MILLROCK TECHNOLOGY

Freeze-Dry Process Monitoring

Gregory A. Sacha Ph D., Senior Research Scientist at Baxter

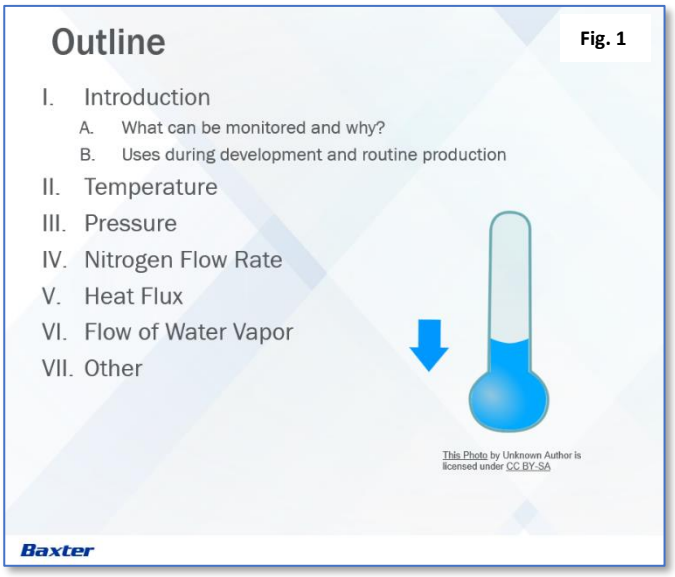

Fig 1 What can we collect during the freeze-drying process and why do we want to? This presentation will discuss just that, and then also take a little more in-depth look at temperature and pressure throughout the process.

We will consider alternate methods of collecting data such as nitrogen flow rate. We will also discuss Heat Flux, which is offered by Millrock Technology. Finally, we will also discuss flow of water vapor, and possible other routes of data collection.

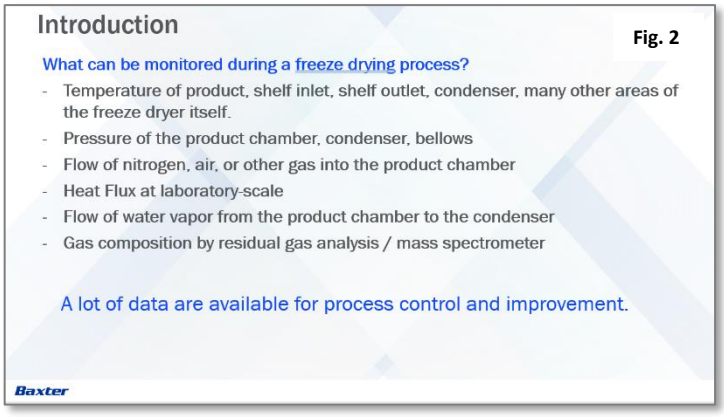

Fig. 2 If you work with lyophilizers in the lab, you may already be familiar with monitoring the temperature of the product, shelf inlet and shelf outlet, and how shelf temperature compares with product temperature. You'll also keep a close eye on condenser temperature, and then many other areas of the freeze-dryer itself. It's complex, how many places within the freeze-dryer that temperature can be measured.

We also look at pressure. Pressure of the product chamber, of the condensers, as well as what is known as “bellows.” Bellows, at full-scale, are those areas that circulate around the stoppering ram. We want to know what the pressure is there, in order to be aware of any possible leaks that could affect our sterility assurance.

We are also interested in the flow of nitrogen, air, or other gas into the product chamber during both primary drying and secondary drying, but more so during primary drying. As we are removing water vapor in the chamber, and as that water vapor flows, there is a large pressure differential created from the vapor pressure of ice in the product chamber, to the vapor pressure of ice in the condenser. The condenser is at a much lower temperature, at around -60 degrees C or lower. That means the vapor

pressure of ice is very low at condenser, and that establishes the flow of water vapor during primary drying.

For that water vapor to flow and to maintain pressure, we must replace it with something. And typically, during full-scale manufacturing, it is nitrogen that we are continually bleeding into the chamber.

At laboratory-scale and at some pilot-scale, a heat flux sensor is available. That heat flux sensor is very similar to what we use for differential scanning of calorimetry, where we can sense the heat flowing from the shelf to the actual vial or vials. In this way, we should be able to determine exactly what is happening in those vials based on the heat flow into or out of the vials. We also can have the option of looking at the flow of water vapor from the product chamber to the condenser.

And finally, we can also analyze the gas composition by either a residual gas analyzer, an RGA, or a mass spectrometer.

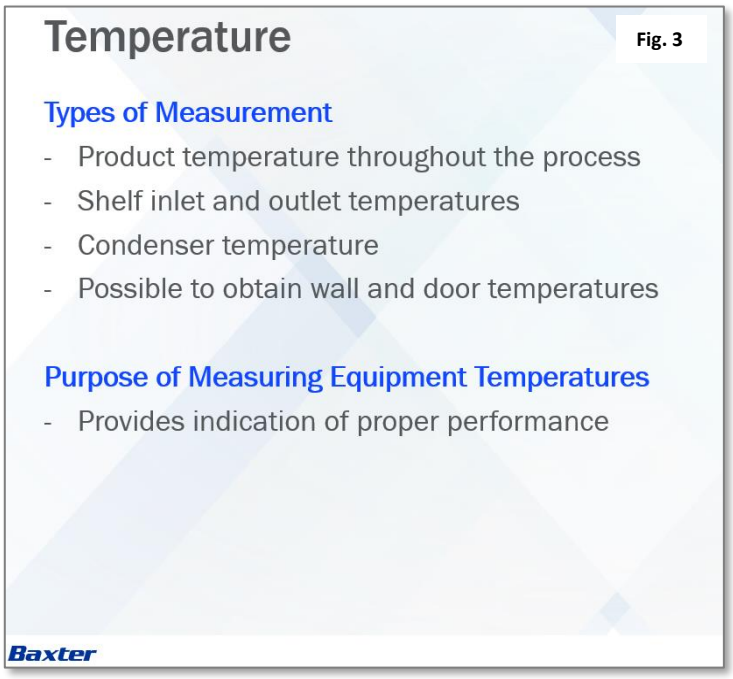

Fig. 3 We can collect product temperature throughout the process, but only on a few vials. We cannot measure the product temperature in every vial. We often do collect shelf inlet and outlet temperatures on all the shelves. We collect the condenser temperature; we need to do that closely because that is what ensures that we have a very low vapor pressure and establishes our flow of water vapor. And it's even possible to obtain our wall and door temperatures.

We want to monitor wall and door temperature because heat can be transferred from the doors or the walls to our product sitting on the shelves. Vials located along the edges of a shelf tend to be slightly warmer than the vials located in the center. That difference is larger, the lower the temperatures at which we work. So, measuring the temperatures outside the product temperature is very useful for monitoring and having an indication of performance of our equipment.

Fig. 4 Product temperature can be tested in a few vials, but we cannot control it directly. In the laboratory, we monitor a few vials with thermocouples, or some type of temperature sensor. The product temperature is completely dependent upon the combination of shelf temperature and chamber pressure. They work together for us to obtain our desired product temperature and maintain it below some critical value.

We must maintain the product temperature below some critical temperature, and that critical temperature we typically identify using thermal analysis such as differential scanning calorimetry to examine what occurs in our frozen solid when we cool it and then slowly warm it. We can also use freeze-dry microscopy, which gives us a good visual analysis of when product failure occurs.

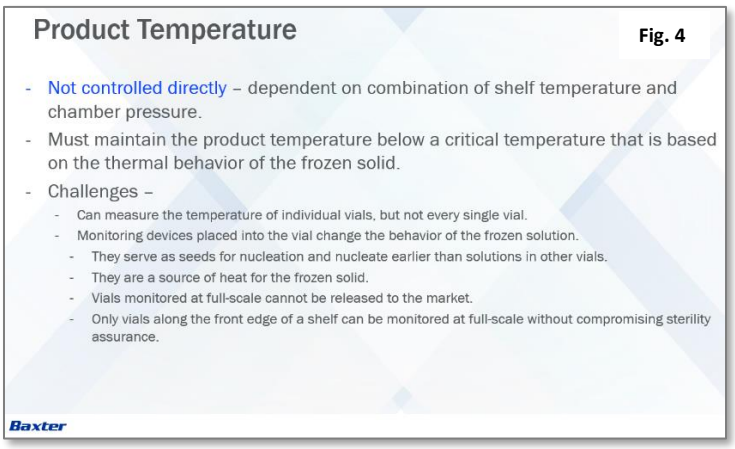

Fig. 4

Together, both can give us a good idea of where our product can fail based on temperature, but a true measure of this is by identifying the failure point during primary drying, which we can do using about five to ten vials filled with our desired fill volume in the specific vial that we intend to use. Product temperature does change based on fill volume and total concentration.

Therefore, we want to see and challenge that product to determine at which combination of shelf temperature and chamber pressure we will reach that critical product temperature.

There are several challenges that we run into during this process. We cannot measure temperature of individual vials, meaning every single vial in our batch. We wouldn't want to because those vials that have that temperature sensor in them are different from every other vial present. This is because the temperature sensors act as seeds for nucleation. If a vial contains a temperature sensor, they will likely nucleate a lot earlier than the other vials, allowing for more time for crystal growth. If there's more time for growth of ice crystals, we have larger ice crystals. And when they begin to be removed, we have larger pores so those vials can now dry faster.

Those temperature sensors are also a source of heat for the frozen solid. This past summer, we had an intern that designed a method of placing cameras inside a lyophilizer focusing on one of the vials containing a temperature sensor. You could easily see how that temperature sensor influenced the behavior of the frozen solid. In some cases, it even began to melt right around that temperature sensor.

Vials that are monitored at full-scale cannot be used at market and are only used for in-process data. An added challenge about product temperature monitoring is that we can only measure those vials located along the front edge (unless we have wireless sensors) otherwise we have to reach over all of the other filled vials to place a sensor, and that would compromise sterility assurance.

Fig. 5 We often use thermocouples to measure temperature, which are point sensors. But we can also use resistance temperature detectors (RTDs), which are more area sensors. And then we'll also look at another sensor called temperature dependent resonance. All three of these must be placed directly into the solution. Therefore, vials equipped with those sensors are completely different than the rest of the batch. We most often use these sensors during laboratory-scale development and then scale up in

engineering batches so that we can compare that data with our laboratory-scale data. However, there are companies that use temperature sensors routinely during manufacturing.

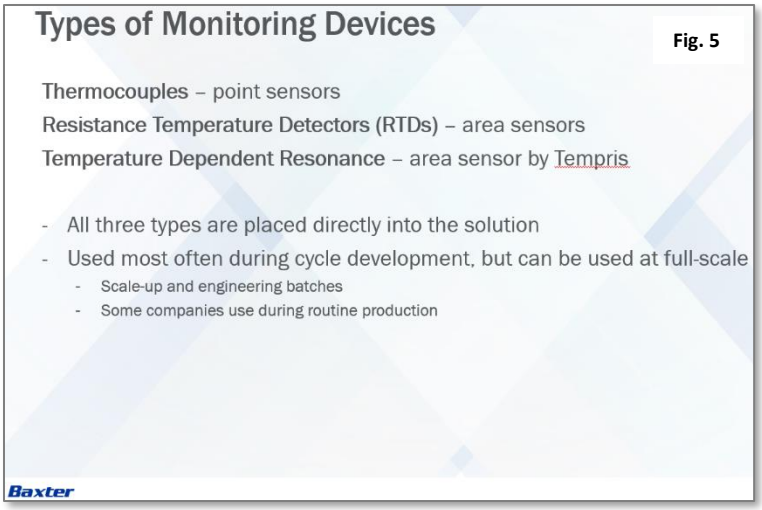

Fig. 5

We must ensure that the sensors are placed and maintained in the vials. Once the lyophilizer door closes you cannot make sure, visually, that the sensor is still in the vial.

Fig. 6 A thermocouple is typically two wires from different metals that are joined at the ends. These ends now create junctions and are placed at two different temperatures. One can either be placed in an ice bath as a reference temperature or, usually, we use electronic simulation. Once these two ends are at different temperatures, a current will flow as a result of the voltage difference. This is known as the Seebeck Effect. Looking at the photo, we see one end on the left, the end of the thermocouple is the point sensor. It measures temperature only at that point. Therefore, when placing this into the solution it's important that we place it directly into the center. Then there is the connection port, this is what we plug into the lyophilizer.

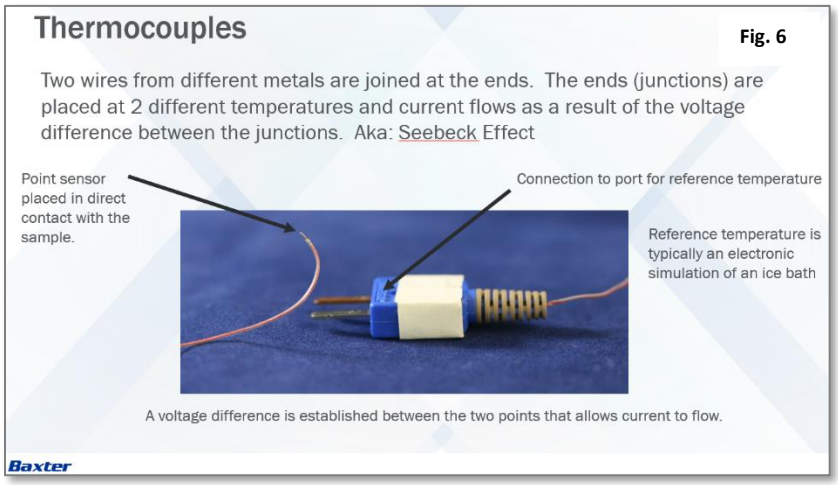

Fig. 6

Fig. 7, 8 Here's an example of a thermocouple, with a type of thermocouple holder. It is very difficult to place a thermocouple and maintain its position in a vial without some type of thermocouple holder. With this thermocouple holder, you thread the thermocouple through the top and pull it through the bottom, where it helps maintain it in the center of the vial.

The desired placement is in the center touching the bottom of the vial. That's the last area in the frozen solution that will have ice, and therefore will tell us when all the ice has been removed, which is when the temperature of the product starts to increase.

Fig. 7

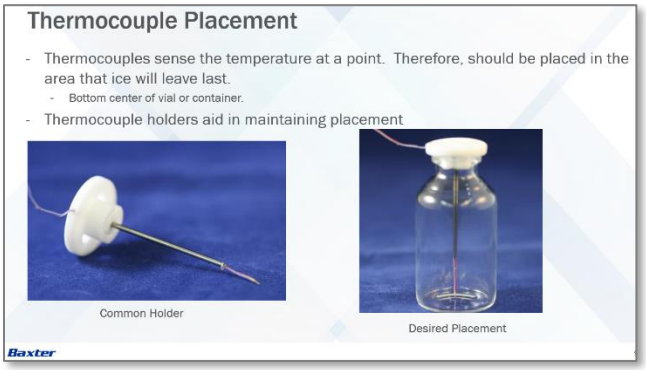

Fig. 8

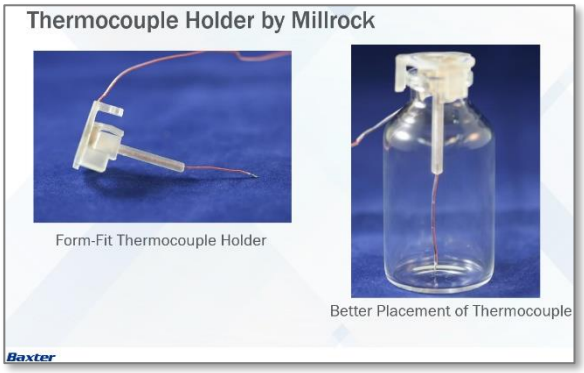

Fig. 9 Another type of temperature sensor is a resistance temperature detector, or RTD. These have a ceramic core that has a metal coil that's wrapped around. This measures the electrical resistance of the metal as the temperature changes. You will notice that the bulky end in the vial is larger than a thermocouple and therefore takes up a lot more room. It also measures an area- total temperature in that area. It's typically used in full-scale manufacturing for fixed point measurements. We can place it to measure shelf temperatures, inlet and outlet, and other temperature sensitive areas within the equipment.

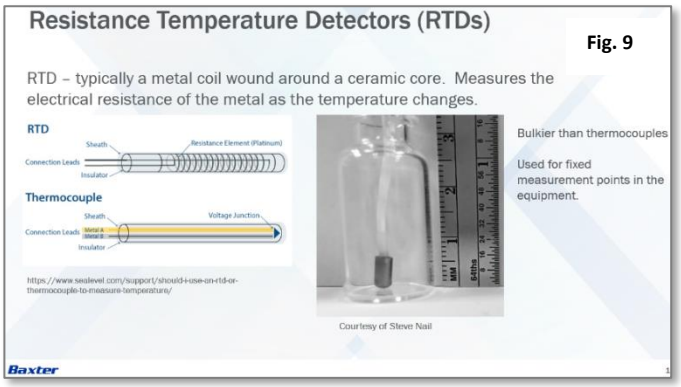

Fig. 10 This is a Tempris System, by iQ-mobil. It's different from a thermocouple or an RTD. In this case, there's a handmade quartz crystal. If you look at the picture there where it says "sensor" and there's an arrow pointing to it. That crystal oscillates in a temperature-dependent frequency. It is excited by modulated microwave signal, and then it's transferred directly to a control and recording instrument that's equipped with an antenna.

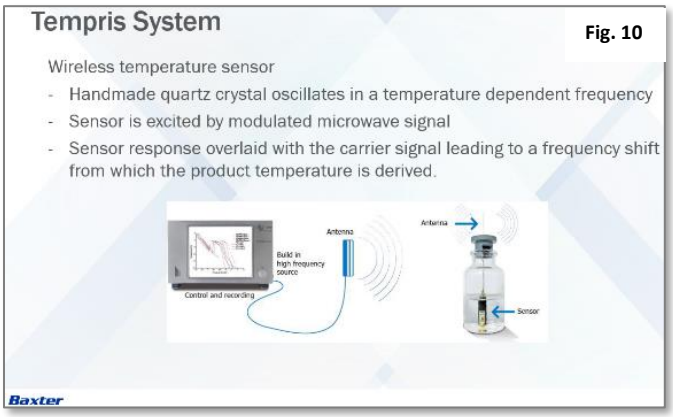

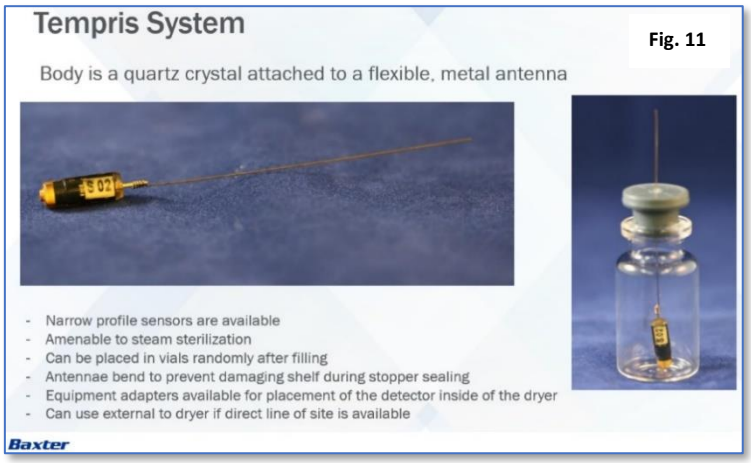

Fig. 11 We can take that sensor, thread it through a stopper, and that piece above it is a flexible antenna. There are also narrow profile sensors available. Even though these look bulky, they really do not drastically affect the volume within the vial, and their temperature recordings match very well with thermocouple data. These are amenable with steam sterilization and what's nice about them is that since you thread them through a stopper and they're wireless, you can now place them along the conveyor as vials are filled. Those antenna bend, and that's important. If they did not bend, the stainless-steel shelves are quite thin, they look robust, but they're sensitive enough to be punctured by something like that.

The challenge here is that we must have a direct line of site to the antenna to record the data. There are some full-scale lyophilizers that have a window port on the side. In some cases, you can place the antenna externally and collect the data. If that doesn't work, there are options of placing the antenna on a moveable pole within the lyophilizer. That antenna will then collect the data and transfer it to your control so you can see data in real-time.

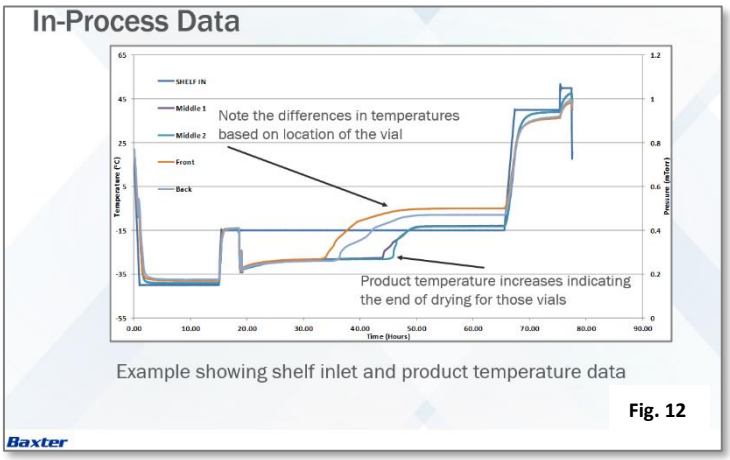

Fig. 12 Here's an example of in-process data. The x-axis is time, y-axis is temperature, secondary y is pressure (we are not showing the pressure graph). And what you see are two thermocouples that were placed in vials located in the front and back of a shelf, and two thermocouples that were placed in the center- vials in the center of the shelf. There's a drastic difference here in the temperature change. You can see those vials located in the front and back start to become warm a lot sooner. That is because of the temperature difference between the edge vial and the center vials. That's why, when measuring the temperature at full-scale, if you're only looking at the front vials, you really are not getting the full picture.

Fig. 13 Vials that are equipped with a sensor are different than the rest of the batch. There is less supercooling, meaning that they'll nucleate at warmer temperatures. So, let's say they may show nucleation at -5 rather than the typical supercooling around -10 or -15. That allows more time for freezing and crystal growth, leading to larger ice crystals. Challenges at full-scale: the sensor can move. We may not be collecting the data we think we are. Also, placing those thermocouples within the lyophilizer can risk sterility assurance.

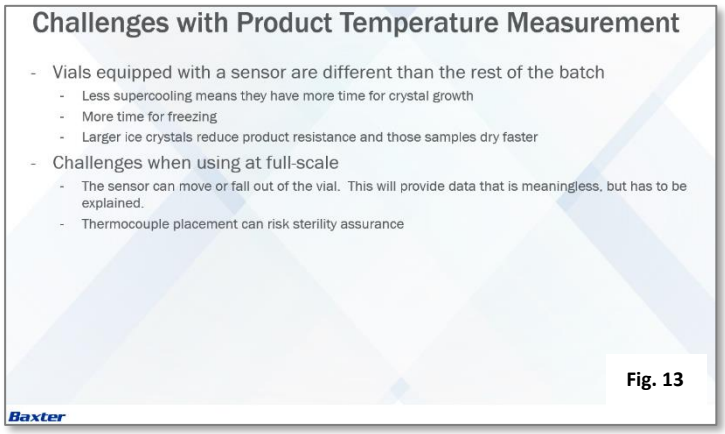

Fig. 14 Let's look at pressure measurement. Common instruments are the capacitance manometer and Pirani gauge. We cannot directly control product temperature. The only way we can control product temperature is by manipulating chamber pressure and shelf temperature. We must work with those together to indirectly control our product temperature.

If we increase the pressure, meaning less vacuum, we'll increase the product temperature. If we decrease the pressure, meaning more vacuum, we'll decrease the product temperature.

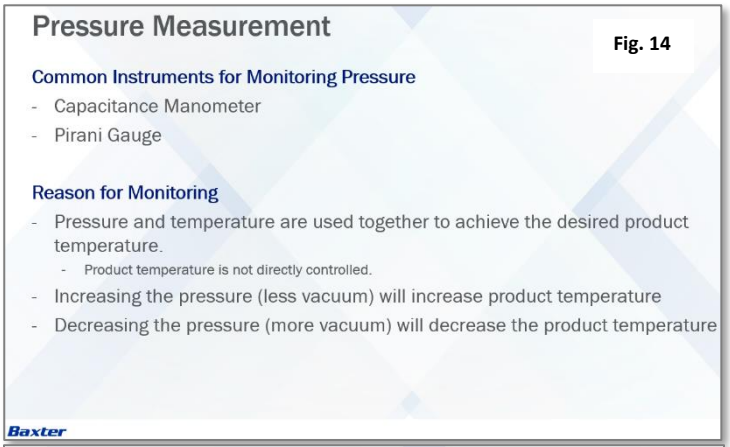

Fig. 15 One method of measuring pressure during the process is using a capacitance manometer. The Capacitance Manometer has a diaphragm that moves based on differences in chamber pressure. The output for this is completely independent of gas phase composition. And it's linear for over four orders of magnitude. What you're looking at in the picture on the right is the actual size of the capacitance manometer placed on the port of a lyophilizer.

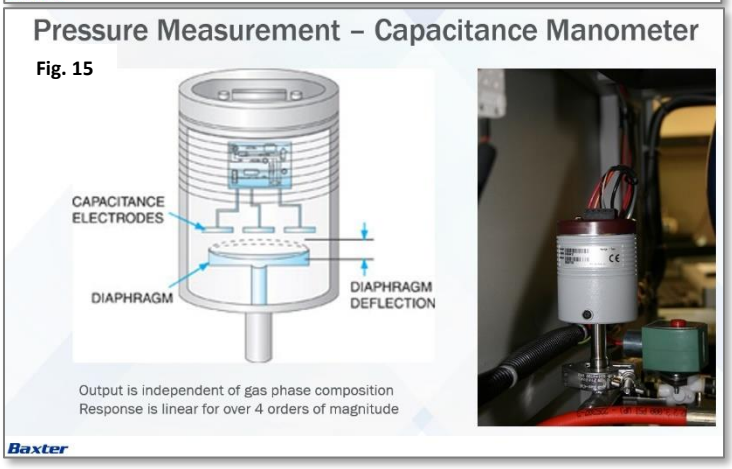

Fig. 16 We often compare the pressure from our Capacitance Manometer with the pressure from our Pirani gauge. The Pirani gauge is a thermal conductivity type of gauge. There are two filaments there that create a Wheatstone Bridge. One filament is at a reference temperature at constant pressure within a constant gas phase composition. The other filament is the measurement filament, it is heated.

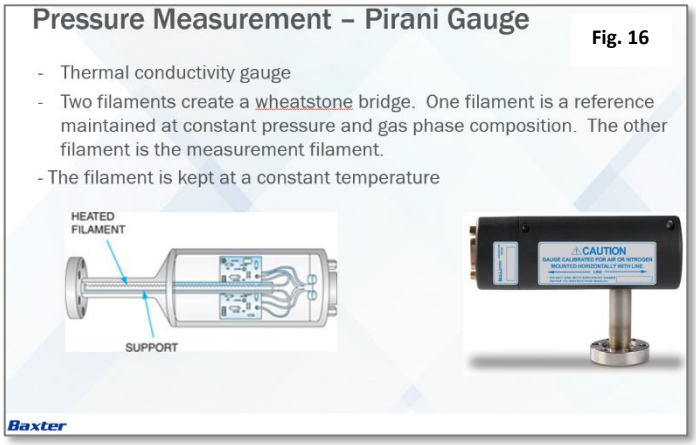

We're trying to keep this filament at a constant temperature, but when water vapor is very high in the chamber during the main portion of primary drying, it decreases the resistance to heat for the filament. Therefore, it reads an artificially high pressure in the chamber. We can now use the comparison of the Pirani gauge with that of the capacitance manometer, which is your true pressure, by looking at when the pressure from the Pirani gauge decreases. That is an indication of when primary drying is complete.

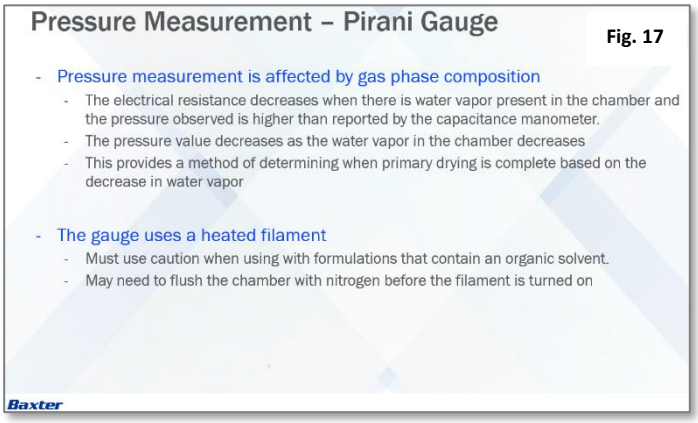

Fig. 17 The electrical resistance decreases when there is water vapor in the chamber. And this allows us to make a direct comparison of the pressure difference between the Pirani gauge and a capacitance manometer. The capacitance manometer displays and records the pressure that we set our process to. If you intend to use 100 millitorr, it will read that true pressure. The challenge with the Pirani gauge is that it does use a heated filament. That means we need to use caution if we intend to use an organic solvent in our formulation. In those cases, we need a method of turning off the Pirani gauge until we have a vacuum initiated and nitrogen within the chamber. At that point, it is now safe to turn the Pirani gauge on.

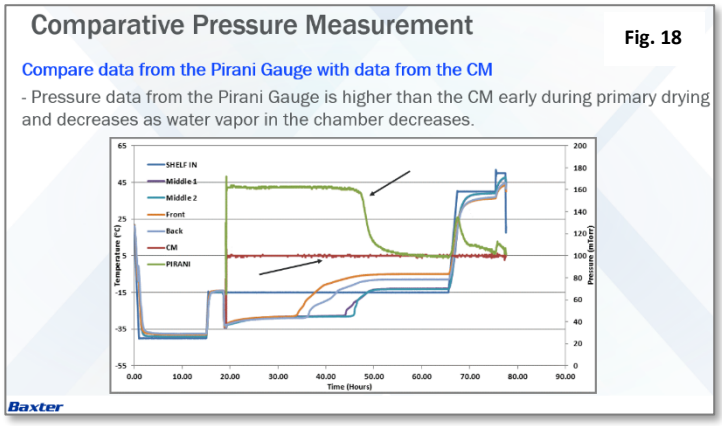

Fig. 18 Here's an example of the data. This is an example of comparative pressure measurement. We see, just like in the graph before, x-axis is time, the thermocouples show front and back, middle, and middle 2. But now we have the red as the capacitance manometer, that's your true pressure. And your Pirani gauge, which, during primary drying, suddenly you see it decrease, and that indicates that water vapor in the chamber is now removed and it's now safe to proceed onto secondary drying. You'll see that for those vials that were monitored with a thermocouple, they tend to end earlier, or what we call "break", when they start to increase in temperature. We say the thermocouples have

broken, meaning that there's no longer ice there, and they can warm to the temperature of the shelf. We'll see that they do that much earlier than the entire batch by which the Pirani gauge is measuring.

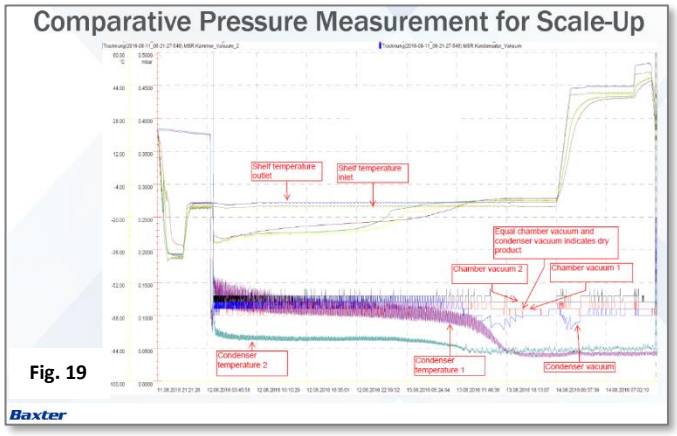

Fig. 19

Fig. 19 We can easily use this now at full-scale. This is an example of our full-scale data from our site in Germany, where, during an engineering batch, it was equipped with thermocouples. There you see the shelf temperature, and it is a bit more difficult to see the chamber pressure here, but we can determine when the Pirani gauge is equal to that of the Capacitance Manometer. This is a great and easy way to compare lot to lot variations as well as a method for identifying and solving excursions during a process. Excursions such as temperature excursions of your shelf or pressure excursions during your process. If we know a little bit about what's going on during that process, we can easily determine if it negatively impacts our process.

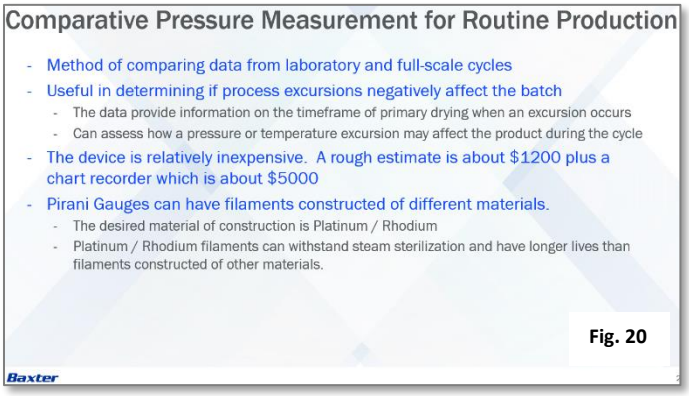

Fig. 20

Fig. 20 Regarding comparative pressure measurement for routine production, it's amazing how many sites do not have this. It's an easy and useful way to determine if process excursions will negatively affect your batch and the really nice thing is that it's inexpensive. Let's say a rough estimate for the device is about $1,200. And a chart recorder, which can record the data electronically, is about $5,000. That's nothing compared to the cost of the equipment or the price of your batch.

One thing to remember is that Pirani gauges have filaments constructed of different materials. It is recommended, when using these at full-scale, that you use a platinum/rhodium blend. Only the platinum/rhodium blend is robust to multiple steam sterilization cycles. There was a lot of discussion a while ago stating that Pirani gauges were not robust during freeze-drying and that they had to be replaced often. When it was investigated further, those sites that did not use platinum/rhodium, had the most challenges with having to replace the Pirani gauges. It's well worth understanding the construction of your gauge.

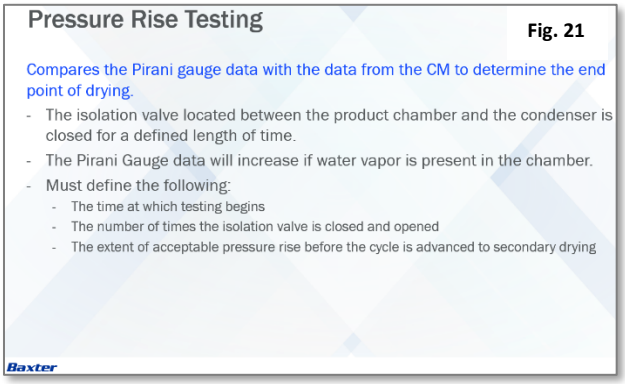

Fig. 21 These Pirani gauges and capacitance manometers can also be used in pressure rise testing. This is where the Pirani gauge data is paired with a capacitance manometer, and it's done in a methodical method to determine the end of drying. What occurs is that you have a product chamber separated by your condenser by some type of valve, in this case, an isolation valve. The isolation valve is closed, and as soon as it closes, you'll see a rise in your pressure recorded by the Pirani gauge. This can be a measure of how much water vapor is still the in your chamber. In order to use this method, you must define a few things. One, you must define the time at which testing begins. At what time during the process do you start experimenting with closing that isolation valve? Next, you must define the number of times the isolation valve can be closed and opened during a process and still be acceptable. Finally, we need to determine the extent of acceptable pressure rise before the cycles advance to secondary drying. There's some work that needs to be done to better understand how to use that data. We need to determine exactly how much that pressure can rise.

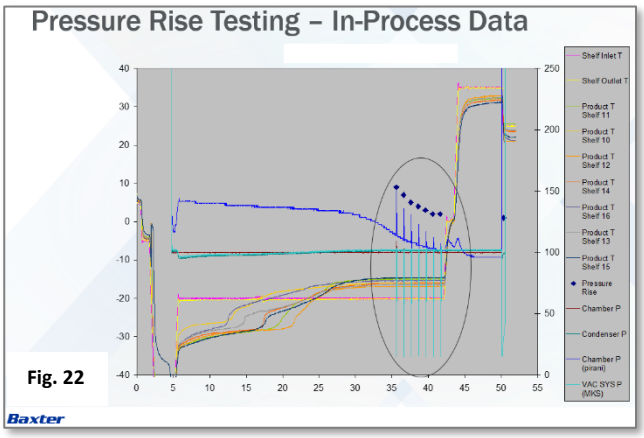

Fig. 22 Here's an example of the data. Right where the oval is, you'll see that blue line start to increase. That's the Pirani gauge. The red line is the capacitance manometer. What you see are spikes in the Pirani gauge when the isolation valve is closed. Multiple experiments or trials were conducted to determine when it was acceptable to advance to secondary drying. You'll see earlier in the process when that Pirani gauge just starts to decrease, you're maybe halfway through, and there's a pretty substantial rise in pressure meaning there was a lot of water vapor present.

And that type of rise will differ between different products.

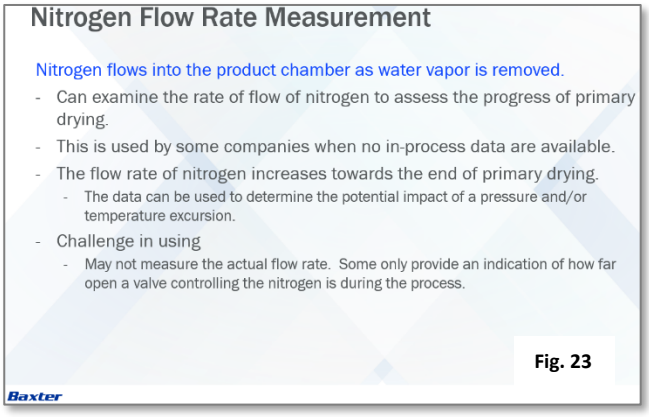

Fig. 23 Regarding Nitrogen Flow Rate Management: During the process, especially during primary drying, there is a valve that opens to allow nitrogen flow into the chamber. And that must occur to maintain our pressure in the chamber. We can examine that data overall. Let's say if there was an excursion during the process. We can go back and look at the flow data of nitrogen to determine at which point during the process did it occur. If it was earlier during primary drying, it may be more of a risk to our product. The challenge to

using this method is that nitrogen flow is not always monitored by a flow valve or flow meter. Instead, there is often just a valve that just opens and closes, and the extent to which it opens is what is measured. It is not always that reliable.

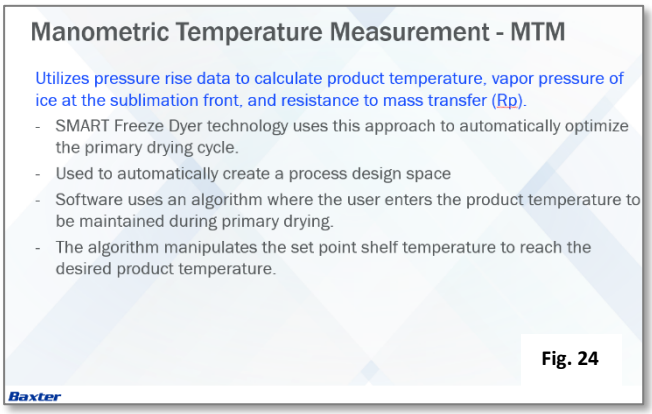

Fig. 24 Next there is something that is called MTM: Manometric Temperature Measurement. It's also based on a pressure rise. Because if we understand how the temperature of our product changes when the pressure rises, we can then do some calculations. These calculations can help automatically determine how to keep our shelf temperature and chamber pressure in line to ensure our critical product temperature, is never exceeded. It is a method of automatically creating your design space. There are some dryers that are equipped with this and it just does everything automatically for you. The challenge here is that the data at small-scale, or laboratory-scale, may not directly translate to full-scale. You must understand the capability of your freeze-dryers at full-scale as well.

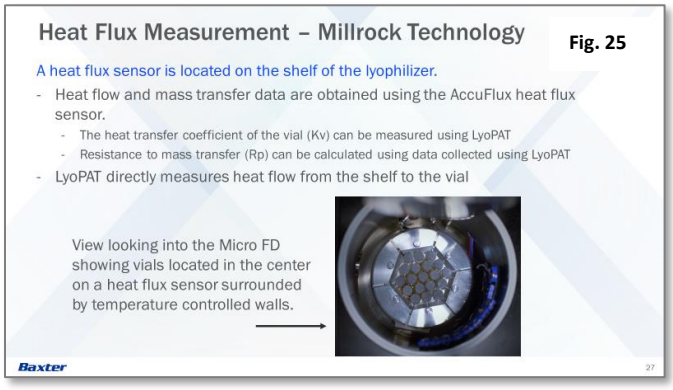

Fig. 25 Next, there's something called a heat flux measurement. Millrock has this developed and this picture is a downward view looking into a micro freeze-dryer. And what you see are these metal walls around those- those walls are blocks that control the temperature of the walls that the vials see. So now, we can adjust the temperature to match and create vials that behave just like center vials. In between those blocks, you'll see all those vials in line, and below, it looks orange. That's because right under there, there is a mat that is a heat flux sensor. That heat flux sensor can then automatically measure what is known as the heat transfer coefficient, your Kv, for your vial. And from there, we can use that data in a series of calculations to calculate the resistance to mass transfer of your formulation during drying. Between both the Kv and the Rp, we can now use those to calculate all possible combinations of chamber pressure and shelf temperature that will keep our product temperature below some critical product temperature.

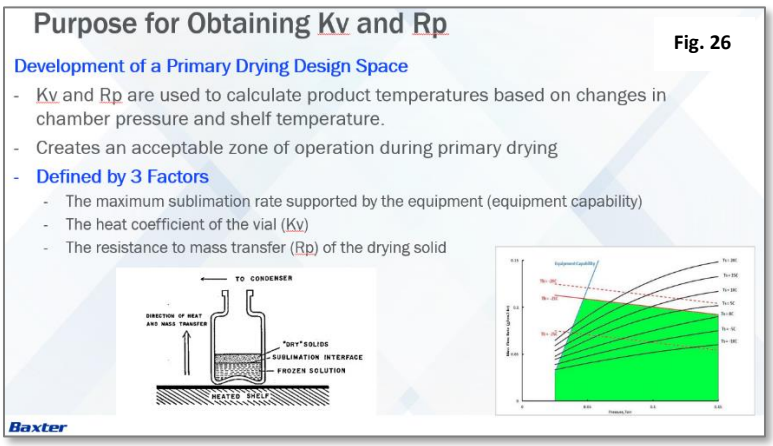

Fig. 26 We can use this to create a design space. Kv and Rp are used to calculate those product temperatures. That's defined by three factors. One, your maximum sublimation rate, your equipment capability, that's the blue line in that graph. The heat transfer coefficient of the vial, the Kv, and the Rp, the resistance to mass transfer of water vapor of your drying solid. Looking at the graph, the black lines are shelf temperature isotherms, the solid red line is the critical product temperature, and anything within the green area is the safe zone of operation. The triangle area is right where the equipment capability and product temperature cross; those are the conditions that lead to the most efficient primary drying cycle.

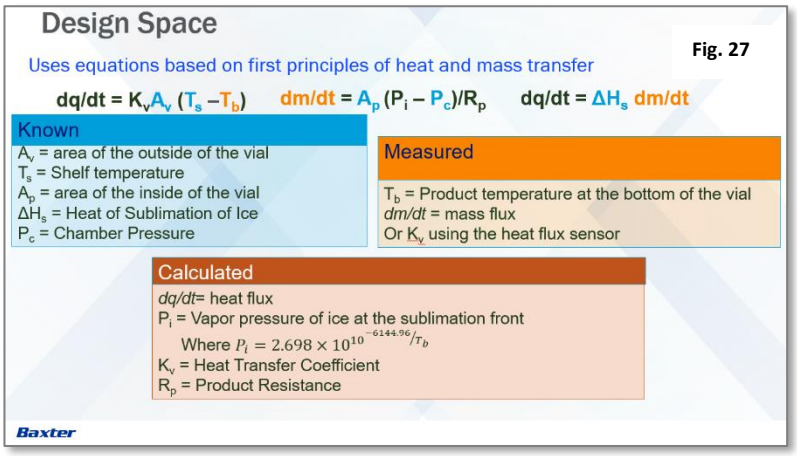

Fig. 27 Here are the calculations that go into this. There are some things that are known, those are in the blue area. We know the area of the outside of the vial, our shelf temperature, our chamber pressure, the area inside of the vial, and the heat of sublimation of ice. Then we can measure the heat flux, as in the flux sensor, we can measure Kv, we can measure the product temperature at the bottom of the vial, and it's also possible to measure mass flux. From there, we use that data to calculate the remaining values.

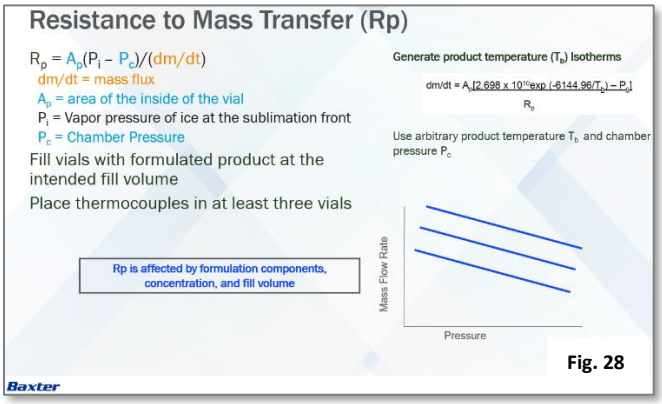

Fig. 28 This is the resistance to mass transfer. If we have mass flux or if we have Kv, we can then alter the equations to obtain the missing data. Here, we can obtain the resistance to mass transfer by filling the vials we intend to use with the intended fill volume of the product, we place thermocouples in the vial, and then we conduct our cycle. We can either measure the heat flux from the vials or the mass flow of water vapor.

From there, we can then get Kv. With Kv we can either measure using heat flux sensor, or we can Measure using vials filled with water and then carry out the cycle by adjusting pressure at a specific shelf temperature to collect the data.

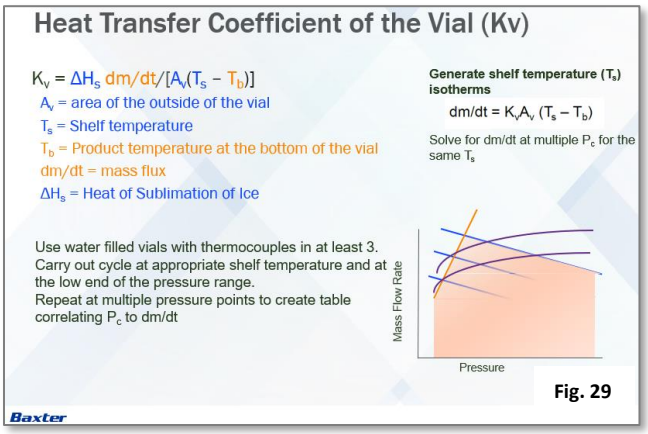

Fig. 29 From there, we can then determine all the possible combinations of shelf temperature and chamber pressure that will maintain our product temperature below that critical product temperature.

Fig. 30 Finally, another method we can use is called tunable diode laser absorption spectroscopy, TDLAS. Here, we place a laser and a detector and a spool piece that's located between the product chamber and the condenser. Fig. 31

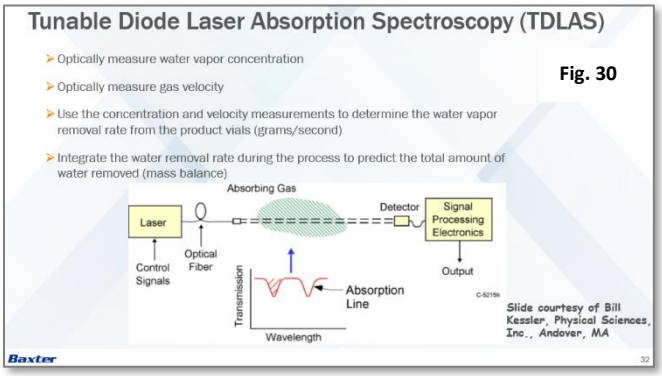

This can automatically measure the flow rate and concentration of water vapor going through that spool piece. The challenge here is that it only works with lyophilizers equipped with the spool piece. And there are many different lyophilizers that do not operate in this manner. Some just have a wall with an opening between the product chamber and the condenser. Others have a platform that raises and falls between that wall.

Fig. 32 Finally, we can also use a Residual Gas Analyzer. It is a type of mass spectrometer that we can connect to a port on the product chamber. Here, one example of use is looking for the loss of organic solvent from a product. It can be important to learn that, so that you can adjust your cycle conditions

and maintain the best appearance. You might also look for loss of other components. There are some formulation components such as buffers that can be lost during the process.

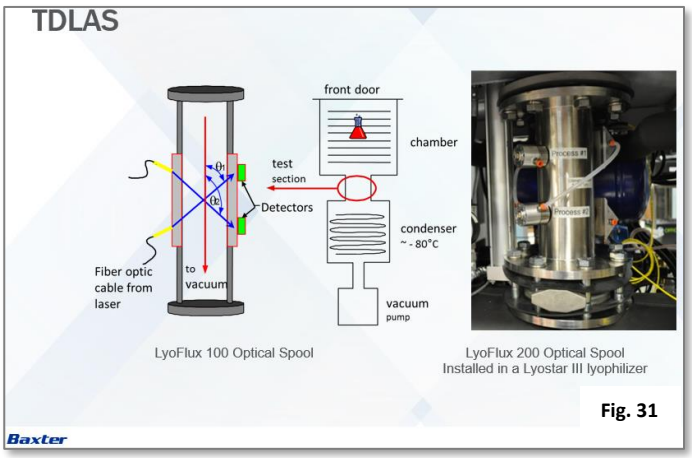

Fig. 31

Fig. 33 You can even use a mass spectrometer connected to the chamber. A couple of companies offer these, there could be more, but there's one by GEA called LyoPlus and Quantum by IMA Life. These are marketed for detecting silicone leaks during the process. The shelves are supplied with silicone fluid or heat transfer fluid through flexible metal tubing. That metal tubing must bend throughout every process and can develop leaks. If it develops a leak, you want to know when that occurs so it does not affect your product. It can also monitor for leaks in a vacuum.

If you start seeing a lot of oxygen, you can then assume that there is a leak somewhere.

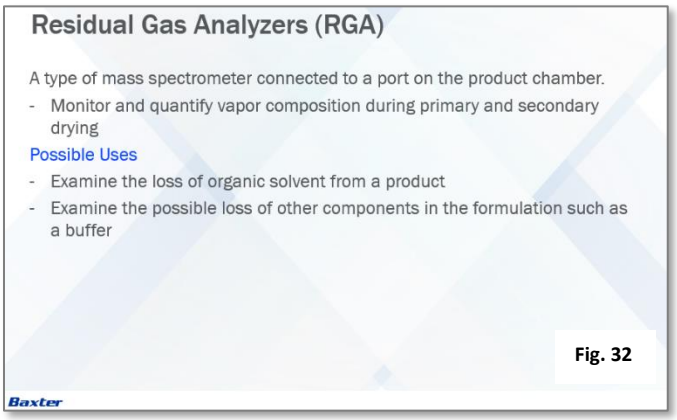

Fig. 32

In summary, we as companies can be using in-process data a lot more than we do at full-scale. That would really make the lives of our transfer group and our quality assurance group a lot easier, because now we can do some homework and calculate when the process is most at risk, and perhaps even calculate a product temperature. There are certain methods here that can tell us when primary drying is complete such as comparing the Pirani gauge measurement with a Capacitance Manometer.

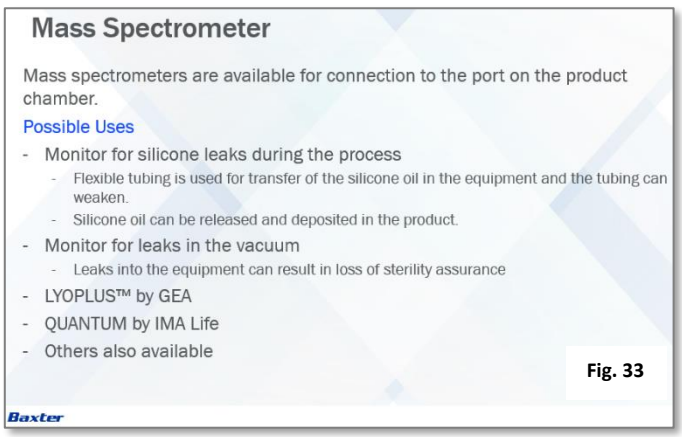

Fig. 33

Product temperature is good for measuring the completion of primary drying at laboratory-scale but can be really challenging at full-scale and may not provide us with the best data. Heat flux measurement can provide us a direct measurement of Kv, which can then be used to solve all those other equations and help us create a design space. Similarly, mass flux can be measured and solved for all those equations to create a design space.