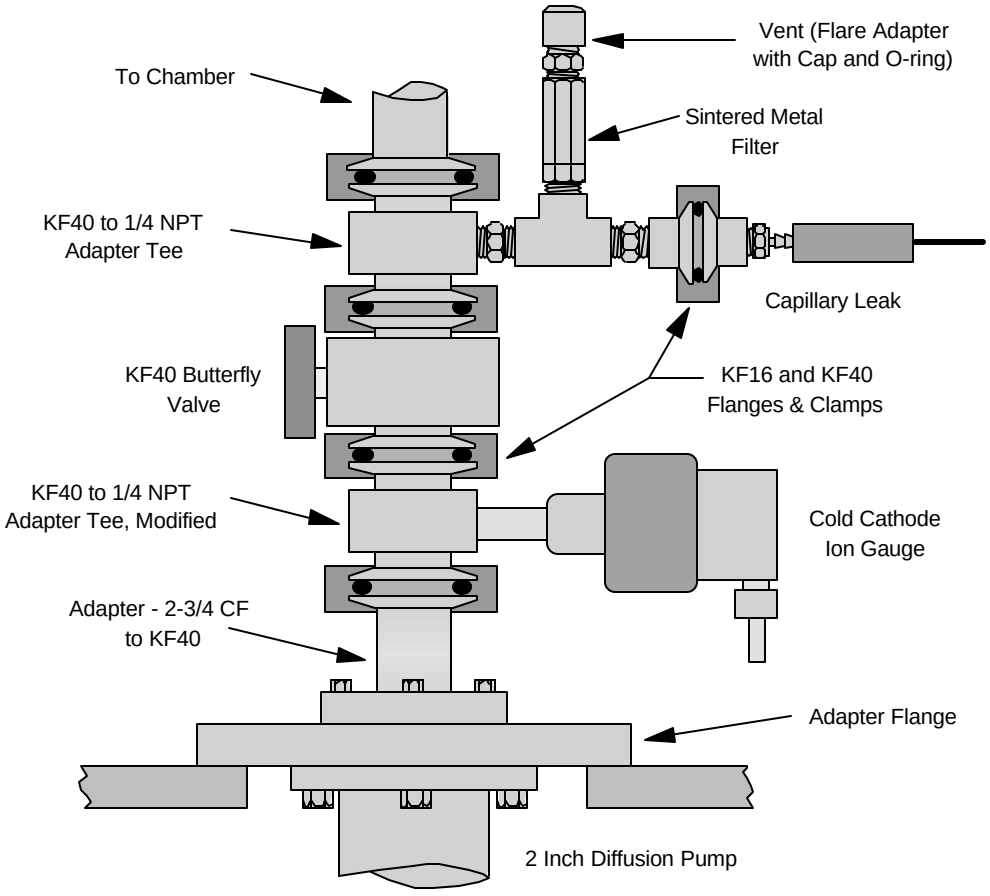

Figure 5 - Pressure Measurement and Control Manifold for a Small Diffusion Pump

