Design and Manufacture of a Modular Continuous Unit Dose Pharmaceutical Lyophilizer

By

Steven Burcat
B.S., Duke University (2018)
S.M., Massachusetts Institute of Technology (2020)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of
DOCTORATE IN MECHANICAL ENGINEERING
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
February 2025

Authored by: Steven Burcat
Mechanical Engineering
January 10, 2025

Certified by: Alexander H. Slocum
Walter M. May and A. Hazel May Professor of Mechanical Engineering
Thesis supervisor

Accepted by: Nicolas Hadjiconstantinou
Professor of Mechanical Engineering
Graduate Officer, Department of Mechanical Engineering

Design and Manufacture of a Modular Continuous Unit Dose Pharmaceutical Lyophilizer

by

Steven Burcat

Submitted to the Department of Mechanical Engineering
on January 10, 2025 in partial fulfillment of the requirements for the degree of

DOCTORATE IN MECHANICAL ENGINEERING

ABSTRACT

Pharmaceutical lyophilization (freeze-drying) enables long term storage and simplified transportation for aqueous vaccines and protein formulations. Modern industrial pharmaceutical freeze-driers rely on large batch and open loop formulation processing, limiting supply chains and resulting in variable quality products. This work describes the design and manufacture of a modular continuous lyophilization machine for pharmaceutical production. Additionally, the scaling and design methodology outlined in this work enables the development of both smaller systems for laboratory testing and larger machines to fit the needs and requirements of individual facilities.

This machine introduces three new technologies to the pharmaceutical freeze-drying process. The first innovation is a continuous flow lyophilization topology which separates the lyophilization steps spatially rather than temporally. This layout allows product to travel through the system in smaller batches for increased product uniformity and quality control. The second innovation is a weight-based sensor for monitoring residual water content. This sensor enables in-situ monitoring of product during sublimation, and it resolves mass measurements as small as 5mg. The third innovation is the implementation of a thermal shock method of inducing controlled nucleation. The convective cooling and spatial non-uniformity within the machine allow vials to experience a 40°C temperature drop in less than 30 seconds. This nucleation front starts on the vial walls, rather than at the top surface of the solution in the vial, potentially increasing the water sublimation rate during drying compared to current nucleation methods.

The machine designed and built for this work integrates into modern factory processes and can be scaled from the lab bench to a production line. The manufactured prototype demonstrates improvements on the production rate, flexibility, and quality of existing machines.

Thesis supervisor: Alexander H. Slocum
Title: Walter M. May and A. Hazel May Professor of Mechanical Engineering

Acknowledgments

This thesis was made possible with support from the FDA.

Thank you Prof. Trout and Prof. Braatz for your insights into lyophilization and your guidance on this work.

Thank you Alex for your incredible mentorship and unwavering support throughout the ups and downs of this project. I have learned so much about both machine design and life from working with you over these years.

Thank you Rohan for working on this project with me from beginning to end. I could not have asked for a better person with whom to have spent late nights debugging vacuum systems and debating what is a hardware versus software problem.

Thank you Lorenzo, Ellen, and Ryan for your time and efforts on this project. I enjoyed working with each of you, and I hope you're happy to see that the full machine has been built and run successfully.

Thank you Josh for your manufacturing insights and for helping us take these designs and build the full machine.

Thank you to all the members of PERG throughout the years. You all welcomed me in when I first came to MIT, and even across two lab generations I always feel re-energized when I come to lab and see you all.

Thank you to all the student groups I have spent time with at MIT. The graduate school experience has always been more than the classroom and lab, and these opportunities and time have been essential to my growth and enjoyment of graduate school.

Thank you to friends for all the adventures, celebrations, commiserations, and shared experiences throughout our time at MIT and beyond. I could not have gotten through grad school without you all.

Thank you to my family for supporting me and pushing me since day 0. I have only been able to get to this point because you all have helped me get here.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

After a pharmaceutical product is synthesized, there is a complex supply chain required to deliver the product to the patients who need it. The product must be kept stable and viable as it moves through this supply chain such that it can provide the necessary medical benefits to patients. Unfortunately, many products do not have long term stability in their synthesized liquid state. These unstable products can be lyophilized to transition them into a stable dry state [1][2]. Many pharmaceutical solutions are not stable at room temperature in their liquid state. There are two primary methods for stabilizing these products. One, the solution can be frozen. When water in an aqueous solution solidifies as ice, the degradation rate of the product decreases [1]. However, this method requires keeping the product at relatively low temperatures, sometimes as low as -90°C, until the product is about to be used. This cooling can be expensive and unwieldy, especially during transportation, and any failure in the cooling system can result in the loss of significant amounts of product. Two example freezers are shown in Figure 1.1.

(a)

(b)

Figure 1.1: Example freezers used for pharmaceuticals. (a) a travel freezer that can be used to transport products [3], and (b) a stand freezer used for on-site storage [4].

The other primary method for stabilizing the product involves removing the liquid entirely from the solution. Without liquid, the product can form a solid cake. Once the liquid is removed, the product can be kept stable at room temperature without requiring significant temperature control. However, the liquid must be almost completely removed to gain this stability [5]. Also, the liquid removal process must not damage the product, which limits the available methods for liquid removal. Additionally, the product must be reconstituted in a liquid solution prior to injection into a patient. This reconstitution process typically only involves adding the correct volume of water or other solvent to the product prior to injection. However, the liquid removal method can influence the effectiveness of cake reconstitution [6]. An example reconstitution process is shown in Figure 1.2. These reconstitution processes can have a high degree of complexity. However, companies have been developing syringe technologies that can perform reconstitution and injection in a single smooth process [7].

While liquid removal has more complexities and initial energy requirements around its implementation than freezing, it can result in significant energy and logistical savings regarding the storage and transportation of pharmaceutical products in the long term. Thus, removing liquid from pharmaceutical solutions is a process that adds significant value to the pharmaceutical production process [6]. This work will primarily focus on the removal of water from aqueous solutions, though the principles apply to liquid formulations which do not use water.

Multiple techniques exist for removing water from solutions. The water can be evaporated or sublimated out of the solution, where the liquid water is converted into vapor leaving behind the solute. Water can be evaporated through inducing boiling by putting the solution in an environment with a low vapor pressure or by heating the solution. Heating the solution to induce boiling also significantly heats the pharmaceutical product, which can denature

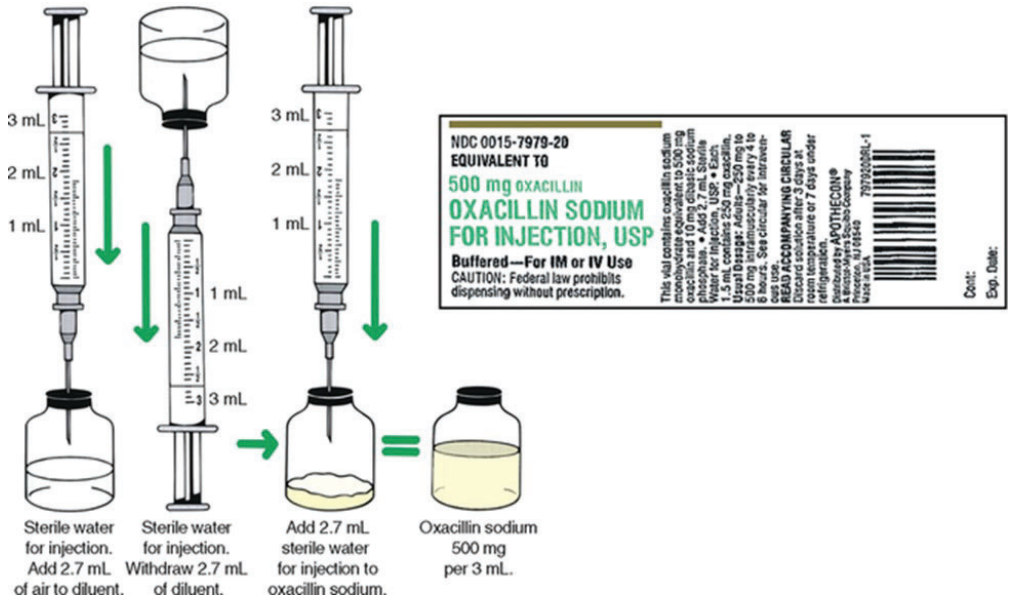

Figure 1.2: An example reconstitution process example, where sterile water is added to a dried formulation for use. The reconstitution process typically involves a diluent injection and mixing with the lyophilized product. Image adapted from [8].

proteins, so boiling by heating is not an acceptable water removal method. If the solution is placed in an environment with a sufficiently low vapor pressure, the water can undergo vacuum induced boiling. This boiling process does not require excessive heating of the pharmaceutical product, but the boiling is a violent process which can also damage the protein chains. Thus, vacuum boiling is also undesirable. The water could evaporate if exposed to extremely dry air, but it can be difficult to maintain a sufficiently dry air flow for the full drying process. Due to these concerns, evaporation is not a widely pursued method for removing water from pharmaceutical solutions.

Sublimation removes water from a solution after the solution is frozen by changing the solid ice directly into water vapor, which leaves behind the solute in a solid cake. This process removes water from the solution that is already in a stable frozen state, reducing the risk that the product destabilizes during the water removal process. Sublimation also does not physically disturb the product in the same way that boiling does. Furthermore, the sublimation process can remove bound water from the product, resulting in a shelf stable product at room temperature. Additionally, the concentration of the product and the cake structure formed by the product around the ice crystals during freezing result in a cake with a high specific surface area, which facilitates rapid and complete dissolution of the solid during reconstitution [9]. The complete process of freezing a solution to prepare for sublimation and removing the water through sublimation is known as lyophilization. While other techniques such as sterile crystallization, spray-drying, microglassification, and powder filling can be used to produce dry product, lyophilization is currently the most common method for drying pharmaceutical solutions [2][10].

The lyophilization process consists of two main processes: freezing and drying. During the freezing process, the product temperature is reduced well below water's freezing point of 0°C to prepare for sublimation [11]. The frozen product is then moved into a vacuum environment for drying, where the low ambient vapor pressure causes the ice to sublimate into water vapor. The water vapor is removed through the vacuum system, leaving behind the solid product mass. This process is shown on a water triple point diagram in Figure 1.3. Typical conditions used to follow this process are shown in Figure 1.4.

Lyophilization has been implemented in the pharmaceutical production process for nearly 100 years [13][11]. Once the pharmaceutical product is synthesized, it is loaded into a lyophilizer. The lyophilization process occurs, leaving behind vials with solid product cakes. These vials are then sealed and delivered to their final destinations, where they are stored until needed. Typical pharmaceutical production lyophilizers operate on a unit vial scale rather than in a bulk product system. A bulk product system would take large amounts of solution in, remove all of the water, and then produce a dried mass of product, which could be subsequently divided into appropriate dosages. The unit vial system keeps pre-metered dosages in vials, sublimating the product in the vial. This process maintains the precision of liquid dosing and the form factor of the vial for use and storage in external facilities. It also helps maintain the sterility of the product during lyophilization and subsequent reconstitution, as the water can be injected into the vial directly through the rubber stopper. Examples of lyophilized cakes in vials are shown in Figure 1.5.

Existing pharmaceutical lyophilization systems suffer from three major limitations. The first limitation is the difficulty in translating processes between research-scale systems and industrial-scale systems [14]. Because research systems and industrial systems have signifi-

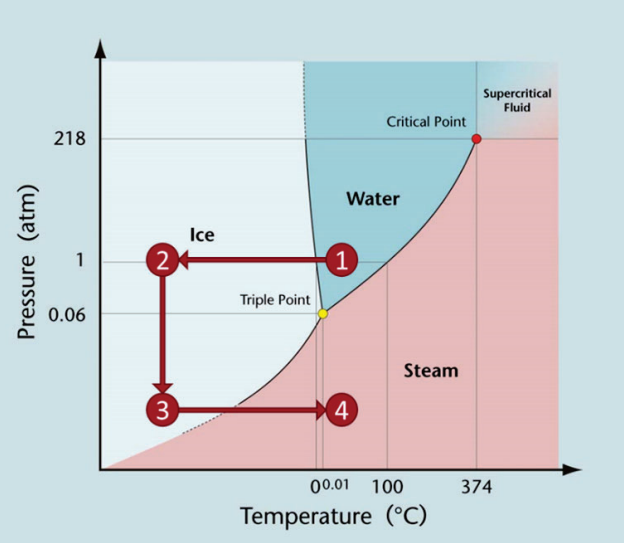

Figure 1.3: During lyophilization, the solution in the vials starts at position 1, with the liquid solution at atmospheric temperature and pressure. The solution is then frozen at atmospheric pressure, moving to position 2. The vial environment is then pumped down to a vacuum, moving to position 3. Energy is provided to the solution through heat, causing the ice to sublimate as the solution moves to position 4. Once the water has finished sublimating, the vial is capped and released back to atmospheric pressure.

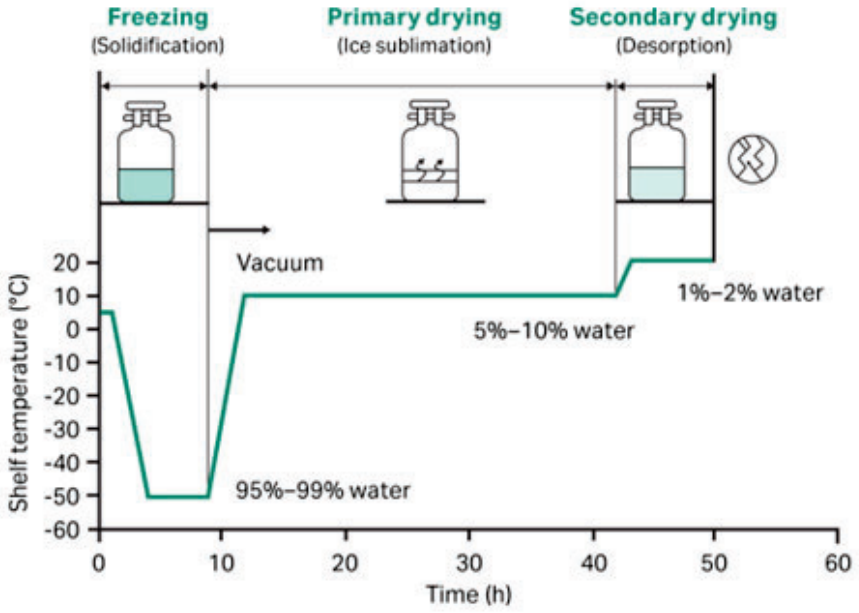

Figure 1.4: A typical lyophilization process on a conventional lyophilizer moves the solution through the freezing and drying stages to remove water from the sample. Image sourced from [12]. The percentages listed are weight percentages for the solution at any given time in the process, not percentages of the initial solution.

Figure 1.5: The representative lyophilized cakes shown are examples of what the product looks like after undergoing the lyophilization process.

cantly differing geometries, the processes developed on research machines do not directly correspond to lyophilization processes that work on industrial equipment. Example of research- and commercial-scale lyophilizers are shown in Figure 1.6 to demonstrate this difference in geometry. Thus, when developing a lyophilization process for new formulations, the process must be developed on both the research equipment and then again on the industrial equipment. This delay can be costly, as it requires the industrial equipment to be taken offline while testing lyophilization parameters for new formulations, and the long cycle time of lyophilization can cause this process to take days [14].

The second limitation is the non-uniformity in vial conditions within an industrial pharmaceutical lyophilizer. The internal volume of these lyophilizers is large enough to create significant variation in the thermal and vacuum conditions between vials [17]. This variation

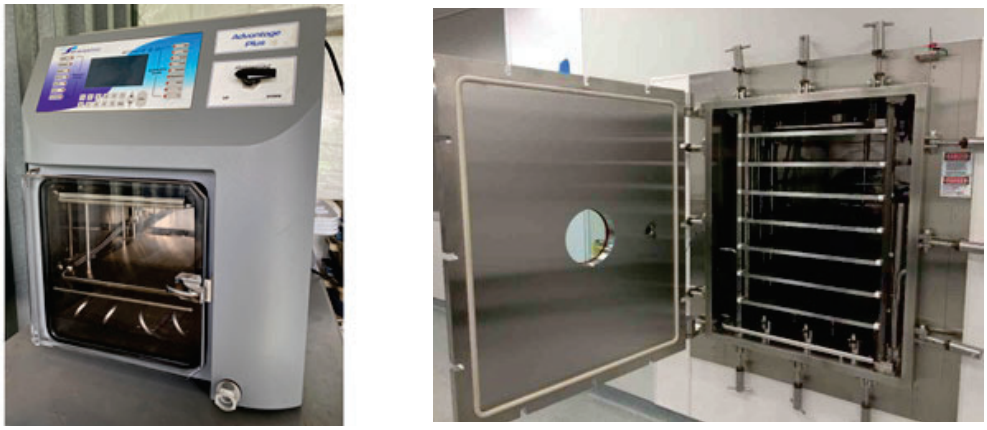

(a) (b)

Figure 1.6: Research lyophilizers (left) are much smaller than commercial lyophilizers (right), and their differing geometries, including the number of shelves and internal volume, contribute to variable performance given the same process parameters. Left image from [15], right image from [16].

in conditions leads to differences in lyophilization time and quality. An example method for delineating different regions within these large internal volumes is shown in Figure 1.7.

This variation can be sufficiently large to cause some vials to not successfully lyophilize. The variation is tied to the distance of the vials from the vacuum system inlet and the proximity of the vials to the walls of the lyophilizer. The farther the vials are from the vacuum system inlet, the higher the local pressure at the vial, which can slow down the sublimation rate. The closer the vials are to the machine walls, the more they are heated by the radiative heat emanating from those walls, increasing their temperature relative to other vials. These higher pressures and temperatures make the lyophilization process more prone to failure. Accordingly, machines can be run with empty vials around the outermost rows of shelves because any product in those vials would not successfully lyophilize [18]. These vials represent lost productivity for this equipment, reducing their effective capacity.

The third limitation is the batched nature of industrial lyophilizers. Pharmaceutical production lines have been evolving into continuous systems for improved production rates, flexibility, and quality [19]. However, the lyophilizers at the end of the process have remained batch systems. To match overall production rates, these machines have large batch sizes, as seen in Figure 1.8. This disconnect leads to a requirement for large buffers between the production line and the lyophilizers, which takes up space and increases the risk of vial contamination as they are waiting to be moved between factory subsystems. There have been preliminary investigations into continuous systems for pharmaceutical lyophilization, but fully functional systems are currently limited [19].

Addressing the third limitation requires a significant topological change to existing lyophilization systems. Continuous lyophilization exists in other industries, such as food production. These lyophilizers often operate on a bulk mass principle, such as the machine shown in Figure 1.9, where material is moved in large quantities through the system. This method does not have as precise control of unit quantities as a unit vial system. In the food industry, for example, this loss of precision is acceptable; however, in the pharmaceutical industry, this potential variation carries higher risks of failed products. Thus, bulk-mass systems are not desirable for a pharmaceutical lyophilizer.

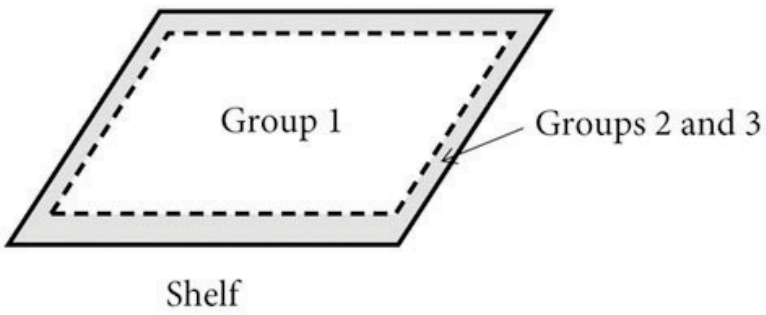

Figure 1.7: The vials placed in the central part of the shelf (Group 1) and the vials placed in the first and second rows along the perimeter of the shelf (Groups 2 and 3) can experience significantly varying sublimation rates due to the relative proximity of the vials to the walls. Figure adapted from [17].

Figure 1.8: A commercial freeze dryer loading one shelf in a freeze drying batch run. These commercial systems include thousands of vials per batch. Image from [20].

Figure 1.9: Bulk lyophilization machines like the one shown here are often used in food freeze drying, where specific mass requirements are not as strict as they can be in pharmaceuticals. Image from [21].

The current work on continuous pharmaceutical lyophilizers is limited. The groundwork analysis for the value and fit of a pharmaceutical lyophilizer has been presented by [19]. That work presents a theoretical concept for a continuous lyophilization system which is realized with this work. A company has created a continuous lyophilization system, but it has not yet demonstrated significant traction within the industry [22]. A concept diagram for this potential lyophilization system is shown in Figure 1.10.

This thesis presents the design, manufacture, and evaluation of a continuous lyophilization system for unit-dose vial production. This system addresses the three existing batch lyophilizer limitations presented above. The translatability between the research system and the industrial system is addressed by using a modular geometry which is maintained between both a research and an industrial scale system, ensuring that processes developed on the research system scale directly to the industrial system. The vial non-uniformity is addressed by using a smaller internal geometry and continuously moving trays, ensuring that each vial is exposed to similar conditions throughout the lyophilization process. The continuous nature of the system addresses existing batch process limitations, enabling more effective interfacing with the continuous nature of modern pharmaceutical production equipment. The system built for this work is shown in Figure 1.11.

Chapter 2 describes the overall system architecture and sizing methodology used to design a continuous lyophilizer using this topology. This chapter starts with initial system sizing, which can be done based on system requirements without specifying the individual components and methods used to perform the lyophilization process. This chapter also highlights the modular nature of the machine and how these modules are assembled together into a unified system, along with the design for the structure on which these modules are placed.

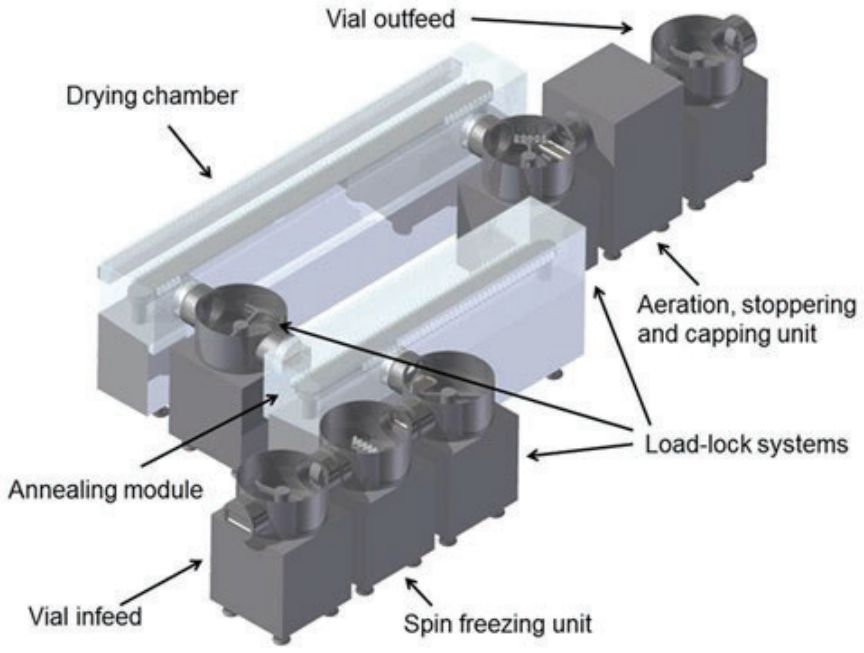

Figure 1.10: A concept for a commercial continuous freeze drier layout developed by Rheavita [22].

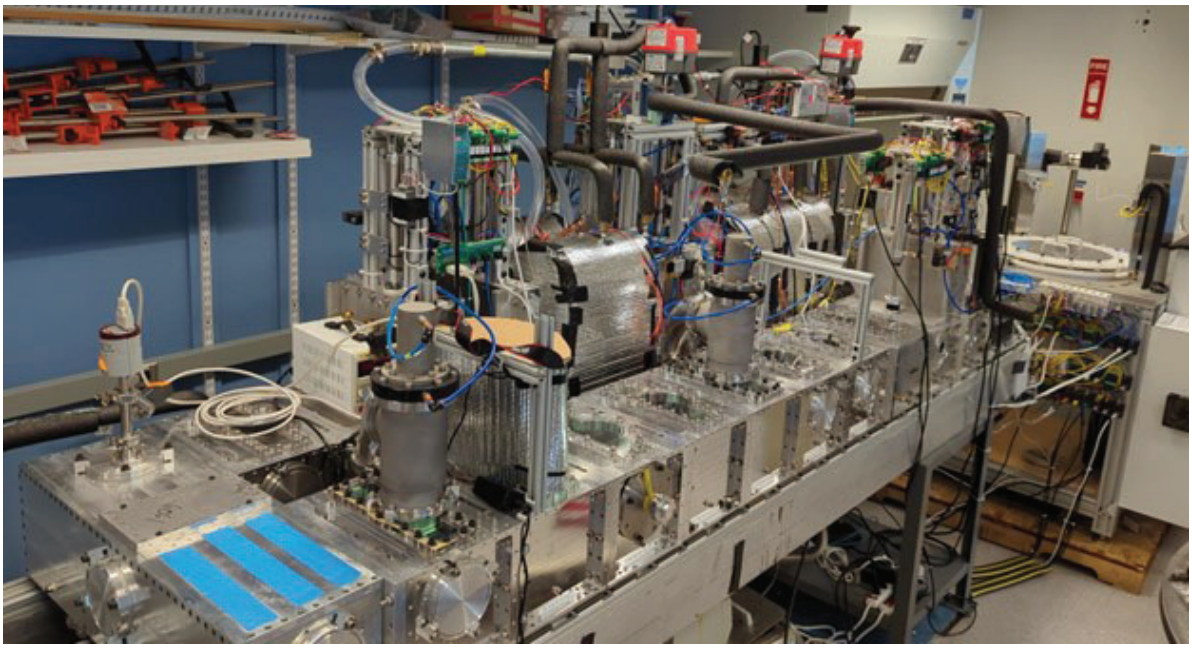

Figure 1.11: An angled view of the continuous lyophilizer manufactured in this work.

The specific motion system selected for this work is described as well.

Chapter 3 describes the specific design of the forced convection freezing system used in this work. This chapter includes the component sizing and a detailed description of how to design and build a similar chamber. The experimental results from using this chamber shows that it performs as well or better than commercial lyophilization shelf-based freezing on metrics of process time and vial condition uniformity. This chapter also details the novel implementation of a thermal quench method for nucleation to provide enhanced control over the vial freezing process.

Chapter 4 describes the design and testing of a novel in-situ unit vial weight sensing system used to monitor the lyophilization process and provide live insight into the end of drying. This weight sensor is based on a deflecting helical spring which causes an optical fiducial to move as the vial mass changes. The change in this optical fiducial position is monitored with a camera outside the vacuum chamber. The weight sensor is demonstrated to achieve mass measurement precision to 10mg on a total initial mass system of 15 grams, including the vial, cap, and solution. While this resolution is sufficient to identify the end of primary drying, it requires further refinement to identify the end of secondary drying.

Chapter 5 describes the design of the machine elements used to create the vacuum system for the drying process. This chapter focuses on the modularity of the chambers to adapt to vacuum, drying, and sensing modalities. Additionally, the capping system used to seal the vials at the end of the process is described in this chapter.

Chapter 6 describes the process of building and assembling the continuous lyophilizer described in this work. This chapter specifies the process for creating individual modules and assembling them together into a larger system. The construction of the machine base structure is also detailed in this chapter, along with the system integration with a commercial vial filling machine.

Chapter 2

System Architecture

A critical aspect of this work is the robust design, manufacture, and testing of a continuous lyophilizer that can be realized at multiple scales with consistent and predictable performance. The lyophilizer is based on a modular design achieved by a deterministic methodology which translates the system requirements to the machine layout and design details. This deterministic process establishes the overall machine architecture without overly constraining the design of individual machine elements. Following this process structure enables a top-down machine design, ensuring that the final machine meets the system requirements while retaining flexibility in module and machine element design. The latter is important to allow for the technological development of elements such as sensors and actuators.

The overall machine sizing is based on the relative time required for each part of the lyophilization process, the desired production rate, and the vials used. This sizing uses a rectangular loop to ensure that the system entrance and exit are connected, enabling continuous motion through the entire geometry. This motion is driven by a magnetic levitation system to reduce the potential contamination that would result from frictional wear on sliding interfaces. The machine itself is separated into individual modules that can be developed for each process function independently while ensuring that they can be assembled together into the larger system. The modules themselves are divided into four categories: loading, load lock, freezing, and drying. Each module consists of three main elements: the machine chamber, the stator, and the base connecting the first two elements. The stator and the base are the same for each module, while the machine chambers include features designated for their specific category and features shared among all modules. The category-specific features are described in the relevant chapters, while the shared features are described in this chapter.

The shared machine chamber features are based on creating a uniform internal geometry and sealing modules both internally and to their neighbors. Each machine chamber is sealed using an injectable sealant to a stator which drives motion through the continuous lyophilizer. Additionally, the machine chambers are sealed to each other using O-rings on consecutive faces. These features are shared by all modules to ensure consistent performance and to enable the machine to be reconfigurable. The modules are assembled on a custom machine table consisting of parallel rails mounted on top of two support points. The use of two support points decouples the table flatness from the surface profile of the floor upon which it is built. The parallel rails include linear sliding elements to make it easier to assemble the system and change module positions.

2.1 Machine Sizing

In a manufacturing line, the system's production rate is related to the number of units within the system and the time spent in the system, as shown by Little's Law in Equation 2.1.

$$L = \gamma W \tag{2.1}$$

Where L is the total number of units in the system, γ is the average rate of units entering the system, and W is the amount of time spent in the system. The time spent in the system corresponds to the process time plus the time required to move between process stations. In a lyophilizer, this variable is relatively fixed, as the time spent in the system is dominated by the process time required to freeze dry samples. Thus, W can initially be treated as a constant for the relationship between L and γ.

In an ideal fully continuous system, the rate of units entering the system remains constant and is equivalent to the average speed of individual elements moving through the system. Given that the samples entering the system must leave the system at the same rate to prevent accumulation, the overall machine production rate is given by γ. Given a desired production rate and total process time, the number of units needed in the system at any given time can be calculated from Little's Law as shown above in Equation 2.1. This calculation can yield the first insight into the overall physical size and layout of the machine, helping ground further analysis and design. So far, no limitations have been placed on what the system will look like or how components will move through the system, yet there are now some basic physical constraints that start to build the framework of how the machine will come together. An example of this initial size estimation appears in Table 2.1. Since this size estimation only accounts for the area covered by the vials themselves, it serves only as a minimum size for a planar system.

2.2 Machine Shape

The units moving through a machine can either enter in one location and exit from a different location, or they can enter and exit from the same location. If the units move through the

Parameter	Name	Value	Value	Value	Units	Formula
Vial size	D_{vial}	20	40	20	mm	
Production Rate	$Prod$	100	100	1000	vial/hr	
Process Time	$t_{process}$	10	10	10	hr	
Total Vials	tot_{vial}	1000	1000	10000	vial	$:= Prod * t_{process}$
Vials total area	A_{vial}	0.35	1.39	3.46	m^2	$:= tot_{vial} \frac{\sqrt{3}}{2} D_{vial}^2$

Table 2.1: Initial estimates for the machine size can be determined based only on the vials used, desired production rate, and overall process time by using Little's Law. The size estimate only accounts for the space taken up by the vials themselves, assuming a hexagonally close-packed assembly. The actual machine is expected to require more space to account for the components that create the required conditions for lyophilization.

system but the system itself is static, like a system with slides or roller conveyors, the entrance and exit can be separated. If the units are carried through the machine by an actuated motion system, then at least the motion system must return to its initial state to receive new units. This return motion can occur without units, such as the bottom side of a conveyor belt, or with units, such as with a circular turntable. In the machine described in this work, the entrance and exit are linked because the pucks that transport units through the system, described in Section 2.4.1, need to return to the machine entrance to receive new units. Though this motion could occur without units, that would be a missed opportunity to use space for process chambers. Accordingly, the system path follows a closed loop, where one section is used for loading and unloading units while the rest is used for lyophilization processes.

Understanding the scaling of system area requirements with respect to desired process parameters helps inform the decision on the overall shape of this closed loop. In general, there are three main ways that systems can move. They can follow rotational paths, they can move linearly, or they can incorporate some hybrid combination of both. As a first pass in comparing a linear system to a rotary system, one can imagine the products moving along the system perimeter. Moving along the perimeter of a simple closed shape involves less direction changes than a snaking path, making it simpler to implement. These calculations can be done in a table similar to Table 2.1, giving the results seen in Table 2.2. The rectangular system covers significantly less area, particularly at higher production rates. Given the importance of floor area in production settings, this trade-off leans heavily in favor of pursuing a rectangular system profile over a circular one.

The analysis in Table 2.2 uses only the perimeter of the system shapes for vial motion, assuming that the vials do not move in a snaking path. Redirecting the product path through this internal area could reduce the relative difference between the spatial efficiency of the two system geometries. However, in a lyophilizer the products must be exposed to significantly

Parameter	Name	Value	Value	Value	Units	Formula
Vial size	D_{vial}	20	40	20	mm	
Production Rate	$Prod$	100	100	1000	vials/hr	
Process Time	$t_{process}$	10	10	10	hr	
Total Vials	tot_{vial}	1000	1000	10000	vials	$:= Prod * t_{process}$
Vial rows	n_{rows}	4	4	4		
Perimeter length	L_p	5	10	50	m	$:= \frac{tot_{vial}}{n_{rows}} D_{vial}$
Circular area	A_c	2	8	199	m^2	$:= \frac{\pi}{4}\left(\frac{L_p}{\pi}\right)^2$
Rectangular area (width 1 m)	A_r	1.5	4	24	m^2	$:= \frac{L_p - 2}{2}$
Area ratio		1.3	2.0	8.3		$:= \frac{A_c}{A_r}$

Table 2.2: Different system layouts can be compared by evaluating the effective area covered by each shape if the vials are arranged along the shape's perimeter. The perimeter is used as an initial estimate for simplicity of path implementation. A rectangular width of 1 m is chosen to normalize the rectangular area calculations.

different temperature and pressure conditions. For a continuous machine, these states must physically exist simultaneously, necessitating that they be sufficiently separated and isolated to not adversely affect vials in a different process stage. These isolating elements, such as chamber walls, prevent maximum density area packing of vials within the entire machine area. Additionally, fully packing the process rows within the system area would limit access for maintenance, assembly, and monitoring. Furthermore, using a large volume for the process chambers leads to non-uniformity within the volume, as described in Chapter 1. This non-uniformity is reduced by decreasing the surface area to volume ratio of the process chamber, such as by using a long tunnel rather than a large cubic chamber. Thus, using the shape perimeter is a reasonable estimate for the overall machine size during this analysis.

A further consideration for improving machine size efficiency relative to production rate involves parallelizing the process elements within the machine. For example, multiple rows of vials can pass through widened chambers, rather than accommodating only a single vial at a time. In a rectangular system, this scaling can be done easily and naturally, as a widened row of units traveling along the perimeter all move at the same speed. However, in a circular system the increased width would result in neighboring rows moving at different speeds or being inefficiently packed. Maintaining a constant angular speed would result in different distances traveled, so the process chambers would need to be shaped differently for each radial row to maintain optimal vial packing within the system. This added complexity does not apply to the rectangular system, making it a simpler geometry to implement. Thus, a rectangular geometry is used for this system.

2.3 Travel Speed

While the overall production rate equals the average speed of the units moving through the system, this speed may not always be the same for all vials in the system. The vial speeds are bound by two extreme conditions. In the slowest case, all vials move simultaneously at the same speed. In the fastest case, only one vial moves at a time, and once it stops the previous vial moves into the newly vacated space. The system can use some combination of these states, where subsets of the vials move together at intermediate speeds, such as in parallel rows or connected on a tray. These bounding speeds, along with example row and tray intermediate options, appear in Table 2.3.

These values show that some states, such as moving only one vial at a time, lead to unreasonably high speed requirements. While having all vials move in this way is not practical, vials moving between process chambers may need to approach these speeds if entering and leaving load-locks. Thus, it is important to recognize the system bounds early to understand where focus may be needed to alleviate extreme values. In these situations, moving multiple vials at a time into and out of load-locks can reduce the speed requirement during that process, either by dividing the vials in the system into multiple rows or grouping them onto trays.

Parameter	Name	Value	Value	Value	Units	Formula
Vial size	D_{vial}	20	40	20	mm	
Production Rate	$Prod$	100	100	1000	vials/hr	
Process Time	$t_{process}$	10	10	10	hr	
Total Vials	tot_{vial}	1000	1000	10000	vials	$:= Prod * t_{process}$
Vials continuous motion speed	$v_{vial_{min}}$	0.56	1.11	5.56	mm/s	$:= D_{vial} * Prod$
Vials single motion speed	$v_{vial_{max}}$	556	1111	55556	mm/s	$:= v_{vial_{min}} * tot_{vial}$
Vial rows	n_{rows}	4	4	4		
Vial row max speed	$v_{row_{max}}$	139	278	13889	mm/s	$:= v_{vial_{max}}/n_{rows}$
Vial tray size	$tray_{vials}$	10	10	10	vials	
Tray single motion speed	$v_{tray_{max}}$	56	111	5556	mm/s	$:= \frac{v_{vial_{max}}}{tray_{vials}}$

Table 2.3: Initial estimates of vial motion speeds inside the system can be determined based primarily on the vials used, desired production rate, and overall process time. The minimum speed assumes all vials moving together continuously, while the maximum speed assumes only one vial per row or one tray moves at a time. Using rows or trays which move vials simultaneously can alleviate the maximum speed bounds.

2.4 Motion System

Given the above analysis, a rectangular machine shape is chosen for this system. Once the layout and motion speed bounds have been determined, the next level of machine detail can be approached. The focus should remain at the level of shared elements between the entire system, such that potential integration risks are not overlooked. The major elements that interact with the whole system that have not been discussed include the system structure and the motion system. The system structure can vary depending on the motion system selected, so it is more practical to focus on how elements move through the system before proceeding further in the system design.

The system which moves units through the machine is a critical element to a continuous manufacturing system. This element is one of the key differentiators between continuous systems and batch systems. In batch processing systems, products can remain stationary within a process element until moved by an external operation, removing any need for internal moving parts. The continuous system's requirement for moving parts necessarily adds complexity that must be addressed and incorporated into the design early on. The motion system requirements primarily center on motion speed, load capacity, positional accuracy, and cleanliness.

The speed requirement can be estimated from the basic system size and throughput estimates above, as seen in Table 2.3. The more vials moving together, the lower the maximum speed required. The load capacity comes from how many vials, along with any other corresponding hardware moving with the vials, are moving through a given part of the system

at a time. This value can be estimated for a given moving system based on the moving element's surface area and the number of vials that would cover that surface area.

The product positioning requirements are connected to the precision required for the other elements in the process. A well-selected/designed motion system will be at least as accurate as the most precise operation in the system. The positioning system precision requirements can be determined during the error budgeting process of each subsystem. These error budgets show how different positional errors affect downstream elements like control and sensing, such that a designer can adjust the precision requirement as the system design evolves. If selecting the motion system before designing the subsystems, as was done for this project, it is best to choose a system likely to have a higher positional precision than expected to be necessary. While developing a new machine, this choice may result in higher costs, but it reduces the risk of errors created by the positioning system that either reduce performance or require replacement of the motion system during development.

A pharmaceutical freeze-drier processes medical-grade products, so it must adhere to a stringent cleanliness standard. This standard is met through periodic cleaning and minimizing potential contamination sources. Pharmaceutical equipment is typically cleaned through a steam sterilization process, such as using an autoclave. Accordingly, all components of the system, including the motion system, must be steam sterilizable. In addition to cleaning, the machine must not introduce undue contaminants into the process that could make the products unsafe for medical use [23]. There are two main potential ways these contaminants could be introduced into the system: 1) the contaminants could enter from the outside environment, or 2) the contaminants could be produced from equipment operation. The first source is addressed primarily through sealing the system, as discussed in Section 2.7.2.1, and it can be reinforced using elements such as a purge chamber at the machine entrance. This contaminant source does not significantly affect the motion system design/selection. The second source can be related to a few different mechanisms: 1) friction between moving elements causing wear, generating particulate matter, 2) materials decomposing in the conditions created within the machine for the freeze-drying process, or 3) materials reacting within the machine, creating by-products released into the system. The second and third mechanisms are addressed by selecting materials which will not be affected by the system conditions, such as materials that minimally off-gas in a vacuum. The first mechanism creates significant concern for a motion system, as many motion systems include elements that contact and move past each other. Conveyor belts, slides, and rotary tables are all common elements in continuous manufacturing systems, but they all include frictional contact between elements which can produce particulates that could contaminate product.

There are a few approaches to motion system selection that can address the particulate generation concern raised by rubbing elements. One such countermeasure involves using materials with low-wear characteristics to reduce the wear particles generated. This method is commonly used in applications with rolling elements, as these systems have high cycle counts and excessive wear would cause premature failure. However, this measure still results in particulate generation and potential contamination, which means that the system must be carefully designed to ensure that particulate generation remains below the risk threshold provided by industrial standards [23]. Another countermeasure involves lubricating the moving elements so that a thin fluid film prevents the parts from rubbing against each other. This method introduces another material into the system, where the lubricant adds

compatibility restrictions and could also contaminate the products. A third countermeasure, often used in conjunction with these first two for clean applications, consists of using shrouds or sealed covers to enclose the moving elements separate from the sensitive volume. These covers generally need to be flexible to accommodate the system motion through elastic deformation, thus preventing them from producing additional particles. Separating the moving elements from the sensitive volume allows particulate generation to occur without contaminating products. This solution tends to be more effective over short distances, where the elastic element does not risk entering a plastic deformation regime. Additionally, cyclic motion can still cause the shroud to fret, producing particulates. A fourth countermeasure follows a different approach, using a non-contact form of motion such as electromagnetic levitation. This motion does not require contact between the moving element and a static reference surface, removing the source of particulate generation. However, these systems can be more limited in their available geometries, motion ranges, and profiles, and they are often more costly than other mechanical motion systems.

2.4.1 Planar Motors

This work implements an electromagnetic levitation system. This system eliminates the risk of particulate generation, ensuring that this system will not violate the standard for acceptable particulate generation [23]. Multiple commercial mag-lev systems which can provide non-contact motion exist; for this work, the motion system developed by Planar Motors Inc. is used [24]. This motion system is selected because it can move at speeds covering the range described in Table 2.3, its pucks have a sufficient carrying capacity to be fully covered by vials, and it is a mag-lev system that can satisfy the cleanliness requirement. Additionally, this system is tested for steam sterilization, making it capable of being cleaned to pharmaceutical requirements. This system provides six degrees of freedom motion control, enabling a wide range of motion profiles. This motion flexibility is particularly useful in sensor development, where the required motion profiles are not initially known. Furthermore, the Planar Motors system is based on modules that can be assembled into reconfigurable geometries, corresponding to the modular design planned for the full lyophilizer.

Using a commercial motion system frees up time to focus on the design of other machine elements. A downside of using a commercial motion system is that there is less flexibility in adapting the motion system to the other elements of the machine developed in this work, as the corresponding components must be designed to interface with an off-the-shelf system. Accordingly, if volumes are high enough, a custom motion system may be worth considering in the future; however, the basic design and operating concept developed using the Planar Motors system would remain.

The Planar Motors system consists of two primary elements – the powered stators (referred to by Planar Motors as "Flyways") and the passive pucks (referred to by Planar Motors as "XBots"), examples of which can be seen in Figure 2.1. The stators can be arranged in a 2D layout, and the pucks can be seamlessly passed off between consecutive stators. This feature makes a stator unit a natural choice as the basis of each machine module. Because the stators can be added and removed from the motion system, building machine modules around these stators ensures that they can be easily integrated and moved within the larger system. This flexibility allows for scaling of the machine size (and throughput), including

upgrading a machine to have a greater throughput as production need grows.

The Planar Motors stators come in a single fixed size, a 240 mm x 240 mm x 70 mm box. The pucks have a variety of available sizes, ranging from 120 mm x 120 mm x 10 mm to 450 mm x 450 mm x 16 mm. In this lyophilizer, each puck serves as a tray that carries a group of vials through the system together. The smaller the group, the more continuous the nature of the system. Thus, the smallest usable puck is the preferred choice. However, the design choices made to ensure a vacuum seal between each module, as described in Section 2.7.2.1, require that the stators have a gap between them in assembly, rather than directly abutting each other. This gap is limited to a maximum of 10 mm based on the Planar Motors system capabilities. Planar Motors' smallest, 120 mm x 120 mm pucks currently cannot cross this 10 mm gap between the stators; the smallest available puck that can cross this gap is a 120 mm x 180 mm puck. These pucks must cross the gap moving along their longest dimension, which does impose some limitations on the system motions. However, this orientation allows two pucks to fit side by side on a single stator. Thus, this 120 mm x 180 mm puck size is used in this system.

Two of the selected pucks can fit side by side on a single stator; accordingly, the system can have two rows of trays moving through it. On each puck, at a maximum packing density, there could be as many as 11 rows of five 20 mm vials, for a total of 55 vials per puck. However, spacing the vials apart from each other by at least one vial diameter creates room for cooling gas flow, as described in Chapter 3, and room for the weight sensing system described in Chapter 4. Based on these spacing needs, a maximum of three rows of four vials with a 20 mm diameter can fit on a single puck, giving a total of twelve vials. Accordingly,

(a) (b)

Figure 2.1: The Planar Motors "Flyway" (a) is a stator which uses electromagnetic coils to control the motion of an "XBot" (b), the puck which has static magnets that interact with the generated magnetic field. The stators are stationary within the system, while their coils drive the puck motions above their surface. The stator is 240 mm x 240 mm x 70 mm, while the pucks can range from 120 mm x 120 mm x 10 mm to 450 mm x 450 mm x 16 mm; the pucks used in this system are 120 mm x 180 mm x 10 mm, like the one pictured here.

further calculations used to determine machine size will use these parameters. These values can be used to update the machine sizing spreadsheet, as shown in Table 2.4. Using the Planar Motors puck geometries and vial spacing listed above for updated size estimations shows about a 50-85% system area increase across the sample vial sizes and production rates listed over the initial estimates based only on dense vial packing area.

2.5 Modularity

This machine uses a modular design to make it adjustable to different user needs. As Table 2.4 shows, the machine size requirements can vary significantly depending on the desired production rate, vial size, and process time. The modular nature of the system makes it reconfigurable as necessary to optimize its layout. This modularity also helps in system development, as the system can be assembled in sub-assemblies to validate performance. This build-up process is described further in Chapter 6.

Each module consists of three primary elements: the Planar Motors stator, the machine body, and a base used to connect the module to the machine structure. An example sublimation module is shown in Figure 2.2. The machine body creates the main vacuum seal between consecutive modules, and it connects the module to auxiliary components as needed. Each machine body includes features that are specific to that module's function and features that are shared between all modules. Because each module connects to a stator, a base, and a consecutive module, all modules can share the features used for these interfaces. These features cover the bottom surface and two sides of each module. The upper surface and two other sides then have features specific to each module's purpose. The specific design details for these shared geometries appear in Section 2.7.

Parameter	Name	Value	Value	Value	Units	Formula
Vial size	D_{vial}	20	40	20	mm	
Production Rate	$Prod$	100	100	1000	vials/hr	
Process Time	$t_{process}$	10	10	10	hr	
Total Vials	tot_{vial}	1000	1000	10000	vials	$:= Prod * t_{process}$
Puck rows	$puck_{rows}$	2	2	2	pucks	
Puck Length	L_{puck}	180	180	180	mm	
Puck width	W_{puck}	120	120	120	mm	
Vials per Puck	$puck_{vials}$	12	6	12	vials/puck	
Puck area	A_{puck}	0.0216	0.0216	0.0216	m^2	$:= L_{puck} * W_{puck}$
Total Pucks	tot_{puck}	84	167	834	pucks	$:= \text{Ceiling}(\frac{tot_{vial}}{puck_{vials}})$
Pucks total area	A_{tot}	1.81	3.61	18.01	m^2	$:= tot_{puck} * A_{puck}$
Perimeter length	L_p	7.56	15.03	75.06	m	$:= \frac{A_{tot}}{A_{puck} * puck_{rows}} * L_{puck}$
Rectangular area (width 1 m)	A_r	2.78	6.515	36.53	m^2	$:= \frac{L_p - 2}{2}$

Table 2.4: Limitations on vial packing density on each puck significantly increase the minimum area required for for the lyophilizer.

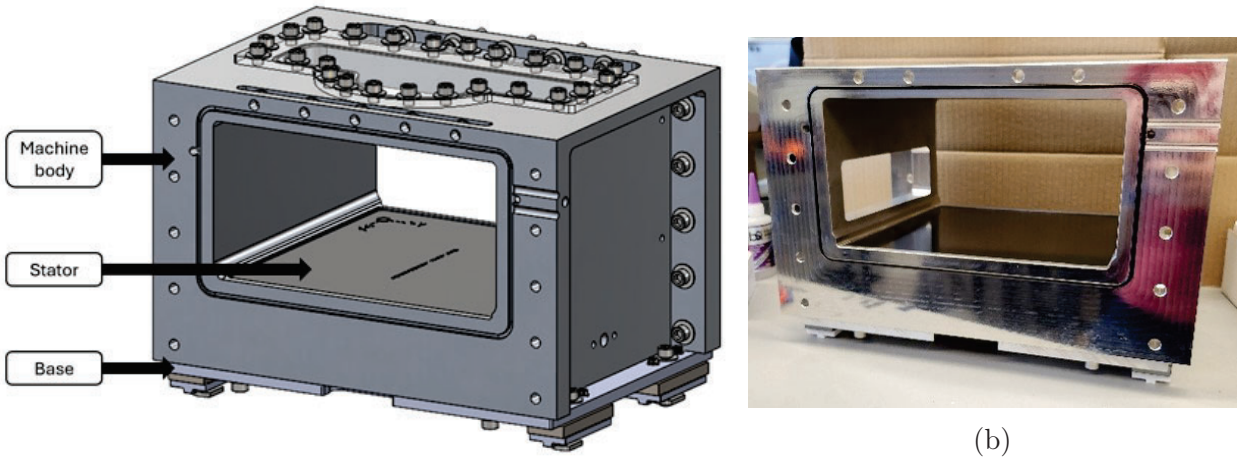

(a) (b)

Figure 2.2: The key elements of each module include the machine body, the stator, and the base.

When setting up the modular system, it is important to establish which elements of the modules are used as constraints during the assembly of the system. As more modules are added to the system, misalignment between consecutive units can accumulate and result in large shifts across the system. Some features must be used as constraints to align consecutive modules to ensure the total misalignment does not prevent full system assembly. These constraints can be purely internal to the modules themselves, or they can interface with features external to the modules that are built into a separate structure on which the modules are assembled. The effects of variation in these alignment features vary depending on which features are used for this alignment. The features that are not used as constraints for alignment will suffer from the tolerance error accumulation of not only themselves but also those of the features used for alignment. Consequently, the features most sensitive to misalignment should be used to reference consecutive modules to each other. However, if many features are used as constraints for alignment, they risk over-constraining the system and preventing successful assembly. Thus, the preference is to use fewer features as constraints for aligning consecutive modules.

The first consideration for a constraining alignment feature is the commercial Planar Motors system. These units must be positioned so that the pucks can successfully traverse between the stator modules without losing levitation and dragging on the surfaces. If the pucks drag along the stator surfaces, then they will produce particulates due to frictional wear, losing the benefit of being a non-contact motion system. The second consideration is the sealing interface between consecutive modules, which is discussed in detail in Section 2.7.2.1. Positional constraints can lead to gaps between sealing surfaces. These gaps must be sufficiently small to ensure the sealing faces can elastically deform if necessary to close these gaps without damaging components, such that the vacuum seal can be maintained between consecutive units. The third consideration is the positioning of the stator and the chamber on their base, and through these connections how the module mounts onto the machine table. As more units are added to the table, any misalignment between the consecutive

unit positioning can stack up and lead to a drift in the overall system shape. This drift can prevent the system from forming a complete loop, as seen in Figure 2.3. Given the overall system geometry which forms a loop, significant drift can lead to a break in the loop which does not allow the system to connect its entrance and exit. These positions must be connected for this continuously operating system to reuse the trays which carry products through the system.

The Planar Motors alignment can be successfully managed by the features that locate it within the machine body, as described in Section 2.7.2.2.1. Thus, the selection is narrowed down to the vacuum faces and the table assembly. This system uses a combination of these two features for assembly, driven primarily by the vacuum faces. The vacuum face parallelism and flatness are specified so that the total error accumulation if all modules are offset would fit within the linear shift allowable by the rail system. The modules are then placed on the rail system during assembly before being bolted together. The parallelism and flatness requirements on the machine body faces ensure that the maximum necessary deflection needed to maintain a vacuum seal between the units is within the elastic regime of the body's bolt flanges.

2.6 Process Division

A unit-dose continuous freeze drying concept was proposed in [19], laying out a conceptual design with separated chambers in which each freeze drying operation occurs. Such a continuous lyophilizer separates each process step in the lyophilization process spatially. This topology differs from existing batch lyophilizers, where all process steps occur within the

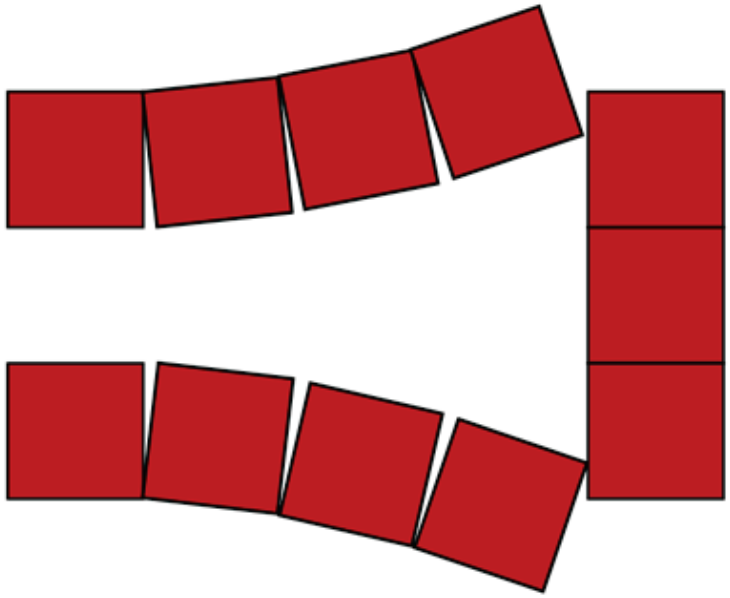

Figure 2.3: If consecutive modules have an accumulation of angular error in their relative positions, they may not be able to match up with the end of the system. Angular error exaggerated for demonstrative purposes.

same volume, but they are separated temporally. Separated process steps in a continuous lyophilizer allow all operations to occur simultaneously, with the products moving through the machine spatially rather than temporally. To ensure product is consistently moving through the system, the process sections are sized relative to the amount of time required to complete each process. This sizing comes from Little's Law, where the total process time is separated into distinct process sections that all share the same rate of units entering and exiting the system. Table 2.5 shows how the relative process time lengths are translated into machine section lengths.

Using these calculated machine section lengths, the overall system layout can be created. In addition to the process sections, the system includes load-locks that separate each process from each other. The design of these load-locks can be found in [25]. Also, one end of the system is used to interface with the equipment which adds and removes vials from the system. These additional sections add length to the overall machine size. A gap is left in the center of the rectangular layout to leave room for process chamber equipment, including piping and chamber bodies. The system layout created using these calculated lengths is shown in Figure 2.4.

This layout includes 6 different types of process stations. The first station, loading, is where the system interfaces with other elements of a production line. Importantly, this station is connected to both ends of the machine. The vials are carried by pucks that do not leave the machine, and these pucks need to return to the beginning of the machine to receive new vials once the lyophilized product is removed. Connecting the entry and exit of the system ensures that the pucks follow a continuous looping motion to maintain operation through this loading and unloading of vials.

Parameter	Name	Value	Units	Formula
Vial diameter	D_{vial}	20	mm	
Vials per tray	num_{vial}	15		
Vial rows per tray	$vial_{rows}$	3		
Vial tray length	L_{tray}	0.180	m	
Production rate	$Prod$	200	vials/hr	
Trays per module	width	$tray_{rows}$	2	
Average tray speed	v_{tray}	0.33	mm/s	$:= (\frac{Prod}{num_{vial}}) * (\frac{L_{tray}}{tray_{rows}})$
Conditioning time	t_{cond}	30	min	
Conditioning length	L_{cond}	0.6	m	$:= v_{tray} * t_{cond}$
Freezing time	t_{freeze}	30	min	
Freezing length	L_{freeze}	0.6	m	$:= v_{tray} * t_{freeze}$
Drying time	t_{dry}	10	hr	
Drying length	L_{dry}	12	m	$:= v_{tray} * t_{dry}$
Total length		13.2	m	$:= sum(L_{cond}, L_{freeze}, L_{dry})$

Table 2.5: The relative amount of time required for each process translates directly into a relative process section length in the lyophilizer based on the relationship given by Little's Law.

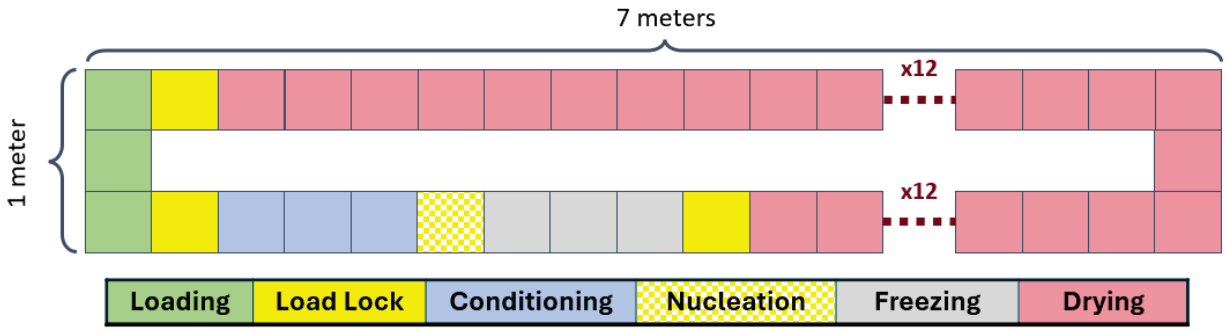

Figure 2.4: The expected machine layout shows the relative lengths of each process section given the required process times. This layout provides an expectation for the overall machine size given a desired 100 vials/hour production rate.

The second type of station is the load lock, which separates each process chamber. These load-locks are needed to isolate the conditions created in each process from their neighboring processes. The load lock design is detailed in [25], and it is outside the scope of this work.

The third station is the conditioning station. In this station, the product is subcooled to a designated temperature just below 0°C. This subcooling prepares the product for controlled nucleation. More details about this process can be found in Chapter 3.

The fourth station is the nucleation station. In this station, controlled nucleation can occur. The different potential mechanisms for this nucleation are described in detail in Chapter 3. In this system, thermal shock is the preferred nucleation method, as described in Section 3.2.2.

The fifth station is the freezing station. In this station, the product is further cooled to a designated temperature, typically between -40 and -60°C. This freezing ensures that the ice crystals fully solidify and that the product is well below the collapse temperature before entering the drying section. More details about this process can be found in Chapter 3.

The sixth station is the drying station. In this station, the water sublimates out of the product, leaving behind the lyophilized cake. This station must maintain a vacuum to drive this sublimation process. More details about this process section can be found in Chapter 5.

An additional process section can be added between the conditioning and freezing sections if a user wants finer control over the ice crystal growth during the freezing process. The ice crystal growth rate depends on the freezing temperature, and if the desired growth temperature is significantly different from the freezing station setpoint, then an intermediary crystal growth station can be added before the freezing section to provide the desired growth temperature conditions.

2.7 Chamber Module Core Design

The overall lyophilizer is built from an assembly of chamber modules. Each module shares a core base geometry while also including corresponding features for its application. The matching core geometry improves manufacturability and helps to ensure proper sealing between any pair of modules. This section will discuss these core geometry features shared

among all modules. Chapter 3 describes the features specific to modules in the conditioning and freezing stations, and Chapter 5 describes the features specific to modules in the drying station. The features specific to the load lock modules can be found in [25], as their design is outside the scope of this work.

The lyophilizer consists of four types of modules, shown in Figure 2.5. The load lock modules require features to accommodate doors moving between their open and closed positions. Because these doors operate in the gap between two modules, they require features on the modules on both sides of the appropriate gap. The load lock doors are angled and face the same direction on both sides of the load lock based on the pressure differential created by their orientation relative to the vacuum environment in the drying chamber. This matched orientation means that the load lock is not symmetric about its middle plane. The asymmetry of the load lock unit results in the need for two types of modules, one for interfacing with each side of the load lock. The most common module seen in the system is only the tunnel structure, and it does not include any door features. Each of these module types has additional features as needed for the specific stations in which they are installed.

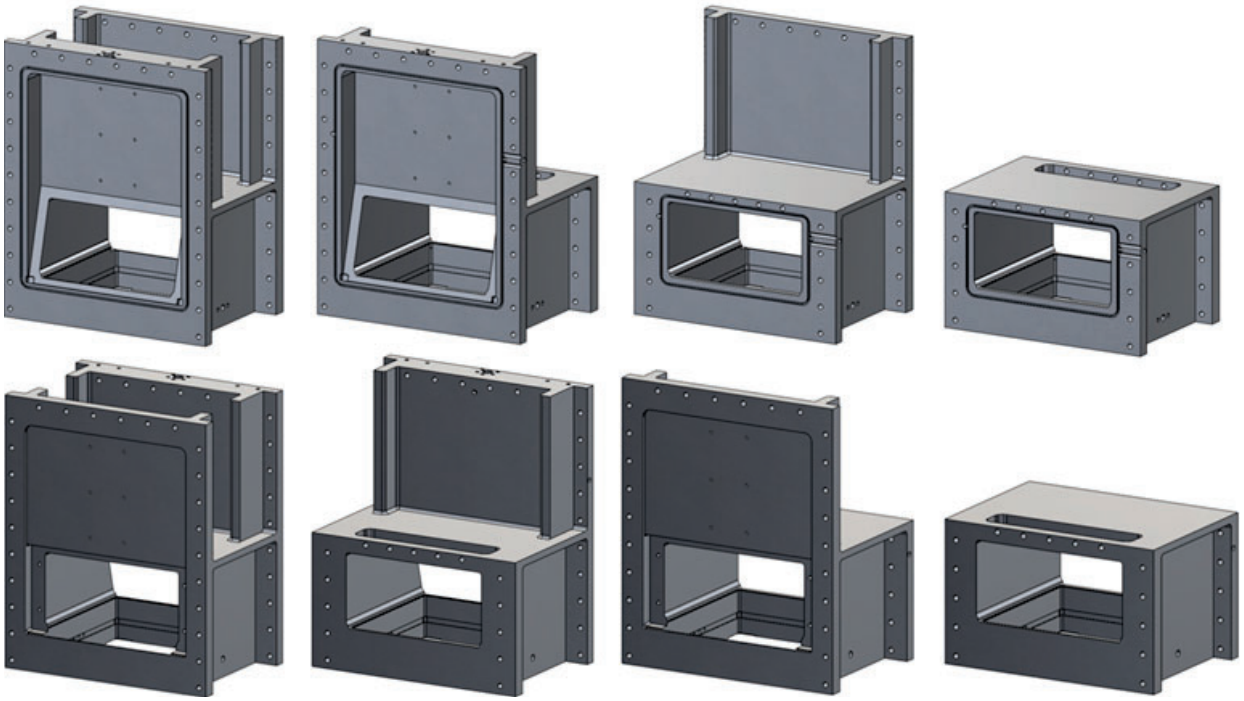

Figure 2.5: The system uses four main types of modules. The first type (first column) accommodates doors on both sides of the module, and it is used for load-lock chambers. The second and third module types include features to mate with the load-lock, either on the load-lock door support (second column) or door (third column) side. The last type of module (fourth column) does not include any door features. This module is the most common, and it makes up most of the tunnel structure.

2.7.1 Tunnel Geometry

The general tunnel geometry defines the internal volume through which the vials move. Keeping this internal volume smaller makes it easier to manufacture because it requires less material and makes it easier to maintain vacuum conditions. However, a smaller internal volume holds fewer vials. The minimum width of the internal tunnel is 240 mm to leave sufficient room for two rows of Planar Motors pucks. The width is extended to 250 mm to provide 5 mm of clearance on each side of the tunnel. The tunnel height is determined with considerations for the weight sensing system. The minimum tunnel height is based on the height of the vials plus the thickness of the Planar Motors pucks, which is a total of 65 mm. However, as described in Chapter 4, the weight sensing system uses vertical sensing arms that extend above the height of the vials. To leave room for the weight sensing components, the internal tunnel height is set at 125 mm. As the weight sensing system is further developed, the tunnel height can be adjusted on future hardware. If the weight sensing system requires more space above the vials, the height can be increased. Similarly, if the weight sensing system requires less space above the vials, then the tunnel height can be decreased. The tunnel wall thickness is set at 20 mm for the walls to provide both resistance to vacuum loads and to provide sufficient depth to bolt the accessory features to these surfaces. The tunnel roof is thicker, set to 25 mm, to handle the weight of additional mounted components such as vacuum valves.

Chambers that require doors to interface with the load-lock system include an additional vertical extension that accommodates these doors. The details of the design for this extension

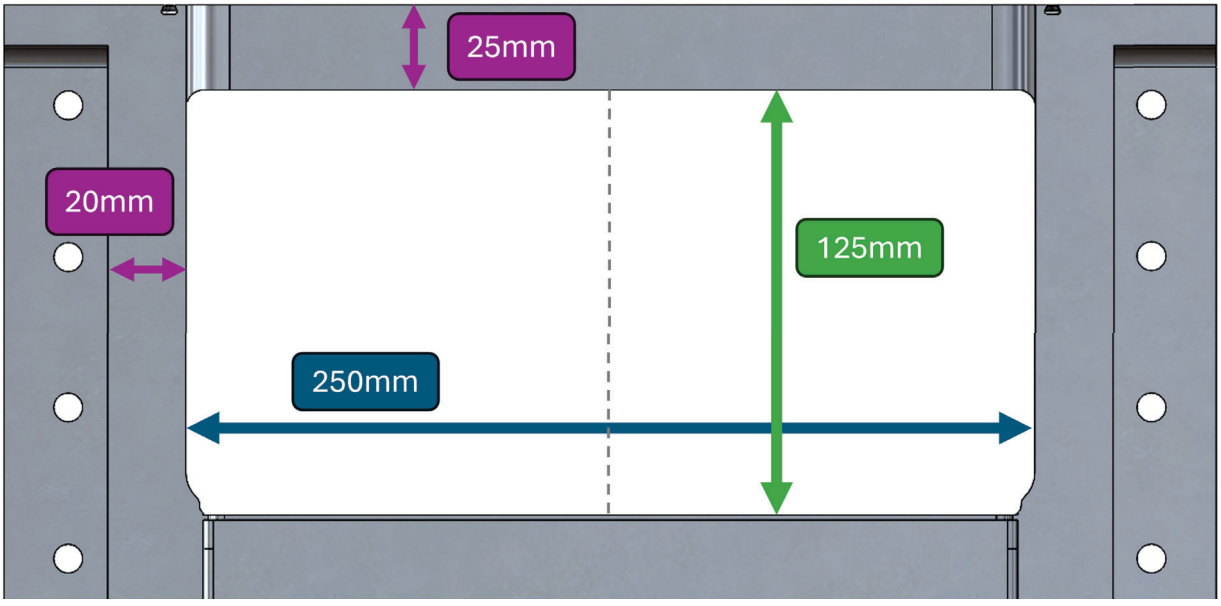

Figure 2.6: The tunnel interior width is sized at 250 mm to accommodate two side-by-side 120 mm wide pucks. The tunnel height of 125 mm provides room for the vials and the weight sensing system described in Chapter 4. The 25 mm roof thickness is greater than the 20 mm wall thickness to accommodate the weight of external components that mount to the roof, such as vacuum system valves.

appear in [25]. These extensions add 236 mm to the height of the chamber, shown in Figure 2.7.

2.7.2 Module Sealing

There are two sealing interfaces shared by every module. Each module has a seal between consecutive modules and another seal between the machine chamber and the stator. Both seals must be robust to a vacuum environment. The locations of these seals are highlighted in Figure 2.8.

2.7.2.1 Consecutive Module Sealing

A critical module feature is the sealing method between consecutive modules. This sealing can be done between only the machine chambers, or it can be shared between the machine chamber and the Planar Motors stators. To simplify the sealing method, it is preferable to avoid having seal paths that cross multiple different components. Sealing paths that cross seams create additional possibilities for leaks which can compromise the vacuum environment needed for the sublimation process. Thus, the machine chambers have separate seals between consecutive modules and their Planar Motors stator. This separation requires the chamber modules to extend into a gap between consecutive stators. As noted in Section 2.4.1, the maximum allowed gap between stators is 10 mm. While a simple choice would be to have each chamber fill half of this gap, the work done in [25] shows that the chambers that include

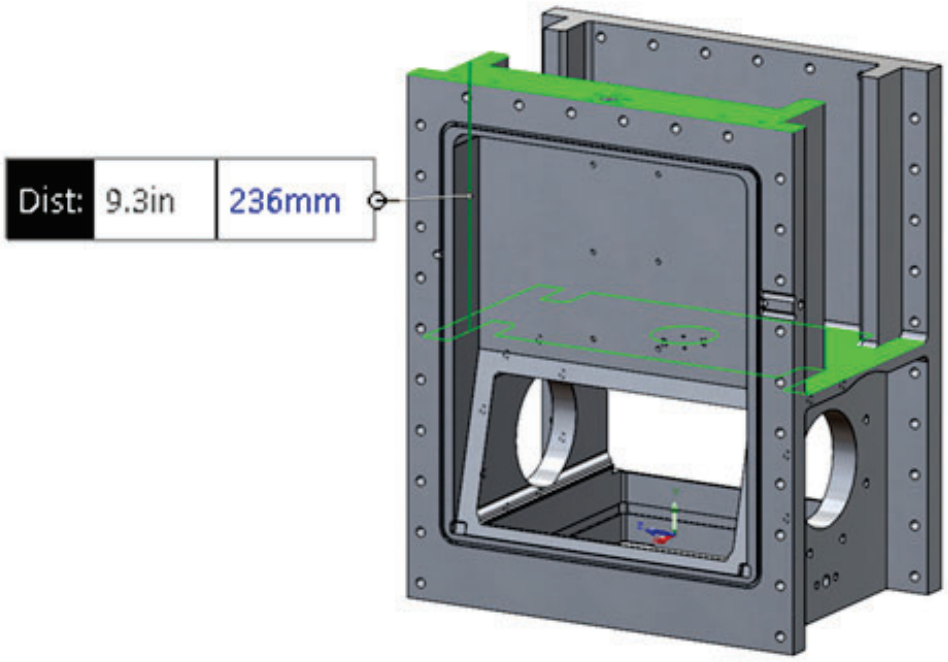

Figure 2.7: The vertical extensions added to the chamber to accommodate the load-lock doors add 236 mm to the overall height of the module.

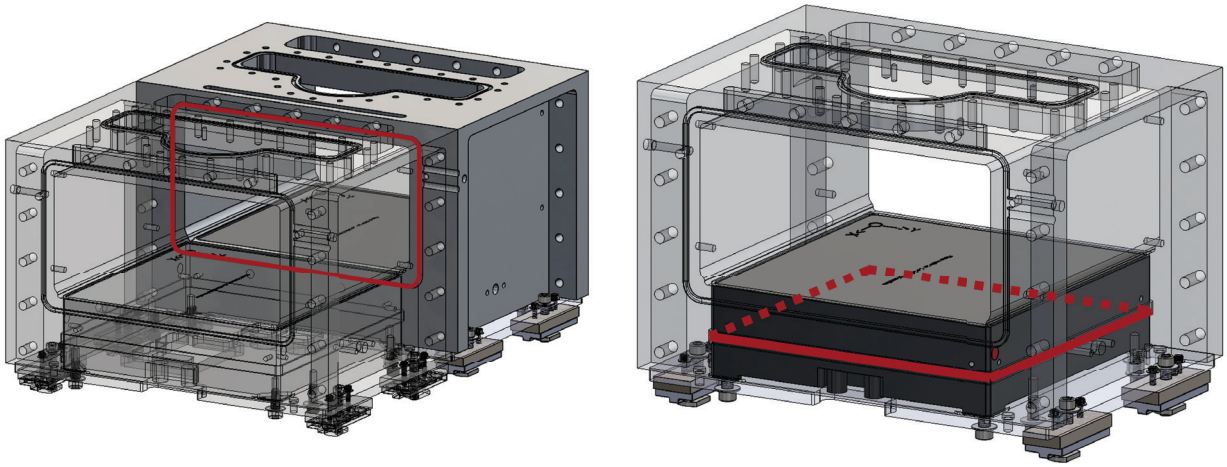

(a) (b)

Figure 2.8: All modules include a sealing interface, highlighted in red, between consecutive modules (a) and between the machine chamber and the stator (b). The sealing methods for each of these interfaces does not need to be the same, but both must be robust to vacuum.

door features need to extend more than 5 mm into this gap. This decision is primarily driven by the need for space to accommodate the load-lock door, which extends down into this gap. The asymmetry also provides space for sealing features to be machined on the thicker side of the chamber extension in the gap, while the thinner side remains flat to interface with the sealing material. This machining choice helps ensure that the minimum thickness of any particular material wall is no less than 1.5 mm.

The interface between consecutive modules is sealing two faces. Two common face seal methods are gaskets and O-rings. Gaskets can have looser surface flatness requirements and can cover large areas. O-rings require specially machined grooves, but they take up a smaller area on the sealing surface. Gaskets are typically cut specifically for their application shape, while O-rings can come in standard sizes or be custom made. O-rings are selected for this face seal because there is enough room on the interfacing faces to fit features for a standard O-ring size, making the sealing material easier and cheaper to source. Additionally, the O-ring groove can be machined using a dovetail tool, which holds the O-ring in place after installation. Having features that keep the sealing material in place during assembly makes the assembly process easier. A nominal 1/8 in O-ring is used because it is the largest size which does not require cutting more than half of the material depth into the material in the gap between stators. The recommended O-ring groove parameters in the Parker handbook [26] are used for the dovetail groove profile.

The Parker handbook specifies the recommended crush for a face sealed O-ring as 20%[26]. The O-ring performance is expected to be maintained as long as the crush remains greater than 15%. The geometric errors which could create a potential gap between consecutive module faces must be specified to remain within limits which ensure that this crush is maintained on the O-ring. These geometric errors can arise from the flatness of the module face and the relative positioning between consecutive faces during assembly. The assembly error can come from deflections in the table structure under load and positioning of the

module on the rails. Given a nominal O-ring crush of 20%, a minimum O-ring crush of 15%, and an O-ring diameter of 0.139 in, the total error allowed is 0.007 in. The flatness of the module face with the O-ring groove is allotted 0.002 in of this total error. Correspondingly, the parallelism of the opposite face is specified to 0.004 in. The parallelism requirement also ensures that the modules' potential offset relative to the rails can be accommodated by the sliders in the t-slotted aluminum. Thus, the positioning of the modules on the rails does not affect O-ring crush. These specifications leave 0.003 in available for the table structure deflection under load. This allocation is used in Table 2.6 to ensure that the table structure is sufficiently stiff to maintain the O-ring vacuum seal. The adjustability provided by the leveling screws described in Section 2.8.3 helps to loosen the requirement for the table structure stiffness.

The normal force used to crush the O-rings comes from the bolts used to connect consecutive chambers. Each chamber has bolt flanges on the sides through which these bolts are connected. These flanges have the same thickness as the tunnel walls. The bolt holes are clearance holes for simplicity and to provide flexibility for minor misalignment between consecutive chamber positioning. The chambers are connected using 5/16-18 bolts. The bolt holes are positioned so that a standard 5/16 in washer has a 2 mm clearance from the wall. The bolt holes are spaced 1.75 in apart from each other.

2.7.2.2 Stator Module Sealing

Because each standard module includes a single stator, the geometry and features that determine the connection between the machine chamber and the stator are part of the core module geometry. This geometry must include proper tolerancing to ensure that any given stator can fit into any given module. Because this system requires numerous modules, attempting to machine each module to match a pre-assigned stator would create significant additional assembly burden. In addition, a vacuum seal must be formed between the stator and the chamber. The top surface of the stator is exposed to a vacuum environment in the sublimation station, while the bottom of the stator is exposed to atmospheric conditions. A seal must be added between these surfaces to ensure that the system can maintain a vacuum environment for sublimation.

2.7.2.2.1 Stator Pocket

The machine chamber includes a pocket in which the stator is inserted. This pocket includes a small lip that contacts the top surface of the stator, shown in Figure 2.9. This lip extends over the stator surface by 1 mm on all sides, and it is used during assembly to ensure that the stator is inserted to a proper depth within the machine chamber before both are bolted to the baseplate. An epoxy shim layer is added to the bottom of the stator on its bolting feet during assembly. This epoxy shim layer ensures that any mismatch in the tolerance directions of the stator and the machined pocket does not prevent either the stator or the machine chamber from bolting to the baseplate. This epoxy shim layer is nominally 0.005 in thick, which is larger than the total potential over-sizing of the stator and under sizing of the pocket depth. An epoxy shim is used instead of a solid-material shim to account for variations in the actual dimensions of each part due to their independent tolerances,

such that assigning specific stators to each pocket remains unnecessary. The bottom of the machine chamber includes overflow grooves to ensure that any extra epoxy added during the shimming process can flow out of the groove, preventing the risk of the epoxy shim layer becoming too thick. Using sealing features between the machine chamber and the stator side walls, such as the injectable sealant, allows the bottom of the stator to remain accessible so no electrical feedthroughs are needed.

The chamber module does not cover the entire stator surface beyond this lip. Because the levitation height for the Planar Motors pucks has a 3 mm maximum, any material between the stator and the puck would need to be less than 3 mm thick. If a thin material sheet is spread over the entire 240 mm x 240 mm area covering the stator and subjected to a vacuum load, it would deflect and interfere with the pucks. This deflection risk can be mitigated by bonding this thin material section to the stator surface. However, this method would permanently connect the stator to the machine chamber. As this machine is still in primary development for this work, it is important to retain the flexibility to remove stators

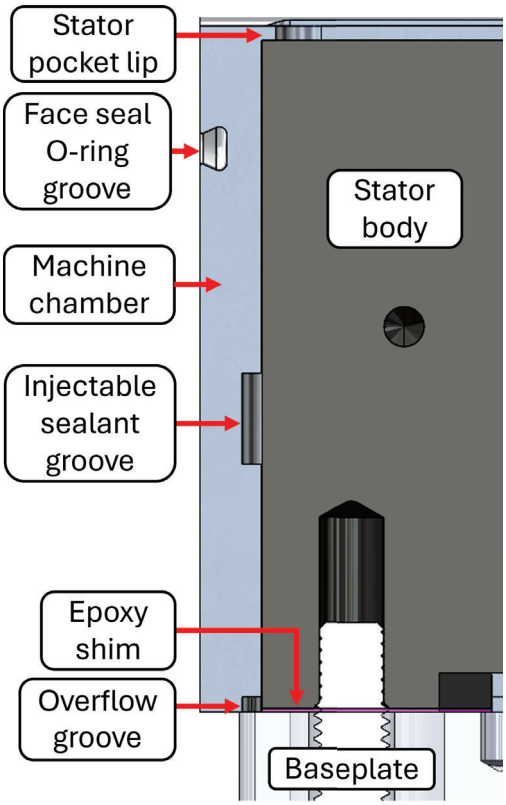

Figure 2.9: Key features for the stator pocket include the pocket lip used to position the stator during assembly, the O-ring groove used to seal consecutive chamber units, the sealant groove for creating a vacuum seal between the stator pocket wall and the stator, the epoxy shim used to ensure the baseplate and the stator are both bolted to the baseplate regardless of dimensional errors within their respective tolerances, and the epoxy shim overflow groove used to ensure that any excess epoxy used in the shimming process can flow away from the stator and baseplate interface rather than create too thick of a shim.

if necessary. After using this system to demonstrate reliable performance, future iterations can fully cover the stator surface and allow the stator to be bonded to that surface. If using this bonding strategy, it is important to ensure that there are minimal air bubbles trapped in the bonding layer, as they can create localized vacuum loads on the machine chamber and the stator.

The stator pocket dimensions are set based on the maximum possible size of the Planar Motors stators. The stators have a nominal length and width of 239.7 mm with a maximum positive tolerance of 0.30 mm [24]. This measurement means that the maximum potential stator dimensions are 240 mm x 240 mm. Accordingly, the minimum dimensions for the machine chamber's stator pocket must be no less than 240 mm x 240 mm. The maximum pocket dimensions determine the possible offset between consecutive stators. Because the pocket will be larger than the stator, the stator will effectively float within the pocket. This means that the stator will likely abut one or two walls, but that final position relative to the chamber is not necessarily known before assembly. If this relative positioning is a critical dimension, then the stator and the chamber can each be referenced off the base plate using features such as dowel pins.

Prior Planar Motors testing indicated that the system can handle at least 1 mm of positional offset in any direction. This value serves as the total positional error that can be budgeted to the features which determine the consecutive stator positions in machine chambers. This positional error can come from three main sources: 1) errors in the position of the stator relative to the pocket, 2) errors in the pocket dimensions and position relative to the machine chamber, 3) errors in the overall machine size, and 4) errors in the assembly of consecutive modules. Because the consecutive module faces serve as the assembly references, the fourth source of error is expected to be negligible. Accordingly, the stator positional error is budgeted to the first three sources.

The first source of error has a minimum value which comes from the tolerancing on the stator dimensions. These units are machined by a commercial manufacturer, so their tolerances must be accepted as they are and designed around in this system. The maximum negative tolerance on the stator dimension is 0.10 mm. This tolerance means that the smallest expected stator would be 239.6 mm x 239.6 mm. If two of these stators were to be placed in consecutive 240 mm x 240 mm pockets, then there could be a maximum error of 0.4 mm between their nominal centers, as each could be offset from their corresponding pocket center by 0.2 mm. Accordingly, only 0.6 mm of positional tolerance remains to allocate to the stator pocket size, the relative pocket position, and the overall module size. This tolerance is evenly allocated to all three error sources, providing 0.2 mm tolerance for each. This tolerance is reflected in the drawings for a tunnel module. For a part of this size, a 0.2 mm tolerance is reasonable to achieve and does not require specialized machining capabilities.

This tolerancing only considers linear offsets, not twisting between the stator and the chamber. The potential for twisting is limited by setting a parallelism and perpendicularity requirement for the stator pocket, using a similar tolerance to the linear tolerances listed above. Using the parallelism and perpendicularity requirements ensures that the effective linear offsets between any points on the stators do not exceed the allowable offsets from the error budget.

2.7.2.2.2 Stator Seal

The sealing path between the stator and the chamber goes around the sides of the stator because its top and bottom surfaces are exposed to different conditions. The seal is not done on the stator's top surface to maximize the available usable space on the stator. The seal is not done on the bottom surface to keep the electrical connections needed to operate the stator available without adding electrical vacuum passthroughs.

O-rings are a standard gland sealing method for plugs inserted into housings. However, when stretched around the side of the plug, the O-rings have a minimum recommended corner radius to avoid overstressing the O-ring before vacuum load compression [26]. This recommended corner radius is at least three times the diameter of the O-ring. The commercial stators have a corner radius of only 3 mm, which means they would require a 1 mm O-ring diameter. The smallest commonly available O-ring diameter is 1.5 mm, which is larger than this specification. Additionally, stator pocket tolerancing used to guarantee that the stators fit in the machine chamber would not be able to guarantee sufficient crush on such a small O-ring diameter to create an effective seal. It is possible to get molded O-rings that can accommodate the tight corner radius requirement, but these are much more expensive. Thus, a different sealing strategy from O-rings is pursued.

Rather than inserting a sealing material like an O-ring and then adding the stator to the pocket, an opposite strategy of inserting the stator to the pocket and adding the sealing material afterwards is implemented. In this method, a groove is machined into the chamber module and filled with an injectable sealing material after the stator and the machine chamber are assembled. The sealing groove is machined into the chamber because the stators are a separate commercial product for which modification is not available. The sealing groove path is located at the point where it can have the largest cross-sectional area to minimize the resistance to sealant flow during injection. The path is positioned to avoid any features on the sides of the stators, as these could compromise the seal integrity or serve as leak points, preventing the injection from completely filling the groove. This positioning is shown on Figure 2.10. The maximum sealing path depth is set by maintaining a minimum thickness of the walls around the stator. This minimum thickness is set at 1/16 in to ensure material integrity during machining and assembly.

The injectable sealant testing and selection process is detailed in [25]. The sealant used is Self-leveling Green, produced by Av-DEC. The sealant groove and the stator are both sandblasted to improve sealant adhesion. The sealant is injected into the groove from one side of the chamber, where it then flows around the groove and exits from the other side of the chamber. This injection fill path is shown in Figure 2.11.

The sealant is injected through a tube which is joined to the chamber wall through a 1/8 in NPT adapter that mounts to a corresponding 1/8 in NPT threaded hole in the side of the chamber. A matching assembly is used on the opposite side of the system for the sealant outlet. This outlet is monitored for air bubbles. During the initial injection, the sealant traps air bubbles as it displaces the air in the sealant groove. As the injection continues, these air bubbles are purged from the system. The outlet tube is oriented upwards relative to gravity to promote this purging process. On the outlet side, there are also two 1/4-20 threaded holes, which are used with corresponding threaded plugs to ensure the sealant fills the groove on both sides of the stator. During the initial injection, the plugs are left open

(a) (b)

Figure 2.10: (a) Two potential injectable sealant path locations highlighted on the stator, where the path would not cross any features on the stator. The path with the larger potential height is used to reduce the sealant flow resistance during injection. (b) The sealant groove thickness, and the resulting wall thickness after the material is removed to create the sealant groove. The resulting wall thickness is maintained at 1/16 in, setting the sealant groove thickness at just over 2 mm.

until the sealant is seen from one side. The corresponding plug on that side is then inserted, increasing the flow resistance on that side and causing the sealant to preferentially flow to the under-filled side. This strategy ensures that both sides of the stator are filled while using only one entrance and exit located opposite each other. Removal of test stator stand-in units that were sealed to test machine chambers showed no notable seam line at this exit interface.

2.7.3 Module Base

The stator and the machine chamber for each module are mounted on a baseplate. This baseplate is shown in Figure 2.13. The baseplate includes two sets of through holes, one for mounting the stator and one for mounting the machine chamber. Dowel pin holes can be included for positioning each of these components if necessary for positional tolerances. The baseplate also includes cutouts for the electrical and cooling channel connections on the stator. Additionally, the baseplate includes threaded holes used for jack screws which can be used to separate a machine chamber from a stator after sealing them together. Bolts can be threaded through these holes, allowing them to push on corresponding cutouts on in the machine chamber to create a shear force on the injected seal described in Section 2.7.2.2.2. The total baseplate length is less than the nominal 250 mm length for each module to ensure that the machine chamber faces will contact each other before the baseplates.

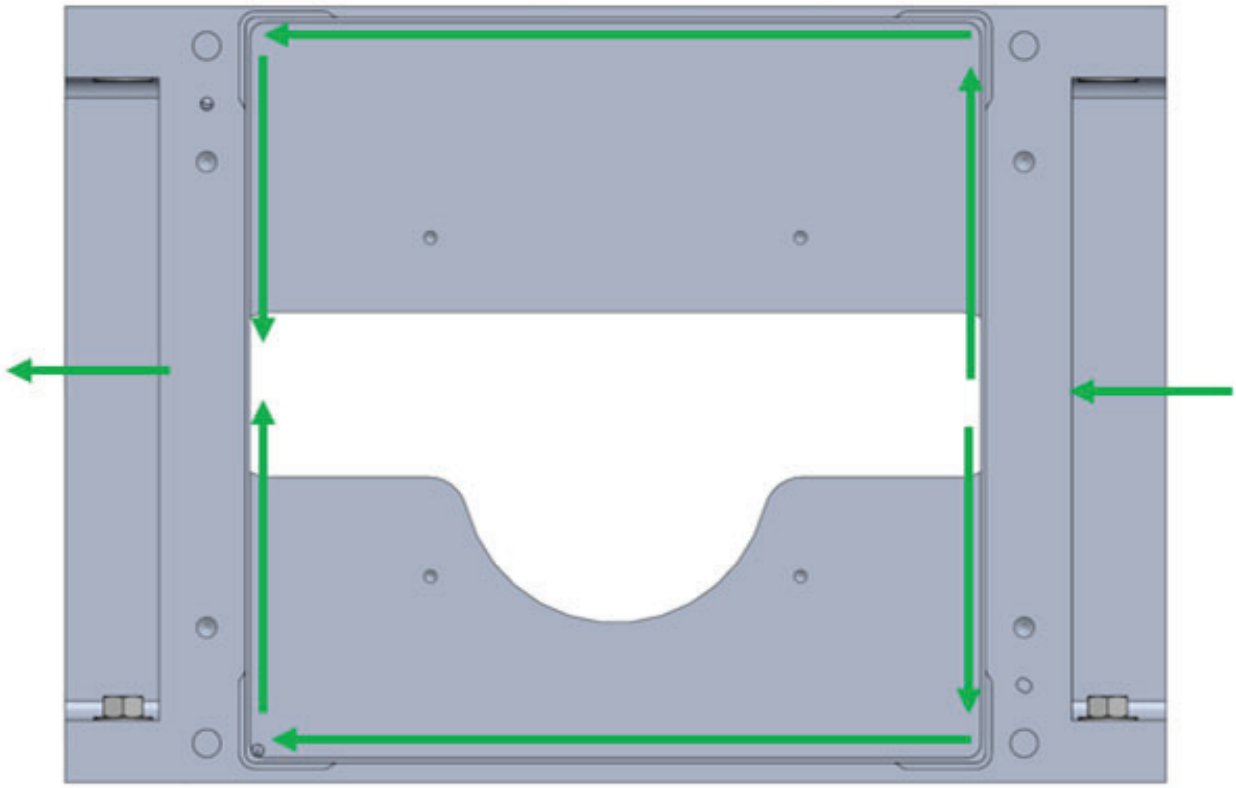

Figure 2.11: The injected sealant used to create a vacuum seal between the machine chamber and the stator flows in from one side of the chamber, around both sides of the stator, then out an exit on the opposite side of the system. The green arrows indicate the sealant flow direction.

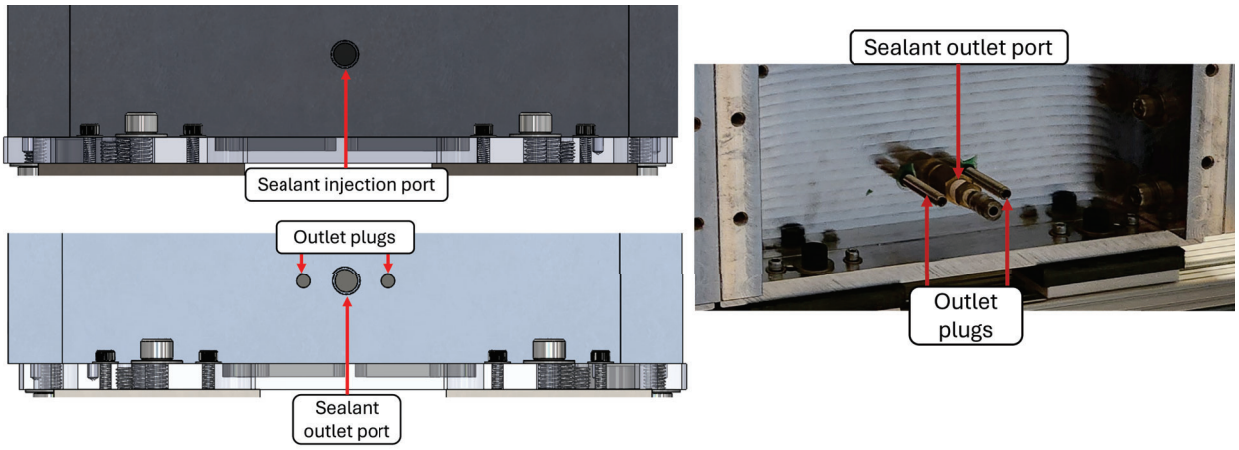

Figure 2.12: The side where sealant is injected includes A) a single 1/8 NPT threaded hole which connects to an adapter used for sealant injection. On the outlet side, B) the same 1/8 NPT threaded hole is used to monitor sealant flow for air bubbles. C) Two $\frac{1}{4}$-20 threaded holes are used with threaded plugs to ensure sealant flows around both sides of the groove.

2.8 Machine Base Structure

The machine base structure serves to position the modules relative to each other and to isolate the system from the ground where it is installed. This structure also creates space for piping, wiring, and the placement of auxiliary components such as vacuum pumps. The base structure must meet deflection requirements to ensure that it will not create gaps in the vacuum seals between units. The base must also accommodate variations in the floor on which this system is built. This second requirement enables the machine to be built on surfaces with looser flatness tolerances, particularly so that specialized flooring is not required to install this system. The base also includes features that allow different modules to be placed in a variety of locations so the potential system arrangements are not limited by the base structure's design.

2.8.1 Table Strategy

There are two general approaches to designing a structure that can meet the maximum deflection requirement. One solution would be to take this requirement as a maximum structural deflection which can be used to calculate the system's required stiffness. An alternative, opposite approach would be to design a system with a high degree of adjustability, such that the structure can be measured and adjusted during assembly. Additionally, these options

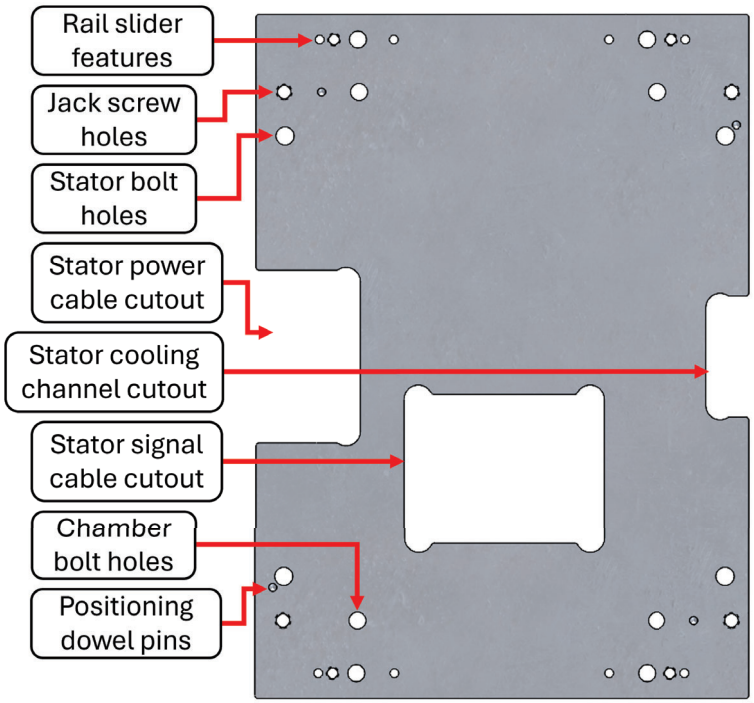

Figure 2.13: The module baseplate connects to both the machine chamber and the Planar Motors stator. The baseplate includes cutouts for the electrical and cooling channel connection on the stator. The baseplate also includes threaded holes which can be used to push the machine chamber if needed to separate it from a stator after sealing them together.

could be hybridized, where a generally stiff structural system ensures that the overall deflection remains low while adjustable elements can correct for any areas that require minor adjustments during assembly.

Evaluating and comparing these strategies requires an understanding of the structural support deflection requirement. This deflection requirement comes from how the structure's deflection affects the system built on top of it, shown in Figure 2.14. When there is a difference in the angular orientation of consecutive modules created by a changing slope on the table surface, this angular difference results in a growing gap between the mating planes of consecutive modules. This gap is largest at the point furthest from the contact location for two consecutive modules, which can be either the top or the bottom of the modules depending on the concavity of the structure's bending shape. If this gap is sufficiently large, it can prevent consecutive modules from mating closely enough to generate sufficient sealing force, which would compromise their vacuum seal.

The machine modules resting on top of the base structure need to seal for vacuum to corresponding faces, as described in Section 2.7.2.1. As shown in Figure 2.14, when the surface on which the modules rest is not level, the difference in slopes under each unit results in a gap between the top or bottom edges of each module. Though the difference in slope may be small, the angular difference is amplified, as described by Abbe errors, creating a leak risk if this angular difference is not managed. The linear gap width created by a change in slope along the deflected table length is given by Equation 2.2.

$$G = H|\theta_1 - \theta_2| \tag{2.2}$$

Where G is the gap length, H is the height of the furthest sealing interface from the modules' contacting edges, and θ is the relative angle to horizontal of each unit. The height

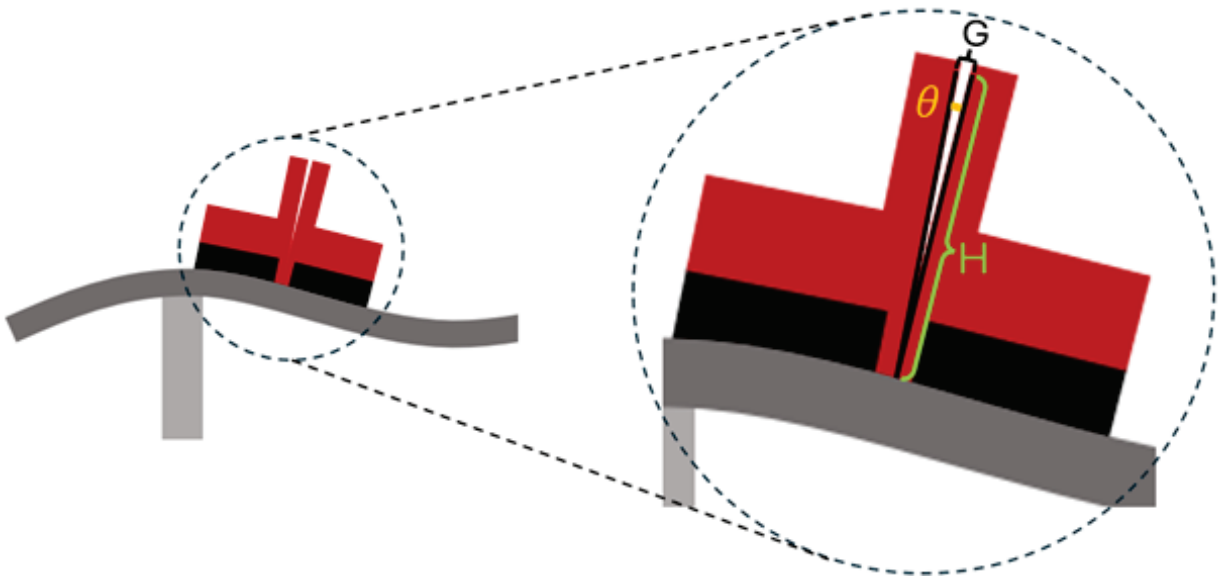

Figure 2.14: When modules are placed on a sloped surface, the angular offset between neighboring modules is amplified. This offset can lead to potential gaps between neighboring modules which would break their vacuum seal.

value comes from the module geometry, which is described in detail in Section 2.7. In this system, the maximum height of the sealing interface from the table is 450 mm. The maximum acceptable gap length comes from the vacuum seal error budgeting, as described in Section 2.7.2.1. This acceptable gap length is 0.003 in. Taking the ratio of these two values gives the maximum acceptable slope difference between any two points on the surface of the machine base structure. For this system, that value is 0.00056 radians.

The small order of magnitude of this requirement indicates that it would be difficult to measure over the course of many adjustments consistently and accurately along the structure surface if using the highly adjustable structure strategy. Thus, the deflection requirement serves as a stiffness requirement for the structure's surface. However, given this small order of magnitude, there is a high risk in only relying on structural stiffness given the strictness it could place on manufacturing tolerances and how the system can change with temperature cycling. Thus, the hybrid strategy of a highly stiff table structure combined with localized adjustable supports is implemented. The stiff table structure design is described below in Section 2.8.2, while the localized adjustable support design is described in Section 2.8.3.

2.8.2 Table Design

The structure deflection profile depends on how the structure is supported. Using many support structures between the top surface and the floor allows the load to be spread out, decreasing the bending length of any given unsupported table section, in turn reducing the required structural stiffness. However, each support needs to fully engage with the floor under the machine to provide any benefit. If the floor is not perfectly flat, then the structure will act as though any support not touching the floor does not exist, and it will bend further until it either engages those lifted supports or reaches a different equilibrium. This observation indicates that having more than two support points will result in a system whose deflection profile also depends on the floor profile on which it is built. The system can account for floor variation either by incorporating adjustable supports or by supporting the structure at only two points. In the latter case, where the structure only has two support points, it is possible for the structure to not be level. For vacuum sealing between modules, this consideration is not a concern, as the seal gap effect is based only on the relative angle between consecutive units rather than the angle of the units relative to the floor. Additionally, if keeping the whole system level is a critical requirement, one of the two support points can be made adjustable to meet the leveling requirement while maintaining the system independence from floor flatness. Assuming a structure can be built to meet the stiffness requirements for a structure with only two support points, this strategy is much simpler to implement than a multi-support option, so it is pursued further.

A key parameter for a 2-support structure is the positions of these two supports. Depending on the support positions, the resultant beam deflection can vary widely, as seen in Figure 2.15.

The optimal support positioning can be determined by modeling the physical state of the structure using singularity loading modeling. This method uses singularity functions to represent changes in the loading conditions for the beam over its length. The structure built for this system is modeled as a beam with a distributed weight load on top representing the system weight and two point loads on the bottom representing the supports. The free body

diagram for this loading appears in Figure 2.16. The point load locations are assumed to be symmetric. While the support locations do not necessarily need to be located symmetrically, this alignment is expected to minimize the maximum slope variation across the structure surface. While the system weight is not necessarily fully symmetric, it is not expected to vary sufficiently to require analyzing the parameter space including asymmetric support locations. Equation 2.3 shows the singularity modeling for this system:

$$q(x) = -w\langle x\rangle^0 + R_1\langle x - L_1\rangle^{-1} + R_2\langle x - L_2\rangle^{-1} \tag{2.3}$$

$$-\int_0^x q(x)\,dx = V(x) = w\langle x\rangle^1 - R_1\langle x - L_1\rangle^0 - R_2\langle x - L_2\rangle^0 \tag{2.4}$$

$$-\int_0^x V(x)\,dx = M(x) = EI\frac{d^2y}{dx^2} = -\frac{w}{2}\langle x\rangle^2 + R_1\langle x - L_1\rangle^1 + R_2\langle x - L_2\rangle^1 \tag{2.5}$$

Assuming the supports are equally spaced and the load of the machine is evenly distributed, the reaction forces at the supports should be equal to each other. This means that each support is expected to bear half of the total machine weight load, as shown in Equation 2.6:

$$R_1 = R_2 = \frac{wL}{2} \tag{2.6}$$

Additionally, because the supports are symmetrically located about the table center, the positions L_1 and L_2 can be written in terms of a single variable, S, representing the spacing between the two supports. The value of S can range from 0 to L, where L is the full length of

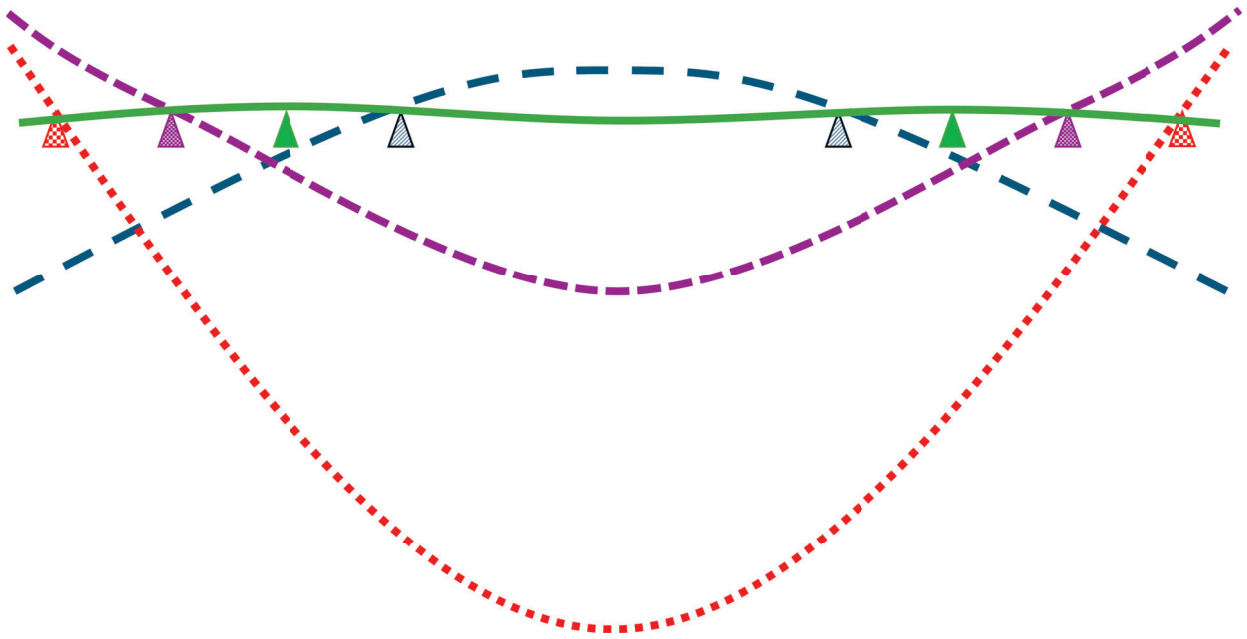

Figure 2.15: Different support placements can result in high variation in the overall beam profile. The range in deflection is exaggerated for demonstrative effect. Figure adapted from [27].

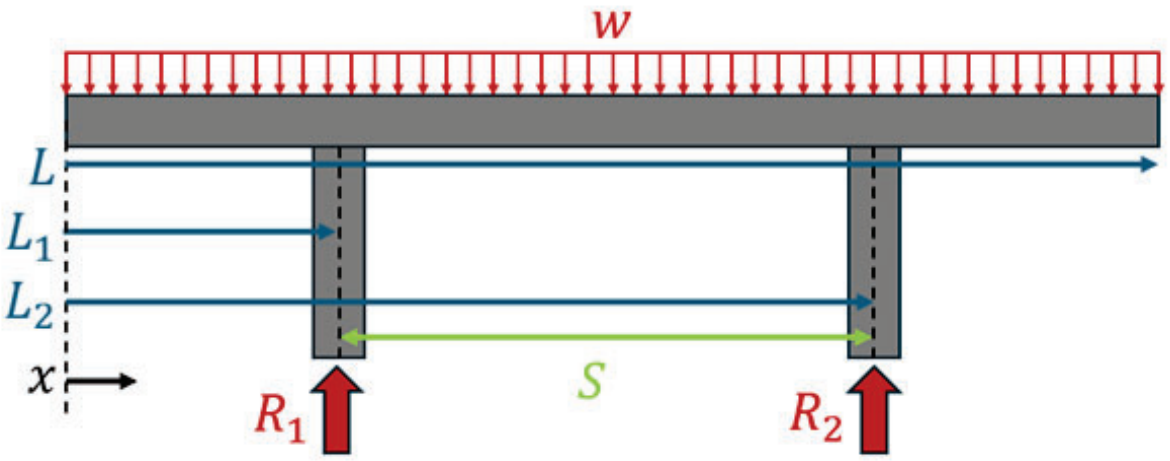

Figure 2.16: A free body diagram showing a table with two supports experiencing uniform loading on its surface.

the system. These boundary conditions represent when there is effectively a single support in the center of the system, $S = 0$, and when the supports are at each end of the system, $S = L$. Using this S term, L_1 and L_2 can be rewritten as the following:

$$L_1 = \frac{L - S}{2}, L_2 = \frac{L + S}{2} \tag{2.7}$$

As expected for symmetrically located supports, the sum of L_1 and L_2 is equal to L. Substituting these relationships into Equation 2.5 above gives the following relationship:

$$-\int_0^x V(x) = M(x) = EI\frac{d^2y}{dx^2} = -\frac{w}{2}\langle x\rangle^2 - \frac{wL}{2}\langle x - \frac{L-S}{2}\rangle^1 - \frac{wL}{2}\langle x - \frac{L+S}{2}\rangle^1 \tag{2.8}$$

$$\frac{d^2y}{dx^2} = \frac{wL}{2EI}\left(-\frac{1}{L}\langle x\rangle^2 + \langle x - \frac{L-S}{2}\rangle^1 + \langle x - \frac{L+S}{2}\rangle^1\right) \tag{2.9}$$

Integrating Equation 2.9 with respect to x gives the equation for the slope at any point across the beam, as shown in Equation 2.10.

$$\frac{dy}{dx} = \theta = \frac{wL}{4EI}\left(-\frac{2}{3L}\langle x\rangle^3 + \langle x - \frac{L-S}{2}\rangle^2 + \langle x - \frac{L+S}{2}\rangle^2\right) + C \tag{2.10}$$

To determine the value of the integration constant, a point with a known slope is needed. For a symmetric beam, the slope at the center will be zero, so the equation can be evaluated at $x = \frac{L}{2}$:

$$0 = \frac{wL}{4EI}\left(-\frac{2}{3L}\langle\frac{L}{2}\rangle^3 + \langle\frac{L}{2} - \frac{L-S}{2}\rangle^2 + \langle\frac{L}{2} - \frac{L+S}{2}\rangle^2\right) + C \tag{2.11}$$

$$-\frac{4EI}{wL}C = -\frac{2}{3L}\left(\frac{L}{2}\right)^3 + \langle\frac{S}{2}\rangle^2 + \langle-\frac{S}{2}\rangle^2 \tag{2.12}$$

$$\frac{4EI}{wL}C = \frac{L^2}{12} - \frac{S^2}{4} \tag{2.13}$$

$$C = \frac{(L^2 - 3S^2)wL}{48EI} \tag{2.14}$$

A reasonable choice of support spacing is to use the Airy points. The Airy points are the support locations that cause the angular deflection at the ends of the beam to be zero [28]. This deflection profile ensures that the system is flat at the end where it is interfacing with other equipment in the production process. The Airy points for a beam supported at two positions are separated by a distance equal to $\frac{1}{\sqrt{3}}L$. To ensure that this support spacing is acceptable, the Airy points spacing can be plugged into Equation 2.8 to set the requirement for beam stiffness:

$$\frac{d^2y}{dx^2} = \frac{wL}{2EI}\left(-\frac{1}{L}\langle x\rangle^2 + \langle x - \frac{1}{2}\left(L - \frac{1}{\sqrt{3}}L\right)\rangle^1 + \langle x - \frac{1}{2}\left(L + \frac{1}{\sqrt{3}}L\right)\rangle^1\right) \tag{2.15}$$

$$\frac{d^2y}{dx^2} = \frac{d\theta}{dx} = \frac{wL}{2EI}\left(-\frac{1}{L}\langle x\rangle^2 + \langle x - \frac{3-\sqrt{3}}{6}L\rangle^1 + \langle x - \frac{3+\sqrt{3}}{6}L\rangle^1\right) \tag{2.16}$$

Due to the singularity functions included in this equation, it is simplest to find the maximum slope change positions graphically. Because most of the parameters other than x serve primarily as scaling parameters, they can all be normalized to 1 to find the value of x which maximizes the slope change. Accordingly, the following Equation 2.17 is plotted in Figure 2.17 to find the maximum slope change position. The value of L is set to 1 for the normalization, but it is kept in Equation 2.17 to maintain correct units.

$$\frac{d\theta}{dx} = -\frac{1}{L}\langle x\rangle^2 + \langle x - \frac{3-\sqrt{3}}{6}L\rangle^1 + \langle x - \frac{3+\sqrt{3}}{6}L\rangle^1 \tag{2.17}$$

The normalized plot shows that the maximum absolute slope change occurs at the support locations. Because the beam is symmetric, only one of these positions is needed for evaluating this maximum slope change. Using the support position for x in Equation 2.15 creates a relationship between the beam stiffness and the deflection requirement.

$$x = L_1 = \frac{L-S}{2} = \frac{L-\frac{\sqrt{3}}{3}L}{2} = \frac{3-\sqrt{3}}{6}L \tag{2.18}$$

$$\max\left|\frac{d\theta}{dx}\right| = \left|\frac{wL}{2EI}\left(-\frac{1}{L}\langle\frac{3-\sqrt{3}}{6}L\rangle^{2} + \langle\frac{3-\sqrt{3}}{6}L - \frac{3-\sqrt{3}}{6}L\rangle^{1} + \langle\frac{3-\sqrt{3}}{6}L - \frac{3+\sqrt{3}}{6}L\rangle^{1}\right)\right| \tag{2.19}$$

$$\max\left|\frac{d\theta}{dx}\right| = \left|\frac{wL}{2EI}\left(-\frac{1}{L}\left(\frac{3-\sqrt{3}}{6}L\right)^{2} + \langle\frac{-\sqrt{3}}{3}L\rangle^{1}\right)\right| \tag{2.20}$$

$$\max\left|\frac{d\theta}{dx}\right| = \left|\frac{wL}{2EI}\left(\frac{2-\sqrt{3}}{6}L\right)\right| \tag{2.21}$$

$$\max\left|\frac{d\theta}{dx}\right| = \left|\frac{(2-\sqrt{3})wL^2}{12EI}\right| \tag{2.22}$$

$$EI = \frac{(2-\sqrt{3})wL^2}{12(\max\left|\frac{d\theta}{dx}\right|} \tag{2.23}$$

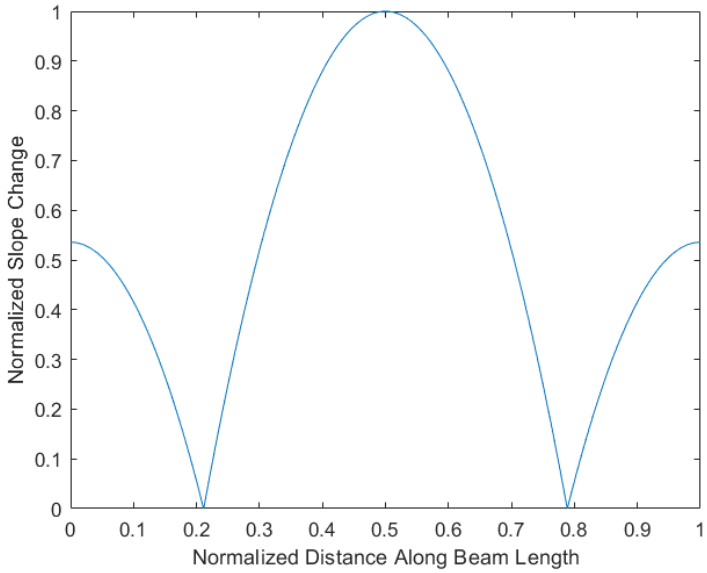

Figure 2.17: The normalized slope change is plotted as a function of the normalized distance along the beam length to find the maximum absolute slope changes. These maximum values match the support locations.

With the general relationship between maximum slope change and table parameters determined, the next step is to start determining the values for these parameters. The first requirement is the maximum allowable slope change, which comes from the gap tolerance between consecutive modules. As discussed above, the maximum gap tolerance is 0.003 in, leading to a maximum slope difference of 0.00056 radians. This value can then be inserted into Equation 2.23, providing the following relationship between table length and required stiffness, shown in Equation 2.24:

$$EI = \frac{(2-\sqrt{3})wL^2}{12(1.7e-4)} \tag{2.24}$$

$$EI = 500(2-\sqrt{3})wL^2 \tag{2.25}$$

The structure length comes from the system layout established in Section 2.6. The distributed load can be estimated by assuming an approximate module volume and preliminary material selection. Using a single module as the basis, given an aluminum structure with 1 in wall thickness, 5 in tall chamber interior height, and a factor of safety of 1.5 to account for additional hardware (such as vacuum and sensing equipment), the chamber contribution to w can be estimated as 728 N/m. The stators weigh approximately 6.6 kg, which over a module unit length of 250 mm contributes 259 N/m. The structure itself will also have some weight which must be included.

The general structural support material selected for the table is extruded aluminum rectangular tubing. Extruded aluminum rectangular tubing has flat and squared sides, a good strength-to-weight ratio, and can be sourced relatively easily. Using tubing rather than solid beams maintains much of the beam's stiffness while removing a significant amount of the beam's weight. Given aluminum rectangular tubing, the beam weight per unit length and the bending stiffness can be estimated as follows:

$$\frac{m_{beam}}{L} = \rho_{Al}(2t_{wall}(h_{beam}+w_{beam}) - 2t_{wall}^2) \tag{2.26}$$

$$I_{beam} = \frac{1}{12}(w_{beam}h_{beam}^3 - (w_{beam}-2t_{wall})(h_{beam}-2t_{wall})^3) \tag{2.27}$$

Given the estimates for system weight and the selection of aluminum extrusion for the table structure, Equation 2.24 can be rewritten to depend on the aluminum extrusion height, width, and wall thickness. Given that this equation will still provide 3 free parameters, there remains flexibility to accommodate other considerations that may affect final design decisions. Given the ability to choose multiple parameters, these relationships are added to a spreadsheet, shown in Table 2.6, to adjust them until acceptable beam stiffness and structural deflection values are achieved. In this system, a 3 in beam width is chosen to match the standard 2x1 t-slotted aluminum extrusion which creates rails for moving modules on top of the built system (see Section 2.8.3). A beam height of 6 in is also chosen to provide a 2:1 ratio of beam height to width. Using these values in the above equations results in a wall thickness of 0.25 in. Accordingly, the structural rails used to create the basis for the machine structure are made from 3 in x 6 in x 0.25 in extruded aluminum beams.

Parameter	Name	Value	Units	Formula
Table structure				
Chamber material		Aluminum		
Chamber density	ρ_{Al}	2700	kg/m^3	
Chamber Width	W_c	350	mm	
Chamber Height	H_c	250	mm	
Chamber Wall Thickness	t_c	25	mm	
Chamber area	A_c	0.028	m^2	$:= (W_c * H_c - (W_c - 2 * t_c) * (H_c - 2 * t_c))$
Chamber mass per unit length	m_c	74.25	kg/m	$:= \rho_{Al} * A_c$
gravity	g	9.81	m/s^2	
Chamber weight per meter	w_c	728	N/m	$:= m_c * g$
Stator mass	m_S	6.6	kg	
Stator effective length	L_S	250	mm	
Stator weight	P_S	65	N	$:= m_S * g$
Stator load contribution	w_S	259	N/m	$:= P_S / L_S$
Safety factor	n	1.5		
Machine load	w_M	1481	N/m	$:= (w_c + w_S) * n$
Structure load				
Structure material		Aluminum		
Structure beams	n_{beams}	2		
Structure height	h_{beam}	150	mm	
Structure width	W_{beam}	75	mm	
Structure thickness	t_{beam}	6	mm	
Structure area	A_{beam}	0.003	m^2	$:= h_{beam} * W_{beam} - (h_{beam} - 2 * t_{beam}) * (W_{beam} - 2 * t_{beam})$
Structure mass per unit length	m_{beam}	6.9	kg/m	$:= \rho_{Al} * A_{beam}$
Structure weight contribution	w_{beam}	68	N/m	$:= m_{beam} * g$
Total beam load	w_{tot}	1549	N/m	$:= w_M + w_{beam}$
Beam Young's Modulus		69	GPa	

Parameter	Name	Value	Units	Formula
	E_{beam}	6.90E+10	Pa	
Beam inertia		7.30E+06	mm^4	$:= 1/12*(W_{beam}*h_{beam}^3)-(1/12*(W_{beam}-2*t_{beam})*(h_{beam}-2*t_{beam})^3)$
	I_{beam}	7.3E-06	m^4	
Beam stiffness	EI_{beam}	5.03E+05	$N\text{-}m^2$	$:= E_{beam} * I_{beam}$
Beam length	L_{beam}	7	m	
slope	θ_{max}	3.37E-03		$:= (2 - SQRT(3)) * w_{tot} * L_{beam}^2/(12 * EI_{beam})$
Sealing height	H_{seal}	450	mm	
gap	gap_{table}	1.51464	mm	$:= \theta_{max} * H_{seal}$
		0.059631	in	
Fine adjustment		0.2	degrees	
	θ_{adj}	0.003	rad	
Gap adjustment	gap_{adj}	1.6	mm	$:= \theta_{adj} * H_{seal}$
Acceptable gap	gap_{rq}	0.003	in	
		0.0762	mm	
Design acceptance		TRUE		$:= gap_{rq} > (gap_{table} - gap_{adj})$

Table 2.6: The relationships derived between the beam parameters and the structure stiffness are added to a design spreadsheet to calculate the resultant maximum expected deflection between two consecutive modules. This spreadsheet allows a user to see how changing the beam geometry affects the resultant deflection, enabling parameter selection. The slope fine adjustment mechanism described in Section 2.8.3 loosens the requirement on the structural deflection while ensuring the overall system will meet the gap requirement.

The system designed in this work is a total of 7 m long, as seen in Table 2.6. Based on this overall sizing, aluminum rectangular tubing of length 12 ft and t-slotted extrusions of 8 ft were used. Each of these lengths are standard sizes that were easily sourced. The assembly of this structure is described in Chapter 6.

2.8.3 Rail Slider System

Because the machine components are aligned relative to each other, constraining them to specific positions on the base structure can create some risks during assembly. If there is any difference between the distance between consecutive module centers due to their own machined geometry compared to alignment features built into the machine base, these constraints will be in conflict. If the constraint based on connecting consecutive modules is stiffer, then the features built into the base could be stressed or deformed, damaging the structure. If the constraint positioning the modules on the base is stiffer, then the modules could either be forced to deform or have gaps between them. If these subsystems have similar

stiffness values, then both the modules and the base structure could be stressed or deformed. These risks can be mitigated by loosening the constraints on one set of features; in this case, the constraint to the base structure is loosened to ensure the vacuum seal between module faces is not compromised.

Additionally, given the modular nature of the system, it is beneficial to include a feature that makes it easier to reconfigure the system. This reconfigurability also makes assembly easier. For this system, the modules are mounted on the machine base using linear rails that allow the modules to slide back and forth. This sliding makes it easy to install and move modules on each side of the system, and it does not strictly constrain the module positions. These rails do not need the precision of commercial linear bearings, as those bearings tightly constrain their slider positions. As established above, including this tight constraint could lead to any error propagation stressing the linear bearings and degrading their life. A lower-cost linear sliding system exists with extruded t-slotted aluminum with Teflon sliders. The Teflon sliders have about a 0.4 mm clearance between the slider width and the corresponding slot width on the extruded t-slotted aluminum, accommodating minor positional offsets between modules.

The modules are mounted to the table using this rail slider system. This system allows the modules to move to their final positions after being placed on the table structure. This system also allows modules to be separated for smaller sample set testing in both development and debugging. The sliders on the rails additionally provide accommodation for system expansion and contraction as it heats and cools without creating thermal stresses on the individual components. The connection between the base plates and the rail slider system can be seen in Figure 2.18.

The slider assembly on the baseplate serves two functions. First, the assembly attaches the linear plastic sliders to the baseplate, allowing the module to mount on the table rail system. This assembly also includes a locking bolt to fix the module in place as necessary. Second, the slider assembly includes an adjustment screw for module orientation adjustments. These adjustments can help compensate for angular misalignment between consecutive modules, whether resulting from the modules themselves after manufacturing or created by the table deflecting under load. These adjustment screws can compensate for up to 0.25 in vertical displacement, which translates to 2 degrees of angular misalignment parallel to the tunnel structure and 1.1 degrees of angular misalignment perpendicular to the tunnel structure. This adjustability in the direction parallel to the tunnel can account for angular misalignment created by deflection in the table, relaxing the requirement on the table stiffness. The 2 degrees of adjustability corresponds to 15.7 mm of gap width at the top of a door unit seal. Because the slope differential may need to be made up over multiple consecutive units, this full width cannot be used to relax the table stiffness requirement. Thus, only 10% of this adjustability is allotted to the stiffness requirement, loosening the gap tightness by 1.57 mm (or 0.06 in). This allocation appears in Table 2.6.

2.9 Turn Module

Most of the machine consists of a long tunnel of consecutive modules, but these modules are only designed for the straight sections of the machine. These modules do not work for the

corners, where the tunnel direction changes. These corners are located at each end of the system. On the loading end of the system, no modules are used because this end is exposed to interface with other equipment such as a vial filling machine. The other end of the system is part of the drying section, so it needs to be sealed for sublimation.

The standard module design does not accommodate corners because the machine chamber extends from the sides of the stators. This extension would interfere with any stators placed directly next to the long tunnel lines. Thus, a special module must be designed to accommodate corners. This module must cover at least two stators in the corner to cover the tunnel change in direction and maintain sealing surfaces like the standard modules. If the tunnel width at the end is at least four total stators long, then a corner specific geometry can be used for the two stators on each corner which can then be bolted together in the direction perpendicular to the long tunnel. If the tunnel width is more than four stators long, then standard modules can be added to the end to complete the loop. If the tunnel width is less than four stators long, then a single module can be used to cover the entirety of the short end. Because the system built for this work has an end length that is only three stators long, it uses a single module to cover these three stators. This module, called the turn module, is shown in Figure 2.19.

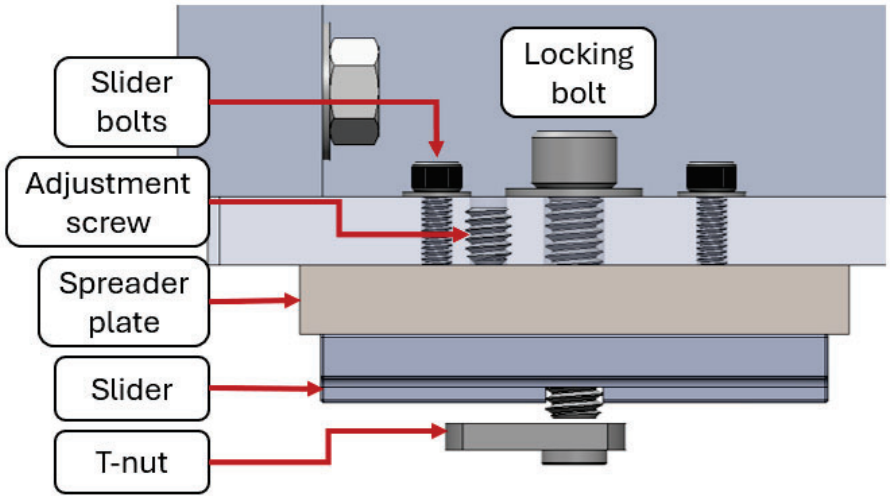

Figure 2.18: The labeled parts are assembled as shown to connect the sliders for the rail system to the module base plate. The slider is a standard t-slotted framing plastic slider, and the t-nut corresponds to the t-slotted framing for locking. The slider bolts connect the slider to the base plate, while the spreader plate spaces the base plate off the slider to create room for bolt heads and electrical connections on the bottom of the baseplate. The adjustment screw can lift the baseplate off the slider for angular adjustments on the rail system. The locking bolt tightens on the t-nut to lock the module's position on the rail system.

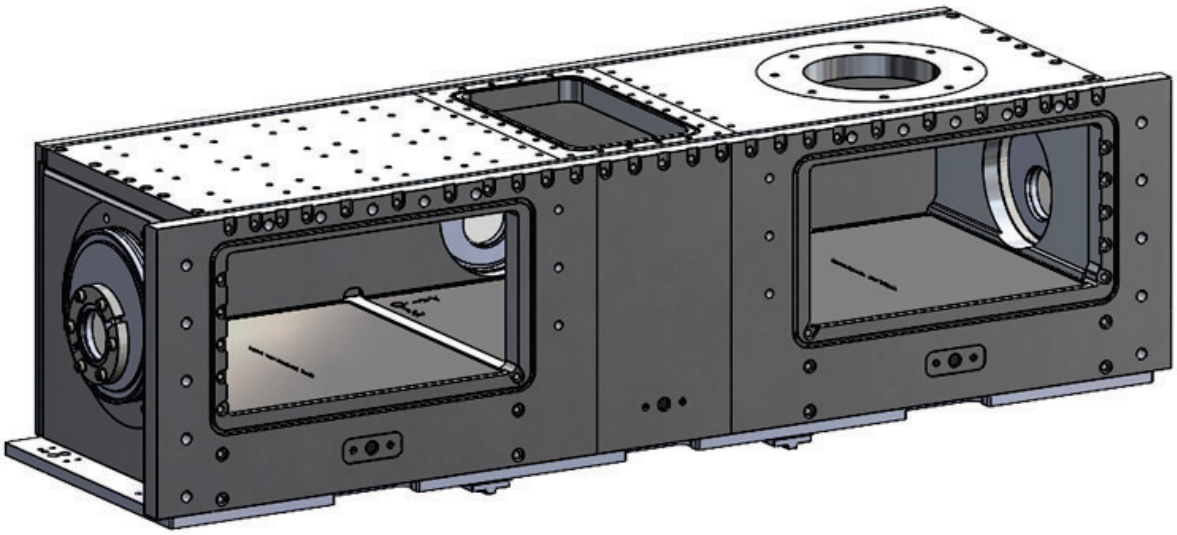

Figure 2.19: The turn module is a custom module which covers the short end of the system in the drying section. Because the bolt flanges would interfere with stators not in the direction of the long tunnel, a custom module is needed for the point on the system where the travel direction changes. This module is assembled from machined plates to reduce waste material in its manufacture.

2.10 Future Work

This work can be further validated by continuing to manufacture and assemble additional modules to make a larger continuous lyophilizer. This expansion can validate the scaling relationships described in this chapter, as well as the tolerancing done to ensure units produced at larger scales will continue to assemble and operate successfully. Additionally, using different formulations with varying required lyophilization profiles can validate the design spreadsheet and its ability to predict the required machine dimensions based on process times. The motion system could also be customized specifically for this application, which could improve the vacuum sealing between the motion system and the machine chamber. This customization could come as a partnership with a commercial manufacturer, such as Planar Motors which was used in this work. A custom motion system could enable changing the width of the machine tunnel, creating a smaller internal volume which could improve internal uniformity. An additional module could be introduced to create a load-lock unit which can hold a vacuum environment with ambient pressure on both sides, enabling more flexibility in system layouts and designs. Minor improvements, such as changing O-ring paths to include larger radius corners and optimizing the bolt patterning for consecutive module face seals, can be implemented based on the performance of the manufactured system.

In an industrial setting, the trade-off between productivity and floor space is crucial. Thus, one of the most critical parameters for this machine is its expected throughput to surface area ratio. This ratio can be improved by increasing the production rate for the existing machine geometry, reducing the machine geometry while maintaining the same pro-

duction rate, or a combination of the two. One improvement method is to stack tunnel systems vertically, as this does not increase floor area coverage while multiplying production rate. These systems can also be operated in parallel lines, removing the complexity of the turn module while keeping stators between pairs of parallel lines to ensure that they can maintain the looping motion required for their reuse. The internal vial density can be improved by decreasing the vial spacing on each tray, either by adding more vials to each tray or by decreasing the tray size while maintaining the same number of vials. This increased vial density given the existing geometry could affect vial uniformity and therefore increase required process times to ensure that every vial on each tray is ready for removal. This uniformity is most strongly influenced by the number of rows of vials across the tunnel width - accordingly, using methods such as decreasing the tunnel width or dividing the surface area into multiple tunnels would enable maintaining the existing system uniformity.

Chapter 3

Freezing Module

The freezing stage of lyophilization sets the basis for the overall process speed and reliability. While it accounts for less than 10% of the total process time, variation in the freezing stage can create variations of as high as 25% [29][30] in the total process time. This variation in process time can lead to large inefficiencies within a system, where some vials finish sublimating much sooner than others. If these vials are grouped together, either on a shelf in large batch lyophilizers or on a smaller tray as used in the system described in this work, then those groups will have to wait for the slowest drying vials before they can exit the system. Thus, it is critical to ensure that the vials undergo a freezing process that is as uniform as possible.

The freezing process also affects product cake stability during drying [31]. During the freezing process, the product is separated from most of the water, as the ice crystallizing separates from the solution, causing the product to concentrate into its cake structure [32]. This separation allows the ice crystals to sublimate when exposed to a vacuum, removing most of the water from the system. These ice crystals also determine the cake structure, as it is formed around the crystallized water. These structural differences can cause variance in drying times, as the relative pore sizes left behind after the ice crystals sublimate influence the cake's resistance to further sublimation.

Given the importance of cooling conditions and uniformity, the freezing system must successfully create and control the thermal conditions required for each stage of the cooling process. There are three stages in this process: conditioning, where the fluid in the vials is supercooled to a few degrees below its freezing temperature; nucleation, where the supercooled solution is exposed to a controlled stimulus to begin crystallization; and freezing, where the crystal growth occurs and the solution is fully solidified in preparation for sublimation. Each of these process steps has different thermal requirements, so they do not necessarily all occur in a single chamber. Additionally, this cooling process must be compatible with the continuous nature of the lyophilizer developed in this work. In particular, the motion associated with a continuous system limits the use of the common shelf-based cooling used in existing batch lyophilizers. This chapter provides an overview of the design and evaluation of a convective gas flow cooling system used to freeze vials in the continuous lyophilizer developed in this work.

This convective gas cooling system is driven by a liquid nitrogen source tank and a series of actuators used to generate the required gas speed. The actuators in this system are muffin

fans. These fans can be placed directly at the chamber inlets and outlets, enabling ductwork to be localized to each cooling chamber. This ductwork includes a recycle system to reduce gas consumption and increase energy efficiency. The recycle system consists of three key elements: two pieces of recycle ductwork, called ram's horns, and a diffuser. The ram's horns direct the gas flow into the chamber inlets, while the diffuser takes gas exiting the chamber and redirects it back into the ram's horns. A mixture of fresh liquid and gaseous nitrogen is injected into the ram's horns to maintain the chamber temperature setpoint, while the diffuser includes an exhaust port to allow gas to exit the system, maintaining constant pressure within the chamber.

The cooling chamber is installed in the continuous lyophilizer in two sections. One section covers the conditioning process, while the other covers the freezing process. Using aqueous surrogate sugar solution samples, the conditioning section operates at a -3°C setpoint, while the freezing process operates at a -40°C setpoint. These setpoints can vary for different formulations. The chambers showed their ability to hold these setpoints within 1°C error, and the cooling times for a tray of vials are shown to be about 20 minutes in each section.

Controlled nucleation is achieved using a thermal quench method. The vial tray is moved from the conditioning section to the freezing section, creating an environmental temperature change of nearly 40°C in approximately 30 seconds. This nucleation is shown to induce freezing in a full set of vials in less than 30 seconds at sample temperatures within 1°C of each other. These combined results demonstrate the cooling chamber's effectiveness and ability to be utilized in a continuous lyophilization system.

3.1 Background

Most freezing in commercial lyophilization is done through shelf temperature control. The vials rest on thermally conductive shelves that follow prescribed cooling profiles, which are typically temperature ramps and holds [30]. Because the continuous lyophilizer developed in this work uses the suspension-based weight sensing system described in Chapter 4, the vials are not in contact with a shelf. Thus, controlling the temperatures of the surfaces below the vials does not have the same direct influence on their heat transfer rates in this system as conduction through the shelf does in traditional lyophilizers. Consequently, other cooling methods for use in this continuous lyophilizer have been investigated and developed.

An alternative, rapid form of freezing is quench freezing using a cold solid or liquid, such as liquid nitrogen. The vials containing the product are partially submerged in the cold material, rapidly freezing the sample. This method freezes the product very quickly, but it also tends to freeze vials heterogeneously [31]. This heterogeneity is likely related to the uncontrolled nucleation that occurs during this rapid freezing process. While this freezing method would significantly reduce the freezing time relative to traditional lyophilizers, it does not provide the necessary freezing uniformity required to ensure consistent product quality. Thus, it is not pursued for this continuous lyophilizer.

Another freezing method is known as spin freezing. This method involves rotating the vial at high speed so that centrifugal forces cause the fluid to spread out and coat the entire interior surface of the vial with a thin film. The spinning vial is then exposed to a cooling environment, such as a cold gas jet, freezing the film inside the vial [33]. A figure depicting

this process is shown in Figure 3.1. Spin freezing requires individualized actuation for each vial to perform the spinning, which adds significant complexity to the system. Additionally, it significantly alters the cake structure from that of traditional lyophilizers. Thus, spin freezing is not used in this system.

Although the full spin freezing process is not desirable due to the spinning element, the cold gas jet element is practical for the continuous lyophilizer. Forced gas convection provides a non-contact cooling method for freezing the vials that is compatible with the Planar Motors magnetic levitation system used to transport vials through the system developed in this work. An inert gas is cooled to the target temperatures for the freezing process and blown through the vials to reduce their temperature through convective cooling. The vials can move freely through this gas flow regime without requiring additional actuation or generating potential contaminants. The gas temperature can be controlled to create cooling profiles similar to those used for existing shelf-based conduction systems, but the gas can also change temperature more quickly because the gas has less thermal mass than the shelf. An example forced gas convection system for lyophilization is described in [35], and a basic diagram is shown in Figure 3.2. Because of the controllability and compatibility of forced gas convection for cooling in the continuous lyophilizer, this method is pursued for this system. Nitrogen is used as the convective gas in this system because it is inert and relatively easy to source.

Radiation could also be investigated as a potential cooling method in a pharmaceutical lyophilizer. However, radiation requires much larger temperature gradients than convection, or it will be a much slower cooling process. To obtain similar cooling rates, radiative cooling would require cooling surfaces much colder than the target temperatures of each freezing step. This difference between target temperature and cooling source temperature would create a notable risk of overcooling, which can be highly problematic during conditioning. If the vials are overcooled during conditioning, the risk of spontaneous nucleation increases. Spontaneous nucleation would introduce freezing heterogeneity, leading to variation in the sublimation results. Thus, radiative cooling is not pursued for this system.

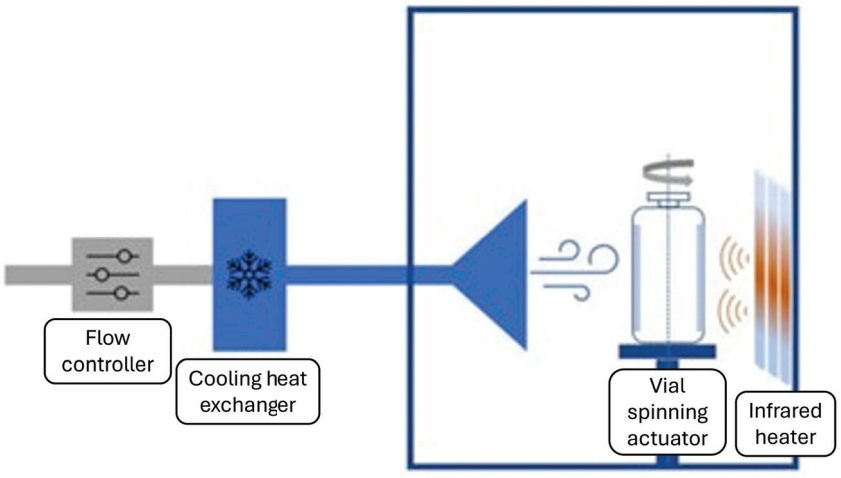

Figure 3.1: The spin freezing process includes a flow controller, a heat exchanger for cooling the gas to the target temperature, the actuator which spins the vial, and an infrared heater for further vial temperature control. Schematic adapted from [34].

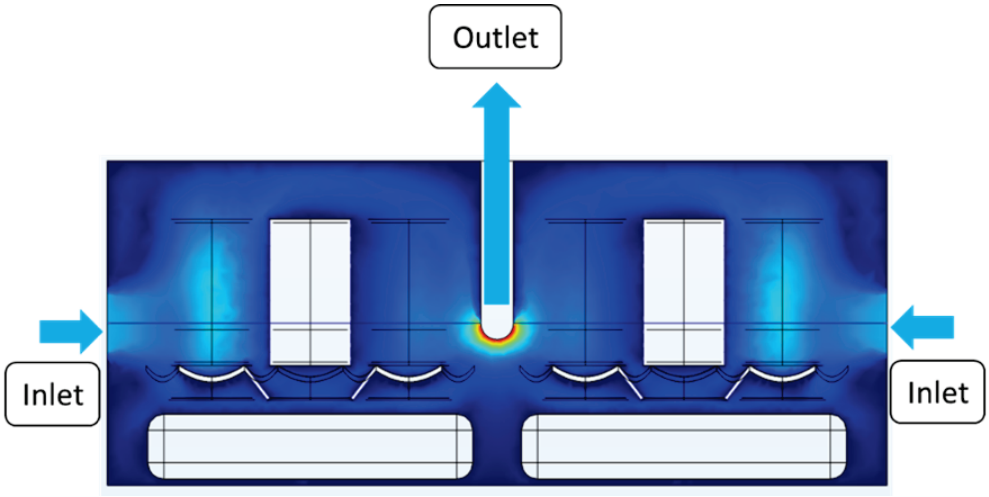

Figure 3.2: A basic forced gas convective chamber diagram. The gas is injected into the chamber through the inlets, travels through the vials to cool them down, then is removed from the chamber through the outlet.

3.2 Freezing Stages

The freezing process is divided into three stages: conditioning, nucleating, and freezing. Each of these stages has its own time and sensitivity requirements. For simplicity in this system, all stages are designed using the same hardware. This decision reduces the complexity of development and production of the parts required to build the full freezing subsystem, but future designs can investigate the use of different cooling methods in these subsections to further optimize the system. Potential heterogeneous design decisions for future development are discussed at the end of this chapter in Section 3.6.

3.2.1 Conditioning Section

The conditioning stage consists of supercooling the product to prepare it for nucleation. This supercooling process aids in controlling the nucleation, as opposed to allowing the product to spontaneously nucleate. The lower the supercooling temperature, the more vulnerable the solution is to spontaneous nucleation, and the more carefully the system must be controlled. However, higher degrees of supercooling can make it easier to induce nucleation once a controlled stimulus is applied to the vials. Thus, the higher the fidelity of the temperature control of the conditioning stage, the lower the supercooling temperature that can be achieved, and the greater the control over the nucleation process. In the supercooled regime, temperature spikes can cause ice nucleation. If the temperature is not sufficiently under control, these spikes could occur before all vials are sufficiently supercooled, negatively affect the rest of the lyophilization process by creating nonuniformity between the ice crystal growth in the vials.

When vials are allowed to spontaneously nucleate, they can be supercooled to temperatures as low as -20°C [36]. When the product in the vials freezes at these low temperatures, it crystallizes quickly, leading to smaller ice crystals and slower primary drying rates [30].

The drying time is reduced because the smaller pores created by the crystals in the resulting dried cake increase the resistance to mass transport to sublimate the ice. Controlled nucleation methods can loosen the subcooling requirement to temperatures as high as -3°C, enabling slower ice crystal growth. This slower crystallization leads to larger ice crystals, which results in faster primary drying rates. Because the lyophilizer developed in this work uses a controlled nucleation method, the conditioning temperature target is expected to be between 0 and -5°C.

3.2.2 Nucleating Section

The nucleating stage consists of initiating the ice crystal formation in the vials. The simplest nucleation method is spontaneous, uncontrolled nucleation. In this method, the vial temperatures are simply driven to the desired freezing temperature, and each vial will spontaneously nucleate at some point as it cools. Because each vial will nucleate at different temperatures and times, they will experience different crystal growth rates and sizes, leading to nonuniform sublimation processes. This variation can introduce inefficiencies because the vials may require different amounts of time in each primary and secondary drying stage, but they are limited in moving between stages based on the vials around them. Additionally, the different pore sizes in the resultant cakes can affect their stability during drying and their resultant ability to be reconstituted after drying [37].

Controlling ice nucleation ensures uniform crystal growth within the product solution across vials. Improving vial-to-vial freezing homogeneity ensures that the vials finish sublimating in similar amounts of time, preventing inefficiencies that arise from the need to leave vials in the system after they have completed the lyophilization process. Additionally, controlled nucleation enables consistent crystal growth at warmer temperatures, leading to larger ice crystals and faster drying times [38]. Many methods are currently used for controlled nucleation in pharmaceutical lyophilizers. One such method involves using vacuum induced surface freezing (VISF). In this process, the pressure in the lyophilizer is briefly lowered to about 100 Pa, leading to ice nucleation on the surface of the solution in the vial, which then propagates through the fluid volume. This process has demonstrated good uniformity in the nucleation temperatures for vials within a lyophilizer [36]. Introducing VISF to a continuous lyophilizer requires adding a vacuum chamber between the conditioning and freezing sections of the system. This chamber addition makes VISF less desirable than other potential nucleation methods, but it remains viable if needed due to its proven robustness on existing lyophilizers.

Another nucleation method is ice fog. The ice fog method involves injecting small ice crystals into the lyophilizer's freezing chamber. These ice crystals will fall into the vials, creating a nucleation point for the fluid [39]. This nucleation method requires the introduction of additional material to the system, and it is still somewhat stochastic as it depends on the ice fog propagating through the lyophilizer volume. Thus, the ice fog nucleation method is not pursued for this continuous lyophilizer.

An additional nucleation method is depressurization induced nucleation. In this method, the pressure within the system is increased during the conditioning stage. Then, to initiate nucleation, the chamber is rapidly depressurized, creating a similar pressure change within the chamber to dropping from ambient pressure to a vacuum environment [39]. Unlike

VISF, this nucleation method does not require generating a vacuum environment, but it does require the ability to pressurize the freezing chamber and to maintain that pressure during the freezing process. This pressurization would hinder the ability of the system to create the forced gas convection flow, as it depends on the chamber pressure being low enough for the gas flow generation system to direct the gas through the chamber and across the vials. Thus, this depressurization nucleation method is not pursued for this continuous lyophilizer.

This work utilizes a nucleation method not commonly used in commercial lyophilization called thermal shock or thermal quench, described by [35]. The ability to rapidly change the gas temperature allows the system to drop the environmental temperature to which the vials are exposed much faster than traditional shelf-based systems. This sharp drop in driving temperature induces spontaneous nucleation at the outer edge of the solution in the vial before thermal convection within the vial can equilibrate the solution temperature. Thus, the bulk solution temperature remains relatively constant, and nucleation occurs quickly. The results of using this thermal shock nucleation appear in Section 3.5. This thermal shock method is like the nucleation induced by quench freezing. However, this method allows for control of the degree of subcooling, while full quench freezing does not.

3.2.3 Freezing Section

The freezing section consists of further lowering the product temperature well below water's triple point. The product temperature must be reduced below the product collapse temperature to ensure that the product remains stable during sublimation [29]. The target temperature for freezing aqueous solutions is typically between –50 and –60°C. The temperature control precision during freezing is not as critical as it is during conditioning, as the primary requirement is to ensure that there is sufficient buffer between the vial temperature and the product collapse temperature to maintain stability during sublimation. If the vial temperature is colder than the target temperature, that condition does not have significant negative consequences, such as the risk of spontaneous nucleation during conditioning.

The freezing section can sometimes include an annealing stage to ensure slow crystal growth in the vials. This annealing process consists of freezing at a temperature similar to the subcooling temperature after inducing nucleation [30]. This annealing can be incorporated into the continuous lyophilizer by adding a specific annealing section between the nucleation and freezing sections.

3.2.4 Freezing Modeling

The vials in forced gas convection can be modeled as cylinders in cross flow, as shown in Figure 3.3.

A collaborator ran numerous simulations on the convective cooling of vials of water in geometries relevant to this work [40]. His work shows that the conditioning and freezing stages should take about 30 minutes each, with a time constant of about 10 minutes. Based on this simulation data, an effective convective heat transfer coefficient can be found using the Equation 3.1.

$$\tau = \frac{\rho V c}{h A_s} \tag{3.1}$$

where ρ is the material density, V is the material volume, c is the material's specific heat, h is the effective heat transfer coefficient, and A_s is the material surface area. The glass vial is assumed to have thin enough walls such that the water in the vial dominates the heat transfer, resulting in an expected heat transfer coefficient of about 23 $W/m^2 - K$. This value is reasonable for this forced convection cooling system.

While the specific heat transfer coefficient varies across the surface of the vial because the gas speed changes, the process modeling can use an effective lumped coefficient to predict cooling behavior. This model is only valid during the periods without phase change in the vial, when the latent heat requirement would prevent temperature change. Thus, the equation can be used to predict cooling times during conditioning and freezing times after crystal growth is complete.

While vials can be modeled individually based on cylinders in crossflow, the vials do not exist independently in the cooling chamber. This modeling can be valid for a single row of vials spaced sufficiently far from each other, but it starts to fail when there are multiple rows of vials. As discussed in Chapters 2 and 4, the lyophilizer designed in this work uses up to six rows of vials across its tunnel. Given multiple rows of vials, the cold gas flow needs to be directed so that it can effectively cool all vials simultaneously.

There exist two choices in setting up the internal geometry for the cooling chamber. The first choice is the vial row arrangement, and the second choice is the gas inlet and outlet positions. Due to the complex fluid mechanics resulting from the flow interacting with multiple vials, combinations of these arrangements and positions are evaluated using simulations, shown in [40]. A brief summary of the available choices, their anticipated effects on system performance, and the key simulations results and their implementation on this system are discussed below.

For vial arrangement, the vials in consecutive rows can be positioned directly in front of each other, or they can be offset by some amount. The vial rows can also be spaced apart from each other by different amounts. In traditional production lyophilizers, the vial rows are hexagonally close-packed to maximize vial density on a shelf. While this arrangement has the highest possible number of vials per unit shelf area, it does not leave space for cold gas to move around the vial to cool it. Thus, the hexagonal close-packing is not used for this system. The simplest vial arrangement would be an evenly spaced grid, creating space

Figure 3.3: The vials in the forced gas convection chamber can be modeled as cylinders in crossflow.

between the vials for gas flow. An alternative arrangement would be to offset consecutive vial rows so that the vials in one row correspond to the gaps in the neighboring rows. This arrangement would be expected to improve cooling performance when the gas flow direction is perpendicular to the row offset direction, as the vials would not be placed directly in the wake of the vials in the row in front of them. These two arrangements are simulated in [40] with different gas inlet and outlet positions to determine which lead to faster and more uniform cooling performance. These simulations showed that offset rows spaced nearly one vial diameter apart from each other demonstrated the fastest and most uniform cooling, so this arrangement is implemented in this system.

For the gas inlet and outlet positions, the available surfaces are limited by the tunnel geometry and the vial motion system discussed in Chapter 2. The tunnel structure limits the use of cooling gas on the inlet or outlet faces of the tunnel for a given module, as features on these faces would block the ability for vials to enter and exit the module. Cooling features also cannot be located on the bottom face of the module, as the Planar Motors motion system is a solid structure that covers the entire module bottom face. Thus, the cooling features can only be placed on the module's side walls and the ceiling surface. Given these surfaces, there are three primary cooling feature arrangements that are evaluated for this system. The first, simplest arrangement uses a gas inlet on one side wall and an outlet on the opposite side wall, where the gas flows unidirectionally across an entire chamber. The second arrangement uses both side walls as gas inlets, where the outlet is then moved to the chamber roof. The third arrangement is the opposite of the second, where the gas inlet is on the chamber roof and the outlets are moved to the side walls.

The simulation work showed that the more vials the cooling gas needed to pass through, the less uniform the cooling rates between the fastest and slowest cooling rows [40]. This result is expected, as with each row the gas flows past, the more energy the gas absorbs from the vials and warms up, decreasing its heat transfer from later rows. This result indicates that the first arrangement of blowing gas across the entire chamber should not be used, as it requires blowing the gas across the highest number of vial rows.

When comparing the second and third arrangements, the initial results showed that putting the gas inlets on the sides of the chamber and the outlet on the chamber roof performed better than having the gas inlet on the chamber roof and the outlets on the chamber walls. This result is reasonable, as the solution to be lyophilized does not completely fill the vials, so the volume to be cooled is closer to the bottom of the chamber than the ceiling. Thus, when the gas is blown from the walls at the height of the fluid in the chamber, it interacts with the fluid in the vials sooner than if it is blown from the ceiling. Accordingly, the cooling chamber designed for this work uses two inlet ports, one on each side of the chamber, and one outlet port located on the chamber roof.

3.2.5 Freezing Control

Many factors require consideration to effectively control the freezing process. First, the machine must maintain the low temperatures used in the freezing process. This temperature control consists of generating and holding consistent low temperatures without compromising the chamber at these conditions. The continuous nature of the machine ensures that the sections responsible for freezing maintain constant temperature setpoints during operation;

however, the components will undergo cycling when the system is shut down for cleaning or maintenance. Thus, the components must survive temperatures ranging from -100∘C, that could be reached during thermal quench nucleation, to 100°C, that could be reached during steam sterilization, in low frequency cycles. This requirement limits the material selection for parts used in the freezing section. This temperature range also raises concerns about the potential for differential thermal expansion between components made of different materials in this system. To mitigate the differential thermal expansion concerns, the structural elements of the freezing system are all made of the same material. The structural elements are made from aluminum in the system built for this work. In further industrial development, the material could remain aluminum, or it could change to stainless steel to match what is commonly seen in pharmaceutical production equipment.

The temperature of the forced convection gas is monitored using thermocouples placed in the gas flow. These thermocouples are placed at the cooling chamber inlets to measure the temperature of the gas as it is exposed to the vials. The gas temperature is controlled through a metered fresh feed of liquid nitrogen. The liquid nitrogen is introduced to the convective gas feed upstream of its injection into the cooling chamber. The liquid nitrogen's latent heat of evaporation rapidly cools the gas, rapidly reducing the cooling gas temperature to the target setpoint. This rapid temperature change allows for a smaller heat exchange volume, shrinking the overall size of the equipment. Additionally, since the convective gas is also nitrogen, the use of liquid nitrogen does not introduce a new material into the system. Because liquid nitrogen is significantly colder than the gas temperatures required in any stage of this lyophilizer, this cooling control method is utilized for all three cooling stages.

3.2.6 Recycle

In a basic forced gas convection layout, the cold gas is only blown past the vials once through the cooling chamber. This process is thermally inefficient because the short time during which the gas and the vials are exposed to each other is not sufficient for both temperatures to equilibrate, particularly at the beginning of the cooling process. Rather than simply ejecting the waste gas, the used gas is reintegrated into the injected gas flow using a recycle system. Using this recycle system ensures that fresh gas does not need to be cooled all the way from ambient temperatures to the system setpoint; instead, gas which is already close to the system setpoint can be reduced back to the target temperature, requiring much less cooling power. The recycle system also reuses the gas mass, which saves on raw material costs. The recycle system includes a metered inlet and a passive outlet. The metered inlet is used to control the cooling provided to the recycled gas flow to maintain the setpoint temperature within the chamber. The passive outlet ensures that mass leaves the system to account for the mass added through the metered inlet, preventing pressure buildup within the chamber.

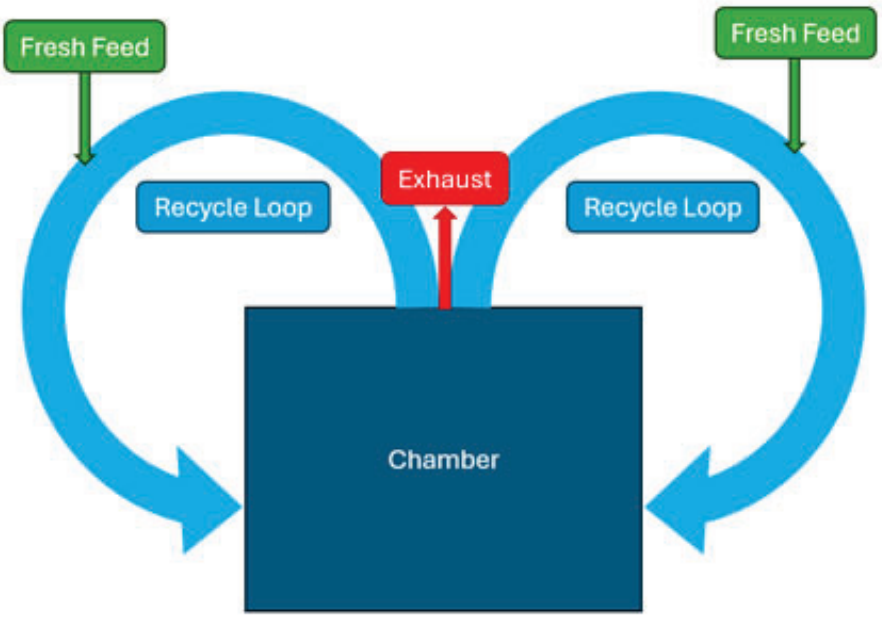

Figure 3.4: The recycle system consists of the cooling chamber, the fresh gas feed, the exhaust, and the recycle loop. Most of the gas which exits the chamber recirculates through the recycle loop to be reused in the system. A small amount of fresh cold gas is fed in to maintain the temperature setpoint, while an equal amount of gas is exhausted to maintain constant mass within the system.

3.3 Freezing Hardware

3.3.1 Module Evolution

From the beginning, the cooling unit was designed with modularity in mind. When building a new process platform, flexible hardware enables exploration of a wider design space than the final unit requires. Accordingly, the cooling module consists of 4 main subsystems: the gas flow generation, the cooling chamber, the recycle system, and the fresh feed injection system. The recycle system consists of the ram's horns which direct flow into the sides of the cooling chamber and the diffuser which splits the outlet flow from the chamber back into the ram's horns on each side of the chamber. Each of these subsystems is built separately before being assembled into a complete module. A complete cooling module is shown in Figure 3.5.

3.3.2 Gas Flow Generation

The forced gas convection system is highly dependent on the mechanism used to generate the gas flow. There are two main methods for using actuators to generate this gas flow. Actuators are preferred for flow generation because they maintain consistent performance over longer periods of time, as opposed to a pressure source that drains over time. The actuated flow can be generated centrally with a large blower which creates a pressure head that drives flow through all the chambers, or the flow can be generated locally with small fans positioned

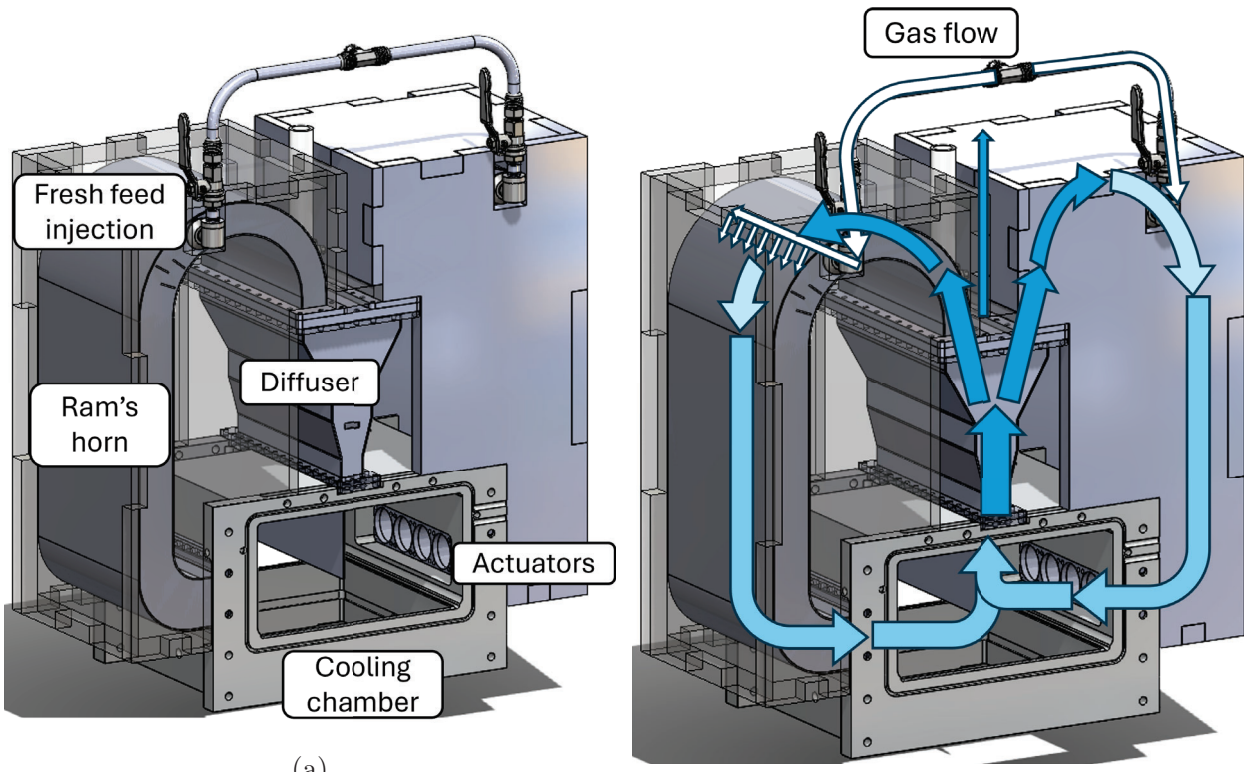

(a) (b)

Figure 3.5: The cooling module consists of 4 main subsystems: the gas flow actuators, the cooling chamber, the recycle system, and the fresh feed injection system. The recycle system consists of the ram's horns which direct flow into the sides of the cooling chamber and the diffuser which splits the outlet flow from the chamber back into the ram's horns on each side of the chamber. The gas flow path is shown in the right image, where lighter colors represent colder gas.

on each chamber. The large blower requires less actuators, but its implementation requires more complex piping to distribute the gas flow among all the chambers. This piping becomes particularly unwieldy when integrated into a system which recycles the used gas, where the flow must be redirected back to the blower from all the chambers. Local flow generation requires many more actuators to generate the flow in each chamber, but the recycle system can be shrunk to operate on each system independently. The local flow generation strategy is used for this system because it lends itself particularly well to modularity, as each recycle system is fully integrated into a single module.

Once the flow generation strategy is chosen, the specific flow generation hardware must be selected. The local flow generation strategy requires small actuators for driving gas flow. Muffin fans are a good choice for this flow actuator because they have been highly developed for cooling in the computer industry. These fans also use sealed bearings which protects them from introducing contaminants to the gas flow.

The fans used in this system are 40 mm x 40 mm x 28 mm muffin fans rated for a volumetric flowrate of 35 cubic feet per minute, shown in Figure 3.6. Over the 40 mm x 40

mm area, this volumetric flowrate translates to about a 10 m/s linear flow velocity. Prior simulation work indicated that to achieve conditioning and freezing times under 30 minutes, a linear flow rate of at least 1 m/s is needed [40]. Choosing fans which could provide an order of magnitude faster flow in ideal conditions increases the likelihood that the fans will maintain required performance when operating in the non-ideal conditions created by the chamber hardware, such as head losses created by the gas recycle ductwork. In situ, the gas speed at the chamber inlet, in front of the fans, measured between 1 and 4 m/s. This measurement indicates that the selected fans perform to the system specification and can be expected to provide the required flow velocity to cool the vials sufficiently quickly.

One risk in using muffin fans positioned directly on the chambers is that they will be directly exposed to temperatures as low as -100°C during the freezing chamber operation. Because the fans are at the chamber, the gas passing through them needs to be as cold as the temperature requirements within the chamber. If the actuators were further from the chamber, such as with a central blower, then the gas could warm up on its way back to the blower before being cooled down again as it returns to the chamber. While this central blower system would be less energy efficient, it would not risk exposing the blower to temperatures outside of its rated operating range. If the actuators become too cold, the lubricant in their bearings could stiffen, causing them to prematurely wear out and fail. However, the coils in the fans also generate heat while they are operating, which can counter the cooling from the gas. This heating is expected to compensate for the heat transfer to the gas so that the fans are not overcooled. During experimental testing and operation, the fans were exposed to temperatures as low as -80°C for nucleation experiments during short (30 s) intervals, shown in Section 3.5.3 and as low as -50°C for extended periods during continuous operation, as shown in Section 3.5.

3.4 Cooling Chamber

The cooling chamber derives much of its design from the requirement to connect to consecutive units and the Planar Motors stators. The feature design and tolerancing used to

Figure 3.6: The muffin fans used to generate the gas flow in the cooling chamber are 40mm x 40mm x 28mm muffin fans. These fans generate gas flow locally at the cooling chamber inlet, localizing the flow generation to each individual module.

meet this requirement is described in Chapter 2. In this chapter, only the unique elements required for the cooling functions are discussed. These elements include the side and top ports, as well as the accommodations needed to implement these features. These features are highlighted in Figure 3.7.

3.4.1 Port Positioning

The chamber has space for ports on the top and sides of the chamber. The bottom of the chamber is occupied by the Planar Motors stator, which is a solid body that cannot be penetrated by gas flow. The front and back faces of the chamber are part of the open tunnel connecting consecutive units, so it also cannot be a source or sink for the gas flow because that would block vial movement through the system. Thus, only the top and sides of the chamber can be used for gas flow ports. As discussed in Section 3.2.4, the side ports are used as inlets and the top ports are used as outlets.

While the ports exist on the chamber faces, the actual position of the inlet and the outlet relative to the vials can be changed by features inside the chamber. If the chamber only uses the ports on the walls and ceiling, the gas will flow diagonally from the side inlets to the outlets, missing some of the vials in the chamber, as shown in Figure 3.8. Instead of isolating the outlet to the top of the chamber, internal walls are extended down from the ceiling to reposition the chamber outlet to the middle of the chamber. This outlet positioning forces the gas to flow past all vials in the chamber before exiting. This expected improvement in cooling performance is validated by the simulation work done in [40].

3.4.2 Outlet Wall Design

The cooling chambers include outlet walls to direct the cooling gas flow for more efficient convective cooling. An example outlet wall is shown in Figure 3.9. Given a box with flow inlets and outlets, the gas will preferentially flow from the inlets to the outlets in the path of

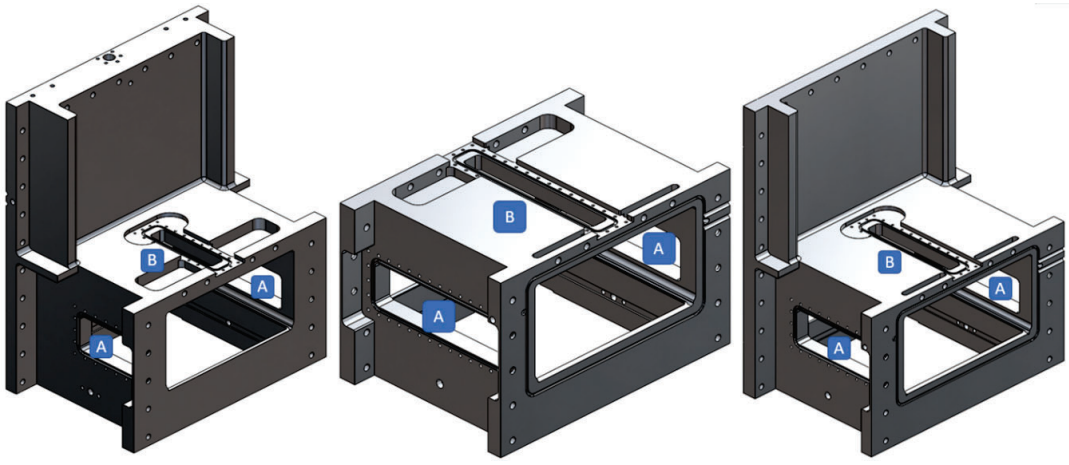

Figure 3.7: Overview of the chambers used in the cooling section. Each unit has two side inlets (A) and one central top outlet (B). The same nominal dovetail O-ring groove and bolt spacing can be used for sealing to each of these interfaces.

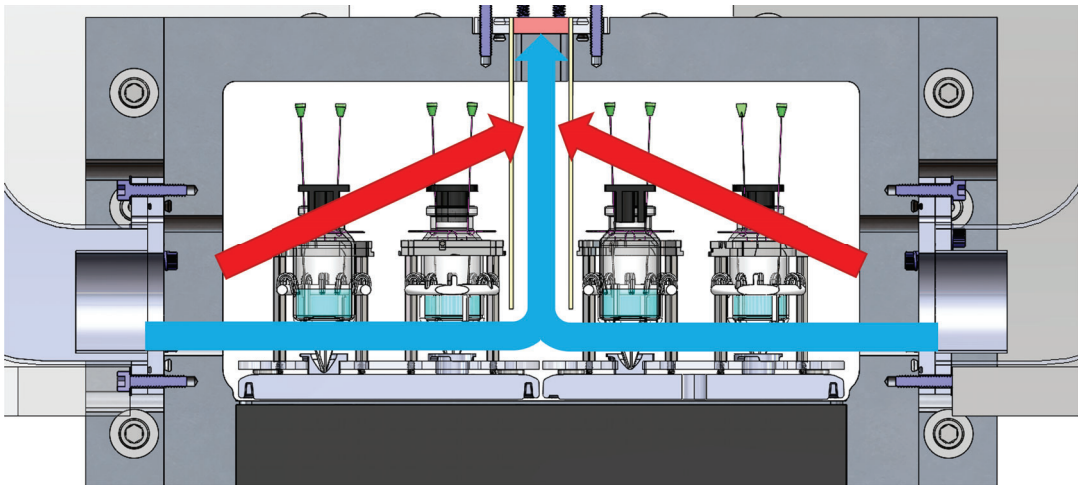

Figure 3.8: If only using the ports on the sides and top of the chamber, the gas flow will follow the diagonal red arrows, which would not effectively cool many of the vials in the chamber. Extending walls down into the chamber, shown in yellow, forces the gas to flow past all the vials in the chamber before exiting through the top outlet, as shown by the blue arrows.

least resistance. In the chamber geometry used in this work, the inlet ports are located near the bottom of the side walls and the outlet port is in the middle of the ceiling, as shown in Figure 3.8. With this layout and no obstructions, the flow will preferentially move vertically towards the outlet before traversing across the bottom portion of the chamber, potentially missing some of the vials. This disparity would only be magnified with the presence of the vials in the chamber, as they create further obstructions along the bottom portion of the chamber, which would direct the flow upwards sooner. With multiple rows of vials, this type of flow would lead to significant non-uniformity between the rows, as any row not directly exposed to the cooling gas front would see much less gas flow, leading to much longer cooling times. Adding walls to the outlet port, extending the outlet entrance down to the level of the fluid in the vials, increases the resistance for gas flow that would bypass the vials by going above them. Rather than seeing a free path to the outlet, the gas would have to be redirected up and then back down through a narrow gap between the vials and the outlet wall. In this configuration, the path between the vials becomes preferential, directing the cold gas flow through the vials and creating more uniform cooling among the rows.

The outlet walls are made thin to decrease the total space they take up in the chamber, allowing for maximizing the width of the outlet given a limited position due to tray geometry. However, if the outlet walls are made too thin, they would deflect and vibrate under the force created by the convective gas flow. For this application, 1/16 in acrylic plates are used for the outlet walls. The 1/16 in thick acrylic provides sufficient thickness to create a stable wall during operation while leaving a sufficiently large outlet to achieve the required linear flow speed. The final outlet wall geometry, shown in Figure 3.9, includes flanges at the top for positioning within the chamber, a cutout to account for the outlet port corner radius, and holes in the bottom corners for spacers.

For each chamber with different outlet port lengths, a different set of outlet walls is re-

quired to match. Most dimensions of the outlet walls and their features are shared among the different outlet walls. The outlet wall length changes to match the length of the corresponding outlet. Extending the length of the outlet walls beyond the outlet length does not provide a significant benefit to controlling the flow path, as the main flow width is dominated by the inlet and outlet port dimensions. Shrinking the outlet wall length would provide direct paths for the flow which would circumvent the outlet, reducing its effectiveness at keeping the maximum gas flow level with the fluid in the vials. When the outlet walls are located in their widest position, the curvature of the outlet port corner fillets would interfere with a straight edged cooling outlet wall. Accordingly, a cutout is added at the top of the cooling outlet walls to ensure the walls will fit in any position on the outlet. The height of this cutout corresponds to the thickness of the top of the cooling chamber, and the depth of the cutout into the cooling outlet wall corresponds to the radius of the outlet corner fillets. These dimensions include an additional half-millimeter nominal spacing to ensure fits within manufacturing tolerances.

The outlet wall height depends on the height of the vials within the chamber. To ensure that the gas flows across the vials rather than over them, as discussed above, the outlet walls must extend below the top of the vials in the chamber. However, the fluid only fills a portion of the space within the vial, and flow interacting with the top, empty portion of the vial is less efficient at cooling the fluid in the vial than flow interacting with the bottom portion where the fluid rests. Accordingly, the outlet wall height is set so that when they are placed in the chamber, the outlet walls extend down to the position of the top of the fluid in the vials as they travel through the chamber. This outlet wall positioning is shown in Figure 3.10.

The tray dimensions can vary depending on the number and size of the vials used within the system. Accordingly, it is beneficial to decouple the outlet wall position from the outlet port machined into the cooling chamber. This decoupling also simplifies the chamber machining, both through removing an operation which would be required to make wall posi-

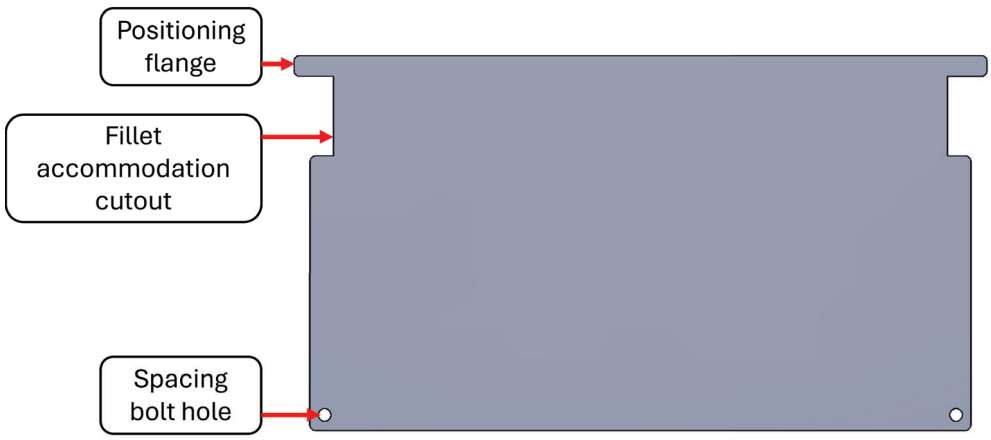

Figure 3.9: The outlet wall includes flanges at the top for positioning the wall in the outlet port, cutouts near the top to accommodate the curvature on the corners of the outlet cutout, and holes in the bottom corners for the spacing bolt used to rigidly attach the bottoms of the outlet walls together.

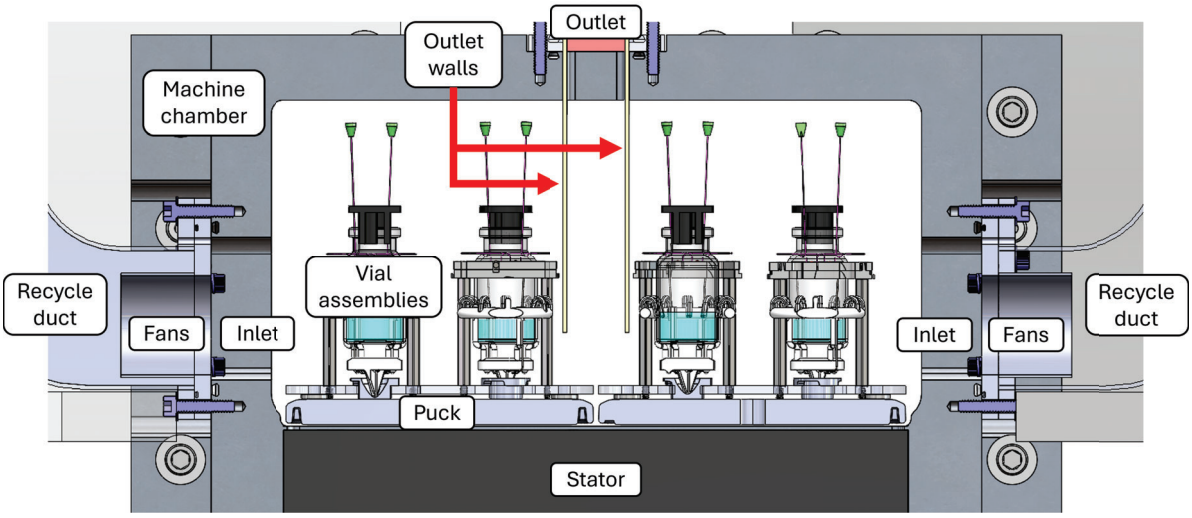

Figure 3.10: The outlet walls extend down to the height of the fluid in the vials to ensure the flow is directed past the vials at the fluid level in the vials. Key elements of the internal geometry of the cooling chamber are labeled.

tioning features, and because this operation would require small tooling to create the features needed to accommodate the thin outlet wall. To separate the outlet wall mounting from the chamber, an additional piece is manufactured separately to hold the outlet walls. This piece connects to one of the plates used in the diffuser assembly, through which it sets the position of the outlet walls relative to the chamber. This piece has two linear cutouts in which flanges of the outlet wall rest. Adjusting the position of these linear cutouts changes the outlet wall positions, widening or shrinking the outlet width, as shown in Figure 3.11. When the outlet wall position is moved to shrink the effective outlet width, additional pieces must be added to block the areas outside of the outlet walls that would lead directly from the chamber into the diffuser. If these areas remain open, the gas in the chamber can escape directly to the diffuser without traveling through the volume between the outlet walls. In this case, the gas going directly to the diffuser would not be forced to travel past the vials, significantly reducing the effective cooling within the chamber.

The outlet walls and their locating piece are supported on the top of the chamber, such that they can be constrained by the chamber and the diffuser assembly without requiring bolting. Using the other components of the chamber to constrain these parts provides degrees of freedom for looser tolerancing on the sizing of the outlet walls and their locating piece. The locating piece is also held in place through a dovetail connection on the top brick plate, which keeps the locating piece in position during assembly. This mounting also leaves the outlet walls free to move when they shrink due to the cold temperatures in the chamber during operation.

During initial testing, the cooling outlet walls deflected under the load of gas blowing from the fans at the inlet. The outlet walls extend down to the height of the bottom of the vials on a tray. Because the outlet walls only have geometric constraints at their tops, they have some freedom to move within the cutouts of the positioning piece. This freedom is necessary to allow the outlet walls pieces to slide in and out of place during assembly without

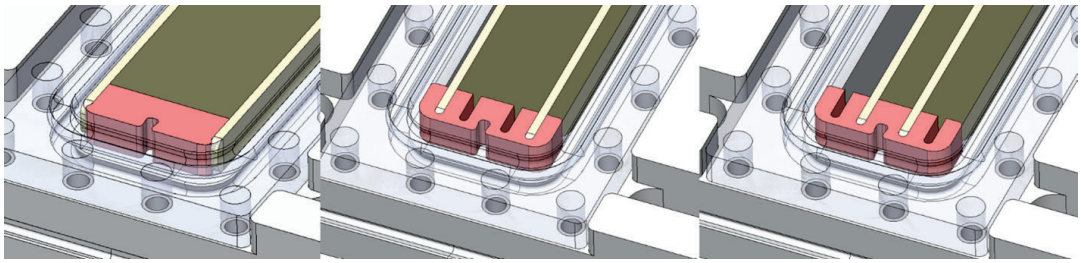

Figure 3.11: Three different outlet wall positions are shown to demonstrate how the effective outlet width can be adjusted by repositioning the outlet walls using a plastic wall holder insert. This insert allows for testing different outlet geometries without requiring a different machined chamber for each geometry. The first insert in the left image is only used for the widest wall positioning, while the insert in the middle and right images has multiple slots for different outlet wall positions.

tightly tolerancing each part; however, the gaps allow the walls to move perpendicularly to their faces, particularly when the cooling gas applies a load to the face of the walls. While the tolerances on the top components could be tightened or bolts could be used to try to fix the top of the outlet walls in place, these solutions would address the motion far from the location where the force is applied. Thus, an alternative measure to resist this loading is implemented near the bottom of the outlet walls.

A spacer placed near the bottom of the cooling outlet walls provides resistance to the motion induced by the gas blown within the chamber. One of the cooling walls has a threaded hole in its bottom corners, while the other cooling wall has countersunk holes on its face. A countersunk bolt is used with a spacer or a pair of nuts to rigidly connect the two outlet walls. This assembly is shown in Figure 3.12.

3.4.3 Outlet

The outlet sizing is determined by the vial trays used in the system, the other features taking up space on the top surface of the chamber, and the fans used in the system.

The width of the outlet depends on the trays used in the system. Because the outlet has walls that extend down into the chamber to ensure that the flow remains fastest at the level of the fluid in the vials, potential interference between the outlet walls and the trays must be prevented. If the outlet walls infringe on the channel through which the trays travel, then the trays would hit these walls rather than travel freely. This geometric restriction creates a trade-off between the potential tray size and the gas flow speed. The smaller the outlet width, the wider the trays that can be used; conversely, the larger the outlet width, the less flow resistance, and the faster the linear gas flow can be generated. To keep the chamber flexible to accommodate a variety of tray geometries for hardware development and differently sized vials, the specific outlet width is decoupled from the machined chamber geometry.

The machined outlet sets the maximum outlet width, as it would create a flow restriction even if the outlet walls were positioned wider than the chamber's outlet cutout. The effective width of the outlet in the chamber can be decreased by moving the position of the outlet

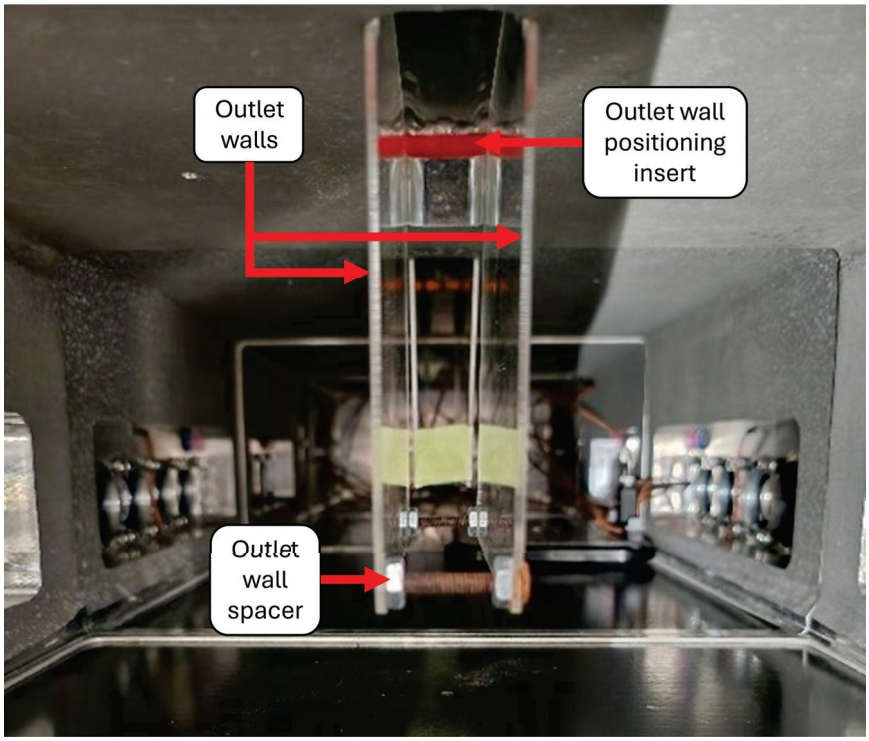

Figure 3.12: The outlet walls include a spacer positioned at the bottom corners of the outlet wall assembly. This spacers helps the outlet walls resist the force applied by the convective gas flow inside the chamber.

walls within this cutout, as described in Section 3.4.2.

The outlet port cutout has curved edges that enable standard machining. For an internal cavity, perfectly square edges are not possible without specialized tooling (such as using a broaching tool). When the outlet walls are in their widest position, their shape needs to account for this curvature. This accommodation is easily achieved by adding a cutout to the outlet wall shape. Because the curvature does not significantly affect the design of other parts, it can be made relatively large so that larger tooling can be used in manufacturing. As such, the outlet wall corners are curved with a 3/16 in radius, allowing the use of tooling with up to a 3/8 in diameter. The outlet wall is shown in Figure 3.9.

The outlet port is centered on the top of the chamber relative to the side inlet ports. For chambers that do not include door features, the outlet port is also centered relative to the chamber's open tunnel faces. Given the dimensioning of the features used in this system, the positioning of the outlet for units without door features results in outlet features spanning the chamber's entire top surface length. For units where the chamber includes door features, the edge of the outlet on the side of the chamber that does not include door features is matched to the location of the outlet when the chamber has no door features. This matching simplifies machining, as it ensures that a subset of features is shared among different types of units, rather than requiring completely new dimensioning and tolerancing for each feature on each unit. This decision also does not affect surface use efficiency because in the no door unit case, the outlet features already span the entire top surface length. These

different outlet geometries are shown in Figure 3.7.

The existence of different door features on the top surface of some chambers restricts the total length of the outlet on those chambers. The closest feature to the wall created by the door features is no less than 0.75 in away, ensuring that there would be room for 1.5 in diameter tooling to machine the required features. The outlets are sized by the largest whole number of fans that would span the outlet (with 0.5 mm gaps between fans to accommodate variance in housing dimensions) without breaking the tool spacing restriction. The resulting outlet lengths are sized for three, four, and five 40 mm x 40 mm fans for the door chamber, door support chamber, and no door chamber respectively. The resulting outlet lengths are 121 mm, 161.5 mm, and 202 mm.

3.4.4 Inlets

The inlet port cutout used for a standard cooling module is shown in Figure 3.13a. This is the maximum width inlet port, as it does not need to accommodate any door features.

The width of the inlet ports is set based on the size and number of fans in the system. This dimension depends on the number of fans used, which varies with the specific chambers. This variation is dictated by the restrictions on the outlet length, as discussed in Section 3.4.3. The side port width is matched to the outlet length to simplify the manufacturing of the recycle system components. Additionally, the features used for the wedge doors between chambers, which limit the machinability of the outlet, similarly limit the area available to be machined for the fan side port, though not quite as severely.

The height of the side ports is primarily determined by the fans used in the system. In this case, because 40 mm x 40 mm fans are used, this port must have an opening at least 40 mm tall. Because muffin fans have mounting holes on their faces within their length and width specifications, the minimum port height remains unaffected by the mounting features. Only a single row of fans is needed at the inlet port because the height of the fluid to be lyophilized in the vials is only about 10 mm. Each type of chamber uses the same fans for simplicity, so their inlet heights are all the same.

When determining the side port length and height, the wiring of the fans must be con-

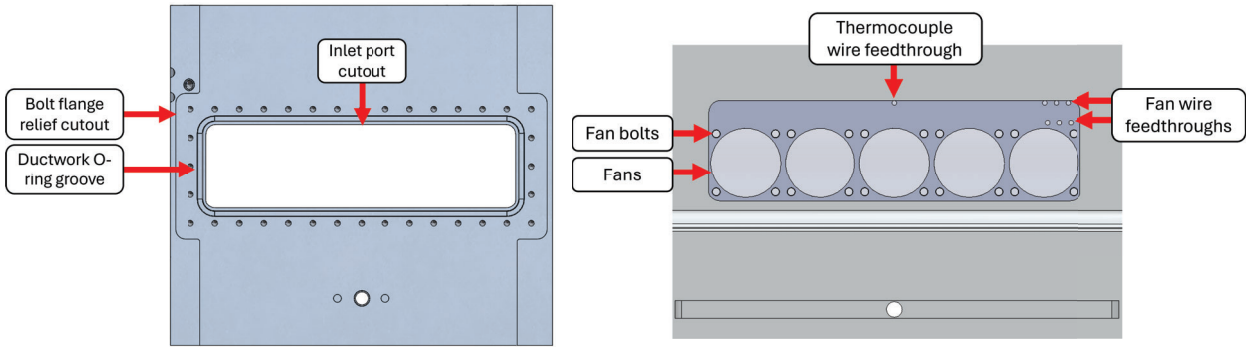

(a) (b)

Figure 3.13: The inlet cutout is located on the side of the cooling chamber. This cutout includes accommodation for the fans, the fan mounting features, and the fan wiring.

sidered. While initial iterations of the cooling chamber included passthroughs for wires in the recycle hardware, these passthroughs proved difficult to assemble and seal. To improve sealing and ease of assembly, the wire passthroughs were moved to the plate on which the fans are mounted. This fan mounting plate is described in Section 3.4.8. This move allows the fans to be mounted independently to the plate, including their wiring, before the fan and plate subassembly is added to the recycle system. To accommodate the physical spacing for these wires, the inlet port must be extended in height and/or width. Given the tight fit on the width of the side port relative to the chamber bolting flange, as shown in Figure 3.13a, extending the height of the inlet port is more viable than extending its width.

Due to the possibility of liquid nitrogen entering the recycle system, it is preferable to put the wire features above the fans rather than below them to reduce the risk of exposing the wires to liquid nitrogen. The wire features require two sets of holes, each with one hole per wire type (with wire types being positive lead, negative lead, and PWM signal). These wire holes must be at least large enough to accommodate one wire per fan, or a single wire large enough to carry the current needed for the total number of fans on that port. For this application, 3 mm diameter holes are used because they can accommodate five of the wires on each fan, which matches the largest number of fans on an inlet. Two sets of wire holes are needed because the wires must go from the fans, which are inside the gas recycle system, to the inside of the chamber, then back out of the chamber and out of the system. Routing the wires through the chamber is a necessity created by feeding the wires through the fan mounting plate, rather than directing them through the recycle system. Because these wire holes go from inside the recycle system to outside it, they must be spaced apart by at least the thickness of the walls of the recycle system ductwork. In this application, the recycle system ductwork uses a 1.6mm wall thickness, as described in Section 3.4.8. With two sets of 3 mm holes and the recycle ductwork thickness, the additional height needed for the side ports is at least 7.6 mm, on top of the 40 mm required for the fans. In this application, the additional side port height is 15 mm, for a total height of 55 mm, to provide a factor of safety on the space needed for wire features. The positions of the fans and the wire holes within the inlet cutout are shown in Figure 3.13b.

The side port corners must be filleted, like the outlet corners. For the side ports, the corner radius must be sized such that it will not interfere with the bolts used to mount the fans to their mounting plate. If the mounting plate is internally threaded such that the bolts do not extend beyond the plate, then this consideration is unnecessary. However, in this application the fans are threaded, and the bolt heads protrude from the plate into the chamber. This orientation allows for shorter, larger diameter bolts, and the bolt heads provide reference features for locating thermal sensors within the chamber. In this application, a 1/8 in diameter fillet is used, allowing the use of up to 1/4 in diameter tooling to machine the corners.

Similarly to how the inlet port length is matched to the outlet port length, the inlet port position along the chamber's length is matched to the outlet port position. This alignment simplifies the recycle ductwork manufacturing, as it keeps the cross-section of the recycle ductwork rectangular.

The inlet port's vertical position relative to the bottom of the chamber is set based on the structure that holds the vials in the chamber. The bottom edge of the inlet port matches the top surface of the bottom plate on the vial holding structure. This positioning focuses

the cooling gas flow on the vials resting on the puck and minimizes the losses due to cooling the puck and the chamber. This vertical positioning on the chamber wall, highlighted in Figure 3.14, is shared among all cooling chamber units, regardless of whether they have door features.

3.4.5 Sealing

The recycle ductwork seals to the chamber through O-ring seals. O-rings are a standard sealing method, and they are used in many sealing interfaces in this lyophilizer. The use of O-rings ensures reliable sealing, uses off-the-shelf components, and provides an interface that can be removed and replaced repeatedly. Additionally, by using similar O-ring specifications as used in the vacuum sections of the machine, these cooling units can be converted to vacuum units if such a need arises. Dovetail O-ring grooves are used to make it easier to load the O-rings into their grooves and to ensure that the O-rings stay in place while the sealing interface is exposed. Although the dovetail groove is slightly more complex to machine, the improvements to the assembly process are worth the increased complexity. Both the inlet and outlet O-ring grooves use a standard nominal 1/8 in dovetail groove profile as recommended in the Parker O-ring Handbook [26]. These O-ring grooves are highlighted in Figure 3.15.

The O-ring groove for sealing around the inlet ports follows a path offset from the inlet port cutout. The offset is set such that the minimum material thickness between the O-ring groove and the port is no less than 1.5 mm. This thickness leaves sufficient material for the groove to be safely machined without piercing the inlet port wall, which would create a potential leak path, and to support the O-ring if a vacuum is applied to the inside of the chamber.

The outlet port uses a similar offset for the O-ring groove along its length for the same reasons listed above. However, at the ends of the outlet port (closest to the open faces of the chamber), the O-ring groove offset is increased beyond this minimum. This extra spacing provides a surface on which the outlet walls and the outlet wall holder rest.

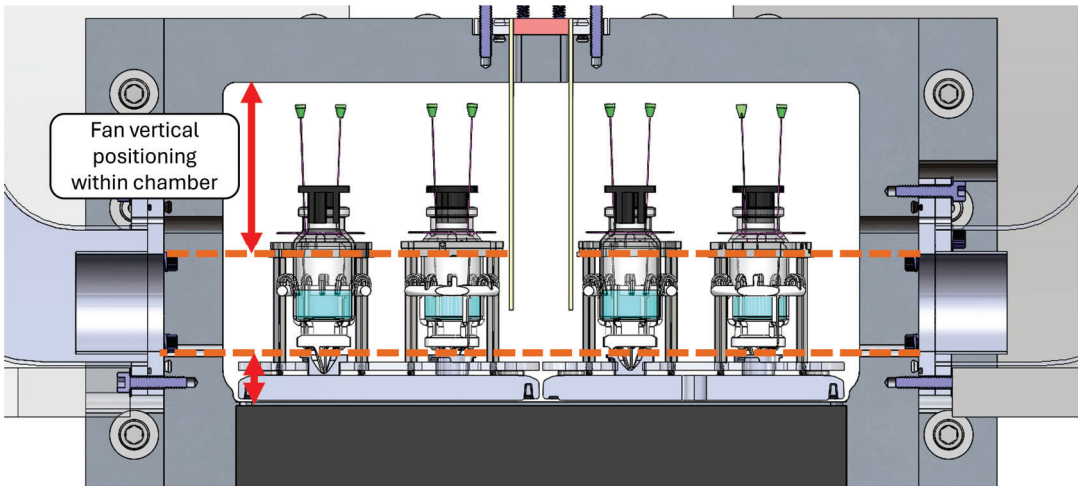

Figure 3.14: The inlet port height is positioned such that the fan face is focused on the vial body to ensure the cooling flow is directed at the vials within the cooling chamber.

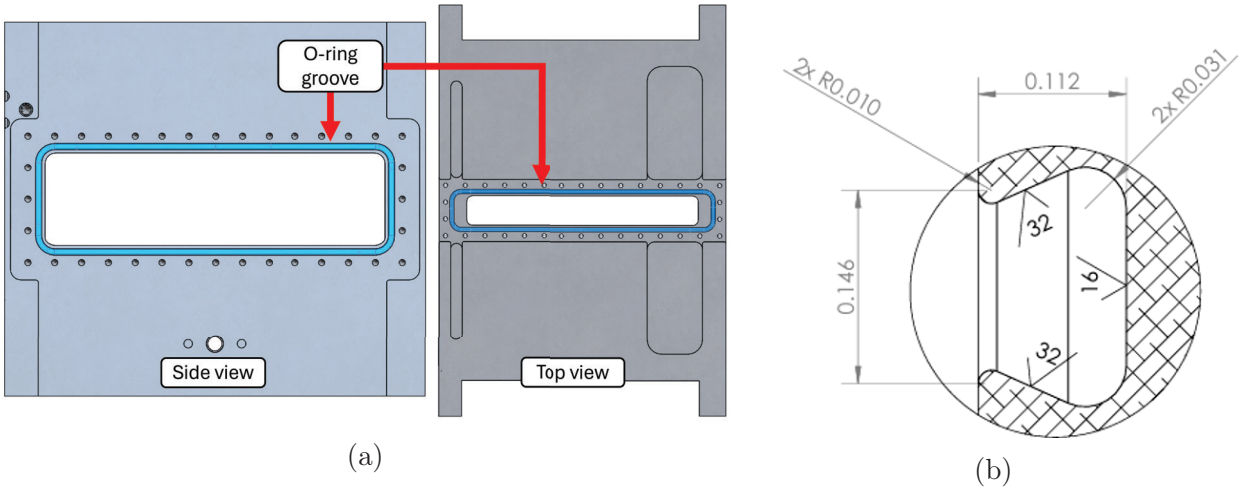

(a) (b)

Figure 3.15: The cooling chamber uses O-ring grooves on the side and top surfaces to seal the inlet and outlet ports. These grooves use the standard nominal 1/8 in O-ring dovetail dimensions (in IPS units) shown in (b).

3.4.6 Bolting

The recycle system ductwork is connected to the cooling chamber through a bolted interface. The bolted interface provides repeatable attachment and removal of the recycle ductwork for assembly and maintenance purposes. To prevent creating additional leak paths between the chamber and the outside environment, the bolt holes in the chamber are blind tapped holes, rather than simpler through holes and fastening with nuts. These tapped holes use a coarse thread pitch rather than a fine thread pitch because the chambers are made from cast metal, as described in Chapter 6. The coarse thread pitch has a larger major diameter, ensuring that the threads engage with a larger material area. This additional area engagement helps minimize the effects of randomly located porosity within the cast material so the bolted connection remains strong enough to provide the necessary clamping force to the O-ring seal.

The bolt diameter can be selected based on the required stress the bolts need to apply to create the bolted connection. This stress can be calculated based on the forces on the bolted components which would cause them to separate and the number of bolts in the system. This calculation can be tedious and require iteration, as it is affected by the bolt spacing, which is in turn affected by the bolt diameters as well. Initial guides for narrowing down the range of diameters include ensuring that at least 3-5 threads are engaged when the bolt is fastened, geometric constraints on maximum bolt diameter, and ease of use of bolts during assembly. Another way to find a starting point is to use a bolt diameter such that the length of the bolt is 3-5 times its diameter. The recycle ductwork connection uses two 0.25 in thick plates, as described in Section 3.4.8, so the required bolt length is 0.5 in plus the thread engagement. The thread engagement is an additional 0.25 in into the chamber wall, resulting in a total bolt length of 0.75 in. Given this total length and thread engagement, 10-24 x 0.75 in machine screws are used to connect the recycle components to the chamber. These bolts engage with 0.25 in of the chamber, ensuring 5-6 threads are engaged. Additionally, given

the bolt diameter of about 0.2 in, the total length is 3.75 times the bolt diameter, meeting the design guidelines.

The bolts both attach the recycle ductwork to the cooling chamber and apply the force needed to apply the preload crush to the O-rings used to seal the components together. To ensure that the O-ring remains compressed along its entire length, the bolts must be spaced sufficiently close to each other such that the recycle ductwork face cannot bend away and create a leak. One way to ensure a secure connection is to ensure that the bolt stress cones overlap. When the bolt stress cones overlap, the bolted joint acts like a fixed connection, and there is no deflection at the interface. However, overlapping stress cones, particularly for relatively thin material like the 0.25 in flanges used for the recycle ductwork, requires many bolts, increasing both cost and time to assemble/disassemble. Thus, it is preferred to space the bolts further apart even though the stress cones do not overlap. The bolt holes are nominally spaced no more than their total length apart to ensure that if the recycle hardware connection design changes, there will be sufficient bolting available to maintain the required preload crush on the sealing O-rings.

The bolt holes form a rectangular pattern outside of the O-ring grooves around the chamber's inlet and outlet ports. While the ports and the O-ring grooves include filleted corners, the bolt pattern remains in a rectangle with sharp corners. This patterning comes from the recycle ductwork, as it is easier to manufacture with rectangular cross sections rather than with rounded corners, as discussed in Section 3.4.8. The bolt holes are offset from the O-ring grooves such that at least 1.5 mm of material will remain between the cut threads in the bolt holes and the O-ring groove. Like the spacing between the O-ring groove and the port cutout, this material thickness ensures a solid wall will remain between the O-ring groove and the threaded bolt holes to prevent the accidental creation of leak paths during machining.

3.4.7 Chamber Bolt Flange Cutouts

The features required to attach the convective gas flow and recycle ductwork to the cooling chambers require a few additional areas of material removal for both fit and ease of assembly. The first material removal area is on the side flanges used to bolt consecutive chambers together, shown in Figure 3.16. When the side inlet ports are wide enough for five fans, the standard flange thickness occupies too much space on the chamber's side face for the O-ring groove, bolt features, and ram's horn flange to fit. Accordingly, a pocket needs to be machined out of the flange to accommodate these features. The pocket also needs to be large enough to fit the nuts and bolt heads used to connect consecutive chambers together. For these bolt holes, hex head bolts and jam nuts are used to accommodate the tight fit, enabling the use of the 40 mm fans in this system. This accommodation allows the fans to span as much of the width of the chamber as they can without compromising the sealing between consecutive modules. The corners of this pocket are filleted for machinability purposes and to remove sharp edges.

Further chamber width coverage could be achieved with smaller O-rings and smaller bolts. This increased coverage would reduce the time that vials are not directly in front of the cooling gas flow. However, smaller O-ring would have reduced sealing performance, and smaller bolts would require an increased total number of bolts, increasing assembly time.

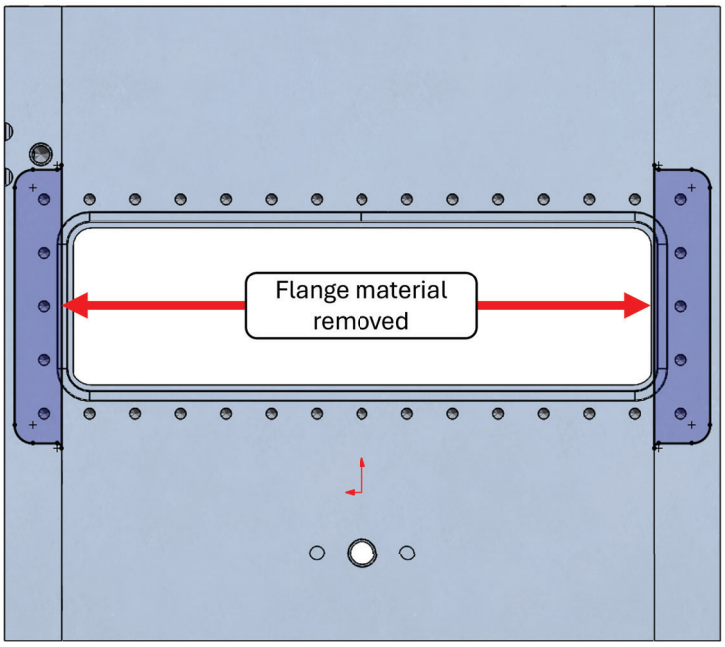

Figure 3.16: Part of the chamber bolt flanges needs to be machined away to accommodate the recycle ductwork attached to the cooling chambers.

The side inlet port is also sized to fit a series of fans positioned consecutively with minimal gaps between them – a wider port would require either differently dimensioned fans or larger gaps between consecutive fans. Alternatively, there can be gaps between the pucks within the tunnel that correspond to these gaps between gas inlet ports. While this strategy ensures the vials are always in front of cooling gas flow, it decreases the vial density inside the tunnel, reducing the overall production rate given the same tunnel length. The effect of increased chamber width coverage by fans and increased gaps between fans requires further study to determine if the potential benefits offset the costs and risks associated with creating a wider inlet port.

The top surface flange on the cooling chamber is also machined away above the areas where the side port inlets are located, as shown in Figure 3.17. Though this material does not directly interfere with the recycle ductwork, it blocks access to the bolts used to attach the ductwork to the chamber. While this top flange provides stiffness for bolting consecutive chambers together, its detrimental effect on assembling the cooling chambers outweighs its benefit. Accordingly, this section of this flange which would be inside the path traced by the recycle ductwork is machined away.

The cooling chambers also have additional machining done on their top faces to facilitate the machining of the features used to interface with the recycle system diffuser, shown in Figure 3.18. This machining is needed to ensure that the recycle ductwork interfaces with a clean machined surface rather than relying on sealing to a rough cast surface. An area slightly wider and longer than the diffuser flange is machined down about 0.25 in from the nominal cast surface where the diffuser flange interfaces with the chamber. On chambers with door features, this cutout includes an extra half circle at the end of the area close to the door features to allow machinists flexibility in the tooling used to machine this area. The

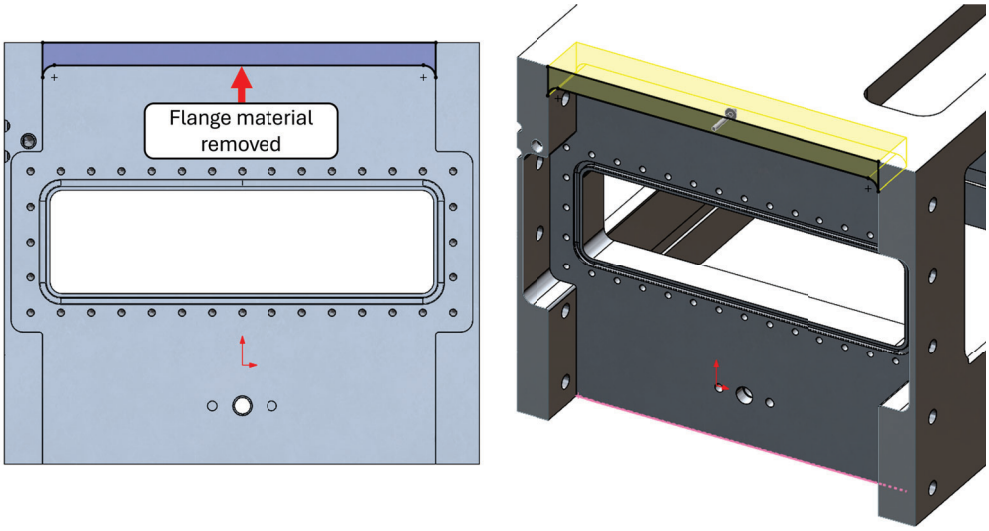

Figure 3.17: The top flange on the cooling chamber is machined away to provide access to the bolts used to attach the recycle ductwork to the cooling chamber.

entire area on which the diffuser flange sits needs to be machined, but machining additional area beyond the flange interface does not negatively affect chamber performance. Thus, the radius of this circle is determined by the machinist based on the tooling they choose to use to machine the top surface. This top surface machining ensures that the sealing interface can achieve the necessary flatness and surface finish to provide a successful seal. This specific area machining is not required on the side faces because those entire faces are already machined as part of the general module manufacturing process. The specific depth of this cut is toleranced from a dimension on the side inlet ports. Dimensioning this cutout from the top surface would create variability induced by the reliance on the cast surface, which has relatively loose tolerances. Additionally, this surface interfaces with the recycle system, which connects back to the chamber at the side ports. Because the side and top features both connect to the interconnected recycle ductwork, it is more useful to tolerance these features relative to each other, minimizing the error between them, than it is to tolerance each of these features to different datums.

3.4.8 Recycle System Design

The final recycle ductwork system, shown in Figure 3.19, consists of three main components: two ram's horns and one diffuser. The fresh gas feed enters the recycle system through manifolds inserted in the ram's horns, and the system maintains constant pressure by expelling gas through an outlet attached to the diffuser. These components are connected to each other using O-ring seals. The gas flow within the recycle system is primarily driven by muffin fans within the system. If the gas flow were primarily driven by the fresh gas feed, then the gas would not re-enter the flow in the ram's horns, preventing the system from effectively recycling the gas flow.

The ram's horns and the diffuser are made from welded 1/16 in aluminum sheets. Each component has a bolt flange welded onto its inlet and outlet which is used to connect the

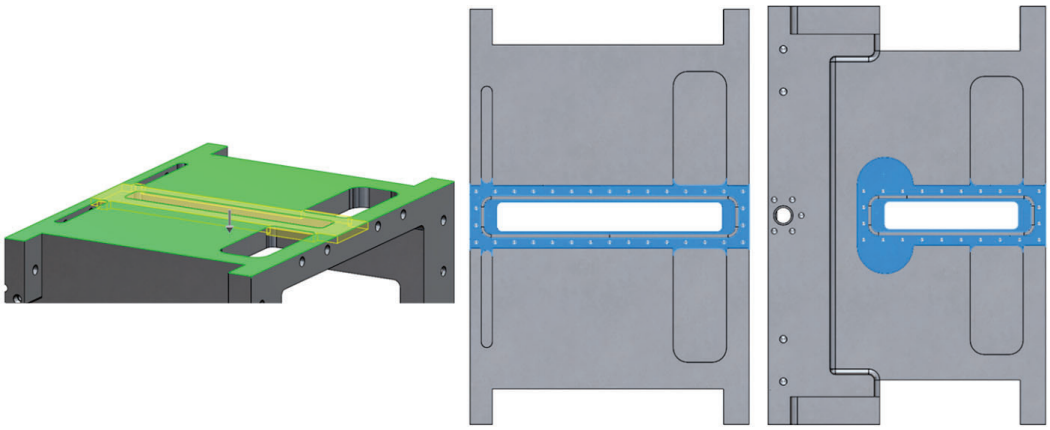

Figure 3.18: A rectangular area on the top surface of the cooling chamber is cut away to provide a clean, machined surface for mounting the recycle ductwork, rather than relying on the cast surface for sealing. On units with door features, an arbitrarily large diameter semi-circular area is added to ensure the full rectangular area for the ductwork can be machined without specifying specific tooling to a manufacturer.

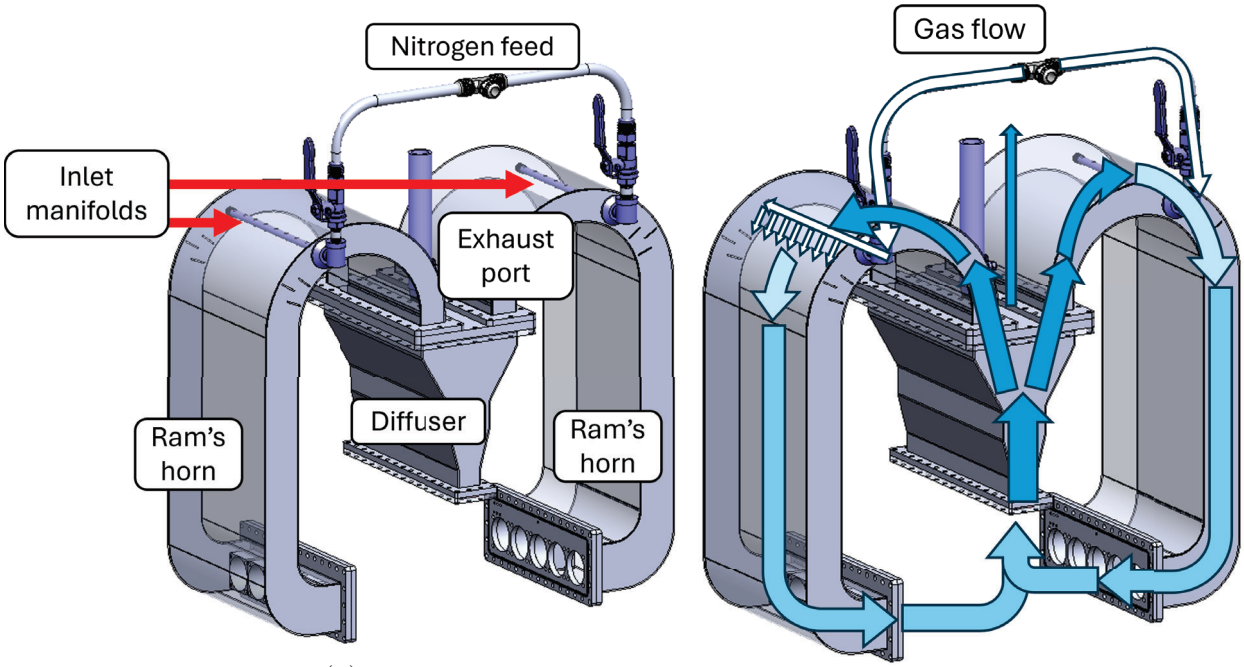

Figure 3.19: The recycle system used in the cooling module consists of two main elements: the side ram's horns which direct flow into the cooling chamber and the diffuser which directs flow from the chamber outlet back into the two ram's horns. The fresh gas feed is injected into the ram's horns through a manifold, and the exhaust gas exits the diffuser through an exhaust tube. The gas flow path is shown in the right image, where lighter colors represent colder gas.

components to each other and to the cooling chamber. This bolt flange is used to compress the O-rings used to seal the components together. The flange is 1/4 in thick to reduce its warping during the welding process. Even with this additional thickness, the bolt flange still needs to be machined flat to address the warping that occurs during welding. This flatness is required for the O-rings to compress effectively and create their seal. The bolt flanges include through holes for the #10 machine screws used to bolt the module components together.

Both the ram's horns and the diffuser have constant widths to ensure their profiles maintain a rectangular cross-section. This geometry reduces the amount of gas compression or expansion that could occur within the ductwork, creating flow head losses. The diffuser is slightly wider than the ram's horn because it needs extra space to accommodate the diffuser fan plate mounting features. The height of the rectangular cross section does need to change for both the ram's horns and the diffuser. These changing areas are guided through the shape of the component side panels, highlighted in Figure 3.20. The diffuser needs to increase the outlet port's face area to match a row of fans, then split the singular chamber outlet into two separate flow regimes for the two ram's horns. The ram's horns use a slow expansion to extend from the diffuser outlet area to match the chamber inlet area. This expansion occurs in the top semicircular path of the ram's horn by using circular paths that are not concentric with each other to create the outer boundaries.

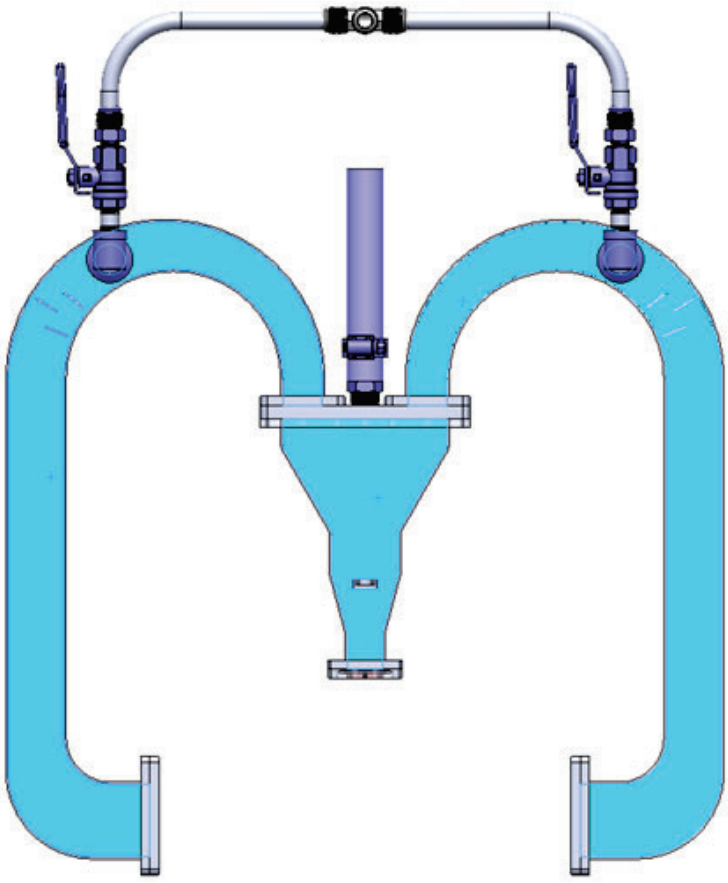

Figure 3.20: The recycle side panels determine the changes in the cross sectional area of the recycle ductwork.

The side profiles of each component are cut with a laser cutter or waterjet to create their irregular geometry, while the other panels are sheared into the desired rectangles and bent into shape. The bolt flanges are also laser cut or waterjet, as all of their features are two-dimensional. During manufacturing, the bolt flanges are attached to a fixture matching the corresponding cooling chamber geometry. The sheet metal parts are then inserted into the bolt flanges, forming the general shape of the completed components. Once in place, the panels are welded to each other and the bolt flanges. This method of welding in place ensures that while the ductwork may not remain perfectly square along its length, the mating faces will match up with the chamber to create a closed system.

The diffuser and the ram's horns are sized differently for each type of chamber to accommodate their different inlet and outlet dimensions. While the diffuser and ram's horn widths need to change to match the changes in the corresponding chamber features, the consistency in the inlet height and outlet width allow some components in these assemblies to remain constant. The side profiles for these features do not change, while the outer panels and bolt flanges adjust their width to match the chamber. The three different diffuser sizes are shown in Figure 3.21.

Each of the ram's horns and the diffuser have an associated set of fans which are mounted to a plate attached to the respective component. There is an extra plate used to bridge the connection between the diffuser and the ram's horns to connect the diffuser to the system outlet and to ensure that no O-ring sealing groove is machined onto a welded plate. The fan plate is mounted to the ram's horns at their outlet, which matches the location of the inlet cutout on the cooling chamber. The fan plate for the diffuser is mounted in the middle of the diffuser. The positions of the fans in each component are shown in Figure 3.22.

The recycle ratio can be measured by comparing the flowrate of gas through the chamber outlet and the velocity of the gas within the chamber. The gas velocity within the chamber is measured in front of the fans, as this measurement is also used as a diagnostic to verify uniform flow generation across the chamber. During conditioning experiments on a five fan width cooling unit, the outlet flowrate was measured at about 70 liters per minute (LPM). During similar tests, the average velocity in front of the inlet fans on one side was measured

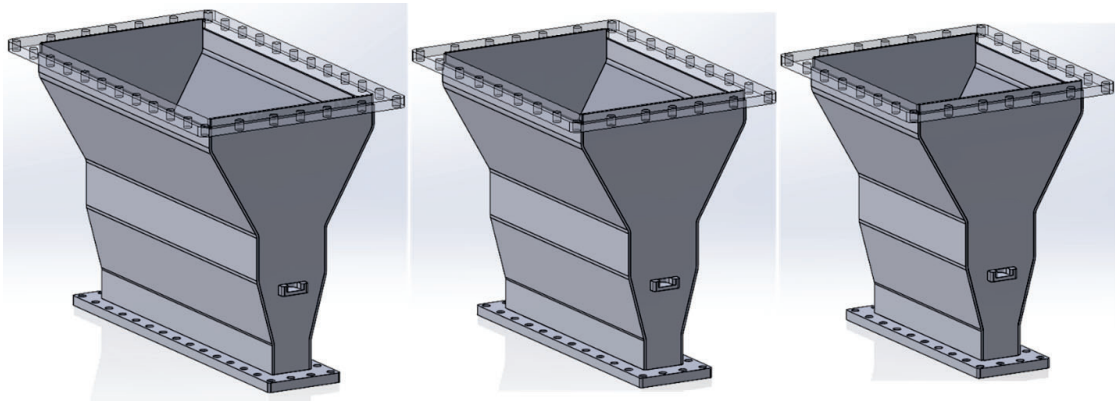

Figure 3.21: The diffuser has three different sizes to match each of the three different cooling chamber geometries. The general shape of the diffuser remains the same across sizes, with only the overall width of the diffuser changing to match the width of the chamber outlet.

to be about 2 m/s. Given a total fan area of 0.016 m^2, this average linear velocity translates to a volumetric flowrate of 1920 LPM. Taking the ratio of the outlet flowrate to chamber flowrate gives a recycle ratio of about 96%, demonstrating that most of the gas used in the system is traveling through the recycle ductwork.

3.4.9 Diffuser

The constraints on the outlet width created by the vial holding structure as described in Section 3.4.3 make it narrower than the fans. Placing the fans directly on the outlet would create a sudden expansion in the flow, which would create significant head losses. These head losses would reduce the gas flow within the chamber and increase the pressure generated within the chamber, both of which are undesirable. Also, when the fan face is partially blocked, its effectiveness at generating gas flow decreases. Thus, the outlet path cross section needs to be expanded before the gas reaches the fans, and this expansion occurs within the diffuser.

The diffuser side profile forms a general "Y" shape as it expands from the relatively narrow cooling chamber outlet to the split between the two ram's horns and the system outlet. The diffuser has five sections within it as its cross-sectional area expands. These sections include straights at the diffuser entrance and exit to facilitate the welded connection to the bolt flanges, a straight for the diffuser fans, and two expansion volumes. The transitions between sections use small radius corners sized to match the bend radius of standard sheet metal tooling. This dimensioning allows the diffuser's outside panels to be bent using a standard punch and die while still matching the side panel profile. These sections are shown in Figure 3.23, and they are described in detail below.

At the entrance and exit of the diffuser (Sections 1 and 5 in Figure 3.23), there are 0.75 in long straight sections. These section provides access for welding the components together

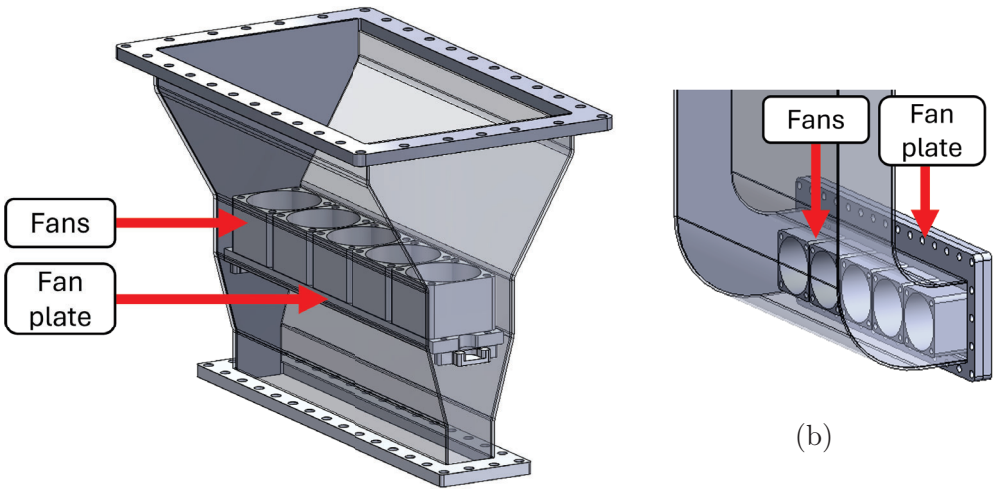

(a) (b)

Figure 3.22: The cooling fans are located in the middle of the diffuser (a) and the bottom part of the ram's horns (b).

and using tools to fasten the bolts attaching the recycle components together. Some of this length (0.25 in) extends into the bolt flanges for welding. These straights can include a small cutout on their corners to help locate them on corresponding features on the bolt flanges to assist during the welding process.

The first expansion volume (Section 2 in Figure 3.23) connects the bottom straight volume with the fans' straight volume. This expansion takes the diffuser width from 25 mm to 48 mm with a 30 degree expansion angle. The maximum recommended angle for a diverging duct is 45 degrees [41], so this angle is expected to provide a sufficiently smooth transition between the two section widths.

The straight fan volume (Section 3 in Figure 3.23) is 28 mm long, matching the length of the fans. This volume is matched closely to the fans to minimize the creation of eddies and to ensure the gas is all pulled through the fans. The width at this section is 8 mm larger than the fans to provide space for the fan wires. While this spacing does create gaps for eddies to form in the flow, it ensures there is sufficient area for the wires to not block the fan faces, as such a blockage would induce greater losses.

The second expansion volume (Section 4 in Figure 3.23) connects the fan volume to the diffuser outlet. This expansion has a total angle of 60 degrees, which exceeds the recommended angle for a diverging duct [41]. However, reducing the angle to the recommendation would nearly double the vertical length of this section, significantly increasing material costs, assembly complexity, and thermal losses. Exceeding the recommend diverging duct angle increases the expected head loss created in the flow and increases the risk of the gas flow separating from the diffuser walls, creating eddy flows. However, given that the system built with this geometry still meets the overall flow requirements (as shown in Section 3.5), these losses are shown to be acceptable.

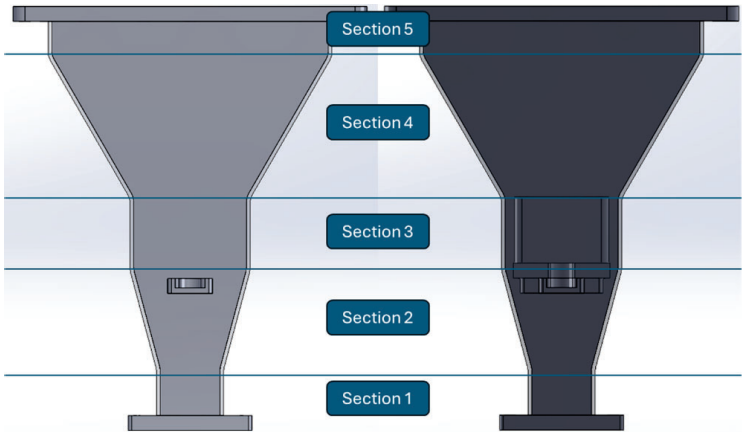

Figure 3.23: The diffuser geometry consists of five sections: 1) the straight section used to connect the bottom of the diffuser body to the bolt flange used to connect to the cooling chamber, 2) the first expansion section used to increase the diffuser width to match the nominal fan dimensions, 3) the straight section which accommodates the length of the fans, 4) the second expansion section used to split the flow to each of the ram's horns, and 5) the straight section used to connect the top of the diffuser body to the bolt flange used to connect the diffuser to the ram's horns.

The diffuser side plate has a rectangular hole for mounting the diffuser fan plate, as shown in Figure 3.24. This hole is centered relative to the angled sides of the diffuser, and it is positioned such that the top edge of this cutout is vertically offset below the straight fan volume by the thickness of the diffuser fan plate. This positioning ensures the fans are centered within their designated volume in the diffuser. The cutout height is set by the thickness of the material used to create the diffuser fan plate's mounting insert. The cutout width is set to three times the height of the cutout to ensure level and stable mounting for the insert that holds the diffuser fan plate.

The bolt flanges on the diffuser inlet and outlet include an outer rectangle of bolt holes and an inner rectangular cutout for allowing gas flow. At the inner corners of the rectangular cutout, small nubs exist to create reference features for the welding process. These flanges do not include any other features because they are part of the welded assembly. The sealing features are located on plates to which these flanges attach because they could warp during the welding process. The diffuser assembly's bottom bolt flange is shown in Figure 3.25.

3.4.10 Diffuser Connection Plates

The diffuser assembly connects to the cooling chamber and ram's horns through intermediary plates. These plates include sealing and positioning features. The plate between the diffuser and the cooling chamber, shown in Figure 3.26, includes rounded dovetail features that hold the outlet wall locating pieces described in Section 3.4.2 This plate has clearance bolt holes for #10 machine screws, as the bolts used to connect the diffuser assembly to the cooling chamber travel through these same holes. This plate also has an O-ring groove which seals

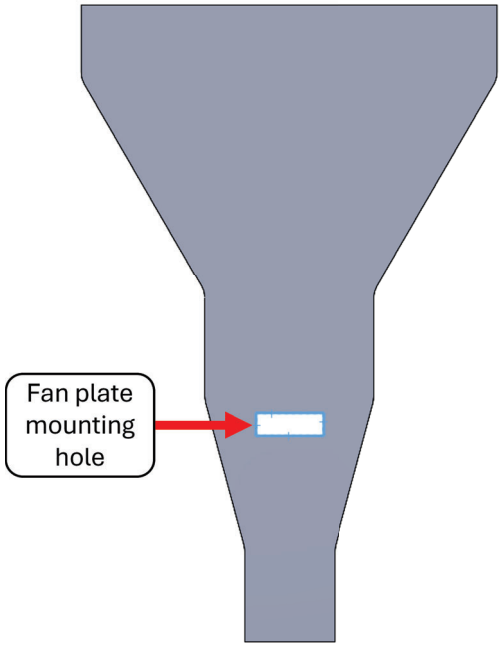

Figure 3.24: The diffuser side panels have a cutout which accommodates an insert that holds the diffuser fan plate.

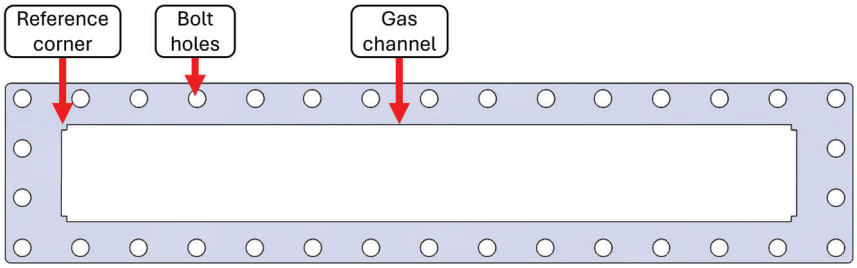

Figure 3.25: The diffuser bolt flange includes through holes matching the bolt pattern on the cooling chamber, a large channel for gas flow, and reference corner features in the corners of the gas flow cutout used to aid in positioning the diffuser panels during welding.

the connection between this plate and the diffuser bolt flange. The O-ring groove around the cooling chamber outlet seals the cooling chamber to this plate. Both O-ring grooves use the standard dovetail profile recommended for 1/8 in nominal O-rings by Parker [26].

The plate connecting the diffuser assembly to the ram's horns, shown in Figure 3.27, requires slightly more complexity. This plate requires multiple O-ring grooves because it connects to welded bolt flanges on both its top and bottom. These O-ring grooves also use a standard 1/8 in nominal O-ring dovetail profile as recommended by Parker [26]. O-ring grooves are not added to the welded bolt flanges before welding the plates to the recycle ductwork because they could warp during the welding process, affecting their ability to properly hold the O-ring used for sealing. Machining the O-ring grooves into the bolt flanges after welding would risk breaking the welds due to the vibrations created during the machining process. To accommodate O-ring grooves on both sides of the diffuser's top intermediary plate, the plate connecting the diffuser to the ram's horns is 50% thicker than the bolt flange plates, giving it a total thickness of 0.375 in. This plate has clearance bolt holes around its exterior to fit the #10 machine screws used to connect the ram's horn to the diffuser through this plate. The plate has cutouts that match the entrance area of the ram's horns on each side, while the central area remains closed to ensure most of the flow goes to the ram's horns. In the center of the plate, a circular hole is cut out and tapped for a 0.5

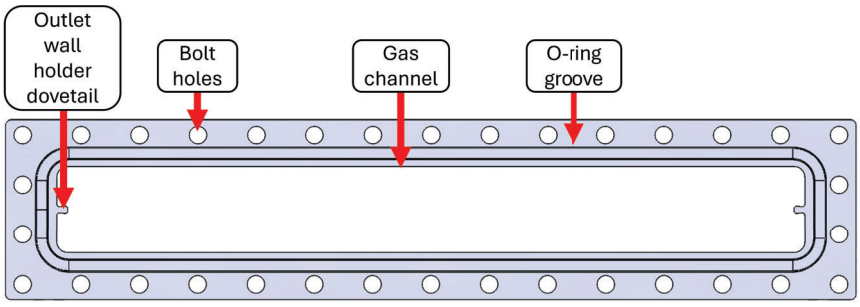

Figure 3.26: The intermediary plate between the diffuser assembly and the cooling chamber includes rounded dovetail features which hold the outlet wall locating pieces described in Section 3.4.2. This plate also includes a dovetail O-ring groove to hold the O-ring which seals this plate to the diffuser assembly.

in NPT fitting which serves as the diffuser (and recycle system) outlet. Used gas is expelled through this outlet to offset the fresh nitrogen feed and maintain balanced pressure within the system. On the plate's bottom face, which contacts the diffuser's outlet bolt flange, the O-ring groove traces a path offset from the inner perimeter of the diffuser bolt flange. On the plate's top face, two mirrored O-ring grooves trace paths offset from the inner perimeters of the ram's horn entrance bolt flanges. A line of bolt holes is partially drilled and tapped into the top face of this plate to enable full bolting of the ram's horn bolt flanges. These holes need to be partially drilled so they do not create leak paths through the chamber; the extra thickness on this plate helps provide sufficient depth for a 10-32 machine screw to have at least 3-5 threads fully tapped in these blind holes.

3.4.11 Ram's Horns

The ram's horns loop around from the top of the diffuser to the sides of the cooling chamber. The ram's horn consists of 5 main sections: straights at its entrance and exit, a half turn and expansion at its top, a vertical straight, and a quarter turn and expansion at its bottom. The ram's horn assembly is shown in Figure 3.28.

Like the diffuser, the ram's horn has straight sections at its entrance and exit (Sections 1 and 5 in Figure 3.28). The entrance straight section serves the same role as the straights in the

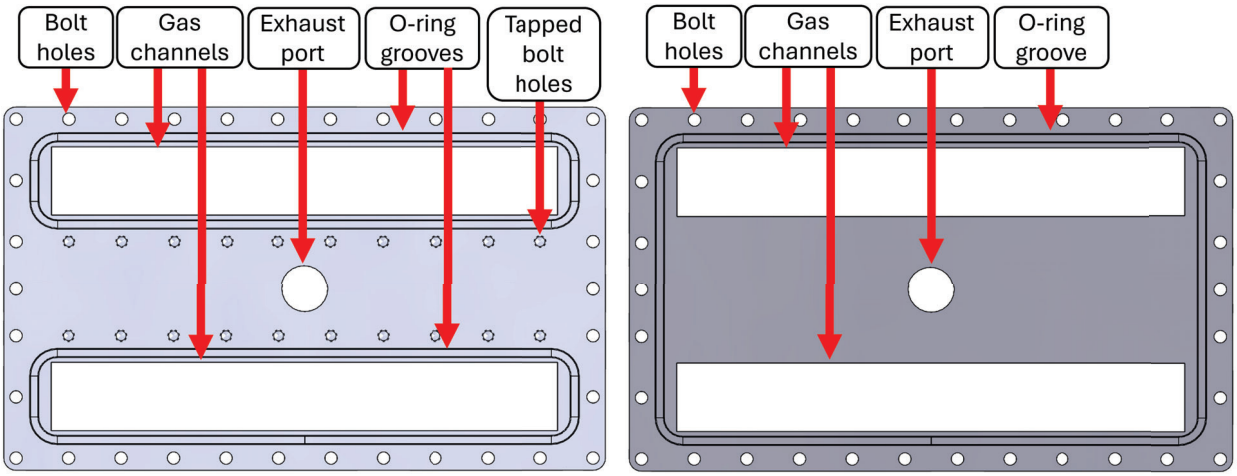

(a) (b)

Figure 3.27: The diffuser's top intermediary plate connects the diffuser to the ram's horns and the cooling chamber exhaust. The top side (a) includes O-rings for sealing the ram's horns to this plate. The bottom side (b) includes an O-ring for sealing the diffuser to this plate. This plate includes gas channel cutouts for the inlets to each of the ram's horns on its top and bottom and a central hole for the exhaust gas flow. The bolting around the outside perimeter of the plate uses through holes to attach the ram's horns to the diffuser. However, bolting is required in the central area to fully seal the connection between the ram's horns and this intermediary plate. The bolt holes which are above the diffuser's central area are blind tapped so they do not create leak paths for gas to flow out of the diffuser separately from its designated outlet ports.

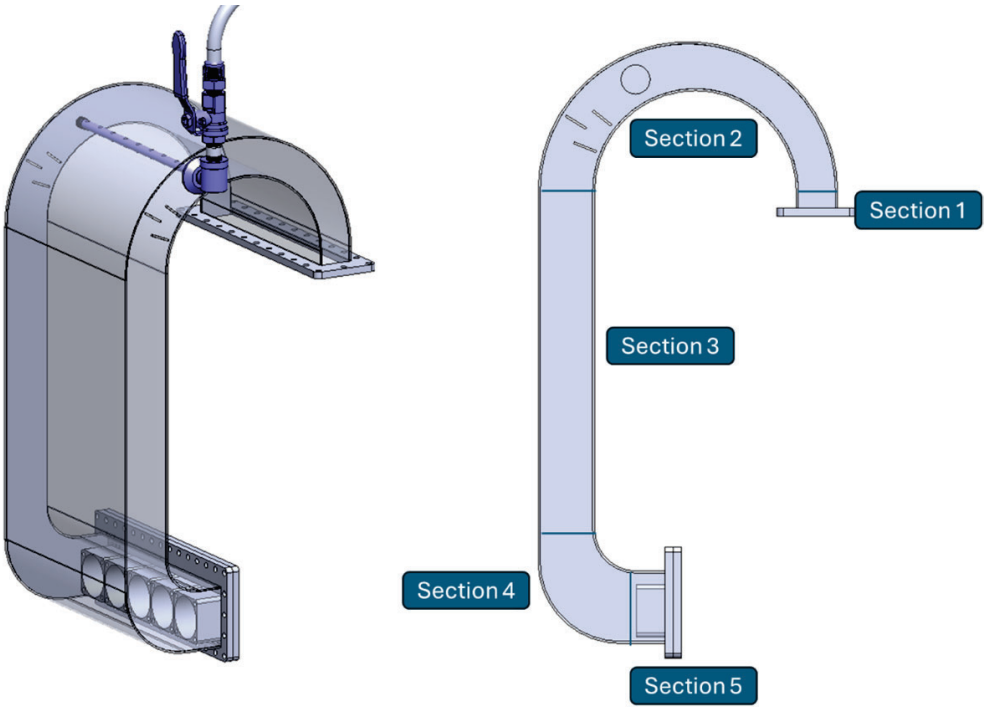

Figure 3.28: The ram's horn includes 5 sections: 1) a straight section used to weld on the bolt plate for connecting the ram's horn to the diffuser, 2) a 180 degree turn section which redirects flow from the diffuser back down towards the cooling chamber inlet and expands the cross-sectional area of the ram's horn to accommodate the cooling fans, 3) a straight section which carries the cooling gas down to the cooling chamber inlet location, 4) a 90 degree turn section to orient the gas flow to match the cooling chamber inlet direction, and 5) a straight section which accommodates the cooling fans and is used to weld on the bolt plate which connects the ram's horn to the cooling chamber. The straight sections at the ram's horn entry and exits and the long 180 degree arc also provide flexibility to the ram's horn to account for thermal expansions and contractions as the system undergoes thermal cycles.

diffuser, providing space for welding and bolting. Since the ram's horn has a constant width, the width at the entrance and exit of the ram's horn ductwork match. The ram's horn's entrance does not have external features that affect its height, so the width is determined based on the diffuser. The diffuser splits the gas flow from a single row of fans, so the area covered by these fans is used to determine the ram's horn entrance shape. The diffuser connects to two ram's horns, so the total area through which the recycled gas flows is double the cross section of a single ram's horns entrance. If the ram's horn entrance area is less than half that of the diffuser fans, then the gas will need to be compressed as it enters the ram's horn. If an individual ram's horn entrance area matches the diffuser fan area, then the gas will need to expand within the limited diffuser volume to double its volume to enter both ram's horns simultaneously. Thus, the height of the ram's horn entrance should be at least half the height of the fans but less than the full fan height to balance the expansion and contraction occurring within the duct. Accordingly, given the fan face length of 40 mm,

the ram's horn entrance width is 31.5 mm.

The half-turn and expansion at the top of the ram's horn (Section 2 in Figure 3.28) redirect the gas from blowing up away from the chamber through the diffuser to back down and towards the chamber again. During this turn, the ram's horn's cross-sectional area expands to match the area required for the fans. This expansion during the curve is created by having the outer curve diameter be slightly larger than the inner curve diameter and having its center slightly offset horizontally from the inner curve center. This top turn section includes the port through which fresh cold nitrogen is added to the chamber. This port is in the second half of the curve, past the curve's apex in the direction of flow. This positioning helps to reduce potential backflow of liquid nitrogen into the diffuser. The side panel has a circular hole cut out of it, where a 1/4 in NPTF threaded insert is welded. The manifold used to inject fresh nitrogen into the chamber is connected through this NPT fitting. The manifold itself and its mounting are described further in Section 3.4.12. This circular hole and corresponding insert are only added to one side of the ram's horn, where the manifold is inserted. The other side of the ram's horn has a solid surface to remain sealed. This section can optionally include features for trapping liquid nitrogen that is injected through the manifold to ensure that this liquid nitrogen does not travel all the way to the cooling chamber while still in liquid form. While these features are not always necessary, they are shown in Figure 3.28 to indicate where they would be positioned if used. These features include three walls which the gas is forced to maneuver around to travel through the ram's horn. While the gas can successfully flow around these obstructions, the heavier liquid nitrogen would fall out of the gas flow and get trapped behind the walls, separating it from the gas flow and preventing the liquid from reaching the cooling chamber.

The ram's horn's straight vertical volume (Section 3 in Figure 3.28) connects the top half-turn to the bottom quarter-turn. This section carries the gas back down from the height at which it exited the diffuser to the chamber's side inlet ports. While the gas travels through this section, the recycled gas mixes with the fresh nitrogen feed to create a uniform thermal profile.

The bottom quarter turn of the ram's horn (Section 4 in Figure 3.28) redirects the gas in the ram's horn back towards the side inlet fans. This curved section also expands the ram's horn cross sectional area to match the final straight sections dimensions. This expansion is done similarly to the top half turn, whereby the outer curve diameter is larger than the inner curve diameter, and the centers are offset. This secondary expansion is required for the ram's horn exit to have enough space to accommodate the fans and their wires, similarly to the diffuser as described in Section 3.4.9.

The ram's horn's bottom straight section (Section 5 in Figure 3.28) also creates space for welding a bolt flange and using tools to fasten the bolts connecting the ram's horn to the cooling chamber. However, it is longer than the ram's horn's entrance straight section because it also accommodates the fan volume. These straights are asymmetric on the top and bottom to match the expansion which occurs in the preceding quarter-turn section. This section determines the width of the ram's horn because it encloses the fans. The width is equal to the total fan width, plus a half millimeter between each fan, plus two millimeters on each side of the fans for fit and wire clearance. The nominal clearance between the ram's horn walls and the fans is larger than between the fans to accommodate the expected greater variance in the final positions of the welded components. The height of this ram's

horn section is set primarily based on the fan height and required space for wiring. The top edge of the ram's horn is nominally placed between the wire holes on the fan mounting plate so that the wires can travel from inside the ram's horn to outside the chamber. The bottom edge is spaced a few millimeters below the nominal fan bottom, both to provide clearance for manufacturing tolerance and to create some space for potential liquid nitrogen within the chamber to pool where it would not directly contact the fans.

The turns in the ram's horn use large radius curves, unlike the small radius curves used in the diffuser parts. Because the turns cover much larger angles, using larger turn radii reduces the losses created by redirecting the flow direction. The large diameters are also sized so that the outside panels can be formed using a roller bending tool.

The bolt flanges on the ram's horn entrance and exit have a rectangular outline of through bolt holes sized for #10 machine screws. The top bolt flange only includes the bolt holes, a rectangular cutout for gas flow, and corner features for welding like the diffuser plates included. The bottom bolt flange, shown in Figure 3.29 includes those same features, plus four holes for fan wires and a thermocouple. These holes are positioned such that they are aligned with the corresponding holes on the fan mounting plate, as described in Section 3.4.13.1.

3.4.12 Fresh gas feed

Fresh nitrogen is injected into the recycle system through manifolds in the ram's horns, as shown in Figure 3.30. These manifolds distribute the gas injection across the ram's horns' long width, leading to uniform temperature gas flow across the vials in front of the fans. The manifolds are made from aluminum tubes with evenly spaced holes drilled into one side. The tubes are threaded on each end. One end has an aluminum cap threaded on it to force the gas to flow through the manifold holes rather than letting it jet through the tube and hit the ram's horn side wall. The other side threads into an NPT elbow used to mount the manifold to the ram's horn. The manifold holes are pointed upwards relative to

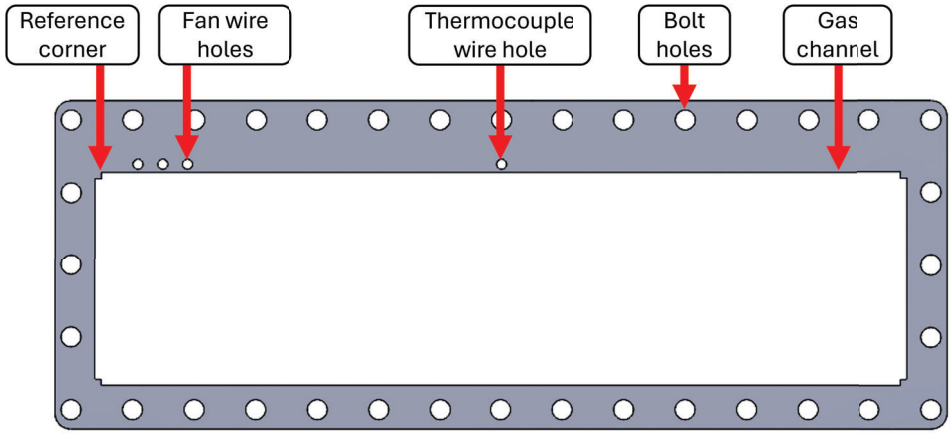

Figure 3.29: The ram's horn's exit bolt flange includes through holes for bolting, a central channel cutout for gas flow, corner features used for alignment during welding, holes for fan wire feedthroughs, and a hole for a thermocouple for measuring the gas flow temperature.

gravity. This orientation forces liquid nitrogen which enters the manifold to be carried by the gaseous nitrogen flow out of the manifold holes, leading the liquid distribution to match the gas distribution. When the manifold holes are oriented downwards, the liquid which enters the manifold is heavy enough to fall out of the flow preferentially through the first few holes, leading to a thermal gradient across the ram's horn. The manifold is also covered with a fine stainless-steel mesh which helps to break up large liquid nitrogen droplets as they exit the manifold. This mesh helps the liquid nitrogen quickly flash and cool down the gas in the ram's horn to reduce the risk of the liquid nitrogen traveling all the way through the ram's horn and into the cooling chamber.

The manifold is mounted to the ram's horn through a male-female aluminum 1/4 in NPT street elbow. The elbow, manifold, and cap are all made from aluminum, matching the ram's horn material. This material matching ensures that the components will not experience relative differences in thermal expansion, which would risk damaging components or compromising seals. The NPT elbow threads into the 1/4 in NPT insert welded into the side of the ram's horn. This connection is sealed with Teflon tape, and enough layers of Teflon tape are applied to ensure that the thread is tight when the elbow is rotated such that the female side faces upwards relative to gravity.

A ball valve is attached to the manifold elbow through a short pipe nipple. This ball valve enables flow balancing between the two ram's horns within the recycle system. If, during operation, one ram's horn within the recycle system consistently measures a colder gas temperature than the other, then the fresh nitrogen feed to that side can be reduced by partially closing the ball valve. This adjustability can account for differential flow preference between the two sides, asymmetric insulation properties, and potential variation in manu-

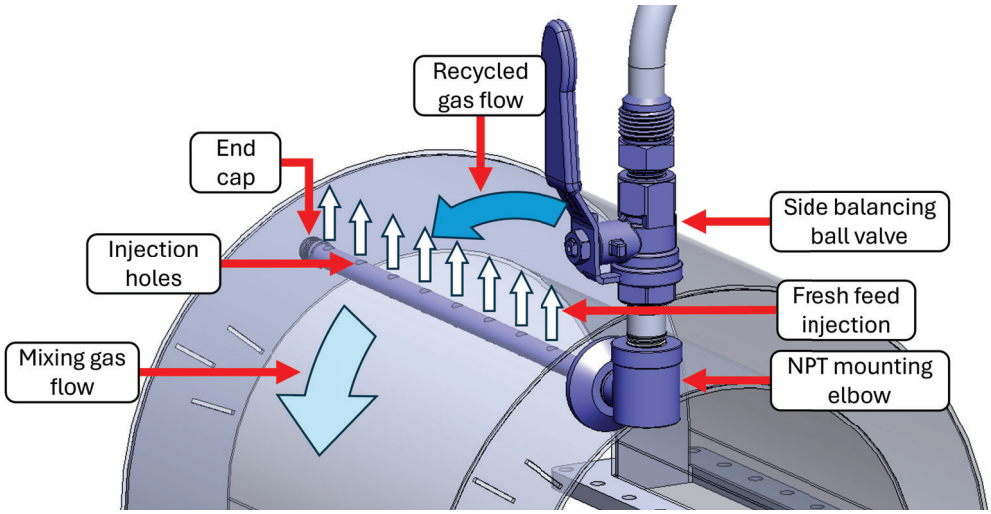

Figure 3.30: The fresh nitrogen feed is injected into the ram's horn through a manifold to evenly distribute it across the width of the ram's horn. This distribution improves the fresh feed and recycled gas mixing, enabling a relatively short mixing length while achieving relatively uniform temperature gas flow inside the cooling chamber. The ball valve upstream of the manifold is used to balance the flow between the ram's horns on each side of a cooling chamber.

facturing. In this application, the ball valves are manually actuated and set during initial startup and tuning, but in future applications they could be actively controlled to adjust to environmental condition changes or to provide an automatic shutoff in cases of pipe failure.

The ball valve inlet has an NPT to flare adapter. Flared piping is used for the nitrogen feed line because it easily connects to custom bent pipe sections, making it simpler to match the piping to the chamber geometry. Also, the flared connections are easy to replace, particularly when connecting to other rigid components, and any given section can be independently removed and evaluated without requiring the rest of the piping to be disassembled. Through these flared connections, straight tubing bent to a 90 degree angle connects the ball valve on each ram's horn to a joining tee. This tee connects to another ball valve, controlling the nitrogen flow into the chamber's recycle system. This ball valve is used for balancing the flow from a single force between multiple cooling chambers, similarly to the side balancing ball valves. From this valve, the recycle system connects to the full system nitrogen feed line. This connection may go through several branches depending on the number of chambers connected to each nitrogen feed line. The nitrogen source line is mixed gas and liquid flow, though at the source the flow may start as only liquid.

3.4.13 Fans

The convective gas flow in the cooling chamber is driven by fans. Fans are located at both the inlet and outlet of the chamber in both the ram's horns and the diffuser. Using fans at both the inlets and the outlet allows the fans to both push the gas into the chamber and pull it out, increasing the overall gas flow through the system. This fan configuration also improves the uniformity of the flow through the chamber, as the fans help reduce edge effects by forcing flow to move over the full face of the inlets and outlet.

The fans are each mounted to plates within the diffuser and ram's horn assemblies. The fan plates used for a standard cooling module without door features are shown in Figure 3.31. Most features and dimensioning described below are kept the same for fan plates for all chambers, with the overall fan plate width scaling with the number of fans. Similarly, fan-specific features numerically match the number of fans that fit on the inlet and outlet of each chamber type. The fan plate for the outlet is entirely within the diffuser assembly, while the fan plates for the inlets are between the ram's horns and the cooling chamber. The side inlet ports are sized based on the fans, so the fans can be placed directly at the inlets without notably affecting the flow. Each type of chamber has its own associated fan plate sized for the number of fans used on that chamber.

Each fan plate has five cutouts per fan – one for the fan's gas flow and four bolt holes for mounting the fan to the plate. These cutouts are patterned with a linear spacing of the fan width plus 0.5 mm. This spacing ensures that the fans will not interfere with each other while nesting closely to each other. The gas flow cutouts are 38 mm diameter holes. This dimension matches the outer diameter of the blades in the 40 mm x 40 mm fans. These 38 mm holes have straight chords at their tops and sides on the outlet fan plate to ensure the plate material thickness is at least 1.5 mm in all areas. The fan mounting holes are sized as standard clearance holes for #8 machine screws. The fans have holes in their corners that are tapped for 8-32 machine screws. The fans are bolted to their corresponding plates with 8-32 machine screws and lock washers. When running, the fans vibrate due to the motion of

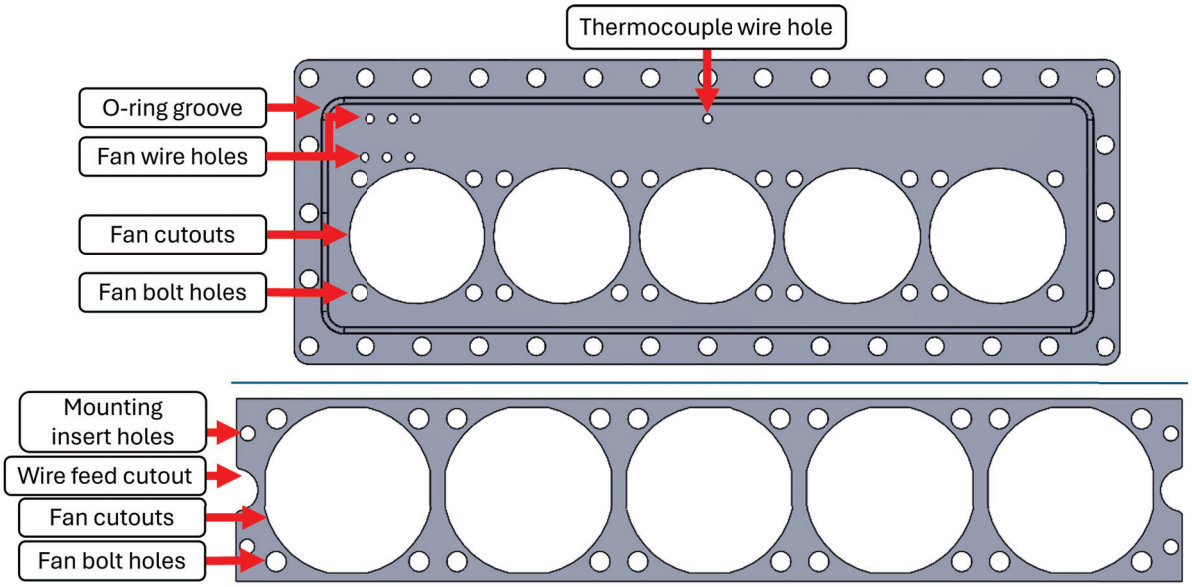

Figure 3.31: The fan plates used in the ram's horn (top) and the diffuser (bottom) include cutouts for the fan faces and through holes for bolting the fans to the plates. The ram's horn fan plate also serves as the intermediary plate for sealing the ram's horn to the cooling chamber, so it includes an O-ring groove, while the diffuser fan plate is located entirely within the diffuser assembly. The ram's horn fan plate includes holes for the fan wires and a thermocouple for measuring inlet gas temperature. The diffuser fan plate includes holes for mounting to the inserts that hold it in place in the diffuser and a cutout for the fan wires to feed out of the sides of the diffuser.

the blades. The lock washers act as preloaded springs, absorbing vibrations to ensure that the fans do not loosen their bolts.

3.4.13.1 Inlet fan plate

The cooling chamber inlet fans are mounted to a simple rectangular plate which connects to both the cooling chamber and the ram's horn, shown in Figure 3.31. This plate has through holes matched to the bolt holes on the cooling chamber for fastening this plate to the cooling chamber. The outside edge of the fan plate is set such that there is 1/16 in of material between the edge of the bolt holes and the fan plate. The fan plate also has an O-ring groove for sealing to the ram's horn's bolt flange. This O-ring is positioned between the bolt holes and the welded section of the ram's horn. Because there is expected variability in the welded surface of the ram's horn bolt flange, the O-ring groove does not overlap with the weld line. Due to the space restrictions imposed by fitting between the bolt holes and the weld line on the ram's horn flange, the O-ring groove is sized for a 1/16 in O-ring instead of a 1/8 in O-ring as used elsewhere in this system. A standard 1/16 in O-ring dovetail groove as recommended by Parker [26] is used for this feature.

The fan mounting features are located based on where they would position the fans within the inlet port. This dimension has some flexibility, allowing for repositioning the fans within the chamber by only changing the fan mounting plate. Based on the vial positions within

the chamber, the fans are located as close to the bottom of the inlet port as possible without risking interference between the fan mounting features and the inlet port, as shown in Figure 3.32. Accordingly, the distance between the fan mounting features on the inlet fan plate and the bolt holes for connecting to the chamber is set such that the bolts and lock washers have about 1 mm of clearance with the nominal bottom edge of the side inlet ports.

The inlet fan plate includes holes for fan wires and a thermocouple. The wire feedthrough holes are required for operating the fans, while the thermocouple feedthrough hole is used to insert the thermocouple used to measure the gas temperature for cooling chamber control. These holes each have a 0.1 in diameter, providing sufficient space for each required feedthrough while remaining tight enough to be sealable. The thermocouple hole is centered horizontally relative to either the middle fan when there is an odd number of fans, or above one of the two middle fans when there is an even number of fans. The middle location is used because it is expected that if the gas temperature across the inlet is not uniform, then the gas will be coldest in the middle of the fan array. For conditioning, the minimum temperature is more critical to measure than the average temperature to prevent spontaneous nucleation, so the measurement focus is on reading this minimum temperature. The thermocouple feedthrough hole is located vertically on the fan plate such that its top edge is nominally tangent with the top edge of the chamber's side inlet port. The thermocouple only passes through the ram's horn bolt flange, not the ram's horn internal volume, so it is positioned as far from potentially interfering with the ram's horn as possible.

Two sets of fan wire holes are used in the inlet fan plate – one set for feeding wires from inside the ram's horn to inside the cooling chamber, and one set for feeding wires from inside the chamber to the outside environment. The first set of holes is positioned just above the nominal fan positions, with about a half millimeter of clearance between the hole edge and the fans. These holes are kept close to the fans to keep them far from the weld line on the ram's horn, allowing the tolerance on that feature to remain loose for easier manufacturing. The second set of wire holes is aligned vertically with the thermocouple hole, as it has the same requirement of feeding wires from the chamber to the environment while minimizing the hole positioning effect on the ram's horn dimensional tolerance. Within each set of wire holes, the holes are spaced 1/4 in from each other so that they are grouped relatively closely but have enough space to avoid tangling during assembly. The fan wires are shown going

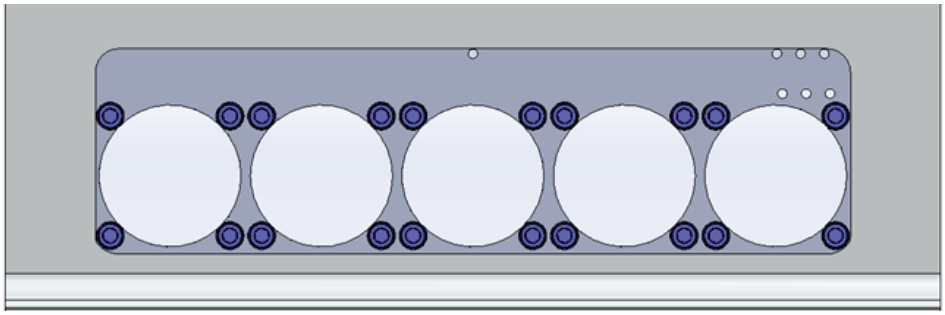

Figure 3.32: The cooling chamber inlet fans are positioned within the cooling chamber inlet cutout such that they are as low as possible within the inlet cutout without creating interference between the fan mounting bolts and the cooling chamber.

through these holes in Figure 3.33.

This wire feed arrangement allows the fan wires to avoid passing through the ram's horn itself, simplifying the parts used in the ram's horn's construction. However, the wires do still need to pass through the ram's horn's bolt flange. These wire holes are positioned inside the seal path created by the O-rings sealing the fan plate to the chamber and sealing the fan plate to the ram's horn's bolt flange. This positioning avoids the risk of gas leaking through wire holes positioned inside the ram's horn volume. However, the wire and thermocouple holes which connect the inside of the cooling chamber volume to the outside environment do create leak paths for gas to exit the chamber. Thus, these wire holes are sealed using a thixoflex sealant. The sealant used for these wire holes is the same material that is used to seal the machine chambers to the Planar Motors stators described in Chapter 2.

3.4.13.2 Diffuser fan plate

The diffuser inlet and outlets are both narrower than the fans used in this system. Thus, the fans are placed inside the diffuser. Placing the fans inside the diffuser rather than splitting the diffuser at the fans reduces the number of required sealing interfaces, however it increases the complexity of mounting the diffuser fans. Because the system sealing directly affects the system's cooling efficiency, the choice to increase mounting complexity to position the fans inside the diffuser is selected. The diffuser fan plate for a standard cooling chamber without door features is shown in Figure 3.31.

Since the outlet fans are placed within the diffuser, their mounting needs to take up as little additional space as possible. Accordingly, the fan plate is only as wide as the fans themselves, plus an additional millimeter to provide a half millimeter clearance on the tops and bottoms of the fans. Due to this tight dimensioning, the fan gas flow cutouts are adjusted

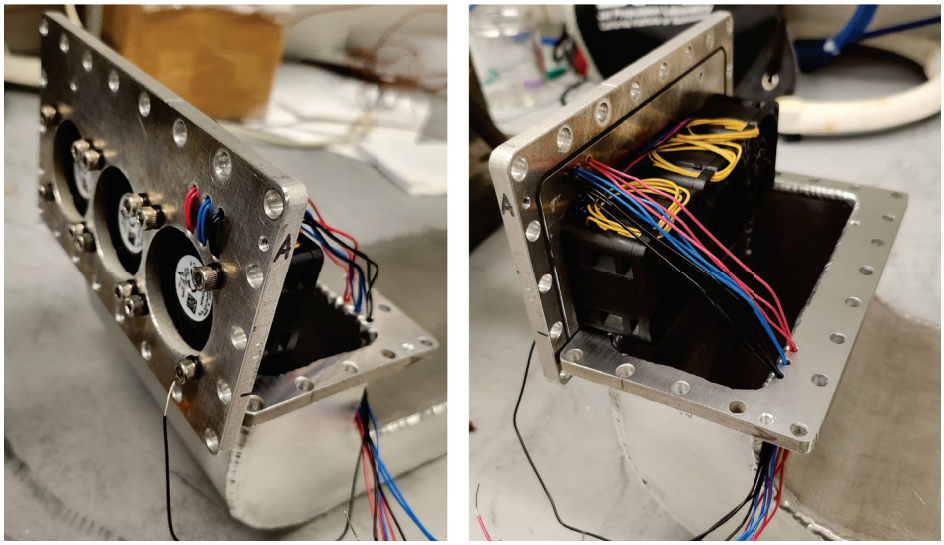

Figure 3.33: The fan wires go through the holes on the fan plate to the side inside the cooling chamber, then back out of the cooling chamber through the ram's horn welded flange. This wire routing strategy allows the fan wiring to be done before attaching the ram's horn to the cooling chamber.

to ensure the plate has at least 2 mm of material at the edges between the cutout and the plate's outer edge. This adjustment fills a minimal portion of the gas flow cutout, creating little effect on the gas flow but strengthening the plate for assembly and operation.

Because the diffuser fan plate rests within the diffuser, it needs features both to mount it within the diffuser and to feed the fan wires out of the diffuser. The fan plate mounting must attach to the sheet metal walls of the diffuser, but it also must be sealable to prevent gas leaks. Also, the fan plate needs to be removable so that the fans can be changed if one of them fails. An additional insert, shown in Figure 3.34, is used to attach the diffuser fan plate to the diffuser walls, meeting these requirements. The diffuser wall has a rectangular cutout into which the insert slots, providing vertical and horizontal support. Two 4-40 machine screws are used to fasten the diffuser fan plate to the insert using corresponding threaded holes in the insert. When the diffuser fan plate is bolted to the inserts on each side, the plate provides a constraint that prevents the inserts from falling out of the cutouts in the diffuser walls. The inserts also have a cutout for the wires to pass through from the inside of the diffuser out to the environment. Once the wires are fed through the insert cutout, the remaining space in the cutout and the gaps between the insert and the diffuser wall are sealed using silicone. These inserts avoid the need to bolt through the diffuser wall, which would require adding out of plane features to the diffuser fan plate.

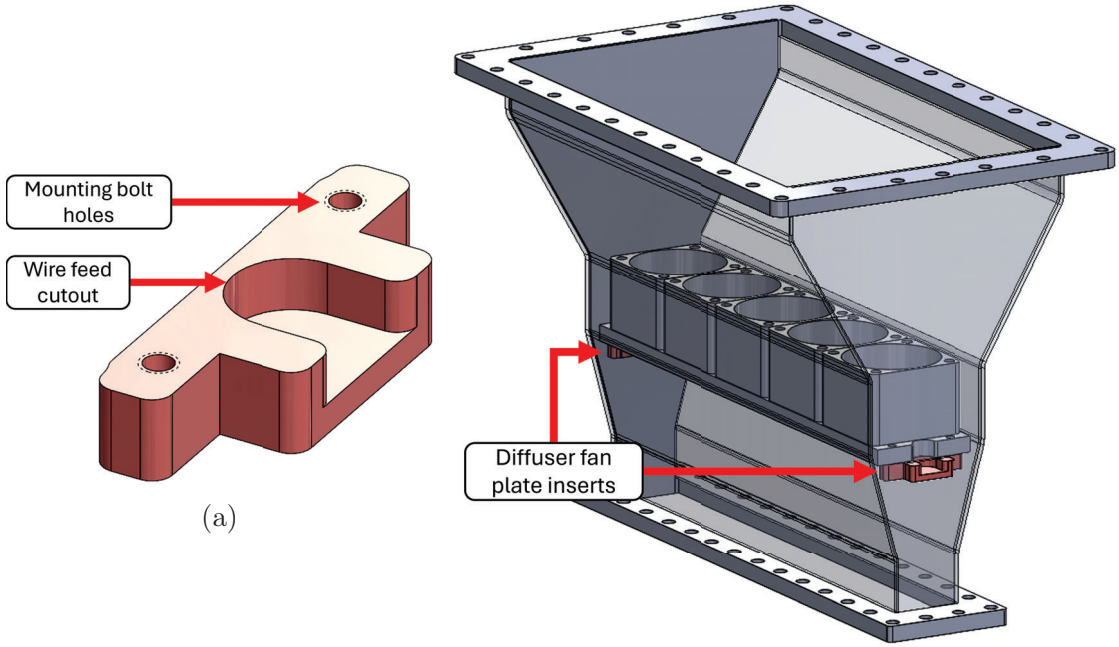

(a)

(b)

Figure 3.34: The diffuser fan plate is mounted inside the diffuser through an additional insert (a). This insert includes tapped bolt holes for mounting the fan plate and an open pocket for creating space to feed the fan wires through the diffuser wall. The diffuser fan plate is mounted on two of these inserts, one on each side of the diffuser.

3.5 Freezing Results

The freezing section of the lyophilizer is evaluated based on the speed at which it can cool vials and the uniformity of the cooling profiles of the vials throughout the system. The vial temperatures are considered sufficiently uniform if they are within 2°C of each other at the end of the cooling process. This metric is based on using a -4°C subcooling target temperature. The maximum 2°C variation allowance ensures all vials will be at least 2°C below freezing when entering nucleation, so they will all be subcooled to some degree. Also, the vials will be no more than 6°C below zero, limiting the risk of spontaneous nucleation. The cooling speed target is 30 minutes for each conditioning and freezing. Nucleation is considered successful if all vials nucleate at a temperature within 1°C of each other.

3.5.1 Uniformity Testing

The uniformity of the freezing system can be evaluated at multiple levels. While the critical parameter is the uniformity of the vial temperatures within the system, the convective gas parameter uniformity can be measured across various elements of the cooling system as a proxy. The first level of uniformity is across the face of one inlet into the chamber. The next level of uniformity is between two sides of the chamber. The final level of uniformity is between different chambers.

The first uniformity metric is an evaluation of the gas temperature and speed distribution across the face of a row of muffin fans. This metric ensures that the vials are receiving a uniform temperature exposure regardless of their position on a puck. This uniformity is most strongly influenced by the recycled gas and fresh feed mixing in the ram's horn before it enters the cooling chamber. This uniformity was improved by using the manifold described in Section 3.4.12, with the key features including evenly spaced holes, a stainless-steel mesh covering to help disperse liquid nitrogen droplets, and orienting the manifold holes up relative to gravity. This uniformity is most important at low production rates, where the vials spend more time in front of the same rows of fans. At higher production rates, the vials spend less time in specific locations within the cooling chamber because they are moving through a larger volume, so variation in cooling conditions across the inlet face are averaged out by the vials moving past each individual chamber inlet face. When operating at lower production rates, the importance of cooling chamber inlet face uniformity can also potentially be mitigated by implementing a carousel puck behavior keeping the pucks moving more constantly through the smaller volume.

The inlet face uniformity is measured by moving a sensor across the row of fans, generating a gas flow rate and temperature profile. The temperature is measured using a thermocouple, while the gas flow rate is measured using a hot wire anemometer. An example set of these measured parameter profiles appears in Figure 3.35. This figure shows a variation of temperature less than 1°C and a gas speed variation of less than 2.5 m/s across the face. The valleys of the gas flow velocity appear to correlate with the center and edges of the fans, which is expected for the sensor measurement level with the middle of the fans. Despite these dips, the gas flow velocity is greater than 1 m/s across the full inlet port, ensuring that all vials will see a sufficiently fast gas flow to meet the heat transfer requirements expected

based on the prior simulation work.

The second level of uniformity is the comparison of gas temperatures and speeds on each side of the system. This uniformity is controlled by adjusting the ball valves on the lines that divide the fresh nitrogen feed into the ram's horns on each side of the cooling chamber. Because the chamber is operating at a steady state condition, the valves can be adjusted in that steady state until the desired uniformity is achieved. A plot showing the temperature tracking of two sides of a single chamber is shown in Figure 3.36. The two sides of the chamber track each other within 0.5°C. The temperature measurement used for generating this plot is taken from the front of the middle fan on the inlet.

Given the levels of uniformity measured for these two levels, the expected total variation in gas temperature across all vials within a chamber is no more than 1.5°C. This total variation meets the target variation maximum of 2°C across the entire chamber.

The third level of uniformity is the comparison of gas temperatures between different chambers in the system. This uniformity is controlled by a pair of ball valves similarly to the side balancing within a chamber. These ball valves can be adjusted once the chambers are in a steady state, where the positions can remain constant because the system operates at constant set points due to its continuous nature.

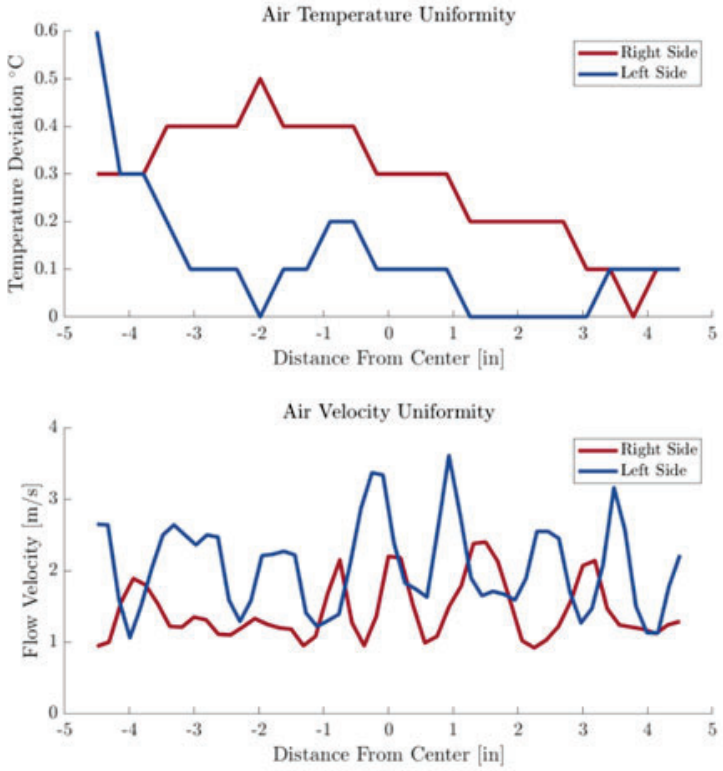

Figure 3.35: The air flow uniformity tests show that the gas temperature in front of the inlet varies by less than 1°C across the cooling chamber inlet face. The gas flow velocity uniformity shows more variation, but this variation results from crossing the borders and centers of the fans. Overall, this level of uniformity is acceptable for the freezing chamber.

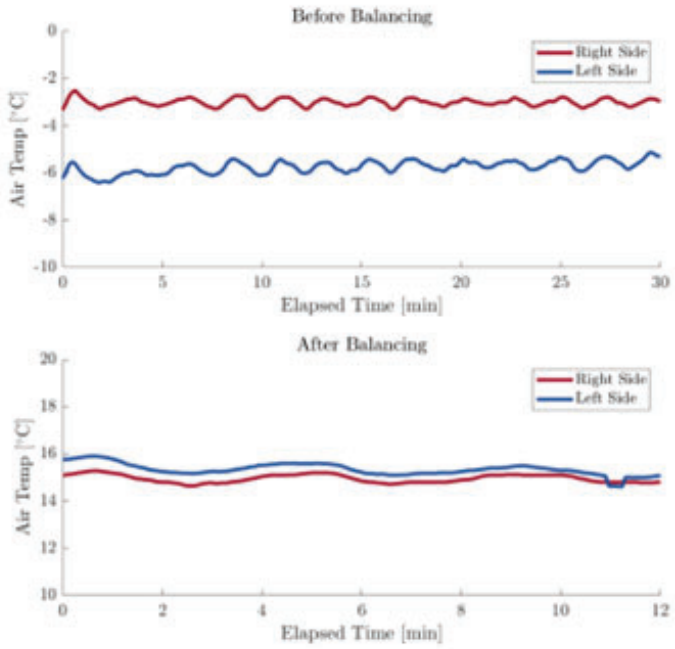

Figure 3.36: The gas flow into each side of the cooling chamber can be adjusted to match the gas temperatures on each side of the chamber to within 0.5°C. This uniformity is achieved through manually adjusting ball valves, and the balanced valve positions remain consistent between system shutdown cycles.

3.5.2 Cooling Rate Testing

Once the chamber uniformity metrics are evaluated, the chamber is then evaluated on how quickly it can cool vials to the required setpoints for conditioning and freezing. This test was performed with thermocouples placed in all 20 vials to directly measure their temperature. The presence of thermocouples in the vials is expected to affect the nucleation behavior of the vials because it provides a nucleation site for ice formation to start. However, it does not affect the bulk vial temperature behavior during conditioning and freezing. Each vial has a 3D printed cap with a hole in the top which holds the thermocouple in place, ensuring that it is located within the fluid in the vial without touching the vial walls. The thermocouple wires go through a KF25 port on the rear wall of the test chamber, then through slits cut into a 0.5 in plastic tube. A secondary section of the same plastic tube is latched over the thermocouple wire slits to seal these slits, preventing warm atmospheric air from leaking into the chamber and slowing down the cooling process. The thermocouple length inside the chamber is much shorter than the thermocouple length outside of it, so finning is not expected to be a significant issue affecting the temperature measurements.

This test is performed separately for conditioning and freezing. Due to the isolated nature of the chamber and the wiring used for the thermocouples, the chamber is cooled down along with the vials. Thus, these tests are more like temperature ramp tests than the step change in thermal conditions that vials experience in the actual system. Using the times derived from these experiments for cooling profiles for vials will then overestimate how long it takes for the vial temperature to settle, because the beginning of the cooling process will have a

larger temperature gradient between the cooling gas and the vials. This extra time provides a buffer for the vial temperatures to more confidently reach the chamber temperature setpoint when not actively monitored with thermocouples. The results of the conditioning test with full vial measurement are shown in Figure 3.37.

These results show that the conditioning stage takes about 30 minutes from the start of conditioning to the last vial settling at its final temperature. The temperature variation between vials at their settled temperatures is 1°C. These results match the expected cooling rate for the vials modeled as cylinders in crossflow as calculated in Section 3.2.4. Additionally, the final range in vial temperatures matches the temperature variation seen across the chamber inlet. This matching confirms the use of gas temperature as a surrogate for expected vial temperature uniformity during initial testing, and it verifies that vial temperature uniformity remains within the design specification of 2°C in the cooling chamber.

The freezing uniformity experiments use the thermal quench method to nucleate the vials after conditioning. The gas temperature is rapidly lowered to -80°C for 10 seconds, then raised back up to -10°C for slow ice crystal growth after nucleation. This slow ice crystal growth results in larger ice crystals, which leads to larger pores for drying, as discussed in Section 3.1. The ice crystal growth temperature is held until the vial temperatures start decreasing again from 0°C, which takes about five minutes. The chamber setpoint is then lowered to -50°C, and the vials undergo the freezing stage. The total freezing stage takes about 30 minutes from the start of freezing to the last vial settling at its final temperature, which matches the conditioning time because they are both saturating exponential processes. The temperature variation between vials at their settled temperatures is again 1°C. The difference in time between the first and last vials reaching their saturation temperatures is one minute. These results also match the expected cooling rate for the vials modeled as cylinders in crossflow as calculated in Section 3.2.4.

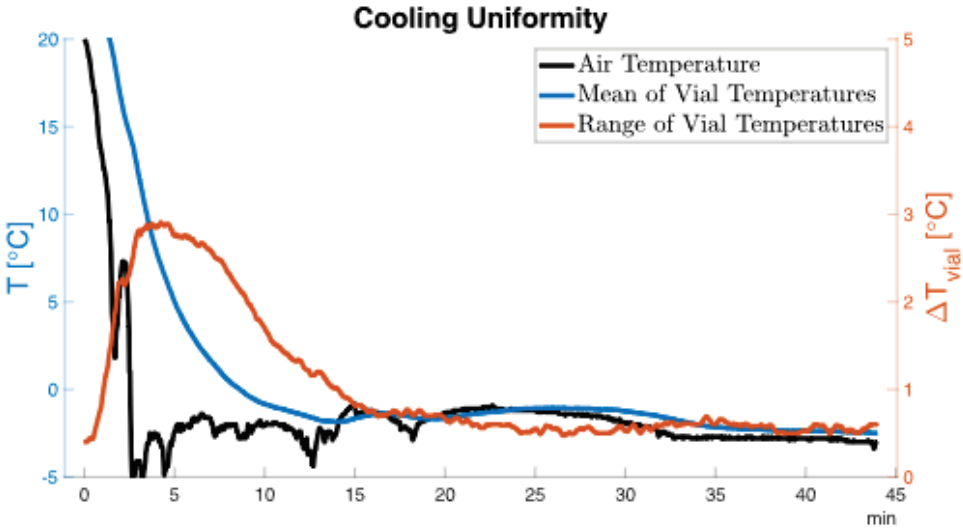

Figure 3.37: The measurement of all vials in the system shows that there is some initial temperature variation between the rows of vials which settles out as the vials reach the target temperature of -2°C after about 30 minutes.

3.5.3 Nucleation Results

The above results show that the designed and manufactured cooling chambers can achieve acceptable levels of uniformity to support controlled freezing for lyophilization. This temperature uniformity is most critical for preparing the vials for nucleation during conditioning. Once this uniformity was validated, the nucleation method tests could proceed without concern of variable vial conditions influencing the ability of different methods to successfully induce nucleation at relatively uniform conditions for the vials.

The nucleation method tested on this hardware is thermal quench. The convective gas cooling enables the vials to experience rapid changes in temperature to create the thermal quench nucleation without adding any additional hardware. Testing other nucleation methods requires additional specialized hardware specifically for that purpose, so they are not tested on this equipment as of this writing. Those capabilities can be added for future validation of the thermal quench nucleation method.

To prepare for the nucleation tests, the vials are filled with 3 mL of water and partially capped. The vials are then placed on a puck and loaded into the cooling chamber. The vials are then supercooled in a chamber with a setpoint of -4°C for 30 minutes. The nucleation method is then tested. There are two methods for measuring the nucleation methods. The first method uses thermocouples in the vials. These thermocouples affect nucleation by serving as ice nucleation sites within the vial. Thus, the results with thermocouples in the vials are not necessarily representative of the nucleation in samples without thermocouples, because the thermocouples make it easier for the solution to nucleate. However, this measurement method provides both nucleation time and temperature, so it can quantify the temperature range over which the vials nucleate. The second measurement method is to optically observe when the vials nucleate and measure the time between seeing the different vials nucleate. The time between first and last observed nucleation from the visual method can then be compared to the results from the thermocouple measurements. If the time between first and last nucleation is similar, then the nucleation temperatures can be expected to be similar as well. Nucleation could also be measured using a thermal camera to monitor the nucleation temperatures optically, but the camera focus could only capture a limited number of vials at a time. The visual method does not require placing anything inside any of the vials, ensuring that it does not interfere with the nucleation in the system. However, it does require having an optically clear interface through which the vials can be monitored.

The results from thermal quench testing with thermocouple measurement are shown in Figure 3.38. These results show that the vials nucleate within 15 seconds of each other, and the temperature variation between vials at their nucleation point is less than 1°C. This uniformity is sufficient to move forward with drying experiments to evaluate whether this level of uniformity is sufficient to create a tight bound on the vial drying times.

The thermal quench nucleation thermocouple results are confirmed by the visual tests. During the visual thermal quench nucleation tests, the vials all nucleate within 20 seconds. Given that this time spread approximately matches the results from the thermocouple experiments, the results from the thermocouple experiments are expected to translate to the case where the vials have only the solution in them.

The above results show that the cooling chambers can successfully condition, nucleate, and freeze the product in vials with low temperature variation. While the uniformity of

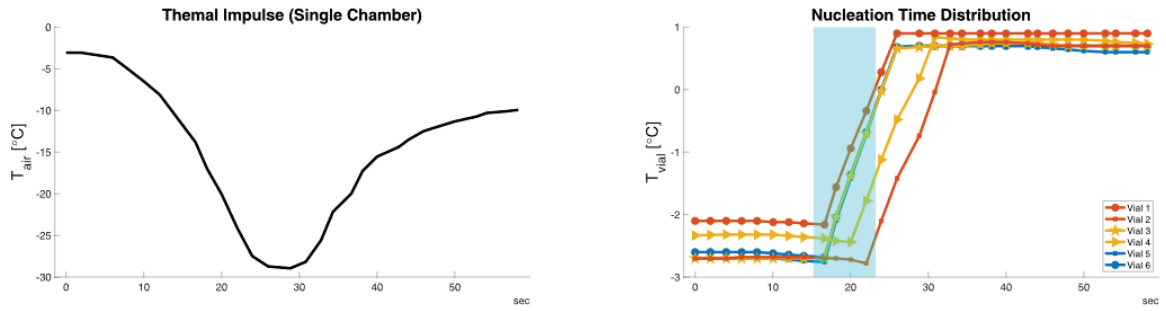

(a) (b)

Figure 3.38: Vials nucleated using thermal quench show nucleation occurs for all vials within a 15 second window, and the temperature range within the vials between nucleation events is less than 1°C.

cooling conditions can be measured during the freezing process, the evaluation of their effectiveness is shown by drying the product after this freezing process. The resultant total drying times and variation in drying times provides insight into whether the cooling process provides sufficient uniformity. The drying experiments performed after going through the freezing protocol established in this chapter are discussed in Chapter 5.

3.6 Future Work

While the cooling chamber developed for this system successfully enables continuous lyophilization, it has room for improvement. One area for future investigation would be looking into using different cooling mechanisms for conditioning and freezing. These two sections have notably different setpoint conditions, so different cooling methods may be appropriate for each section. For example, a flood cooling based method for conditioning could provide faster and more uniform cooling than the forced gas convection used in this work, and the fluid used for the flood cooling system would not need to achieve the temperatures seen during the freezing stage. However, such a method creates more complex sealing and piping concerns, as well as risks for contamination due to fluid evaporation. Another area for future investigation is redesigning the recycle ductwork and its interface with the cooling chamber. The recycle ductwork had a tendency to warp significantly during the welding process, making it difficult to ensure that the ductwork would not leak once it is assembled with the cooling chamber. This ductwork could be manufactured using extruded aluminum tubing with sharp corners to create a stiffer assembly for welding. The losses induced by using sharp corners to redirect the gas flow could be mitigated by installing turning vanes in these corners. In larger volume manufacturing, the recycle system components could be cast with flanged faces and machined for sealing, similar to the process used to make engine intake manifolds.

Furthermore, the cooling chambers could be linked, such that the exhaust gas flow from the freezing section can serve as the fresh cold source for the conditioning section. This linking would require using a controlled splitting valve assembly to allow for some of the freezing section exhaust to still vent to prevent overcooling the conditioning section, but

it could significantly reduce overall nitrogen consumption. Additionally, the cooling chambers could be coated with either additional insulation or a hydrophobic material to prevent condensation buildup between operations, as this condensation can risk damaging electronic components.

Chapter 4

Vial State Sensing System

In-situ product state measurement during lyophilization processes enables the use of a release criteria to optimize the time material spends in a freeze dryer. This specification increases the machine's production rate and, with sufficiently accurate state measurements, ensures product quality. This measurement must occur noninvasively, within a vacuum, and without introducing contaminants to the lyophilized product.

The vials have two key state variables of interest: temperature and residual moisture. The temperature measurement provides insight into the status of the vial during the lyophilization process. This insight includes when the vial is ready to transition between process states and verification that the vial has not exceeded its collapse temperature during primary drying. The residual moisture measurement indicates how far along in the drying process the vial has progressed and when it has finished the lyophilization process.

Current lyophilizers can include indirect thermal and in-situ residual moisture measurement methods. The thermal methods can provide individual vial data, but the current in-situ moisture measurements provide information based on all vials within a system, such as through laser measurements to measure moisture in the lyophilizer headspace [42]. While system-wide residual moisture measurements can serve as release criteria in batch lyophilization, a continuous lyophilizer requires localized information to distinguish finished products ready to exit from new products entering each process chamber. This work proposes a measurement system that directly records individual vial product mass as a direct proxy for residual moisture. An additional measurement system tracks vial temperature during the sublimation process. The temperature monitoring system implemented in this work is similar to the current state of the art, while the residual moisture sensing system is novel.

The mass measurement system uses a deflecting spring to identify changes in the product mass. This spring's deflecting element is a curved wire portion of a helical spring that experiences both bending and torsional stresses. This geometry enables the spring system to take advantage of the deflection efficiency of torsional stress, which acts along the entire length of the deflecting element, as opposed to pure bending deflection where the stress is localized at the beam's root. The curved wire is modeled as a cantilevered curved beam experiencing an out-of-plane loading force. As the force changes, such as during a change in mass during sublimation, the resultant deflection is expected to similarly change. This deflection is amplified through an extended sensing arm which also creates motion in the plane parallel to the tops of the vials, making the measurement visible from the top of the

vials.

The mass measurement system consists of four main elements. The spring wire is the key deflecting element which moves as the mass in the vial changes. Two springs are used in axially symmetric positions so that the vial can be positioned between them, leaving the vial accessible from above for other operations such as loading and capping. These wires are mounted in an aluminum base using a ball press fit to create a secure connection. The tops of the sensing arms on the wires have AprilTags attached which enable high fidelity position tracking using a video camera. The vial itself is placed in a basket that is suspended from the sensing wires. This suspension positions the vial system's center of mass below the system's center of stiffness, providing stability for the sensor as it moves through the lyophilizer. This measurement system is modular and sized for individual vials, so it can be stacked together in an array to create a full tray of vials.

The measurement system underwent validation testing through adding mass to the system while placed on a lab scale, providing a known mass change for the measured distance changes. When these measurements were compared to the linear model derived from the curved cantilever beam deflection, the residual errors were an order of magnitude larger than the desired sensitivity. These residuals were structured, not random, indicating some unmodeled phenomena contributing to the error. After relaxing some linearity assumptions in the modeling did not address this residual error, a polynomial fit was used to match the structured residual and added to the linear model. This combined model provides measurement reliability to 10 mg accuracy for a 3 g starting solution. While this measurement does not meet the desired 3 mg accuracy for a 3 g starting solution, it does provide direct insight into the sublimation process on a unit-vial basis which current systems do not capture. This insight can be shown through tracking mass change trajectories for every vial on a tray and using these mass measurements to determine when to release the tray from the lyophilizer. Releasing the vials immediately upon the sensor no longer resolving a vial mass decrease resulted in Karl Fischer titration residual mass measurements consistent with the end of primary drying but an incomplete secondary drying phase, confirming that further sensor resolution is needed to achieve precise identification of the end of the sublimation process.

4.1 Current Systems

Current lyophilization systems either measure residual moisture content after a completed lyophilization cycle, interrupt the cycle to remove vials with a sample thief to test representative samples, or measure the moisture in the system that is sublimating out of the samples [43]. These systems do not provide individual vial, in-situ residual moisture measurements. As such, lyophilization processes measured with existing methods lack the ability to differentiate the process state of individual vials in their systems, delaying the process to accommodate the slowest-drying units. This distinction has less of an effect in existing batch lyophilizers, where the next process step for the batch cannot start until all units in the batch are finished. However, in this continuous system, these measurements would either be insufficient to determine the state of specific vials or limit the production rate of the machine.

Standard direct residual moisture content measurement methods include Karl Fischer

(KF) titration, thermal gravimetric analysis, moisture analysis devices, near-IR spectroscopy, a pressure differential measurement, and tunable diode laser absorption spectroscopy (TDLAS) [44]. KF titration is a destructive process, where the sample goes through an oxidation reaction to measure moisture. Thermal gravimetric analysis measures mass loss while heating the sample to drive out the remaining moisture, where the mass lost represents the remaining moisture [44]. A moisture analysis device similarly heats the sample to evaporate water, then replaces the moist gas with a dry gas and uses a sensor to measure moisture loss. These methods both require sample destruction, like KF titration.

Near-IR spectroscopy measures the water signal in the infrared spectrum. The pressure differential measurement takes advantage of the water sensitivity of a Pirani pressure gauge. The Pirani gauge measurement depends on the gas composition, while a capacitance manometer measures the inert pressure. When these two gauges read the same pressure value, it is an indication that there is no longer water vapor in the chamber. This measured lack of water vapor means that sublimation has finished, as the vials are no longer producing water vapor through sublimation. TDLAS measures the concentration of various chemicals in a gas mixture, including water vapor, which can be used to identify the water content in the lyophilizer's vacuum system. When this water vapor measurement is sufficiently low, it is indicative that the sublimation process has finished. These measurements from these methods all combine information from all vials in the system, so they cannot differentiate the state of individual vials from each other.

A common method for measuring temperatures in lyophilization systems is thermal imaging, using cameras such as the FLIR A35 shown in Figure 4.1. Thermal imaging enables non-contact vial temperature measurement from outside of the system. In large commercial lyophilizers, this measurement is limited to only vials that are visible at the outside edges of the system. These vials may not be representative of the entire shelf because the vials at the shelf edges and shelf interior can experience notably different conditions, as described in Chapter 1. Additionally, while thermal cameras work well for measuring objects which are at much higher temperatures than their environment, the vials spend most of the time at a temperature lower than their environment. Thus, the thermal camera measurements must be carefully calibrated and tuned for accurate measurements.

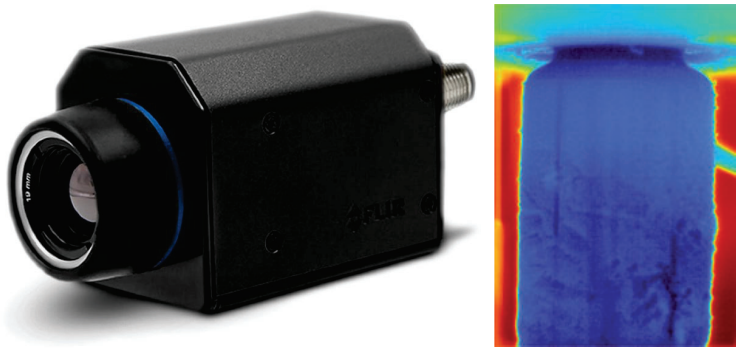

(a) (b)

Figure 4.1: The FLIR A35 thermal camera is used for temperature monitoring in this work. An example vial measured with this camera is shown on the right.

4.2 Vial Mass Measurement System Functional Requirements

A unit vial based, in-situ residual moisture sensing method is needed to enable optimal product release on continuous lyophilizers. Measuring the weight of each vial to directly measure residual moisture content based on water mass removal during sublimation is an obvious choice, but devising the method by which such measurements could be taken proved to be a significant challenge. A viable solution was achieved by suspending each vial on compliant wires such that an external camera could view optical fiducials (for this work, AprilTags are used) attached to the wires to assess wire deflection and hence vial weight.

This vial mass measurement system has the following functional requirements: 1) The sensor system must operate in a vacuum environment, as this is necessary for the drying phase to occur through sublimation. 2) The sensor must provide individual vial data. Systems which measure results at the machine level do not provide sufficient information to distinguish which vials are ready to move to subsequent stages from those which are not. 3) The sensor must resolve mass changes smaller than 3 mg. This metric is based on detecting 2% residual moisture of a 5 wt% 3 mL sucrose-mannitol lyophilized cake. This particular target is used for developing the specific sensor built in this work, but the corresponding design spreadsheet allows for changing the system parameters to meet mass change resolution requirements that may differ for other formulations. 4) The sensor must be vial agnostic. To enable use with a variety of potential formulations and to ease integration with pharmaceutical production, the sensing system must not depend on a specific vial type or size. 5) The sensor must be nondestructive, and it cannot contaminate the products. This requirement rules out invasive sensors and systems that incorporate excessive moving parts which could introduce particulate contaminates to the system.

Of the existing residual moisture sensing methods described above, the gravimetric analysis element of thermogravimetric analysis shows the most promise. Mass measurement meets the above requirements, and its simplicity provides design flexibility for integration. The mass removal element of thermogravimetric analysis occurs naturally within the lyophilizer, though the moisture is removed through sublimation rather than evaporation. The other methods either do not meet all the requirements or require highly specialized and expensive equipment. This weight measurement can be implemented by either periodically moving the vials to a scale within the system, or the scale can be integrated into the vial tray. Moving the vials to a scale within the system requires additional actuation, and the motion creates another potential source of particulate generation if it involves sliding interfaces. Integrating the mass measurement system into the tray requires space for this system, which reduces vial density. However, that reduced vial density is already required for convective gas cooling, as described in Chapter 3. Thus, the measurement system is built into the tray holding the vials. The tray integration is done on a modular unit vial basis, allowing for an assembly of these vial measurement units regardless of the tray size. This modularity provides flexibility in the tray design and geometry for different system configurations.

4.3 Sensitivity

This weight sensor uses an optical measurement system to record vial mass information from the vacuum chamber. The optical measurement does not require any specialized vacuum passthroughs, as only an optical viewing window is necessary. Additionally, this measurement does not interfere with the lyophilization process, nor does it introduce potential contamination sources. This system includes two optical fiducials which move based on the system weight. These reference markers are 3 mm edge length AprilTags. AprilTags are standard optical fiducials that are easily tracked in a camera image. Given a camera field of view covering one vial, the characteristic optical window dimension is expected to be around 40 mm. Using a standard camera with HD resolution of 1920 x 1080 pixels, the camera can resolve distances at about 21 μm/pixel. To achieve a resolution of 1 pixel per the 3 mg minimum resolution target, the desired motion of the optical measurement datum is 7 μm/mg, or 7 mm/g. This requirement can be relaxed by using a higher resolution camera or video tracking to achieve greater centroid tracking confidence.

A key goal for this weight sensing system is to create a linear relationship with applied load. A linear relationship has lower risks of over-fitting data to stochastic information. Cantilevered beams in deflection and torsion have a linear relationship between their load and motion, making them a desired geometry for investigation.

4.4 Spring Geometries

A wide variety of spring deflection systems were prototyped to try to achieve this resolution, utilizing combinations of cantilevered and torsional beams to function as spring elements. Systems that incorporate torsional motion tend to deflect more than those that rely primarily on bending deflections. This difference comes from the efficiency of stress distribution in the beam. In bending, the stress is concentrated near the root of the beam, which limits most of the deflection to that region. However, in torsion, the stress is experienced over the full length of the beam, allowing more of the beam to contribute to the deflection.

Adding an extended sensing arm is a common method for amplifying motion for measurement. In these deflecting systems, the sensing arm is typically implemented by adding a straight beam extending from the end of the loaded beam. When the end of the beam deflects, the tip moves to an angle relative to the fixed base. This angle, multiplied by the length of the straight beam extension, moves proportionally to the cantilevered loading motion. In torsional deflection, the motion at the tip is only angular rotation, so the extended sensing arm changes that twist into linear motion.

The final spring wire mass measurement system uses a portion of a helical spring as the primary deflecting arm, and it includes an extended sensing arm to amplify and change the plane of the motion for the visual sensor. The curved wire of the helical spring takes advantage of torsional loading efficiency to increase system sensitivity while maintaining a small footprint within the system. It is also simple to manufacture using conventional wire bending techniques. The expected deflection for this system is predicted using a model for the deflection of a curved beam experiencing out of plane bending, as derived by Timoshenko [45]. The problem setup appears in Figure 4.2.

Timoshenko shows that the strain energy for the curved beam can be found using the relationship in Equation 4.1:

$$U = \int_0^x \left(\frac{M_x^2}{2EI_x} + \frac{M_z^2}{2GI_p} \right) r d\psi \tag{4.1}$$

Where M_x is the bending moment, E is the beam's Young's Modulus, I_x is the beam's second moment of inertia, M_z is the torsional moment, G is the beam's shear modulus, and I_p is the beam's polar moment of inertia. These moment loads are functions of the distance between the root and the end of the beam, and they can be expressed by the relationships shown in Equation 4.2.

$$M_x = -Prsin(\alpha - \psi), M_z = Pr[1 - cos(\alpha - \psi)] \tag{4.2}$$

The strain energy can be separated into its bending (U_x) and torsional (U_z) components, shown in Equation 4.3.

$$U_x = \int_0^\alpha \frac{M_x^2 r}{2EI_x} d\psi, U_z = \int_0^\alpha \frac{M_z^2 r}{2GI_p} d\psi \tag{4.3}$$

Taking the partial derivative of strain energy with respect to the load, P, gives Equation 4.4:

$$\delta = \frac{\partial U}{\partial P} = \frac{Pr^3}{EI_x} \int_0^\alpha sin^2(\alpha - \psi) + \frac{EI_x}{GI_p}[1 - cos(\alpha - \psi)]^2 d\psi \tag{4.4}$$

The explicit bending and torsional deflections can be split up as before, as seen in Equation 4.5:

$$\delta_x = \frac{\partial U_x}{\partial P} = \frac{Pr^3}{EI_x} \int_0^\alpha sin^2(\alpha - \psi) d\psi, \delta_z = \frac{\partial U_z}{\partial P} = \frac{Pr^3}{GI_p} \int_0^\alpha [1 - cos(\alpha - \psi)]^2 d\psi \tag{4.5}$$

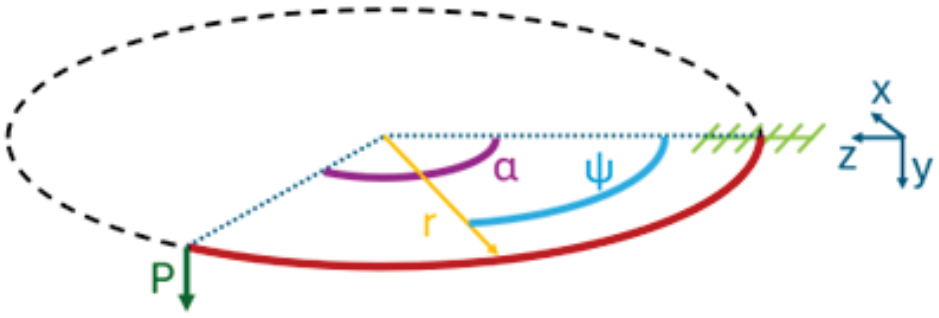

Figure 4.2: A diagram showing a curved cantilevered beam loaded perpendicular to the curvature plane. P is the out of plane force, D is the diameter of the cantilevered beam (assuming a circular cross section), α is the total angle of the wire arc, and r is the radius of curvature of the wire arc.

Using the choice of a circular wire geometry, the moments of inertia can be substituted into Equation 4.5 to get the final deflection relationships shown in Equation 4.6:

$$\delta_x = \frac{32Pr^3}{\pi E d^4}(\alpha - sin(\alpha)cos(\alpha)), \delta_z = \frac{8Pr^3}{\pi G d^4}(6\alpha - 8sin(\alpha) + sin(2\alpha)) \tag{4.6}$$

Where r is the curvature radius of the beam, E is the beam's elastic modulus, d is the wire diameter, G is the beam's shear modulus, P is the out of plane load on the end of the beam, and α is the total angle of the curved beam. These equations provide valuable insight into the key variables which determine the system performance by showing which variables have exponential effects on deflection.

These equations demonstrate that the system deflection has a power relationship with the wire diameter and the arc of curvature bent into the spring wire shape. Thus, these parameters serve as the main driving tools in designing the wires for these springs. Decreasing the wire diameter creates the largest potential increase in deflection, but it is limited by keeping the wire in its elastic deflection regime. Additionally, a thinner wire can be more difficult to mount rigidly into a fixed base, which is required to create the cantilevered beam condition. The wire chosen for this system has a .014 in diameter to balance this preference for a thinner wire and the requirement for a thick enough wire for mounting system flexibility.

The wire deflection and stress relationships can be put into a design spreadsheet to select wire parameters for manufacture. The spreadsheet used for this system appears in Table 4.1.

The spring wires primarily use heat-treated stainless-steel for its strength and elasticity, which allows the wires to deflect repeatably under load while maintaining enough strength to prevent yield. The spring wire tested in this work is a .014 in diameter CH900 stainless steel wire with a 140 degree wire bend arc, 16 mm wire arc radius of curvature. This wire is predicted to have a linear vertical deflection at the tip of 1.68 mm for a 3 g load split between two wires. This deflection results in a vertical deflection rate at the wire tip of 0.56 mm/g. This deflection rate does not meet the system functional requirements, so some form of amplification is required. This amplification can be done using extended sensing arms. As the tip of the curved wire deflects downward, it also experiences a slope change which is amplified by the perpendicular sensing arm. These bending and twisting slopes at the end of the wire can be explicitly calculated as shown in Equation 4.7 from [46], where the specific definitions for coefficients C_1 and C_2 can be found in Equation 4.9. The addition of this calculation to the design spreadsheet is shown in Table 4.2. These two slopes create linear motions that are perpendicular to each other. The bending slope creates linear motion that is parallel to the end of the curved wire, while the torsional slope creates linear motion that is perpendicular to the end of the curved wire. Thus, their total linear motion is equal to the square root of the sum of the squares of each of these motions.

$$\theta_{\text{torsion}} = \frac{Pr^2}{EI}(C_1 \cos(\alpha) + C_2 \sin(\alpha)) \tag{4.7}$$

$$\theta_{\text{bending}} = \frac{Pr^2}{EI}(C_2 \cos(\alpha) - C_1 \sin(\alpha)) \tag{4.8}$$

Parameter	Name	Value	Units	Formula
Wire Arc Radius	R_w	16	mm	
Number of Wires	N_w	2		
Total Vial mass	m_v	3	g	
Force per wire	P	0.01	N	$:= m_v * 9.8/1000/N_w$
Wire Elastic Modulus	E_w	194000	N/mm^2	
Wire Poisson Ratio	v_w	0.3		
Wire Shear modulus	G_w	74600	N/mm^2	
Wire diameter	D_w	0.014	in	
		0.3556	mm	
Wire arc	arc_w	140	degrees	
		2.4	rad	
Wire tip bending deflection	δ_B	0.58	mm	$:= 32 * P * R_w^3/(\pi * E_w * D_w^4) * (arc_w - SIN(arc_w) * COS(arc_w))$
Wire tip torsion deflection	δ_T	1.10	mm	$:= 8*P*R_w^3/(\pi*G_w*D_w^4)*(6*arc_w - 8*SIN(arc_w) + sin(2*arc_w))$
Wire tip total deflection	δ_z	1.68	mm	$:= \delta_B + \delta_T$

Table 4.1: This table shows the z deflection calculation based on the theoretical deflection of a curved cantilever beam loaded perpendicularly to the plane of its curvature using the parameters used in this work. This calculation is based on 3 g of mass change, representative of the 3 g of water that is sublimated from a vial used during the development of the lyophilizer in this work.

$$C_1 = \frac{1+\beta}{2}(\alpha)\sin(\alpha) - \beta[1-\cos(\alpha)] \tag{4.9}$$

$$C_2 = \frac{1+\beta}{2}[(\alpha)\cos(\alpha) - \sin(\alpha)] \tag{4.10}$$

$$\beta = \frac{EI_x}{GI_p} \tag{4.11}$$

The extended sensing arm is 50 mm long, which provides about a 2x amplification to the measured motion, increasing the expected deflection rate to 1.21 mm/g. The system is assembled such that there are two wires bearing the mass, and these wires are moving in opposite directions. Thus, the effective motion measured by the optical system is doubled, resulting in a deflection rate of 2.42 mm/g. While this result is about less than the required motion described above, that calculation is based on using a FHD camera with a 1920 x 1080 resolution to capture images. Using video captures rather than single images allows for averaging of the optical fiducial centroid over multiple frames. This averaging enables a higher precision on the fiducial centroid than individual pixels. Given a video length of 120 s, the AprilTag positional accuracy improved to 0.1 px, reducing the motion requirement

to 0.7 mm/g. Combining the video accuracy improvements with the increased deflection rate from the sensing arm results in a system with a factor of safety of 3 on the expected

Parameter	Name	Value	Units	Formula
Wire Arc Radius	R_w	16	mm	
Number of Wires	N_w	2		
Total Vial mass	m_v	3	g	
Force per wire	P	0.01	N	$:= m_v * 9.8/1000/N_w$
Wire Elastic Modulus	E_w	194000	N/mm^2	
Wire Shear modulus	G_w	74600	N/mm^2	
Wire diameter	D_w	0.014	in	
		0.3556	mm	
Wire bending moment of Inertia	I_w	0.0008	mm^4	$:= \pi * D_w/64$
Wire torsion moment of Inertia	J_w	0.0016	mm^4	$:= \pi * D_w/32$
Wire stiffness ratio	β	1.30		$:= (E_w * I_w)/(G_w * J_w)$
Wire arc	arc_w	140	degrees	
		2.44	rad	
Wire tip bending deflection	δ_B	0.580	mm	$:= 32 * P * R_w^3/(\pi * E_w * D_w^4) * (arc_w - SIN(arc_w) * COS(arc_w))$
Wire tip torsion deflection	δ_T	1.10	mm	$:= 8*P*R_w^3/(\pi*G_w*D_w^4)*(6*arc_w-8*SIN(arc_w)+SIN(2*arc_w))$
Wire tip total deflection	δ_z	1.68	mm	$:= \delta_B + \delta_T$
Slope Constant 1	C_1	-2.89		$:= (1 + \beta)/2 * ((arc_w) * COS(arc_w) - SIN(arc_w))$
Slope Constant 2	C_2	-0.49		$:= (1 + beta)/2 * (arc_w) * SIN(arc_w) - \beta*(1-COS(arc_w))$
Wire tip bending slope term	θ_B	0.047	rad	$:= (P * R_w^2)/(E_w * I_w) * (C_2 * COS(arc_w) - C_1 * sin(arc_w))$
Wire tip torsion slope term	θ_T	0.055	rad	$:= (P * R_w^2)/(E_w * I_w) * (C_1 * COS(arc_w) + C_2 * sin(arc_w))$
Sensing arm length	L_{arm}	50	mm	
Bending slope motion	δ_{s_B}	2.349	mm	$:= \theta_B * L_{arm}$
Torsion slope motion	δ_{s_T}	2.761	mm	$:= \theta_T * L_{arm}$
Total tip motion	δ_s	3.625	mm	$:= SQRT(\delta_{s_B}^2 + \delta_{s_T}^2)$

Table 4.2: Adding a 50 mm sensing arm provides about a 2x amplification in the motion from the wire tip deflection to the sensing arm tip motion. The bending and torsional contributions to the slope at the wire tip are much closer in magnitude than their relative contributions to the wire tip deflection.

resolution compared to the requirement.

4.5 Scale System Structure

The final scale system consists of three main components: the spring wire, the vial basket, and the sensor base. These components are shown in an example scale system assembly in Figure 4.3. Each of these components must be designed to meet its sub-requirements, and they must interact accordingly with the other sensor components as needed. The spring wire is rigidly attached to the wire base, and the vial basket connects the spring wire to the measurement subject – in this case, vials of product to be lyophilized. This sensor is attached to the moving tray through an assembly called the "puck furniture," which is detailed in Section 4.6. Decoupling the scale from the mounting hardware attaching it to a puck allows for each to develop independently and minimizes the changes to interfaces between different components.

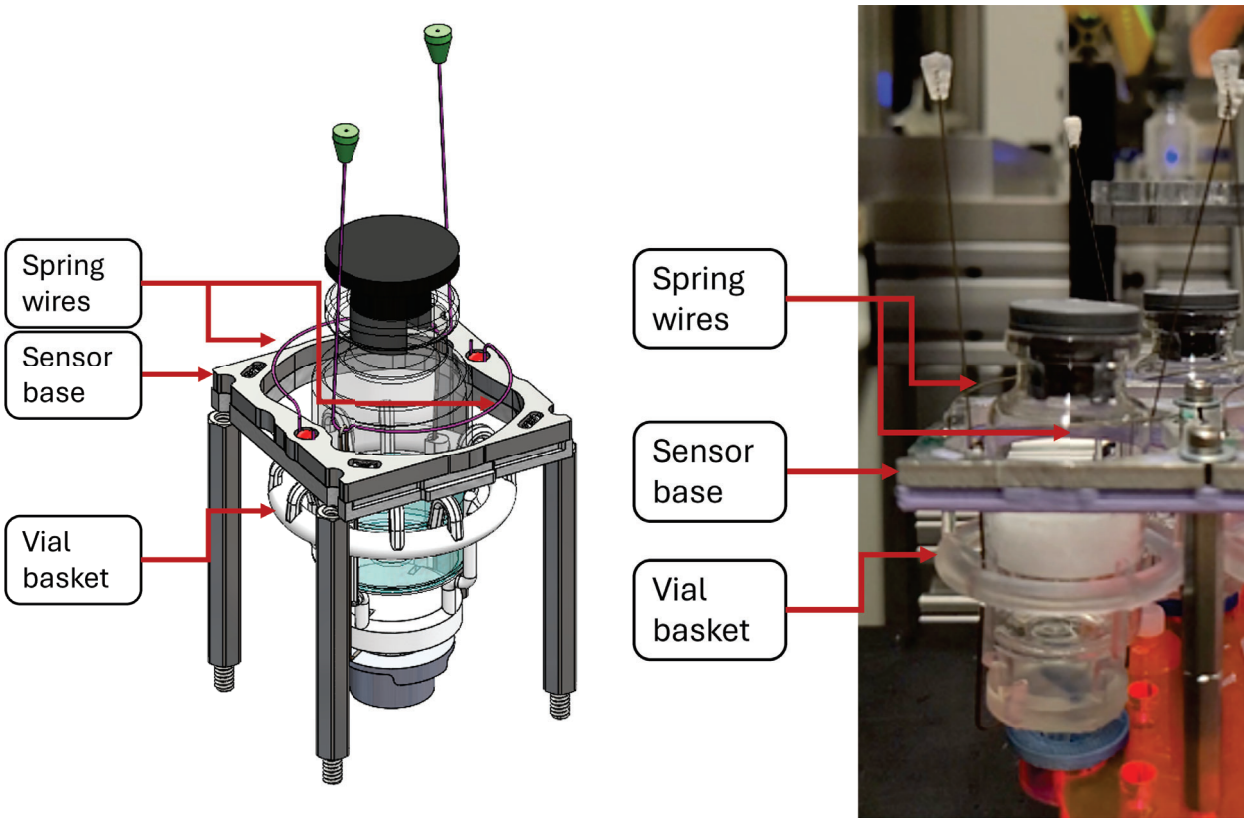

(a) (b)

Figure 4.3: The full weight sensing system assembly consists of the spring wires, the sensor base in which the wires are mounted, and the vial basket which holds the vial.

4.5.1 Spring Wire

Once the spring geometry has been calculated, it needs to be integrated into an actual wire that can be installed in the lyophilizer. This final spring wire, shown in Figure 4.4, includes four key features. The first feature is the deflecting arm which incorporates the curved wire geometry calculated in Section 4.4. This arc serves as the main deflecting element in the wire to create the motion used to measure the mass changes. The second element is the sensing arm which is used to amplify and change the plane of the deflection motion for measurement. As discussed in Chapter 2, the system includes multiple rows of vials, which means that not every vial can be seen from the side. The motion of the tip of the sensing arm is perpendicular to the vial's vertical deflection because the sensing arm is oriented perpendicular to the curved beam plane. This shift ensures that the sensing tip motion will be parallel to the top of the vial so that the motion can be imaged from above the vials. In this plane, the measurements for neighboring vials will not block each other.

The third element is the mounting crook used to attach the wire to the sensor base. This mounting is described in Section 4.5.2.1. This crook is bent in a plane that is offset from perpendicular to the plane of the deflecting arm arc, as shown in Figure 4.5a. The full

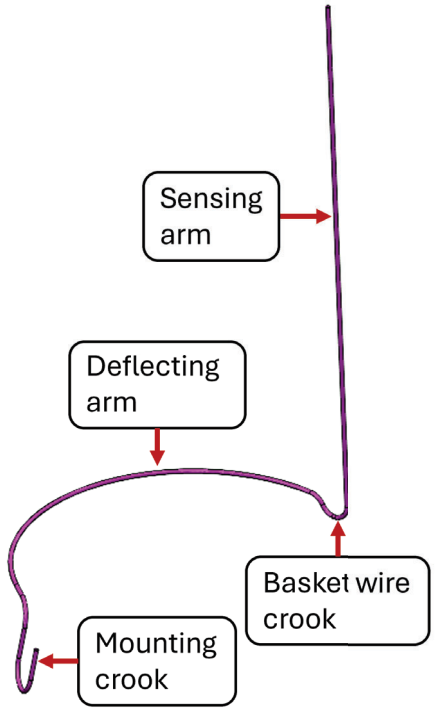

Figure 4.4: The final spring wire geometry consists of the deflecting arm of the wire, which has a 16 mm bend radius over a 140 degree arc, the 50 mm long sensing arm used to amplify the spring deflection and change the motion plane to be parallel to the top of the chamber, the mounting crook used to mount the spring wire to the wire base, which consists of a 4 mm nominal diameter semicircle and two 3 mm straight wire sections, and the basket wire crook, which enables rolling contact between the basket wire and the spring wire, preventing sliding which would reduce repeatability and risk particulate generation.

mass of the vial plus the product causes the beam to experience a large deflection, which is necessary to achieve the desired product mass sensitivity. Because the sensing range is only a subset of the total mass applied to the system, it is useful to preload the wire so that it is deflecting in its most sensitive motion regime during measurement. The spring wire sensor is expected to be most sensitive when the deflecting arm of the wire is parallel to the ground, as that is when the loading at the tip has the longest moment arm and will create the most deflection, as shown in Figure 4.5b. The most precise measurement is needed when the water has sublimated from the vial, leaving behind only the product. At this point in the process, the system mass is dominated by the vial and its cap, so only those masses are used in calculating the wire preload angle. Using a combined mass of 12 g for these components and following the same calculations shown in Table 4.2 show an expected tip deflection of 6.7 mm from its neutral position. The tip of the curved arc is 28.26 mm from the mounting features. Using these two measurements, the angle between the deflected tip position and the mounting position is 13.4 degrees. Thus, a 13.4 degree tilt is applied to the wire as the θ_0 shown in Figure 4.5a so that the wire deflects to a flat position with an empty vial and cap loaded.

The fourth key wire element is the basket wire crook from which the wire basket hangs to load the spring wire. This crook is rounded so that as the tip of the spring wire deflects, the basket wire will roll rather than slide along this crook. This rolling motion is designed to prevent the system from sliding at this interface, which would introduce nonuniformity and particulate generation.

At least two curved beams are needed to keep the vial level while it is suspended in the vial basket. If only one beam is used, then the slope created by the beam's deflection would also cause the vial to tilt. This tilt could be mitigated by hanging the vial directly under the wire crook, like a bird feeder, but doing so would limit access to the top of the vial. The top of the vial should not be blocked because the vials need to be capped while under vacuum in the system, and this is easiest to achieve by keeping the area above the vial clear. As long as at least two wires are used and the vial is supported between them, then the deflection will not necessarily tilt the vial. While more wires can improve vial stability, they also distribute the load of the vial and will each deflect less. Thus, only two wires are used to meet the minimum number of wires needed for balance. These wires are positioned in axially symmetric positions to avoid biasing the vial in any direction.

4.5.2 Sensor base

The sensor base is the spring wire mounting fixture. This base provides the resistant forces and moments needed to keep one end of the wire fixed while the other end deflects under the vial and product weight. A diagram showing the wire base with key features called out is shown in Figure 4.6. All of the base features are designed in 2D to enable rapid and large scale manufacturing with methods like laser cutting. These bases are 3 mm thick to ensure the press fit on the 3.5 mm ball is fully engaged with the wire while maintaining a thickness that is easily available for laser cutting.

The sensor bases used in this system include two types of features. One type is used directly for the spring wire deflection system, while the other enables a modular assembly of these sensor bases during development. The former features would carry through to an

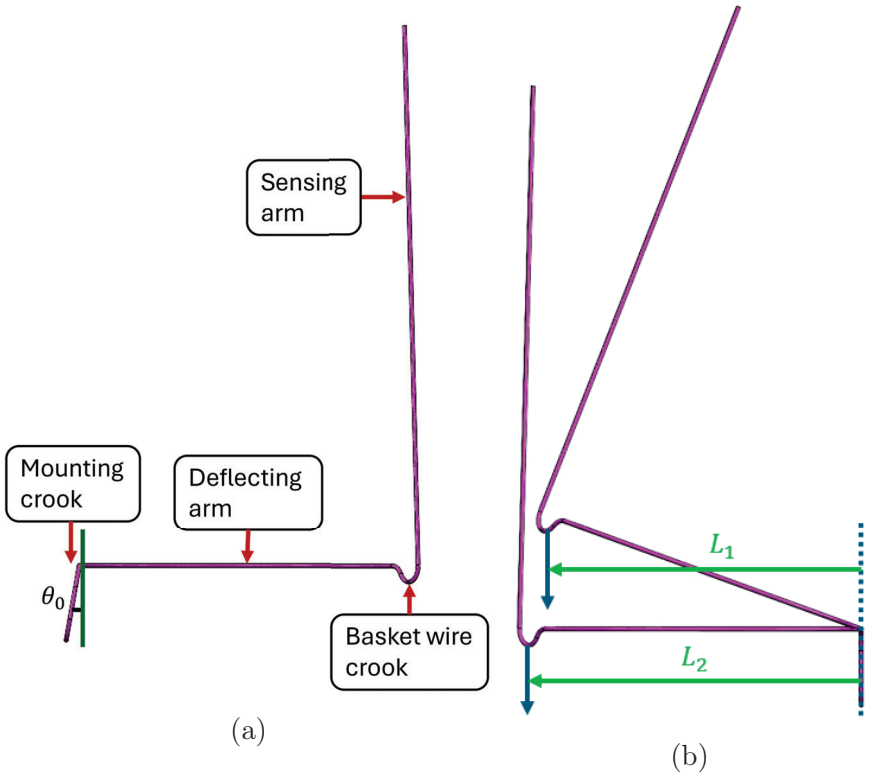

Figure 4.5: The wire mounting features are bent to an angle offset from perpendicular to the curved beam deflecting arm (a). This offset sets a preload angle for the wire, which is calculated based on the predicted deflection in Section 4.4, The initial angle is set such that the preload mass of an empty vial and cap loads the deflecting arm to its horizontal position, where the measurement is most sensitive. When the deflecting arm is angled upwards or downwards, the moment arm is shortened (b), decreasing the system sensitivity to load. The θ_0 calculated for the wires used in this system based on the nominal 3 mL filled 10R vials is 13.4 degrees.

industrial design, while the latter features might be removed or simplified at larger assembly scales. The features used for the spring wire deflection system include the cutouts for the ball press fit used to hold the wire and the central body cutout which creates space for the vial basket assembly to deflect vertically through the base. The features used to enable modular assembly include the assembly mounting cutouts which correspond to tabs on a separate assembly base and the securing bolts cutouts used to clamp the wire bases to the separate assembly base. In an industrialized system, these wire bases and the separate assembly base could be combined into a single part to reduce the number of components and assembly steps.

The overall size of the sensor base is constrained by the spring wire geometry and the Planar Motors puck size which is discussed in Chapter 2. The spring wire geometry sets a minimum unit size based on the deflecting arm's 16 mm bend arc radius. Given that geometry, the base cannot have a characteristic length less than 32 mm. Since the Planar Motors puck used in this system has a characteristic width of 120 mm, only three bases with a

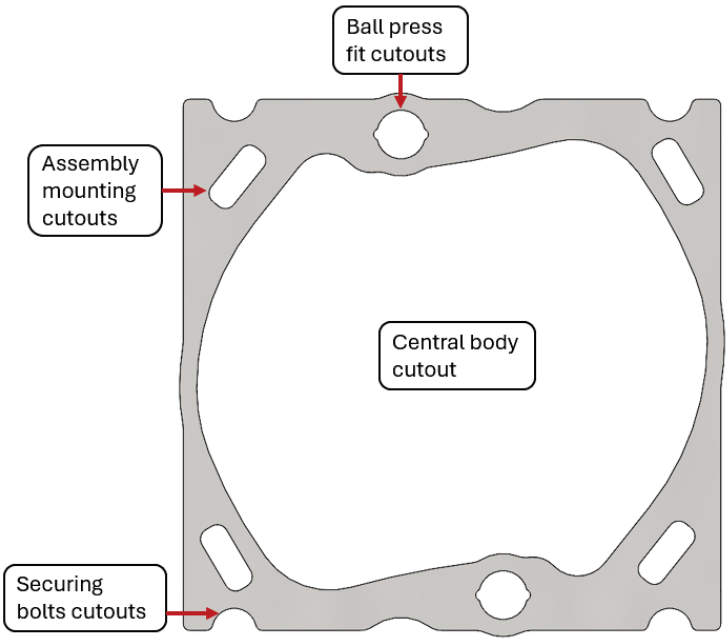

Figure 4.6: A top view of the wire base. The wire base only includes 2D features to improve manufacturability. Key features include the cutout for the ball plug press fit which holds the spring wires in the base, the central body cutout through which the vial basket assembly deflects vertically, assembly mounting cutouts used to attach individual wire base modules to a larger assembly, particularly during system development, and securing bolts cutouts which are shared between consecutive base modules.

32 mm characteristic dimension would fit across a puck. The mass measurement system must be contained within this dimension because it corresponds to the width of the tunnel system. If this width is exceeded, then the puck would not be able to enter the tunnel. Accordingly, the nominal characteristic dimension for the base is scaled up to 40 mm to maximize the usage of this available space. While this scaling causes the units to overhang the puck in its long dimension (180 mm), this overhang is acceptable because it extends in the tunnel direction. The potential layouts based on a weight sensing system characteristic dimension of 32 mm and 40 mm are shown in Figure 4.7. For manufacturability, the nominal dimension is scaled down to 39.5 mm to provide 0.5 mm tolerancing on the actual unit dimensions. System symmetry sets the unit length to match the unit width, setting its maximum at 39.5 mm as well. To maximize space for needed features, a 39.5 mm x 39.5 mm square area is allocated to the wire base.

The final wire mounting uses a press-fit ball mount to secure the wire to a sensor base, as shown in Figure 4.8. The feature on the sensor base consists of a circular hole for the ball and a slot to accommodate the wire. The slot is slightly undersized relative to the wire's nominal dimensions to ensure that the wire will elastically deform the base, creating the clamping force on the wire. The base must have a lower elastic modulus than the wire to ensure it elastically deforms to create the clamping force. Since the wire is made from steel, the base is made from aluminum. The slot width was determined experimentally by cutting a series of different width slots and pressing the ball and wire into each to determine the fit.

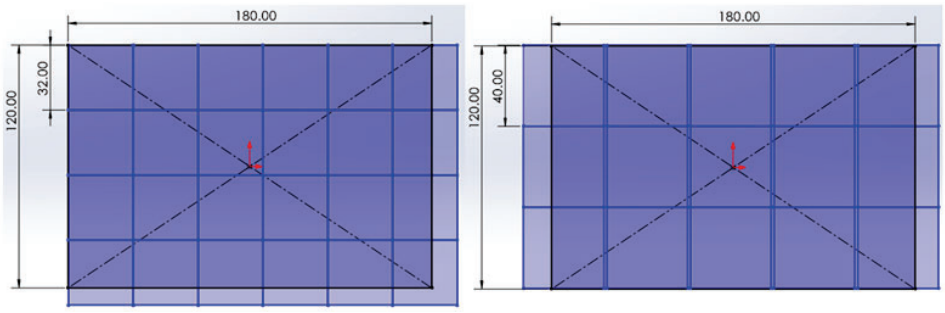

Figure 4.7: Four rows of potential sensor bases with a 32 mm characteristic dimension would overhang the 120 mm width of a puck (left), preventing the assembly from fitting in the width of the tunnel system described in Chapter 2. Since the maximum number of rows on these pucks is thus limited to three, the characteristic sensor dimension can be scaled up to 40 mm. This 40 mm sensor size would overhang the puck's 180 mm length (right), but this overhang is acceptable because it is in the same direction as the tunnel.

The optimal overlap between the ball and the wire was found to be 0.14 mm on each side of the ball. The circular hole for the 3.5 mm ball is slightly oversized to 3.6 mm to ensure that the majority of base deformation occurs where it engages with the wire.

The mounting features are positioned such that the distance between the basket hanger crooks in the spring wires is equal to the diameter of the deflecting arm's curved arc, and the line connecting the basket interface points is perpendicular to the basket hanger crook. These considerations ensure the basket hanger is positioned to achieve the designed rolling contact at this interface, rather than stochastic sliding. The mounting features are also positioned axisymmetrically about the center of the wire base, as shown in Figure 4.9, to ensure the vial is located centrally in the base.

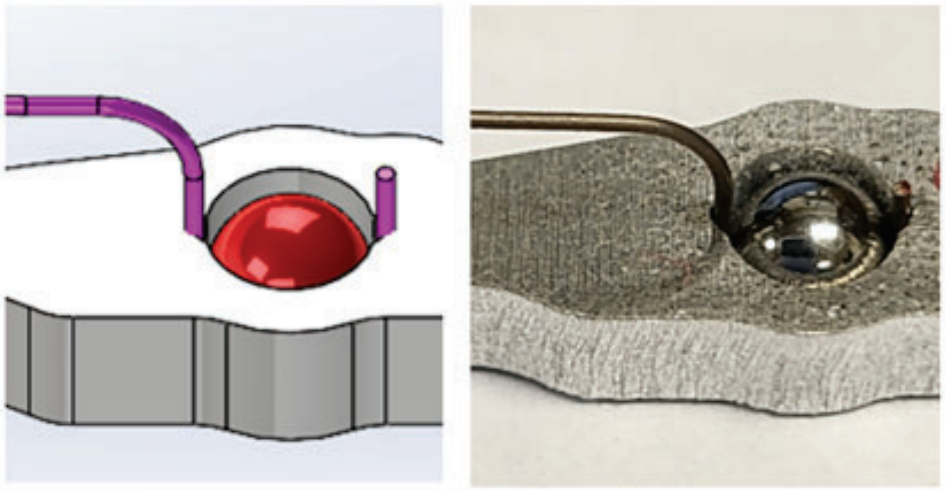

Figure 4.8: The model (left) and as built (right) ball plug press fit mounting used to mount the spring wires to the wire base. This mounting provides a rigid support for the deflecting spring wire.

The central body cutout on the sensor base is sized to ensure that the spring wires can freely move into and out of the cutout as they deflect. As discussed in Section 4.4, the spring wires' deflecting arm has a 16 mm radius. The center of this arc is offset from the sensor base center by 2 mm based on the positioning of the wire mounting features on the wire base. Accordingly, the central hole is based on an 18.5 mm radius circle to provide clearance between the wire and the walls. At the sensing end of the spring wire, the basket wire crook extends a straight section tangent to the 16 mm radius curve. On each point in the sensor base where this hook would be located, there is an additional cutout to accommodate this extra straight length. At the mounting crook, the ball press fit features on the sensor base infringe on the 18.5 mm radius cutout. Because the wire deflection is minimal closest to the mounting point, the cutout is not needed in this area because the wire will not bend low enough to contact this part of the base. These constraints form the overall shape of the wire basket accommodation cutout, which is visible with the wires in Figure 4.9.

To maintain flexibility in developing these spring wire systems, the sensor bases include cutouts that enable each sensor unit to connect to a larger sensor base assembly component. This sensor base assembly component holds a full row's worth of vials and includes locating features for the assembly. Separating this piece from the wire bases allows the sensor bases

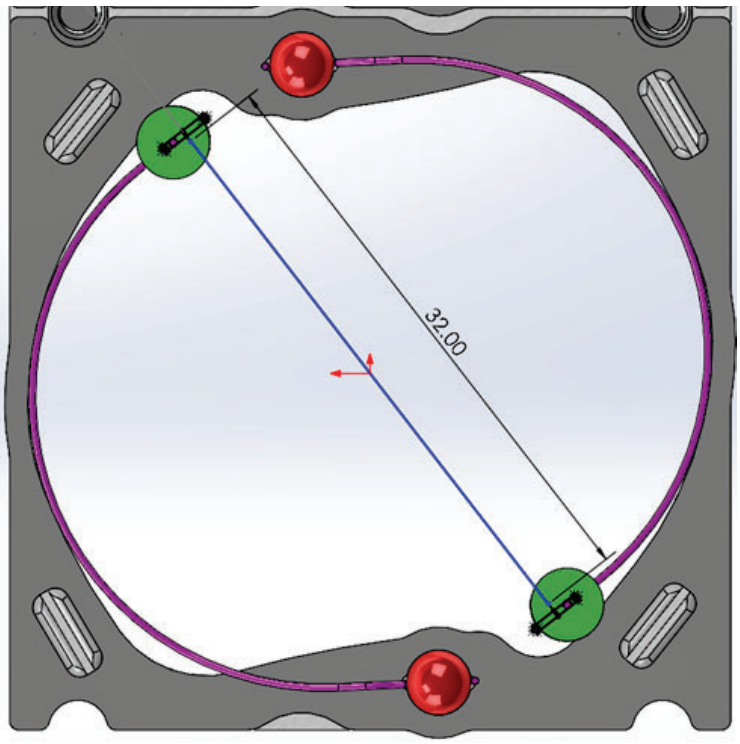

Figure 4.9: The top view of the sensor base assembled with sensing wires shows their axisymmetric positioning. The ball press fits are located such that the opposite wire ends are perpendicular to each other. This perpendicularity is needed for the basket wire to have rolling contact where it connects to the sensing wire. The distance between these basket wire crooks with this mounting is 32 mm to match the deflecting arm's 16 mm radius of curvature.

to be designed at the unit vial level, and the number used in a larger assembly can match the size of any puck geometry. Additionally, the sensor bases can be swapped out of the assembly if they are damaged or later improved without affecting the rest of the system. Four assembly mounting cutouts, one in each corner of the square sensor base, are used to attach the sensor base to the puck furniture. These cutouts are sized based on the philosophy of maintaining wall thicknesses of at least 1.5 mm while maximizing their size such that the corresponding mating features can be as large as possible for improved manufacturability. These assembly mounting cutout slots are also angled towards each other to provide appropriate constraint such that elastic averaging between these mounting points can consistently locate the wire base. Ideally, these slots could be arranged in perpendicular pairs, but the combination of other features and geometry constraints limit the angles such that they are not quite perpendicular.

The sensor base also includes four semicircular securing bolts cutouts near the corners for fasteners. Each of these semicircles matches with the corresponding features on neighboring sensor bases to create a clearance hole for the bolts used to fasten the wire base to the rest of the puck furniture.

4.5.2.1 Wire Mounting

Mounting the wires in the sensor base requires a fixture to reduce the variation in the unloaded wire positions. If the unloaded wire positions deviate too much, then the resting basket position will be tilted off the central base axis. If this tilt is too large, it can create several problems in the system. The cone system described in Section 4.6 is designed to prevent wire basket tilting motions from affecting other system processes, such as capping and loading vials, by ensuring that the vial is properly aligned with external hardware during those interactions. However, if the vial tilts too much when only loaded with a vial and not acted on by other external forces, the cone system components will interact during weight measurement. This interaction brings friction forces into the weight sensor system, preventing the wires from freely deflecting and moving as the vial mass changes, reducing the measurement reliability. Using a fixture for mounting the spring wires into the sensor base ensures that the wires are consistently loaded into their bases, reducing this tilt risk. The wire mounting fixture used for this work includes five components, which can be seen in Figure 4.10.

The first component is the arbor press used to push the ball into place, creating the press fit between the wires and the sensor base. A small arbor press is used for this work. The second component is the arbor press adapter, a 3D printed component used to convert the relatively large flat face of the arbor press into a surface that can press down on just the 3.5 mm ball plug. If this surface is not narrowed, then the press would also crush part of the spring wire, preventing it from performing as required for measurement. This adapter narrows pyramidally to the press point sp that it will not buckle under the press force load. The press point itself is a #6 machine screw threaded into the 3D printed adapter through the tip of the pyramid. The 3D printed material, particularly when narrowed to match the profile of the 3.5mm ball, buckles under the press load. The machine screw can bear this load without buckling, and the bottom of the machine screw is slightly concave which helps it locate itself on the ball during the press operation.

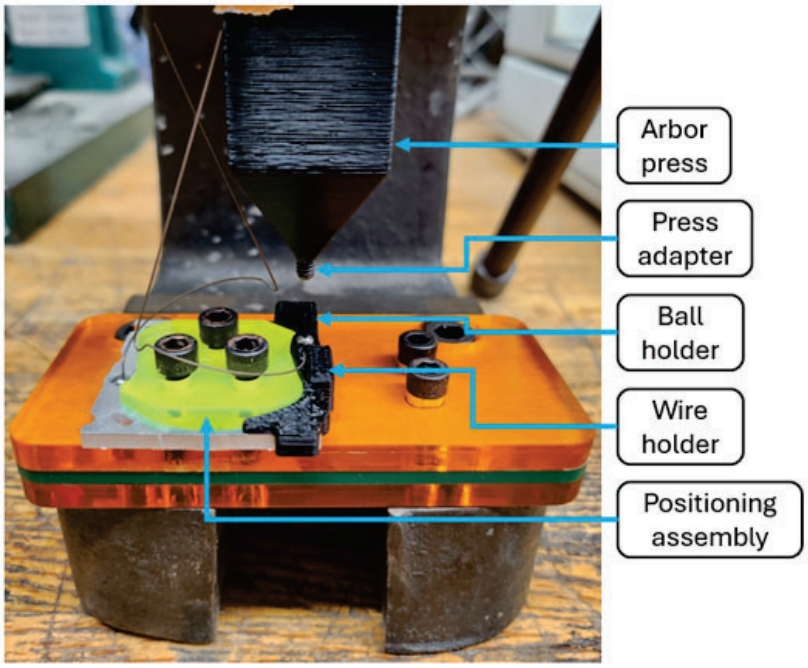

Figure 4.10: The assembly fixture consists of five major components: 1) The arbor press applies the force needed to push the ball into the press fit. 2) The press adapter converts the general shape of the arbor press to match the diameter of the ball. 3) The ball holder ensures the ball stays concentric with the press fit hole during the insertion. 4) The wire holder ensures the wire position is consistent between press operations and does not move during the ball insertion. 5) The positioning assembly ensures the wire base is positioned such that the press fit hole is concentric with the arbor press adapter such that the ball is successfully inserted into the press fit hole.

The third weight sensor assembly component is the positioning assembly which orients the wire base relative to the arbor press. This component ensures that the press is vertically aligned with the ball and the corresponding hole without requiring careful positioning for each individual operation. The positioning assembly is adjusted once, then it is locked into place such that wire bases can be placed and removed with repeatable positioning. The positioning assembly consists of a set of sliding plates and a raised profile which matches the interior cutout on the wire bases. The sliding plates provide vertical and horizontal adjustability so that the sensor base can be correctly aligned to the press location. Once the sensor base is aligned, locking bolts secure the sliding plates in position such that further positioning assembly adjustment is not required. The sliding plates include holes for two positioning bolts which connect the sliding plates to corresponding holes in the arbor press. This connection allows the positioning assembly to be removed and replaced on the arbor press while maintaining alignment between uses.

The remaining assembly fixture components are used to hold the ball and wire in place on the base during the press operation. These components are 3D printed and include alignment features which interface with the sensor base. The fourth component supports the wire, while the fifth component holds the ball plug. The wire holder supports the spring

wire on its bottom and sides, leaving space on top for installing the wire and removing the wire holder after the press operation. The holder is shown in Figure 4.11b. The degrees of freedom which this fixture constrains are shown in Figure 4.11a. This holder removes three degrees of freedom from the wire, which greatly improves the repeatability of the wire's positioning. Two other degrees of freedom are limited by the cutout in the wire base while the wire is resting in position, even before inserting the ball plug. The final degree of freedom is only semi-constrained. When the wire rests on the wire holder, it remains aligned rotationally, but it is not fully constrained until the ball plug's insertion is complete. A future improvement to this fixture would include a full constraint on this rotational degree of freedom during installation which does not limit the ability to remove the wire support after the press operation.

The ball holder, shown in Figure 4.12, simply resembles a claw with a profile matching the 3.5 mm ball's 2D profile, keeping the ball aligned with the wire base cutout during the press operation. Without this ball holder, the ball plug tends to shift to one of the sides of the wire. This shift adds a lateral force during the press operation which tilts the wire to an inconsistent position. Adding the ball holder constraint removed this source of variability, improving the reliability of the assembly process.

The rotational symmetry of the sensor base allows the mounting fixture parts to remain fixed for both ball plug insertions on each sensor base. After the first spring wire is mounted, the sensor base is rotated 180 degrees to reposition the corresponding cutout under the arbor press for each operation. This repositioning holds for individual wire bases and for combined rows of wire bases, as the fixture features are located to avoid interference with longer wire base sets.

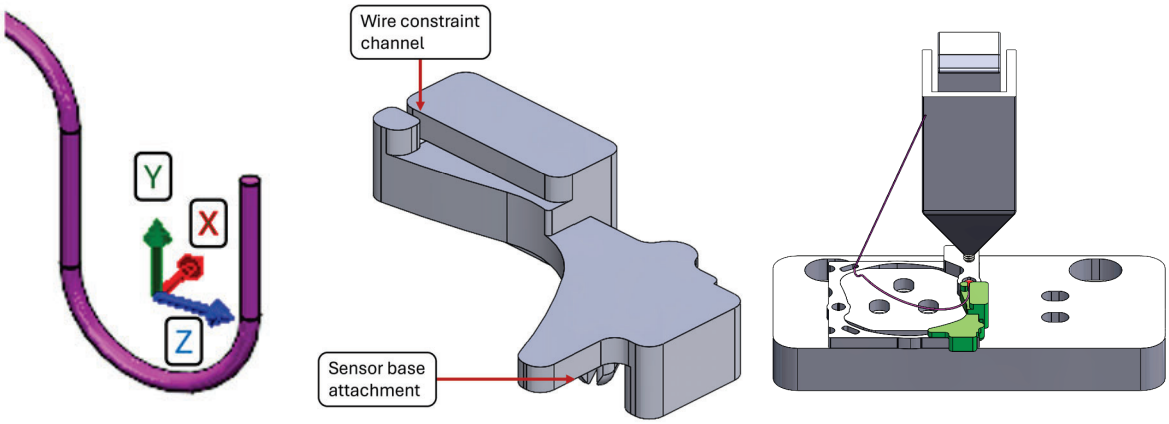

Figure 4.11: The wire has six degrees of freedom (a) that need to be constrained at the mounting interface. The linear motion in the X and Y directions and rotation about the X and Y axes are constrained by the wire constraint channel on the wire holder (b). Linear motion in the Z direction and rotation about the Z axis are constrained by the wire base cutout. The wire holder position in the assembly fixture is shown in (c).

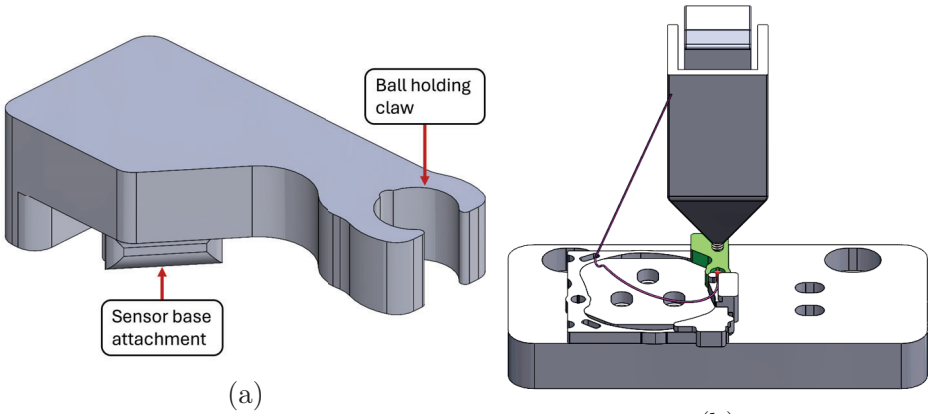

(a) (b)

Figure 4.12: The ball holder ensures the ball plug stays in its proper position during the press operation. Without this holder, the ball has a tendency to shift towards one side of the bottom wire crook during the press operation, tilting the wire assembly and altering the effective preload angle.

4.5.3 Wire Motion Tracking

The wire motion is measured by a camera capturing the sensing arm tip position change during the sublimation process. The reliability of the wire position measurements depends on the optical sensor's ability to reliably measure the sensing arm tip position. Measuring positions in images is a problem heavily discussed in computer vision literature. AprilTags are a simple pattern developed to address this position measurement problem. An example set of AprilTags is shown in Figure 4.13.

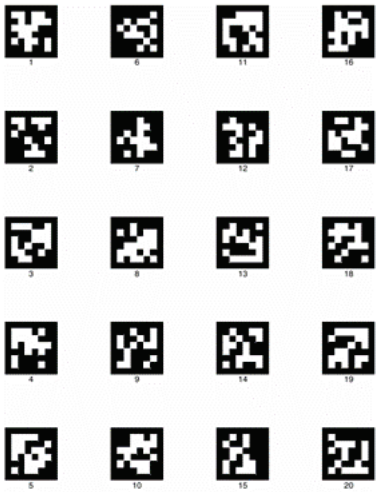

Figure 4.13: AprilTags are a common tool used to provide easy targets for position measurement in images. These optical fiducials are used on the sensing arms to improve the reliability in measuring the wire position changes during sublimation. These AprilTags also enable differentiation of different weight sensors on a single puck.

An AprilTag is mounted on the end of each sensing arm, as shown in Figure 4.14. A printed plastic nub is glued onto the end of the wire, and the AprilTag is glued on top of this nub. The nub has a conical shape to help prevent it from catching on the robotic arm used during the vial loading and unloading operations described in Chapter 6. The AprilTags improve the camera's ability to track the wire movement accurately and locate the sensing arm tip's position in space. If the error in the wire position measurement is too large, it can drown out the actual measurement of interest, thus decreasing the system's overall precision.

4.5.4 Vial Basket

The vial basket holds the vials in the spring wire weight sensing system. This basket needs to securely hold the vial in place while the vial moves through the machine, and it needs to withstand the loads created during each machine operation. These loads include both mechanical (loading, unloading, capping, blowing gas) and thermal (freezing) forces. The basket is additionally used to position the vial for loading, unloading and capping. The basket itself is made from two parts – a plastic basket and a stainless steel basket wire.

The vial basket hangs below the sensor base to improve system stability. The axial symmetry of the vial basket assembly ensures that its center of mass is located on its central axis, and the system is upright when this axis is aligned vertically. The pivot point for the vial basket system is the interface between the vial basket and the spring wires at the

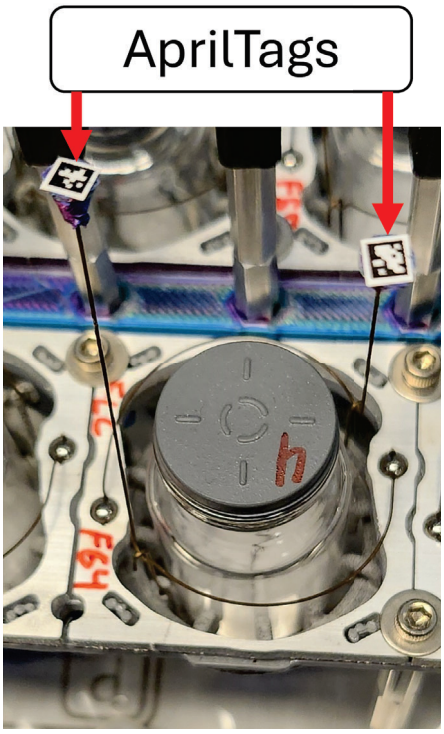

Figure 4.14: The AprilTags are placed on the ends of the spring wire sensing arms. The motion of the AprilTags is tracked by a camera to determine the mass change in the vial during sublimation.

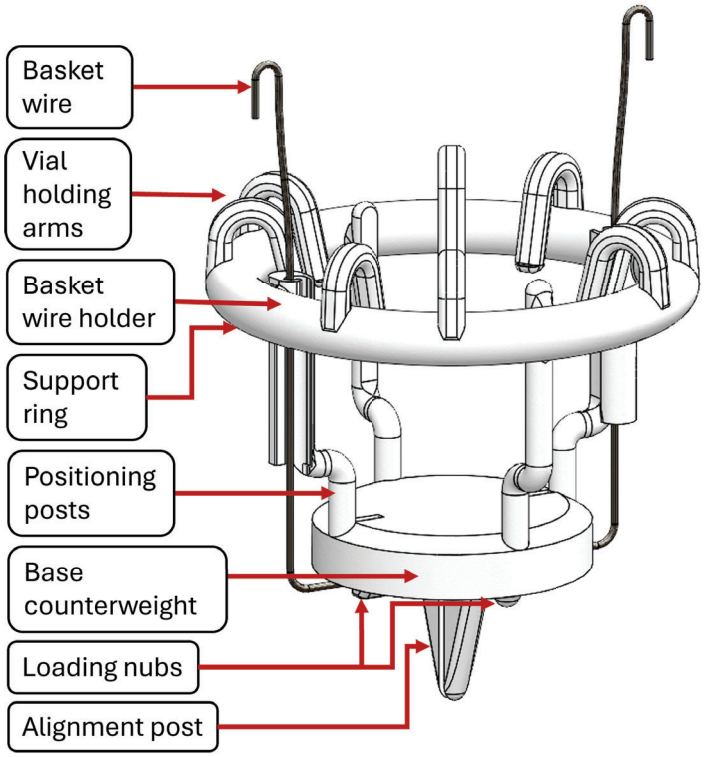

Figure 4.15: The wire basket assembly is used to hold vials in the spring wire system. The basket wire connects the vial basket to the sensing wires. Elastically averaging vial holding arms position the vial in the center of the basket. The support ring helps the system resist the forces created by vial loading and capping. The basket wire holders locate the basket wire in the vial basket. The positioning posts locate the vial in the vial basket. The basket base counterweight connects the other main parts of the basket together, and it is used to position the sensing arms during vial loading and unloading. The loading nubs ensure the basket stays balanced when it is pressed down during capping and vial loading. The alignment post ensures the vials do not move significantly during freezing and that the vial is oriented and aligned correctly during the vial capping and loading processes.

basket wire crook. The line connecting these two points serves as an axis of rotation for the basket assembly. The vial basket assembly is stable if its center of mass is located vertically below this axis, as illustrated in Figure 4.16. The basket assembly is marginally stable at best if the center of mass is located vertically above this axis. If the basket assembly center of mass is located above the rotational axis and it is not directly above the axis, then gravity will induce a moment on the system about this axis causing the system to tip over. If the assembly center of mass is below the rotational axis and it is not directly below the axis, then gravity will induce a moment on the system which tips the assembly back into a vertical state. This hanging setup also ensures that the system's center of mass is located below its center of stiffness. In previous weight sensing system iterations, such as those described in [47], positioning the center of mass above the center of stiffness lead to significant tilting motion and system instability. This motion results from the required low stiffness of the sensing compliant elements. As the center of mass moves, if it does not have

a restoring force re-centering it within the assembly, the vial quickly continues to tilt and fall over. Keeping the center of mass below the system center of stiffness ensures the compliant element's restoring force stabilizes the system similarly to gravity.

The plastic part of the basket includes seven major features: vial holding arms, a support ring, basket wire holders, vial positioning posts, the base counterweight, loading nubs, and an alignment post. These features are shown in Figure 4.15. The basket base counterweight has a variable thickness to accommodate the weight needed to ensure the basket assembly's center of mass is below the rotational axis at the top of the basket wire. The top of the base counterweight is chamfered to facilitate the interaction between the basket and the basket pulling trunk used for vial loading/unloading. On the bottom of this base, there are two loading nubs and two wire holding features, as well as the alignment post. The alignment post has a triangular profile to provide stability and a cutout down the center to allow for basket wire insertion. When the basket wire is inserted in this slot, it has one degree of freedom fully constrained. The two wire holding features on the bottom of the base are opposite facing clips which constrain an additional three degrees of freedom for the basket wire, such that it can only spin and move side to side. The tops of these clips, along with the additional two loading nubs, provide contact points for the basket base to remain stable and upright when the vial is pushed down during loading and capping.

The vial positioning posts extend upwards from the top of the base counterweight. These posts have a step built into them on which the vial rests. This offset between the bottom of the vial and the base counterweight provides space for the cooling gas flow to directly contact the vial's bottom surface, improving cooling process efficiency. This offset feature

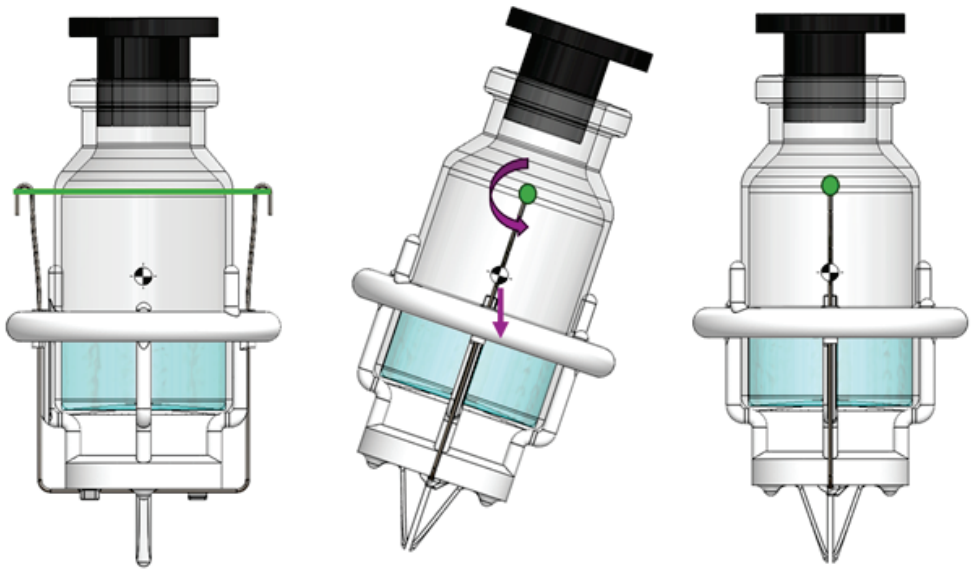

Figure 4.16: The vial basket has a rotational axis which goes through the mounting points on the basket wire. Designing the basket so that the center of mass of the vial basket assembly is below this axis ensures that when the basket tilts, gravity creates a restoring moment which returns the assembly to an upright state.

increases the vial's exposed surface area which is also in contact with the product solution during cooling by about 60% compared to covering the vial's bottom surface. As convective heat transfer scales linearly with surface area, the feature similarly improves expected cooling performance by a corresponding 60%. The empty volume between the basket base and the vial also provides the space needed for the basket puller trunk, a component used during the loading and unloading operation described in Chapter 6, to interface with the vial basket. Two of the positioning posts include a vertical clip for holding the basket wire. Similarly to the clips on the basket base, these clips are oriented opposite each other to constrain the basket wire. These clips also include a slot for width, which helps them accommodate both manufacturing variation and wire thermal contraction.

The support ring connects the tops of the vial positioning posts, providing support for the posts during the loading and capping operations. When the vial is pushed down into the basket, the filleted bottom of the vial interacting with the step in the vial posts leads to a radial force pushing these posts out from the basket center. If this support ring is not included, the posts snap off at the base due to this load. Including the support ring changes the posts from acting like a cantilevered beam to a beam supported on both ends, increasing their stiffness eightfold.

Extending from the support ring are eight vial holding arms. These vial holding arms are bent inwards towards the vial, such that the inner diameter of their contact points is a 21.5 mm circle, which is smaller than the 24 mm nominal vial diameter. This difference ensures that the vial holding arms will engage with the vial. When the vial is inserted into the basket, these holding arms are bent back towards the support ring, generating an elastic deflection that creates secure hold on the vial. While using eight contact points would over-constrain the vial position, the arms are sufficiently flexible that all of them can deflect evenly, and this deflection averages out to center the vial in the basket. Because this holding method is elastic rather than rigid, it can also easily accommodate manufacturing variation in the vial diameter. The eight arms are positioned 36 degrees apart from each other, which would match a circle with twelve evenly spaced divisions. Four of these divisions are occupied by the vial positioning posts, and the other eight are occupied by these vial holding arms.

The wire baskets were prototyped using SLA printing and produced in larger quantities with multijet fusion printing. The SLA printed prototypes were made on an available Form 3 printer which allowed for testing multiple different geometries and evolving the features needed for a fully functional wire basket. Once these features were validated, the model was outsourced to a commercial manufacturer for production using multijet fusion printing. At larger scales, these baskets could be injection molded, and all features have been designed such that a two part mold could effectively create all of the features.

4.6 Puck Furniture

The puck furniture, shown in Figure 4.17, connects the weight sensors to the pucks that carry the vials through the lyophilizer. This furniture also sets the relative positioning of the vials to each other, and it includes features which aid in controlling vial positioning during multiple operations during the vial life cycle.

The base of the wire furniture is the puck plate which serves as an adapter between the

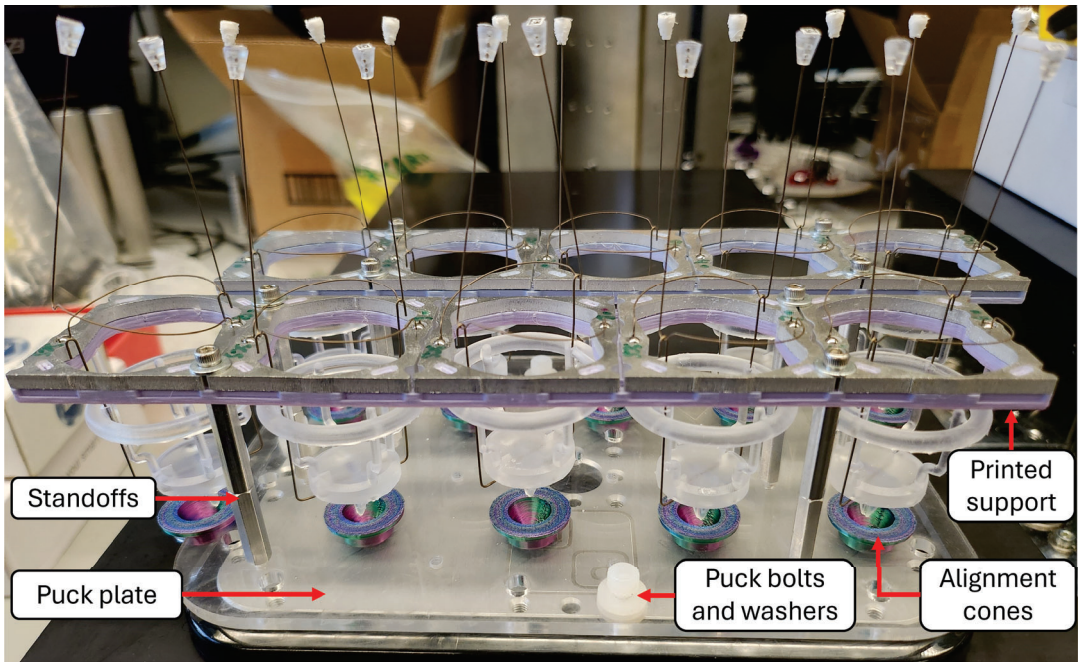

Figure 4.17: The puck furniture connects the spring wire weight sensing system to the Planar Motors pucks. The puck plate interfaces between the mounting pattern needed for the sensor bases and the puck mounting bolt pattern. An optional metal sheet can be added between the puck plate and the puck to serve as a heat shield between the puck and the vials. The puck plate is attached with plastic bolts and washers to avoid interference with the permanent magnets in the puck. The standoffs provide the vertical displacement between the sensor bases and the puck plate for the vial basket's vertical motion. The sensor bases are attached to this furniture through a 3D printed support. Alignment cones placed in the puck plate ensure accurate vial positioning during vial loading and capping.

bolt holes on the puck and the mounting positions used on the furniture. This plate also holds the furniture alignment cones which correspond to the alignment posts on the vial baskets. These furniture cones have a slope matched to the angle on the vial basket alignment post, such that when the vial basket is pushed down during loading and capping, the whole basket naturally centers itself in the furniture cone. This centering, shown in Figure 4.18, ensures the vials are aligned and vertical for the loading, unloading, and capping operations. These operations can fail when the vial is too tilted/off-center, so this alignment is necessary for reliability.

The puck plate is spaced off the puck by nylon washers. The furniture can include a sheet metal heat shield that would be supported between the puck plate and the puck to reduce radiative heat transfer between the puck and the vials. The puck plate and the potential heat shield are fastened to the puck using nylon bolts. The nylon bolts and washers are not affected by the permanent magnets within the puck, which helps keep them from affecting the stainless steel basket wires used in the vial basket assembly. Bolts made from ferromagnetic materials, such as steel, can become magnetized by the puck's permanent magnets, causing them to exert a force on the basket wire which would interfere with the gravitational mass

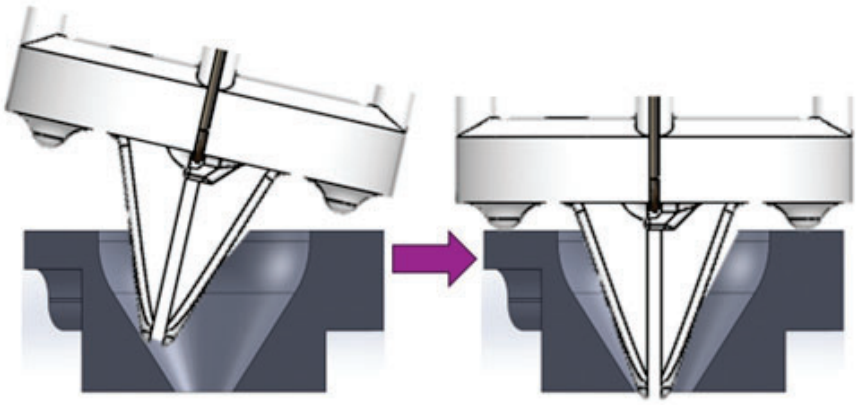

Figure 4.18: The furniture cone serves to center the vial basket assembly during operations which involve vertically pushing down on the vial in the basket, such as capping and loading. When the centering post engages with the conical surface in the furniture cone, the vertical force causes it to slide into a centrally located position.

measurement.

Four standoffs connect each vial row to the puck plate. These standoffs are located between vial bases so they do not interfere with the weight sensing or the lyophilization process. For a row of five bases, the pairs of standoffs are placed between the first and second sensor bases and between the fourth and fifth sensor bases. The standoffs are made from aluminum to provide stiffness and strength while remaining lightweight and non-magnetic. These standoffs are sized for 4-40 bolts because they can fit between each individual sensor base without interfering with any of their features. The 4-40 bolts used to connect the sensor bases to the standoffs also use washers to increase their surface area engagement, stabilizing the connection between the sensor bases and the rest of the puck furniture.

Each row of sensor bases is connected by a 3D printed support, shown in Figure 4.19. This support includes prongs that connect to the assembly mounting cutouts on the sensor bases to hold them in place. The printed support also includes through holes for the 4-40 bolts that connect the row of sensor bases to the puck furniture standoffs. The supports used in this system hold five sensor bases each, as the length of each row of vials on a puck is five vials long. The support includes a large central body cutout at the mounting point for each vial base to allow the vial basket to freely move vertically. This support structure enables replacement of the sensor bases in cases of poor performance or damage that might occur during system development. At larger scales, the printed support and wire bases can be combined into a single part, where the metal sensor bases are manufactured as a single part which is as many units long as is required in the system.

4.7 Results

An automated offline testing setup, shown in Figure 4.20, was built to evaluate the spring wire weight sensing system before implementing it on the lyophilization hardware. In this test setup, the spring wire system is placed on an analytical balance to directly measure

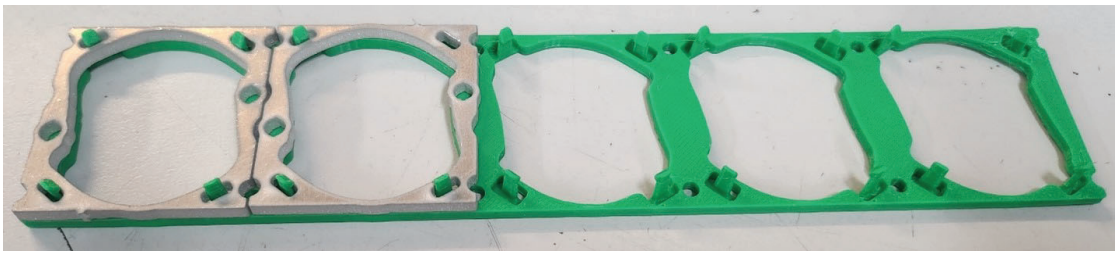

Figure 4.19: The 3D printed support is used to assemble a set of five sensor bases. The support includes prongs which connect to the corresponding assembly mounting cutouts in the sensor bases, through holes for the bolts used to connect the sensor bases to the puck furniture standoffs, and large central body cutouts to allow the vial basket assembly to deflect vertically during sensing.

the mass changes in the system. A peristaltic pump is used to periodically add mass to the system, and a Canon Rebel T7i camera with an 18-55mm lens is used to capture the tag positions during the tests. Because the springs are deforming elastically, the motion when adding or subtracting mass is expected to be symmetric, so adding mass is a reasonable method for running calibration experiments. Removing mass mechanically has a high risk of artificially agitating the vial, which can affect the tag motion measurements. Water's rate of evaporation at atmospheric pressure and room temperature is very slow, which would lead to extremely long tests. Using a fluid like liquid nitrogen which can vaporize more quickly due to boiling is also undesirable, as this boiling process can also induce additional vibrations in the system which may obfuscate measurements.

During a calibration test, water is added to the vial in 500 mg increments. The water is pumped using the peristaltic pump and dispensed through a needle which is positioned over the top of the vial. After each water addition, the system is left to settle for four minutes. The impact of the droplets dispensed by the needle impacting the fluid in the vial causes large vibrations in the system. The system does not have much damping because it is designed to have minimal friction. These impacts are an artifact of the test setup with the water dispensing, so it is appropriate to remove their effects from the measurements. However, in the lyophilizer these spring systems will be subjected to inertial loads as they move around, so they will still have some vibration in-situ. Accordingly, the AprilTag positions are recorded in a 10 second video, and the average wire tip position is reported. This test is performed with 15 mass additions, providing a system response over a total mass change range of 7.5 g. While the mass change in the vials used for this system during lyophilization is expected to only be 3 g, this extended test range verifies whether the performance stays consistent enough to handle variation in vial, cap, and solution fill masses.

The spring system linearity is evaluated by plotting the tag motion against the mass added during an experiment. An example dataset is plotted in Figure 4.21a. If the system is perfectly linear, then the data points should fall on a best-fit line. At the plotted scale, the results appear to be highly linear. However, the residuals plotted in Figure 4.21b show concavity. This result indicates a consistent nonlinear behavior, as the data follows a relatively smooth curve which is offset from the linear fit line.

Multiple potential sources of nonlinearity have been investigated to try to identify the

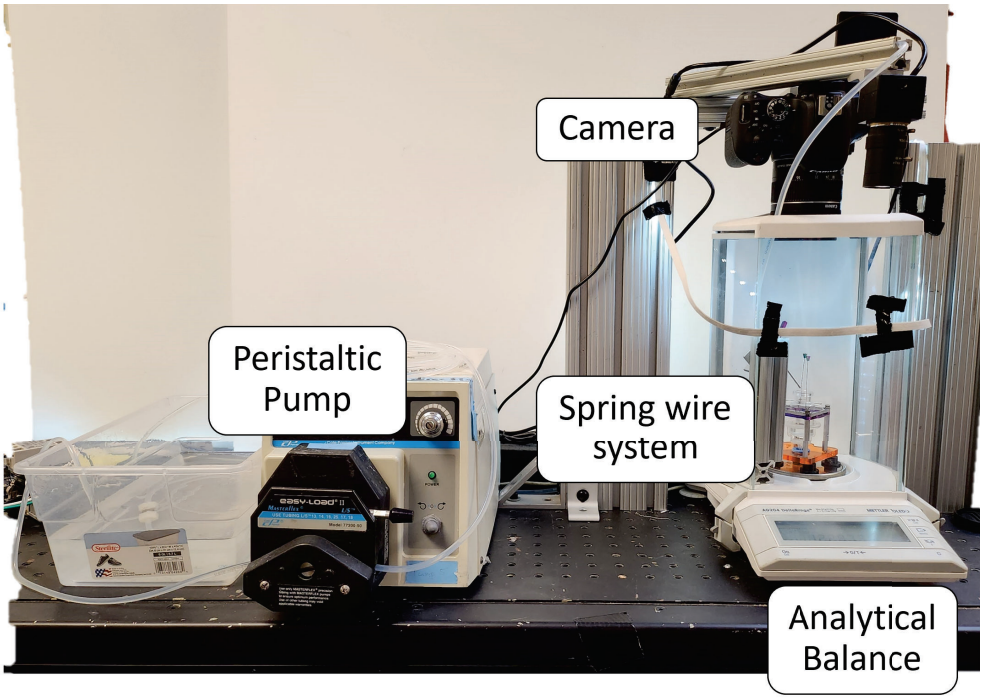

Figure 4.20: The test setup for calibrating the spring wire system uses an analytical balance to provide the ground truth measurements of mass changes. The mounted camera records the AprilTag positions on the spring wire sensing arms as the peristaltic pump intermittently adds water to change the mass in the vial.

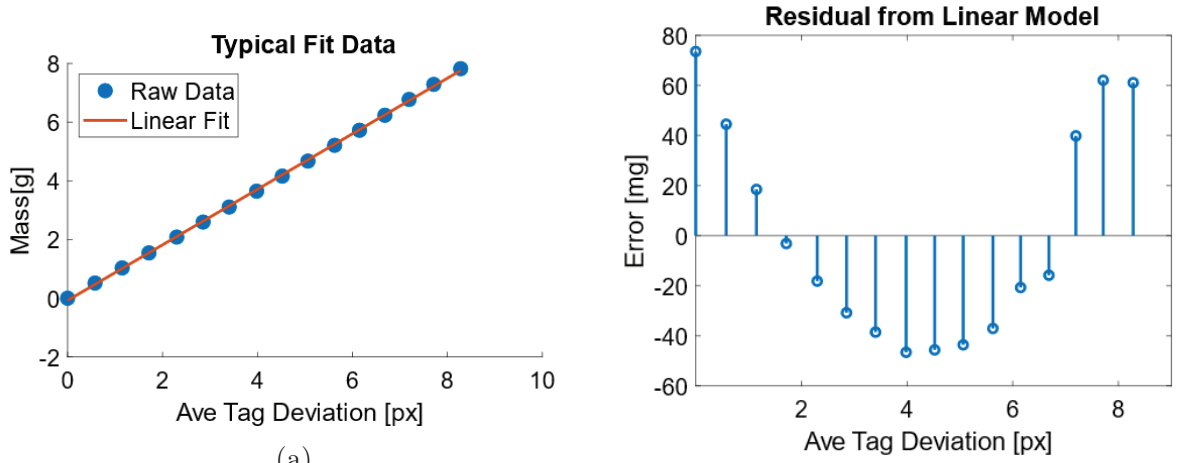

(a) (b)

Figure 4.21: (a) The AprilTag motion plotted against the mass addition from a sample calibration test shows the linearity of the spring system response. The difference between the data and the linear fit line, highlighted by (b) the ordered residual plot, shows a consistent, mildly nonlinear behavior in the response. This nonlinearity error ($\mathcal{O}(10\text{mg})$) is greater than the required system precision ($\mathcal{O}(1\text{mg})$), so it cannot be neglected.

source of this nonlinearity. One source could be sliding at the basket wire and spring wire interface. This motion could be addressed by gluing the wires together, removing system compliance from this interface, or by adding an additional link to the system, such that any sliding motion would instead result in motion of this link. Neither gluing the sensing and basket wires together nor adding an additional link to the system showed any improvement in the system linearity.

Another source of potential nonlinearity is deflection in the parts of the spring wire not currently included in the deflection analysis, highlighted in Figure 4.22. Thus, the deflection created by these wire components is also modeled. This wire section can be treated as a combined beam in torsion and a cantilevered beam with a moment load applied to its end using Equation 4.12.

$$\delta = L_1 cos\left(\frac{PL_3L_1}{EI}\right) - L_2 sin\left(\frac{PL_3L_1}{EI} - \theta_0\right) \tag{4.12}$$

L_1, L_2, L_3, and P are shown in Figure 4.22. E and I are the elastic modulus and second moment of inertia of the spring wire, respectively, and θ_0 is the preload angle bent in the spring wire shown in Figure 4.5a. Adding this deflection modeling to the predicted displacement also does not account for the nonlinearity seen in the calibration data.

Another approach to address the observed nonlinearity involves looking at the 2D motion of the sensing arm as seen from the side. As the deflecting arm moves, the AprilTag not only moves sideways, but down. This motion means that as the deflection angle increases, the corresponding motion seen by the camera can change. A diagram illustrating this motion appears in Figure 4.23. The circular motion transformation between the deflection of the spring wire at the basket wire crook and the motion at the sensing arm tip is shown in Equation 4.13.

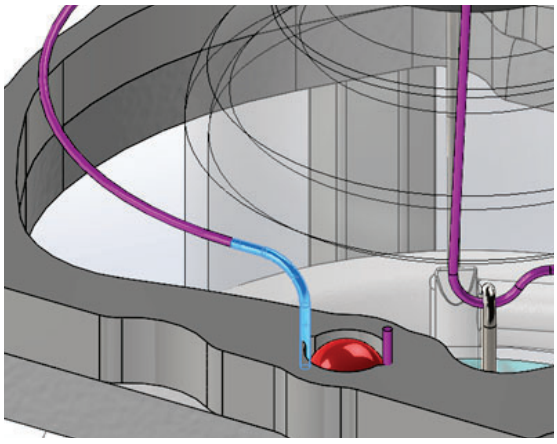

(a)

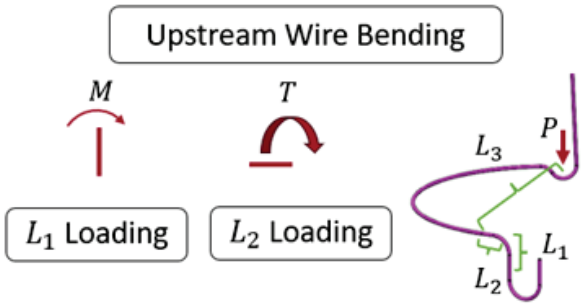

(b)

Figure 4.22: Additional elements of the spring wire which could contribute to bending deflection that are not captured in the curved cantilever beam deflection modeling. When modeled, these elements do contribute to the overall spring motion, but they are also expected to deflect linearly with mass changes in the vial.

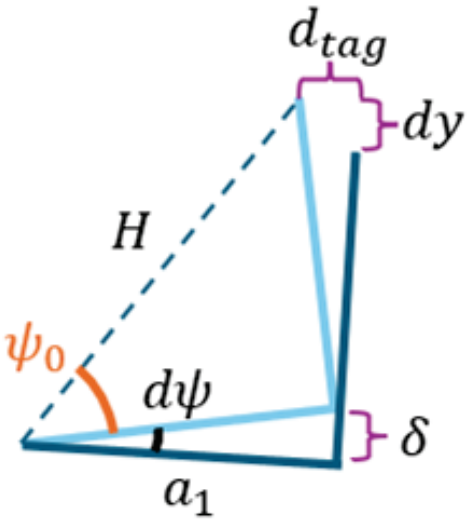

Figure 4.23: As the spring wire deflects, the sensing arm moves both sideways and down. This motion traces a circle, which means the measured tag distance changes as the wire deflects further.

$$\delta = a_1 \sin\left(\psi_0 - \cos^{-1}\left(\frac{d_{tag}}{H} - \cos(\psi_0)\right)\right) \tag{4.13}$$

This transformation also does not account for the nonlinearity seen in the calibration data, as the angles are sufficiently small for these results to remain linear.

Given that the previous modeling and hardware adjustment approaches do not address the curvature seen in the calibration data, a higher order model is used to fit the calibration data. When compared to the data within an experiment, these higher order models show very good fits, as the model can predict the mass measured in the vial based on the tag deflections with less than 5 mg of error. These results are shown in Figure 4.24.

While these models show strong performance on their own calibration data sets, the nonlinear nature of these models increases the difficulty of applying them to data where the mass is not known. During the lyophilization process, these weight sensors are the only form of mass feedback, so the models will not have multiple known data points which can be used to ensure strong performance. To test the model transferability, the coefficients calculated from one experiment using a sample spring wire assembly are used in combination with the deflection results from a second, separate experiment on the same assembly. These predicted results are then compared to the measured results from that experiment and plotted in Figure 4.25. These results show much higher errors, greater than 25 mg, which are well above the target precision of 3 mg.

Though these results do not show acceptable performance, they are based on the assumption of having no known deflection and mass measurements which can be used to calibrate the system in-situ. However, there is one known data point that can be used for in-situ calibration. Before the vial is first inserted into the weight sensor, the mass of the vial, product, and cap are measured, providing a known mass value. Once the vial is placed into the weight sensor, the resultant deflection can be measured. This combination provides a

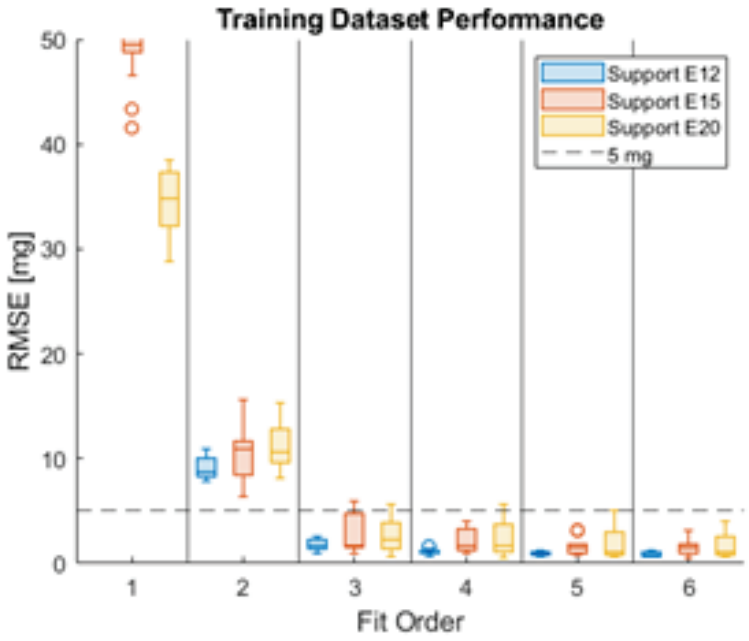

Figure 4.24: The root mean squared error is shown for the mass predictions based on measured deflection from calibration data on three sample spring wire assemblies (E12, E15, and E20). While linear and second order model fits do not capture the data well, third order and higher models can reproduce the mass data with less than 5 mg of error.

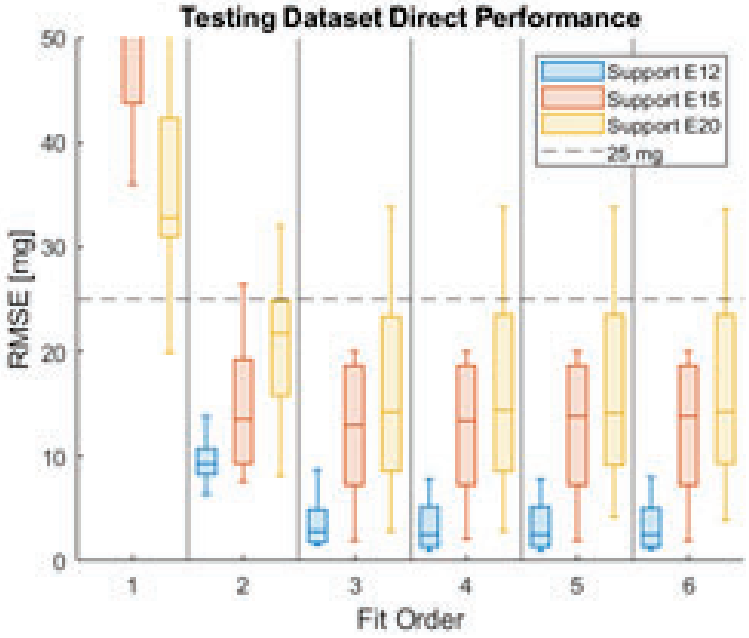

Figure 4.25: The errors in predictions used from a model generated by the calibration data from one spring wire weight sensing assembly applied to additional deflection and mass data measured on the same assembly far exceed the target precision of 3 mg.

single known data point for mass and deflection in the system. Using this data point to update the higher order model used to predict the system's mass versus tag motion curve improves the system repeatability to 15 mg or less, as shown in Figure 4.26.

This repeatability does not meet the 3 mg target required to use the residual moisture measurement as a vial quality certification. Future work on the spring weight sensor design can investigate methods for further refining the measurement accuracy, particularly regarding improving the linearity of the deflection response to system mass changes.

While the system does not yet meet the accuracy requirements needed to certify the residual moisture content in vials, it can identify when the vial has completed the sublimation process in the lyophilizer. When vials supported in the spring weight sensor system are in the drying section of the lyophilizer, they demonstrate a clear leveling phenomena corresponding to the end of the sublimation process. This result means that the system can detect when the residual moisture in the vial is no longer changing significantly, even if it cannot identify exactly how much water has been removed. The vial weight no longer changing significantly signals that the product can be removed from the system, as leaving it in the drying section longer will not significantly change the product's state. While this system does not yet enable in-line certification, it can serve as an in-line release criterion and vials can then be examined externally without affecting throughput of the machine. This release criterion improves system efficiency by ensuring vials are not left in the system longer than necessary. Example vial weight trajectories and their ability to identify the end of the drying process appear in Figure 4.27.

Vials removed from the lyophilizer after the spring mass sensing system indicated that drying is complete had their residual moisture content measured using a Karl-Fisher titration.

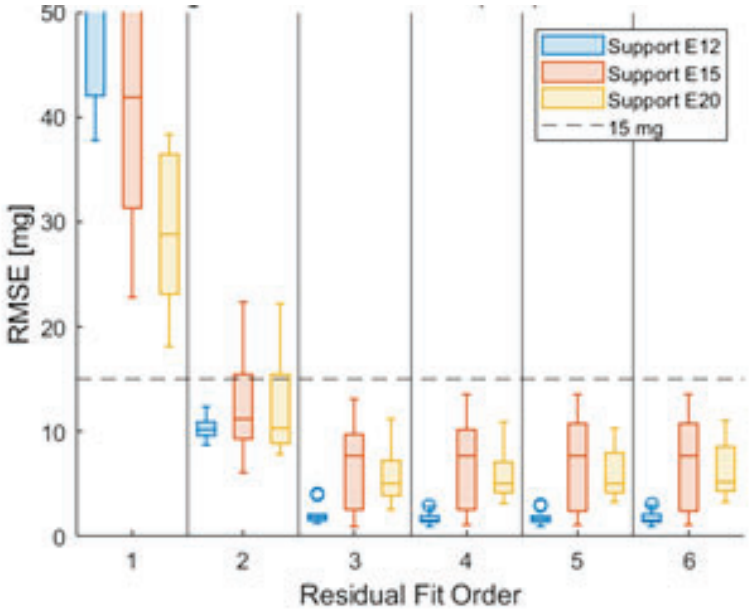

Figure 4.26: Using the full vial mass as a known data point for model calibration, the predicted mass values in situ are expected to have no more than 15 mg error relative to the true mass. Using higher order fits beyond third order does not improve the system performance.

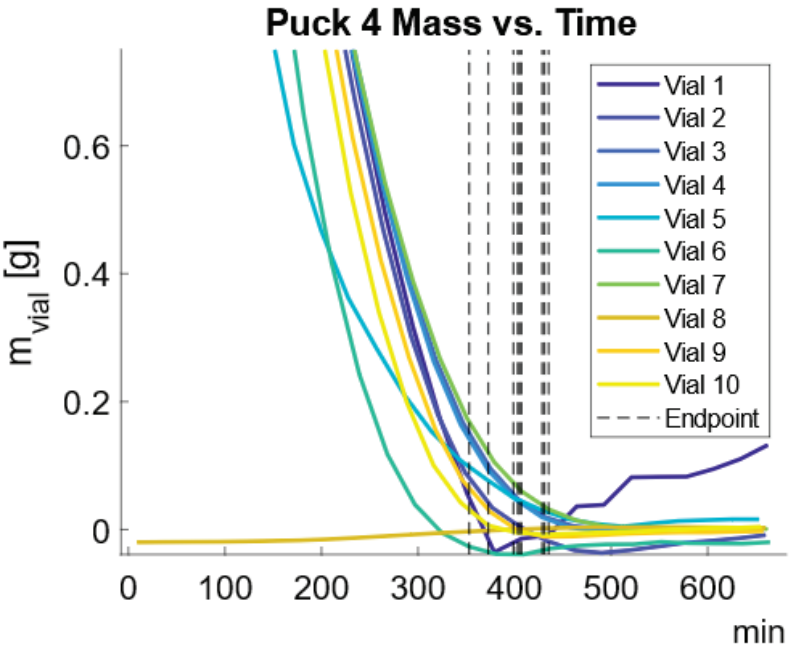

Figure 4.27: The vial estimated mass flatlines when water is no longer sublimating out of the vial. Measuring the difference between when the first and last vial finishes drying within a puck can serve as an indicator as to how uniformly the vials went through the lyophilization process. Once the last vial on a puck finishes drying, the puck is ready to be released from the system. Vial 9 starts flat because it is used as an empty internal standard vial to compare to the other drying vials. The increasing measured mass of vial 1 after its estimated drying endpoint is likely due to an error in the AprilTag position estimation.

These measurements showed less than 2% residual moisture content, meeting the desired specification. Thus, these results show that the system can accurately identify when vials have completed sublimation in the lyophilizer.

4.8 Future Work

Further work on the weight sensing system is required to achieve the resolution needed to certify the end of secondary drying, and thus the end of the lyophilization process. One way to achieve this resolution is to increase the sensing arm length, further amplifying the measurement motion. While the sensing arm length could be directly increased, this length change would require increasing the lyophilizer's internal volume, which in turn would decrease its vacuum performance. Alternatively, the sensing arm could be effectively lengthened by replacing the straight vertical section of the spring wire with an optical lever system, like those used in other precision measurement systems. An optical lever system could utilize a much shorter sensing arm that has a mirror mounted on its tip instead of an AprilTag. A laser from outside of the system could then be directed at the mirror, and its reflected image could be measured to track the spring motion. In this case, the sensing arm length becomes the total laser path length, which is decoupled from the chamber geometry. Thus, this system can

significantly increase the sensing arm amplification and resultant measurement resolution.

The wire mounting system can also be further refined. One alternative mounting method would be to use a dowel pin rather than a bearing ball for the press fit. The pin would provide a longer contact line than the ball to increase the support provided to the spring wire. The wires themselves can also have tightened tolerancing based on sensitive dimensions revealed during the assembly and testing of many of these wires. The angle between the mounting crook plane and the curved deflecting arm plane and the curved sensing arc radius proved to have a strong influence on the resultant wire deflection. When paired wires used for a support had mismatches in these dimensions, the wire basket would sit tilted rather than balanced vertically. Alternatively, the vial and basket assembly center of mass could be further lowered below the weight sensor center of stiffness and the basket rotational axis to increase the restoring force for re-centering the vial in the weight sensing system.

The weight sensing system was only built and tested with a single vial diameter for this work, but manufacturing sensors for larger and smaller vials would demonstrate its ability to adapt to multiple vial sizes. The differently sized vials would have different masses around which they would require the most sensitivity which would require using the design spreadsheet to create an appropriate wire geometry for those vials.

Chapter 5

Drying

The drying process describes the part of lyophilization where the liquid in a solution is removed via sublimation. This machine section requires control of vacuum and thermal conditions to ensure that the sublimation process can successfully be completed. The primary considerations when designing the drying chambers are their vacuum system, condensers for removing water vapor before it reaches the vacuum pumps, and a capping system for sealing the vials under vacuum after they finish sublimating.

The vacuum system consists of the vacuum ports on the chambers, the condensers, the vacuum pumps, and the valve assemblies used to isolate the condensers for regeneration. The vacuum ports on the chambers are based on ISO100 geometry to integrate the machine with standard vacuum equipment. These ports are simple standard ISO100 bulkhead connections on the sides of the chambers and multifunctional ISO100 or sensing window ports positioned on the chamber roofs. The side ISO100 ports can be used with ISO100 blanks that have window ports machined into them for viewing the sides of the vials for thermal measurement and general monitoring. The ISO100 size is used to provide a large enough port for manual access for maintenance and to reduce the frictional losses in the vacuum piping system.

The condensers remove water vapor from the system before it can reach the vacuum pumps, where it could cause damage by degrading the pump oil. The condensers in this system use liquid nitrogen to cool the walls of the condenser body, creating a sufficiently low temperature surface for the water vapor to condense into ice. The condenser body rests in a dewar that holds liquid nitrogen, and it includes level sensors so the liquid nitrogen level can be automatically maintained. The condenser has an ISO100 port at its entrance to interface with the chamber piping and a KF25 port at its exit to feed to the vacuum pump. Once the water vapor has been removed from the system, the vacuum piping diameter can be reduced without creating significant performance loss.

The capping system consists of a pneumatic piston and a compliant pusher that presses the vial stoppers into a fully capped state. This piston uses the same dynamic rod seal as is used on the load-lock doors, described in [25], in a machined aluminum cylinder which is mounted on top of the final chamber before the exit load-lock. The compliant pusher allows the piston to act as a force-controlled actuator to ensure that the vials are fully capped regardless of variation in the vial and cap heights. This system was successfully used to cap hundreds of vials consecutively without failure.

5.1 Overview and Requirements

The drying section of the lyophilizer covers the sublimation process during freeze-drying. In this machine section, the frozen product in the vials is exposed to a vacuum environment to sublimate the solid ice directly into water vapor. This water vapor is then removed from the lyophilizer through the vacuum system, leaving behind the solid product cake. An example dried product cake is shown in Figure 5.1.

The drying section has two phases: primary and secondary drying. During primary drying, the bulk water is removed from the product cake. This water consists of pure ice crystals that have separated from the product during the freezing process. Secondary drying consists of the removal of water that remains bound to the product even after the freezing process.

Most of the water removal occurs during primary drying [2], and this process represents most of the total drying time. During primary drying, the product needs to be kept below its collapse temperature to prevent cake failure. The collapse temperature is a property specific to a given formulation, often determined experimentally, representing the glass transition temperature for the frozen solution. If the product temperature exceeds the collapse temperature during primary drying, then the product softens, allowing for mobility that can trap moisture within the crystalline solid [48]. If this occurs, then the water may not be able to sublimate out, which can lead to cake collapse after it exits the lyophilizer. This cake collapse occurs when sufficient water remains in the cake to partially redissolve the solid. This collapse leaves the product in an unstable state, indicating a failure of the lyophilization process.

During secondary drying, the vial temperature can be elevated to supply more energy to remove the bound water in the product. Since the bulk ice has been removed from the system, the destabilization risk from the solution exceeding its collapse temperature is removed. Figure 5.2 shows an example of a typical primary and secondary drying process in a conventional pharmaceutical lyophilizer. Conventional lyophilizers distinguish between the primary and secondary drying processes by changing the temperature setpoint for the vial shelves.

Figure 5.1: Once a solution has gone through the drying process of lyophilization, the liquid is removed from the product to leave behind a solid crystalline cake.

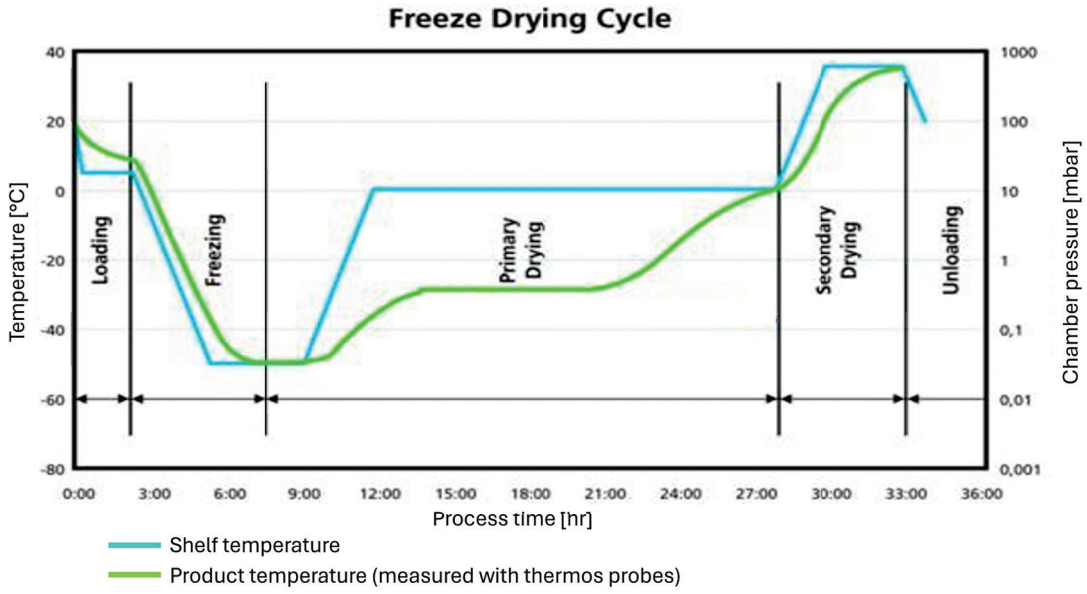

Figure 5.2: Primary drying and secondary drying are separated by shelf temperature setpoints in conventional lyophilizers. The shelf temperature setpoint is higher during secondary drying because the removal of bulk water removes the risk of failure due to exceeding the product's collapse temperature. Figure adapted from [2].

5.1.1 Temperature Control

In the continuous lyophilizer developed for this work, the shelf temperature is not a driving control parameter. Thus, an alternative heat transfer method from conduction is needed. While forced gas convection works in the cooling chamber for heat transfer, convection is not an available energy delivery method for drying because the vacuum environment requires a lack of gas. Thus, the drying chambers must use radiation for power delivery. This radiative power delivery can be controlled by changing the temperature of the vial environment. This environmental change can be driven by either the temperature of the Planar Motors stator, the machine chamber, or both. The view factor between the vials and the stator is blocked by the puck, so it is more efficient to control power delivery by changing the machine chamber temperature.

The collapse temperature can vary between different formulations [48]. Thus, the explicit temperature maximum for the vials during primary drying can vary depending on the product formulation. There are two ways to ensure that the product temperature does not exceed the collapse temperature. First, the primary drying environment temperature can be maintained below the collapse temperature. This method is the safest option because it prevents enough energy from being transferred to the vials to allow them to exceed their collapse temperatures. The second method is to limit the power delivery to the product below the energy required to sublimate the ice in the vial. Sublimation is an endothermic process, so it requires energy input to occur. If the energy delivered to the vials is less than the energy required for sublimation, the sublimation rate will be lower than its potential maximum. However, if the energy delivered to the vials exceeds the energy required for sublimation, the excess energy

will heat up the product. This heating can cause the temperature to approach the collapse temperature, so it must be monitored and/or controlled.

Actively cooling the machine chamber below the collapse temperature requires significant energy investment and would slow down the sublimation process, while relying purely on the passive machine chamber temperature creates an uncontrolled reliance on the environmental conditions around the lyophilizer. Thus, the machine chamber must have some level of temperature control. In the system developed in this work, this temperature control is achieved through using chillers to control the module temperatures in the drying section. The stators serve as an active heat source that consistently warms up the machine chamber. This heating is countered with a chiller that can be turned on and off based on the desired machine chamber temperature.

For this strategy to be effective, the power delivered to the vials by the system must be low enough that the module chilling creates a functional, controllable range within which the power delivery is less than the vial can take for sublimation. If operating the chillers on the modules at full power still results in too much power delivery to the vials and heating of the ice before primary drying ends, then more active cooling of the chamber or radiation limiting tools such as heat shields are required. Implementing chillers to control the machine chamber temperatures in the lyophilizer sufficiently limits the heat transfer to the vials, but on the lyophilizer built for this work two chillers were needed to service the full drying section.

The product temperature during secondary drying is not as sensitive. Because the bulk water has been removed from the product, the risk of partial ice melting is mitigated. Excess heat that does not contribute to sublimating the bound water will heat up the solid mass in the product cake, but this heating does not damage the product as long as it does not denature the proteins. Thus, a higher power delivery primarily serves to increase the sublimation rate during secondary drying. The machine chambers at the end of the drying section can be operated without chillers to achieve this increased power delivery.

5.2 Drying System Elements

The drying section consists of three major elements: the vacuum system, the drying chambers, and the condensers. The vacuum system needs to connect the vacuum pump to the drying chambers to allow them to create the vacuum environment in the drying section. The drying chambers are required to seal sufficiently to maintain the vacuum conditions required for sublimation. The condensers must remove water vapor produced by the lyophilizing product from the drying section so that it does not contaminate the pump oil, leading to pump damage and subsequent failure.

5.2.1 Vacuum System

The drying section must be kept under vacuum for two purposes. First, the vacuum environment drives the sublimation process. Sublimation occurs when the ambient vapor pressure is lower than the vapor pressure of the ice in the vial. This pressure difference enables the ice to sublimate directly into water vapor when the ice is at a temperature below water's triple point. The lower the ambient vapor pressure, the faster the sublimation due to the larger

driving force created by the pressure differential. However, as sublimation occurs, it releases water vapor into the drying section headspace, increasing the ambient vapor pressure. The vapor pressure equilibrium is based on the vacuum system's ability to remove water vapor from the drying section. This water removal rate is limited by the motion of water vapor through the vacuum system up to the condensers, where the vapor is removed from the vacuum system as it is condensed into ice on the condenser walls.

The second purpose of the vacuum system is to limit the heat transfer to the vials due to convection. The strategy of limiting power delivery to the vials to prevent the temperature from exceeding the formulation collapse temperature relies on having direct control of the power delivery and being able to sufficiently limit power delivery to the vials. Convective heat transfer significantly increases the power delivery to the vials as it is a much more efficient form of heat transfer than radiation. Thus, this heat transfer must be limited. A lower inert vacuum pressure is an indication of less free gas molecules in the system that could transmit thermal energy through convection. The vacuum system in the drying section needs to maintain an inert pressure below 13 Pa to sufficiently limit potential convection effects [49].

5.2.2 Drying Chambers

The drying section chambers include two primary features. These features include ISO100 ports on the chamber side walls and a multifunctional port for windows or ISO100 connections on the chamber top surface. The side ISO100 ports include a 102 mm diameter hole sized to accommodate an ISO100 centering ring and eight bolt holes positioned to match an ISO100 bolted flange. The top port shape is based on an ISO100 port at its center with extensions to each side to cover the full internal width of the modular tunnel. These features are shown in Figure 5.3.

Basing these drying chamber ports on standard ISO100 geometry provides flexibility to the chambers while simplifying manufacturing. Each chamber is manufactured with the same features regardless of which role it fulfills in the drying section. The use of standard features provides a direct interface connecting the drying chambers to standard components with proven vacuum sealing hardware. These ports can then be connected to standard ISO100 hardware, or custom parts can be made that interface with the same style of features. The flexibility of this port geometry is shown in Figure 5.4, where each different functional drying module assembly is shown. Each of these module assemblies uses the same drying machine chamber design, and the ports are simply connected differently depending on the functional uses required. This modularity includes a chamber for connecting to the vacuum system, a chamber for weight sensing, and a chamber for temperature sensing.

5.2.3 Side Ports

Part of the drying module flexibility is enabled through custom machined ISO100 blanks. The side ports on each module can support various uses through different ISO100 blanks. Each blank can be placed on any of the machined chambers, quickly adapting the base drying module to the needed purpose. There are three main types of ISO100 blanks used in this system, as shown in Figure 5.5.

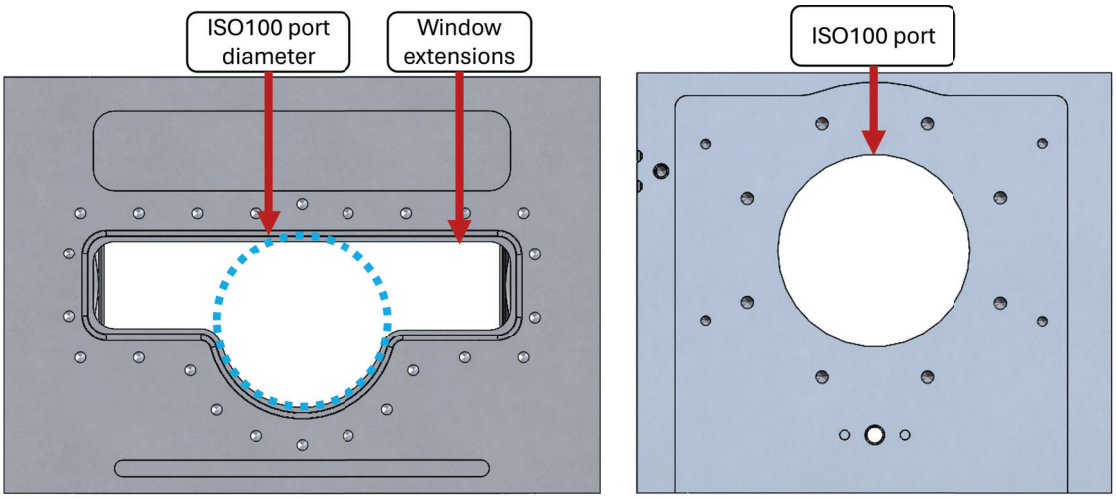

(a) (b)

Figure 5.3: The drying machine chambers include (a) a top port which can serve as a sensing window or an ISO100 connection and (b) side ports sized for ISO100 fittings. This standard port sizing makes use of commercial, proven vacuum sealing equipment to use these ports for multiple functions, such as product monitoring and connecting modules to the vacuum system. The top port requires an adapter plate to connect to ISO100 fittings.

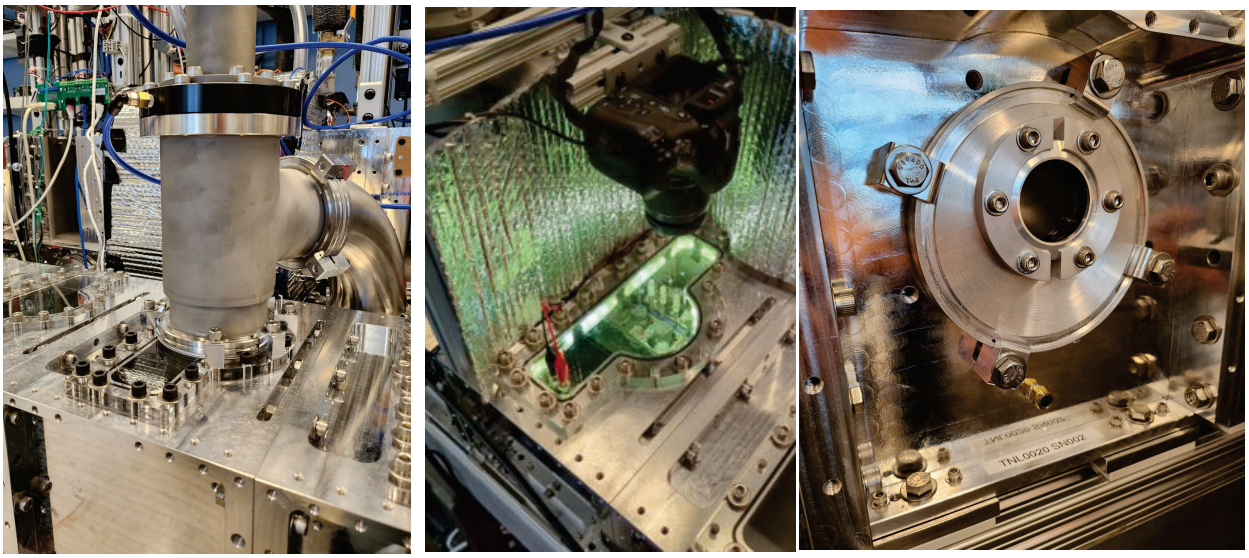

(a) (b) (c)

Figure 5.4: The drying module can be assembled with additional hardware to make it accommodate the different needs required in the drying system. In the system built for this work, the primary chamber applications include (a) connecting to the vacuum system, (b) providing visibility for the weight sensing system, and (c) providing a window for the temperature sensing system.

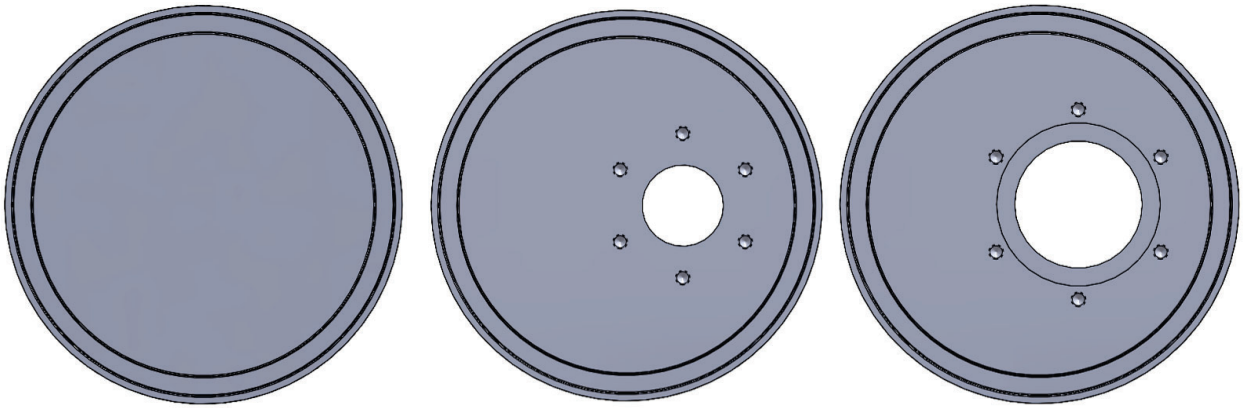

(a) (b) (c)

Figure 5.5: The three different blanks used in the system. (a) The first blank is just a standard blank used to seal the side of a standard drying module. (b) The second blank includes features for a standard bulkhead clamped connection to KF25 vacuum fittings. (c) The third blank includes a 52 mm diameter pocket and a 40 mm hole surrounded by the bolt pattern for KF40 bulkhead clamps used for mounting windows to the sides of the drying modules.

The first type of ISO100 blank (Figure 5.5a) used in this system is the standard blank, which is used to simply seal the side ISO100 port. While the modules include the capabilities to be used for multiple different purposes, they do not always need various windows and vacuum connections. Thus, when a module does not require additional attachments, its side ports are covered by standard blanks. At larger scales, some modules can be produced without the side ports to save components, but during system development it is more helpful to include the flexibility to place windows and vacuum attachments on a wider variety of units as needed.

The second blank (Figure 5.5b) includes KF25 bulkhead clamping features. These features are used to connect the blank to standard KF25 components. The load-lock chambers use these blanks to connect to their vacuum pump. These blanks can also be used with reducers to adapt to other KF sizes, such as KF16 for pressure sensors. This second type of blank is shown in Figure 5.6 attached to a vacuum valve assembly used to pumping down and isolating a load-lock chamber.

These blanks with KF25 features can be made to work with windows, but the windows will stick out of the blank, making it easier for them to bump into or get dirty. Also, the window viewing area would only have a 25 mm diameter, limiting its visibility. Additionally, the windows needed would have to be a custom size to include the 25 mm diameter viewing area plus the additional diameter used for sealing. This additional window diameter works out to about 10 mm, as the O-ring used to seal KF fittings has a 5 mm diameter. For thermally transparent windows, such as the zinc-sulfide windows used in this system, the standard window sizes are 25 mm, 40 mm, and 50 mm. The 40 mm window is a little large for the KF25 fitting, but the 50 mm window is a good fit for KF40 fittings, as the 10 mm diameter addition required for the O-ring corresponds to the difference between the potential

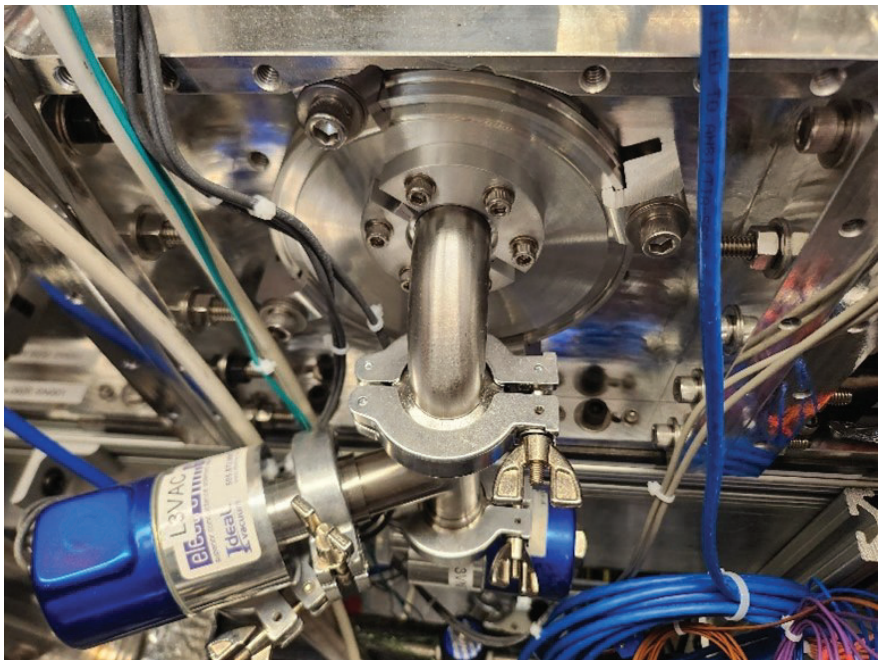

Figure 5.6: The ISO100 blank can have features machined into it to accommodate KF25 bulkhead clamping, enabling the attachment of vacuum connections to the sides of the drying chambers.

window diameter and the nominal KF fitting size. Thus, a third style of machined ISO100 blank is used for these 50 mm windows.

The third ISO100 blank is machined for a 50 mm IR transparent glass window. The window is sandwiched between two O-rings, as shown in Figure 5.7, to maintain the seal on the vacuum chamber on which they are placed. The O-ring between the window and the blank provides the seal for the vacuum environment. The O-ring on the opposite side of the window provides a compliant interface between the window and the clamp used to apply a preload to the vacuum sealing O-ring. This compliant interface ensures that the window is not scratched during installation and removal. The preload force is applied through a machined KF40 blank. The outer diameter for a standard KF40 O-ring is just over 50 mm, making it a good fit for sealing the 50 mm window. These O-rings are used without centering rings because the KF centering rings include an extrusion that would interfere with the window. Although removing this centering ring provides freedom for the O-ring to shrink and be sucked in through the port on the ISO100 blank, the preload force is sufficient to hold the O-ring in place with the vacuum load. If this preload force were not sufficient, or if it needed to be so large that it would damage the window, then the centering rings could be used with their flange on one side machined away. The flange on the opposite side would position the centering ring in either the opening in the ISO100 blank or the KF40 blank used to apply the preload force to the window sealing O-rings.

Using a KF40 blank to apply the preload force to seal the vacuum O-ring allows for the use of bulkhead clamping to provide that preload force. The KF40 blank has a 38 mm diameter circle machined out of its center to create a viewing window to look through

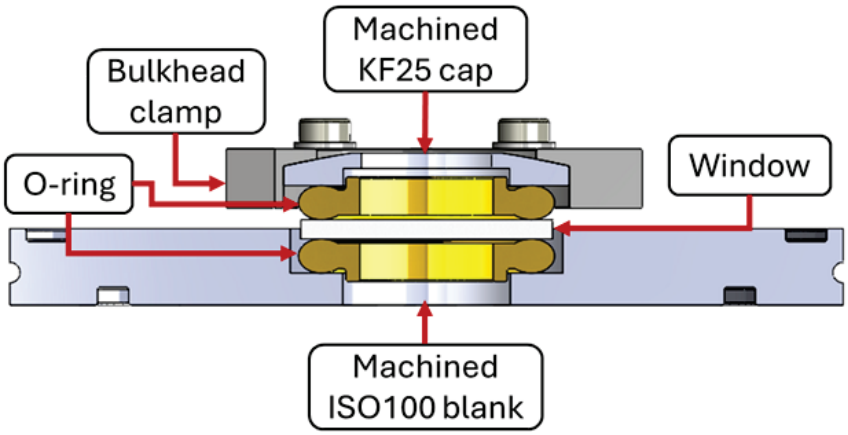

Figure 5.7: The ISO100 blank designed for windows is machined with a pocket for an IR transparent window for measuring the temperature of the product during the drying process. These blank also include bolting for a KF40 bulkhead clamp to provide the preload needed to compress the O-rings used to seal the window to the blank for vacuum. The O-ring between the bulkhead clamp and the window provides a soft interface to prevent scratching on the window during installation and removal.

the IR transparent window. The remaining areas on the blank include the angled surface that interfaces with the bulkhead clamp and the flat surface that interfaces with the O-ring preloading the window. The pocket machined into the ISO100 blank accommodates the depth required for the window and the second O-ring. The pocket is offset from the center of the ISO100 blank so that its position can be adjusted by rotating the ISO100 blank. This adjustability ensures that the window height relative to the machine chamber interior can be adjusted to match the thermal camera viewing position with the vial positions inside the chamber.

5.2.4 Top Window

The top window port is shown in Figure 5.8 with a vacuum valve attached. The central component of this port's geometry is an ISO100 port located along the central axis of the module tunnel, highlighted in Figure 5.3a. This port is shifted from the center of the corresponding stator relative to the chamber's tunnel faces to accommodate the bolt pockets used to bolt consecutive chambers together. Since the assembled modules act as a long tunnel, shifting the port position by small amounts along the length of the tunnel does not significantly affect the performance of parts that interface with this port.

The second element to the top window is the extended rectangle that cuts across the ISO 100 port to create an opening with a width that matches the internal tunnel width. This rectangle is positioned such that the one edge is tangent with the ISO100 port cutout. This positioning minimizes the number of ISO100 bolt holes lost to the widened cutout while keeping the widened cutout within the same tunnel section as the ISO100 port. The widened cutout height is based on the width of a puck and the camera aspect ratio used for measuring the weight sensing system. The camera lens used in this work has a 16:9 ratio. The puck

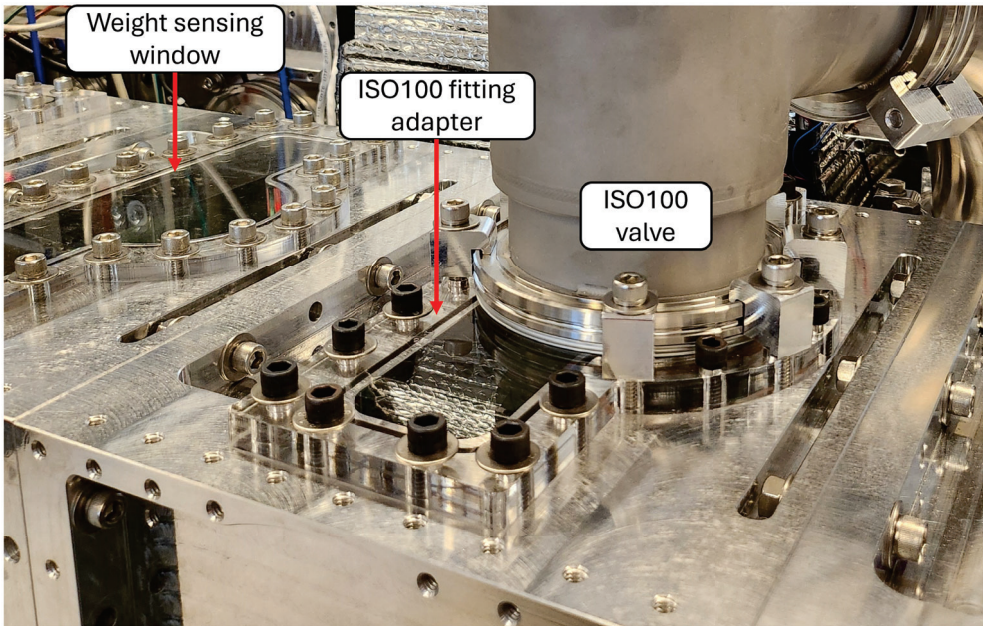

Figure 5.8: The top port on the drying units can be used for viewing the weight sensors described in Chapter 4 or for connecting the chamber to the vacuum system.

width is 120 mm. Thus, the cutout height is 67.5 mm. This cutout height is large enough for the camera to see at least two individual vial weight sensing systems at a time. When combined with the ISO100 cutout, there are positions for the puck which allow measurement of three vials at a time. An image showing the camera's view of the vials through this top window is shown in Figure 5.9.

The top window port has a standard 1/8 in O-ring dovetail groove offset from the cutout edge by 3 mm. A 2-273 O-ring is used in this groove to provide the vacuum seal between the window and the chamber. The bolt pattern around the window starts with the standard pattern for an ISO100 boosted connection. The machine chamber surface can only include six of the standard eight bolt holes used in this pattern due to the extended cutout for the window taking up the space of two of the standard bolt holes. Around the circular arc, additional bolt holes for the window are evenly positioned between the holes for the standard ISO100 pattern. These bolt holes are necessary to ensure that the O-ring is properly preloaded when using an acrylic adapter to connect the drying chamber to ISO100 hardware. The bolts around the rectangular cutout are spaced 38 mm apart from each other.

5.3 Condenser

The drying section requires condensers to remove water vapor from the vacuum system. Without condensers, the water vapor will collect in the vacuum pump oil, leading to pump damage. The condenser, shown with the condensed ice from operation in Figure 5.10, pulls water vapor out of the vacuum system by exposing the vapor to surface temperatures well below the freezing temperature of water at the vacuum system pressure. The vacuum system piping between the condenser and the drying chambers needs relatively large diameter tubing

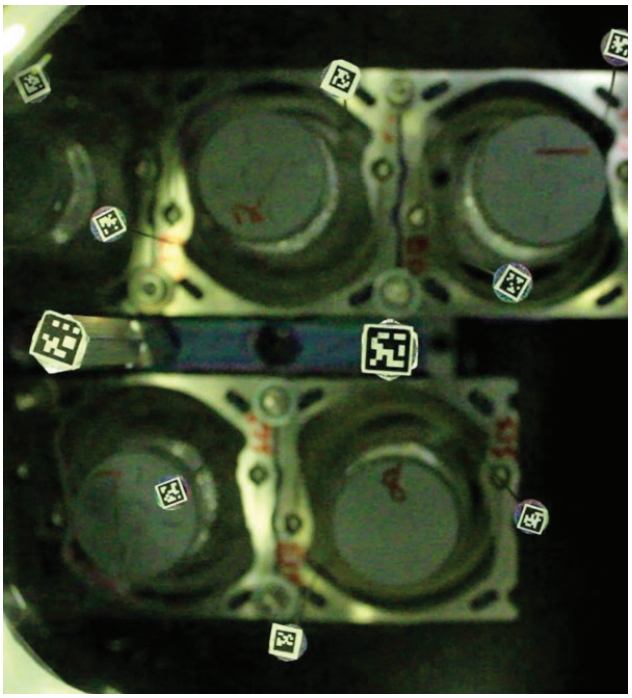

Figure 5.9: The weight sensing system is viewed through the top window on the drying modules. The puck positioning under the window can capture more than one weight sensing system at a time, decreasing the total time required to record weight data for a full puck.

to maintain a low head loss in the water vapor. After the condenser, smaller diameter tubing can be used to transport the inert gas to the vacuum pump.

The condensers in this system use liquid nitrogen to provide the low temperature thermal sink necessary to condense the water vapor out of the vacuum system. The condenser body, shown in Figure 5.11, consists of an extended ISO100 fitting that connects directly to the ISO100 vacuum piping coming from the top of the drying modules. A KF25 half nipple is welded onto the side of the condenser body near the top ISO100 flange which serves as the condenser outlet. This KF25 fitting connects the condenser to the vacuum piping that connects to the vacuum pump. Inside the condenser, a small inner tube extends from the outlet down to near the bottom of the condenser's internal volume. This small tube forces gas to travel down the full length of the condenser body before exiting the system, which increases the usable volume of the condenser for water to freeze.

The condenser body rests inside a standard liquid nitrogen dewar. The volume between the condenser and the liquid nitrogen dewar is filled with liquid nitrogen during operation to cool down the condenser body to operating temperatures. The standard liquid nitrogen dewar lid is replaced with a cut piece of foam that has a hole for the condenser body. A gasket cut from a silicone sheet is placed in the lid's dewar rim connection groove. This gasket improves sealing between the dewar and the lid to prevent liquid nitrogen from splashing out when the dewar is being filled. The condenser body hole is located off-center from the dewar to provide space for the additional features used for the condenser operation. These features include holes for the high- and low-level liquid nitrogen sensors, a hole for the liquid nitrogen inlet, and a groove for mounting to the liquid nitrogen dewar.

(a) (b)

Figure 5.10: The condenser is used in lyophilization systems to separate the water vapor sublimated from the vials from the vacuum system. The water vapor in the vacuum system is condensed into ice which forms on the inner walls of the condenser body. The small tube inside the condenser body forces the water vapor to travel through the length of the condenser, increasing its capacity.

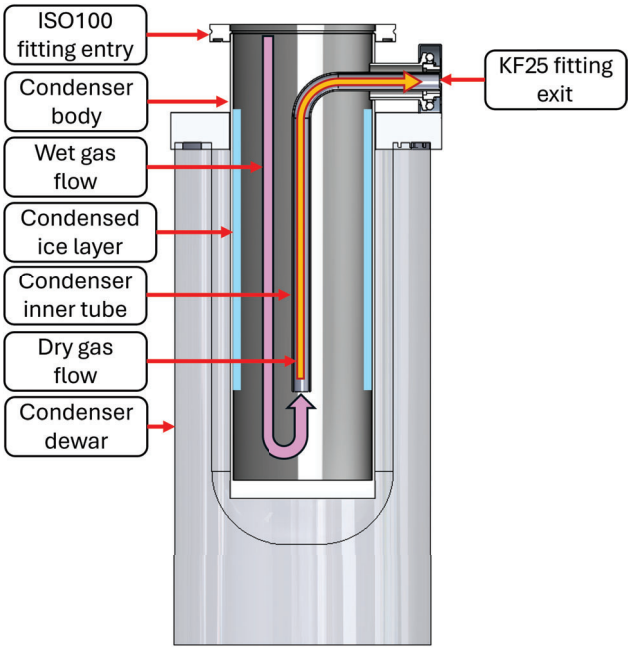

Figure 5.11: The water vapor enters the condenser from the top. The gas is forced to move through the condenser body to enter the condenser inner tube, allowing the water vapor to condense into an ice layer on the condenser body walls. The dry gas then flows up through the inner tube and out of the condenser exit. The condenser dewar is filled with liquid nitrogen to maintain the condenser body wall's low temperature.

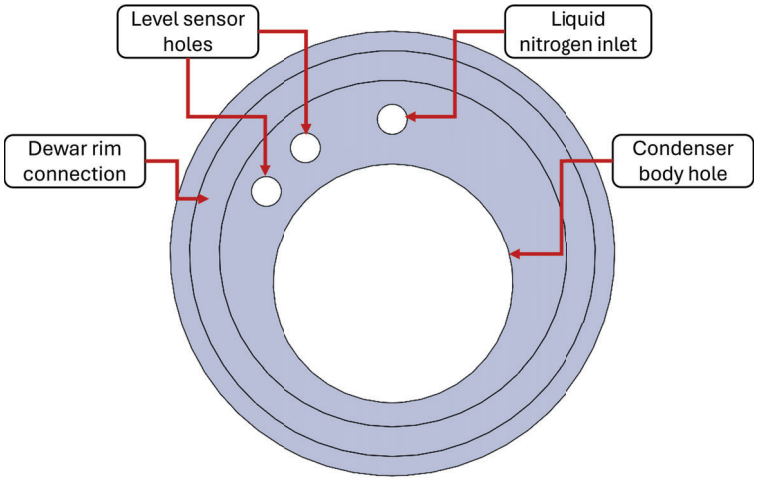

Figure 5.12: The condenser lid includes cutouts for the condenser body, the high- and low-level sensors, and the liquid nitrogen fill inlet.

The condenser has high- and low-level sensors for monitoring the liquid nitrogen level. The low-level sensor identifies when liquid nitrogen should be refilled to maintain condenser performance. The high-level sensor identifies when the liquid nitrogen level is sufficiently high to end the active fill, preventing the liquid nitrogen from overfilling which would create a spill risk. Both sensors are thermocouples which are placed in foam insulation inside a stilling well to isolate the measurements from the splashing induced by the liquid nitrogen fills. The stilling wells are 12 mm aluminum tubing that is flared at the top to rest on the top of the holes cut for the level sensors.

The condenser have a limited amount of water vapor they can remove from the vacuum system before they need to be replaced or regenerated. The capacity limit comes from the growth of the ice layer inside the condenser. As the ice layer grows, it insulates its surface from the condenser walls, increasing the temperature at the ice layer surface. Once this insulation layer grows thick enough for its surface temperature to be above the temperature required to condense the water vapor out of the vacuum line, the water vapor will travel through the condenser and into the vacuum pump. The condenser capacity scales with the interior surface area of the condenser, and the insulative properties of ice determine the resultant thickness that can be sustained before the condenser reaches capacity. The condensers used in this system are sized so that given the 100 mm condenser body diameter, the condenser body length supports sufficient ice build up to have a 12 hour condenser saturation time when running at full capacity.

Once the condenser reaches capacity, it needs to be regenerated for the system to remain functional. This regeneration is done by removing the condenser from the system and replacing it with an empty, dry one. The condenser is separated from the vacuum system by closing the valves that connect it to both the chamber and the vacuum pump. The condenser is then detached from the vacuum line and replaced with a fresh condenser. Once the full condenser is removed from the system, it can be heated to melt the ice, which is then dumped down a drain as liquid water. The condenser interior is then dried off, and it is set aside until it is needed to replace the next condenser that reaches full capacity.

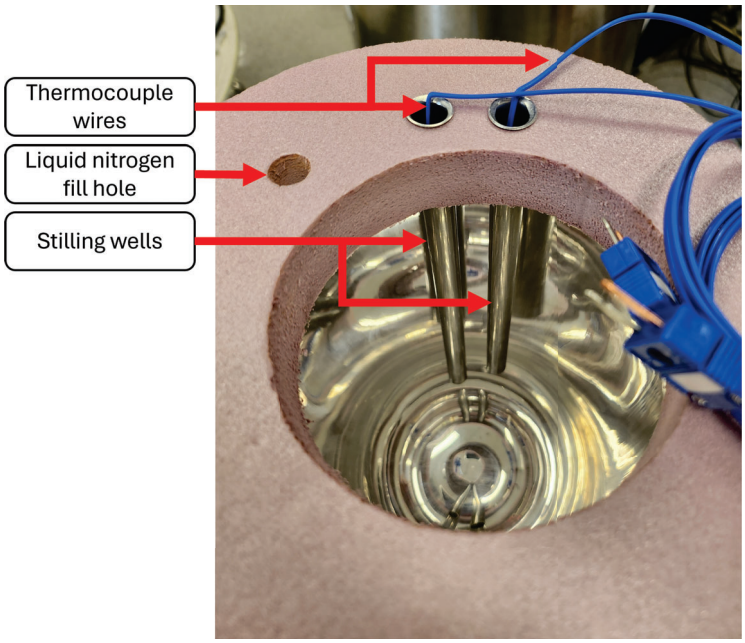

Figure 5.13: The thermocouples used for sensing the liquid nitrogen level are placed within stilling wells to ensure their measurements are not influenced by splashing within the condenser dewar when it is being refilled with liquid nitrogen.

5.4 Vacuum Piping

The vacuum system is connected to the drying section through ISO100 elbow valves mounted to the tops of standard drying modules. These drying modules have an adapter plate which covers the widened section of the top window cutout while keeping a central port open for an ISO100 fitting. This adapter plate seals to the chamber using the top cutout's O-ring. The adapter seals to the ISO100 fitting using a standard ISO100 centering ring. The adapter plate is bolted to the chamber using all available bolt holes except those positioned for the standard ISO100 bolt pattern. The elbow valve is positioned on top of the ISO100 centering ring on the adapter, and half claw clamps are bolted into the remaining bolt holes on top of the chamber to provide the preload for the ISO100 centering ring O-ring. A standard ISO100 90 degree elbow is attached to the angle valve to redirect the vacuum line down through the center of the machine. A 450 mm long ISO100 nipple is then attached to the elbow to move the vacuum connection below the machine table for the interface with the condenser.

The condenser is attached to the ISO100 nipple below the machine table, where the connected interface is accessible for removing and swapping condenser bodies. The condenser attaches to the condenser valve assembly through a 100 mm long KF25 bellows. This bellows provides a compliant connection to account for small variations in the positioning of the condenser between regeneration cycles. The bellows connects to a tee that splits between two valves, the condenser vent valve and the condenser vacuum valve. The combination of the angle valve at the top of the chamber and this valve assembly allows the condenser to be isolated from the machine and the vacuum line without breaking the vacuum environment in the chamber. This isolation is required to vent the condenser in order to swap it with a

fresh one for regeneration.

The condenser valve assembly is attached to a central KF25 vacuum tube which feeds back to the system's vacuum pumps. Each condenser valve assembly connects to this line through a tee. This system can easily be extended for the number of condensers and vacuum pumps required to maintain the vacuum conditions as the system production rate and corresponding length increase, as described in Chapter 2.

5.5 Drying Control

The drying process is controlled by adjusting the power delivered to the vials during drying. In traditional batch lyophilizers, this power is delivered through the vial shelves. However, this system includes magnetically levitated pucks that carry the vials instead of shelves. Since conductive heat transfer is not viable in this system, the power for driving sublimation is primarily delivered through radiation from the machine chamber's internal surfaces. This radiative power delivery depends on the temperature difference between the chamber surfaces and the vials. Thus, the power delivery can be controlled by controlling the machine chamber temperature.

The first decision in setting the chamber temperature control method is to determine whether the system needs heating, cooling, or a combination of both. If the chamber without any temperature control can produce vials that collapse due to temperatures exceeding their collapse temperature, then the system primarily requires cooling. If the vial temperatures remain well below their collapse temperature, then the system primarily requires heating. If the vial temperatures remain close to their collapse temperature during sublimation, then the system would benefit from a combination of heating and cooling. During initial testing, when operated at high capacity and without temperature control, the lyophilizer was able to produce failed cakes due to collapse. This high power delivery is primarily driven by the heat produced by the Planar Motors stators which heat up the machine chambers during operation. Thus, cooling is the primary temperature control implemented in this system.

The drying chamber is cooled through water blocks installed on the outside walls of the machine chambers. While the stators include cooling channels and could be cooled instead, running cold water through the stators resulted in a loss of vacuum performance. This vacuum performance loss could be due to the stator cooling down more than the machine chamber, causing it to shrink away from the machine chamber's walls. This shrinkage can stretch the vacuum seal between the stator and the machine chamber, potentially creating leak paths into the vacuum environment. Cooling the machine chamber instead of the chamber would cause it to shrink around the stator, further compressing the vacuum seal instead of stretching it.

5.6 Capping

The vials go through the lyophilization process with their caps partially depressed. Lyophilization vial caps include a cutout through which water vapor can escape during sublimation. Once the sublimation process is complete, the vials must be fully capped while still within

the vacuum environment. This capping ensures that the product in the vials does not reabsorb moisture from the air after it exits the lyophilizer. After the vials are capped, they are sealed and ready for packaging and distribution.

This capping operation requires the actuation of components to ensure the cap is fully seated in the top of the vial. A preferable choice would be to use the electromagnetic motion system, as it would not require adding an additional actuator to the system. Adding components adds complexity and increases the risk of leaks or contamination, so they are undesirable if not necessary. However, using a planar motion with the mover and a tapered surface to gradually push the caps into the vials requires dragging surfaces against each other, increasing the risk of particulate generation. The Planar Motors pucks also do not have enough vertical travel to fully seat the vial caps, as they are limited to 3 mm of vertical motion while the caps need to move about 6 mm. Thus, another system is required to perform the capping operation.

The lyophilizer includes other elements that already incorporate vertical motion in the load lock chambers. The load lock doors travel vertically, as described in [25]. This motion is driven by a pneumatic piston which uses a dry dynamic rod seal to maintain the vacuum environment. This same piston and seal geometry is reused to create a passthrough for the capping mechanism. Using the pneumatic piston provides a moving component with more than sufficient travel to successfully cap the vials.

The pneumatic piston moves between two set positions, fully opened and fully closed. The pistons used in the load-lock doors have more travel than is required for the vial capping. The load lock doors need to move about 125 mm vertically, while the capping piston only needs to travel about 6 mm. If the piston is allowed to move over its full travel length when mounted to the roof of a chamber, it would either exit the chamber when retracting or drive through the vial during extension. Exiting the chamber would break the vacuum seal, and driving through the vial would shatter the vial and ruin the product, so the piston travel length needs to be shortened. This shortening is achieved by attaching a shaft collar to the piston, shown in Figure 5.14a, limiting its retraction distance. When the piston retracts, the shaft collar contacts the piston body, preventing the piston from fully retracting into its cylinder. The piston is mounted such that its full extension corresponds to the displacement required to fully cap the vial, so its motion is only limited by the shaft collar in one direction.

This displacement-based capping strategy includes a risk inherent in depending on the precision of the piston displacement used to cap the vials. If the piston displacement is too large, then the piston could risk breaking the weight sensing system or the vial itself. If the piston displacement is too small, then the vial caps will not be fully seated, which would prevent them from sealing the product inside the vials. This risk is addressed by adding a compliant element to the end of the capping piston. This compliant element is a short spring with a pusher connected at the bottom. As the piston extends, the pusher engages with the cap, and the compression of the spring provides the force required to cap the vial. The spring provides some flexibility in the specific piston displacement, loosening the strictness on its specific mounting height. The pusher is a 3D printed piece with a textured surface and an aluminum foil coating. The textured surface and foil coating are used to ensure that the pusher does not stick to the vial cap during the capping operation. Initial testing showed that a flat printed surface on the pusher could sometimes generate suction adhesion to the vial cap when it pressed down, causing the system to pull the cap back up off the vial during

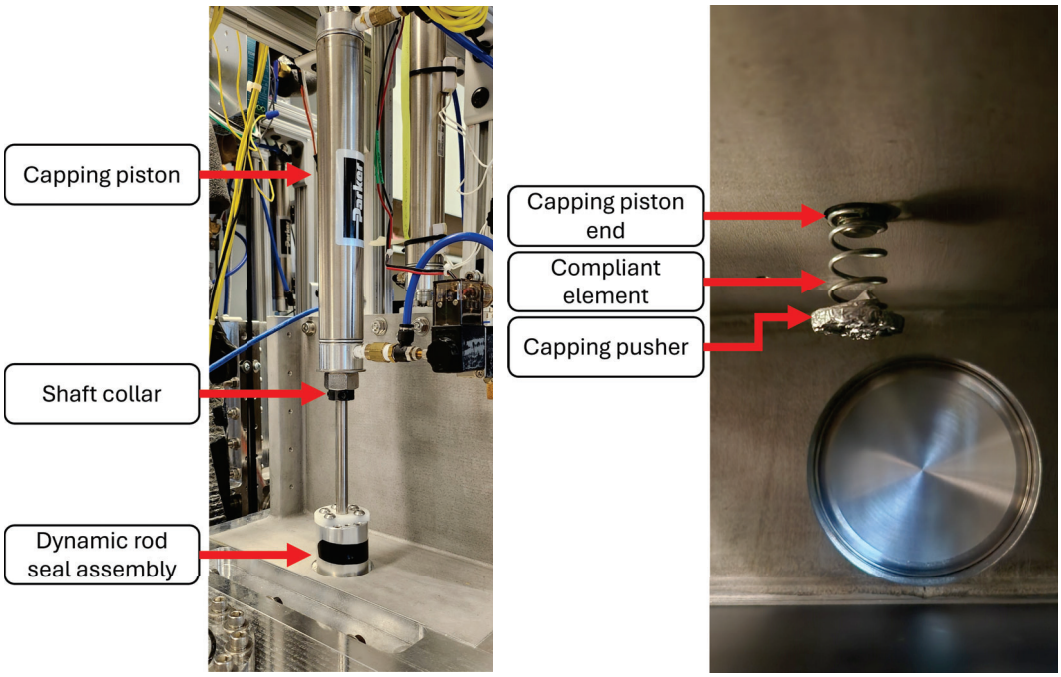

(a) (b)

Figure 5.14: The vials are capped by a pneumatic cylinder pushing down on the vial caps at the end of the drying section. (a) shows the cylinder outside of the chamber with its dynamic rod seal mounted to the top of the chamber. (b) shows the compliant capping plunger used to press the vial caps into the vials. The compliant plunger decreases the risk of underpressing on the cap and not sealing the vial or overpessing on the cap and potentially shattering a vial. This capping process ensures the vials are sealed and ready for release from the lyophilizer.

the capping operation. Once the surface texture and foil coating were added, the pusher no longer stuck to the vial cap. The capping system demonstrates high reliability, successfully capping hundreds of vials without capping or vacuum environment failure.

The current capping system caps vials individually. While this operation is fast enough for the production rate of the continuous lyophilizer built for this work, there are ways to speed it up if needed. One option would be to use an attachment on the piston that caps more than one vial at a time. This attachment would divide the piston force between multiple vials, which can limit the maximum number of vials that can be capped simultaneously. Also, the attachment would need to avoid interference with any of the weight sensing system components. Furthermore, a multi-vial attachment would need to evenly press down on all caps that it engages with to ensure that each vial is fully capped. Another method for improving the capping speed would be to add additional pistons such that multiple vials or pucks could go through the capping process simultaneously. This addition would require more passthroughs, adding potential leak paths to the system.

5.7 Results

The most critical performance validation for the drying system is its ability to maintain vacuum during the lyophilization process. Figure 5.15 shows a typical pressure trace for the system over a 12 hour drying cycle. The pressure in the drying section stays below 5 Pa throughout the drying process, which is sufficiently low to drive sublimation and ensure that convective heat transfer is minimized. This result demonstrates that the drying section can successfully perform the sublimation process during lyophilization.

5.8 Future Work

The drying chambers used in this work provide an environment for the successful sublimation of water during the lyophilization of product in vials. These chambers could be further improved in future development of the lyophilizer. One area of future work would be to investigate methods for power delivery beyond simply using the radiative power of the chamber walls. One such method involves using microwaves to accelerate the drying process [50][51]. A microwave enabled drying chamber could speed up the drying process, increasing the system production rate with less drying modules. Another future change could be to integrate the cooling channels into the machine chambers. This integration would provide improved cooling to the system components, providing better temperature and thus radiative power delivery control. A further area for development is altering the capping system to cap more than one vial simultaneously to meet faster production rates.

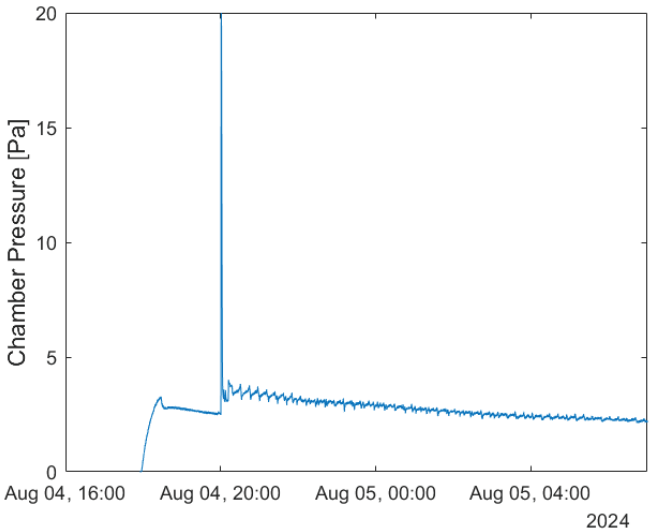

Figure 5.15: The pressure in the drying section stays below 5 Pa during a typical drying run (started at 20:00), which is sufficiently low to both drive sublimation and ensure that convective heat transfer is minimal. The initial spike before the start of drying is an artifact of the load-lock operation

Chapter 6

Manufacture

The machine design described in the previous chapters must be validated by building the actual system. The manufacturability of the system is critical to consider when finalizing the design. A system that is too difficult or costly to create will not be built, which prevents it from achieving the benefits it is designed to create. These manufacturability concerns are discussed in the prior chapters, and the specific details of the decisions made to manufacture the machine built for this work are discussed in this chapter. The complete assembled continuous lyophilizer is shown in Figure 6.1.

The machine module components described in Chapter 2 are manufactured using different methods. The machine chambers which make up the tunnel structure of the lyophilizer are cast and machined. Casting these machine chambers provides a scalable and economical method to produce the general unit geometry, while the machining process ensures that each unit meets the required tolerances for successful assembly. The module baseplates are

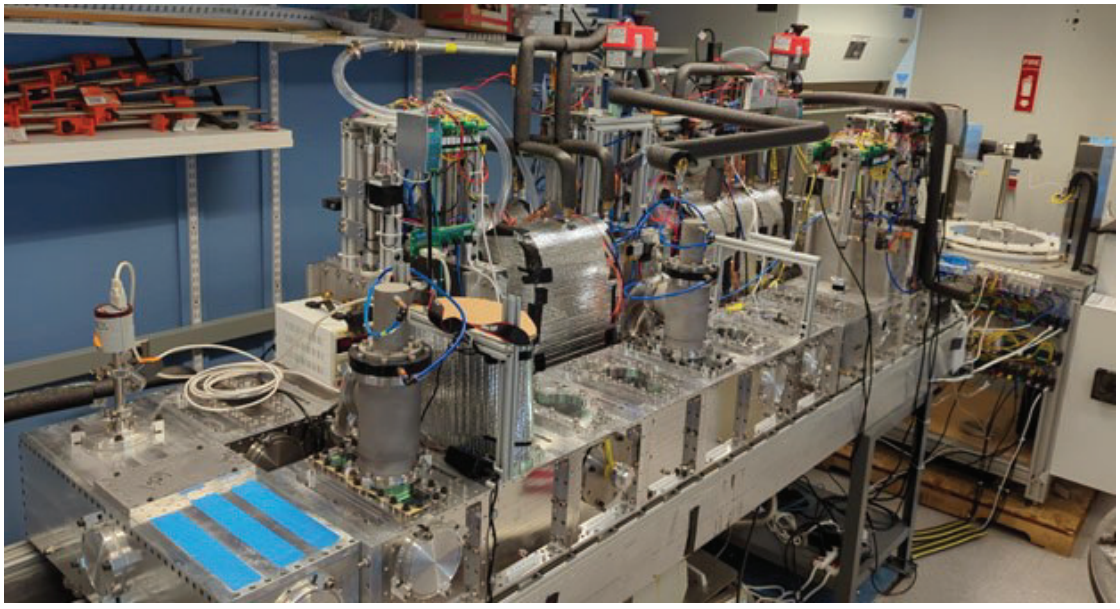

Figure 6.1: The continuous lyophilizer built at MIT for this work. This machine is designed to continuously lyophilize vials at a rate of 30 vials/hour.

machined from plate material, though they could be manufactured using other 2D geometry methods such as waterjetting or wire EDM cutting. The stators are acquired directly from Planar Motors. The modules are assembled by inserting the stators into the machine chambers and then injecting the vacuum sealant into its groove, as described in Chapter 2. Once the sealant is injected, the baseplate is added to rigidly connect the two components. Once the injected sealant cures, the module is ready for assembly into the larger system.

The machine table described in Chapter 2 has two steps for assembly: creating the sliding rail beams and mounting them on the table legs. The sliding rail beams are made by epoxying t-slotted aluminum extrusion to rectangular aluminum tubing. The table legs are welded steel machine tables sized to support the weight of the lyophilizer. These tables are placed on top of an aluminum sheet to spread the load over the floor surface and a rubber sheet to damp vibration which could be transmitted through the floor. Once the beams are epoxied together, they are placed on top of the table legs at their nominal spacing. Assembled modules are then placed on the sliding rails to self-align the rails, reducing the risk of binding in the rails when modules are slid into position.

During this work, the lyophilizer was assembled in smaller portions to enable integration testing during module manufacture. The first assembly was a minimum-sized unit consisting of four modules, one each for loading, cooling, drying, and a load-lock. This unit enables basic process testing and serves as a blueprint for a research-scale system that uses the same internal geometries as the larger production continuous lyophilizer. The continuous system is then built up in multiple stages, with each stage adding a level of integration and complexity, until the fully looping system is in place. This stepwise assembly provides robust testing for each interface so that problems can be identified during development and addressed without requiring disassembly of the fully integrated system.

This continuous lyophilizer is tested with a commercial vial filling machine to demonstrate its ability to interface with standard factory equipment. The vial filling machine interfaces with the lyophilizer at the loading and unloading section of the system. This interface is primarily connected to the motion system and the weight sensing system. A cone is added to the vial filling machine gripper to prevent it from interfering with the sensing arms used in the vial weight sensing system. An additional pulling element further moves the sensing arms away from the gripper to prevent interference. This system was shown to successfully load and unload hundreds of vials from the system on multiple tray assemblies.

6.1 Built System Parameters

The nominal scale used to build the test system in this work is a minimum scale that includes the necessary sections for a full continuous lyophilization cycle. These sections include loading and unloading, conditioning, freezing, and drying, along with load-locks separating each section. The layout used in this system appears in Figure 6.2.

The system must include process chambers for freezing and drying to perform a complete lyophilization cycle. These process chambers operate at relatively stable, constant conditions as the vials move through them. The freezing section must have independent process chambers for conditioning and freezing operations to achieve this constant operational condition. The conditioning section subcools the vials with a temperature setpoint around -5°C, while

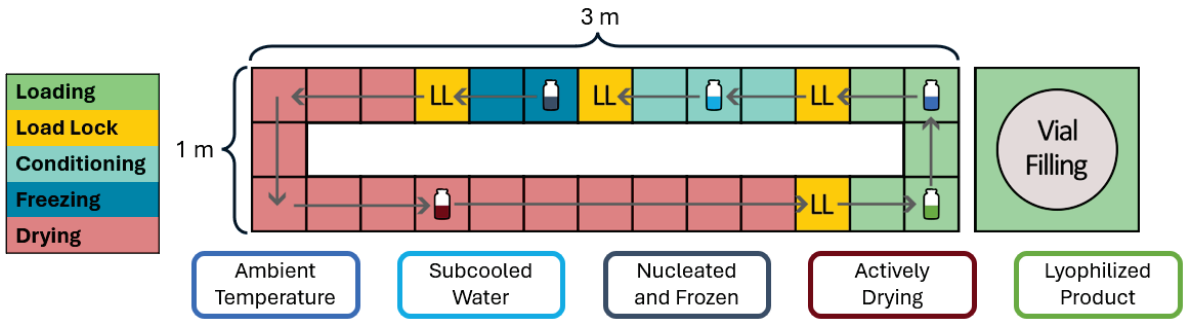

Figure 6.2: The system built to validate the design work in this thesis is sized based on the minimum necessary units to create a fully continuously looping lyophilization cycle. The vial starts in the loading section in the top right at ambient conditions. They then move into the conditioning section where the water in the vial is subcooled. Next the vials move into the freezing chamber for nucleation and solidification. The vials then move into the drying section, where they are actively drying as water sublimates out of the vials. Finally, the lyophilized product returns to the loading section for the finished vials to be removed and replaced with fresh vials ready to be lyophilized. The loop at this size is expected to have a production rate of 30 vials/hour.

the freezing section solidifies the ice crystals in the vials with a temperature setpoint closer to -50°C. These cooling sections each require door units at their entry and exit to interface with the load locks that separate the process chambers. The conditioning section includes an additional standard cooling unit without door features. This chamber provides room in the conditioning section to test freezing profiles that include annealing, as described in Chapter 3. While proper annealing installation would involve adding a dedicated machine section, this additional section would increase the number of doors required to separate the chambers, significantly increasing the overall system size. Instead, the vials can be moved into the freezing section to induce nucleation via thermal quench, then moved back into the final conditioning chamber for slow ice crystallization to facilitate large ice crystal growth. The vials can then be moved back into the freezing section for full ice solidification. These conditioning and freezing sections are kept on the same side of the lyophilizer loop to enable this motion and to leave the other side of the loop available for drying units. If the conditioning and freezing sections are placed on opposite sides of the loop, they would take up the entire loop and leave no space for drying chambers. Because these chambers are placed on the same side of the loop, they effectively set the minimum length of the fully continuous lyophilizer.

After accounting for the modules required for the freezing process, the loop to create the continuous lyophilizer is completed with drying modules. The total freezing process length includes three load locks, three conditioning chambers, and two freezing chambers. One door unit is required to transition from the load lock chamber exiting the freezing chamber to enter the drying section. A turn module is required to move vials from one side of the lyophilizer to the other side. At the lyophilizer exit, a door unit and a load lock are required for exiting the system. This accounting results in eight drying stators on the return tunnel within the lyophilizer. Combining these drying modules with the three drying units in the

turn module and the drying door unit at the entry to the drying section, this lyophilizer will include twelve total drying modules. These twelve drying modules create an effective 3 m of drying tunnel length, given the 250 mm module length. With the puck furniture described in Chapter 4, each tray covers about 200 mm of tunnel length. Given the 3 m total drying length, fifteen pairs of trays could fit in the drying section, giving a total of 30 trays. The vial trays each carry ten vials, so a total of 300 vials can be in the drying section at a time. Given an estimated drying process time of ten hours, the drying section limits the lyophilization rate to 30 vials per hour. Comparatively, the freezing chamber can accommodate only two pairs of trays at a time, limiting the freezing section to 40 total vials. However, the freezing process only takes 30 minutes, so the freezing section only limits the lyophilization rate to 80 vials/hour. Accordingly, the drying section serves as the limiting factor on the production rate for this continuous lyophilizer.

At this scale, the system includes three condensers, two mass sensing stations, and two thermal measurement stations. The number of condensers is based on how many are needed to maintain the required vacuum conditions within the system, as described in [40]. The number of mass sensing and thermal measurement stations is limited to only two to minimize complexity while validating the performance of independent stations. As needed, the number of these stations can be increased for greater data density. The required data density can be determined by keeping a tray at a sensing station for its entire drying period, then sampling that data at lower frequencies corresponding to the frequency at which a puck would go through measurement stations in a balanced measurement system. The pucks can be moved in a carousel style motion to enable multiple mass measurements from a single puck while it is in the drying section. As the drying section expands, each weight station can have an independent carousel, and the station measurements can serve as gates for allowing each puck to progress to the next area within drying. This separation would allow for distinguishing conditions within the drying chamber for primary and secondary drying, where higher thermal input can be tolerated during secondary drying than primary drying. The sublimation process is not affected by having vials at different stages in the drying process next to each other, particularly because each vial can be monitored independently by the weight sensing system described in Chapter 4.

The layout used for the system built in this work requires a loop consisting of 22 Planar Motors stators. Of these 22 stators, five are left exposed without chambers for loading and unloading. An additional three stators are dedicated to the turn module, which is located at the end of the loop to allow pucks to return to the start. Based on the required Planar Motors stator spacing, it is most effective to make a unique unit for this turn module. Only one turn module is needed for the full system because the loading section stators need to remain exposed for interaction with the vial filling machine. This allocation leaves 14 stators to be assembled with machined individual chambers.

6.2 Module Manufacturing

The module elements are split into those which must be custom built and those which can be assembled using off-the-shelf components. Each module consists of at least three key components: the Planar Motors stator, the machine body, and the module base, as described in

Chapter 2. The Planar Motors stator is a commercial part, while the machine body and the module base are custom parts for this machine. These custom components must be manufactured for this specific application. While the minimum unit system only needs 14 machine bodies, this sample set can validate the manufacturing method for much larger quantities. Potential manufacturing options include machining the modules from solid material, assembling the modules from machined plates, and having the machine module base geometry cast and then machined. Machining the modules from solid material creates a significant amount of waste, as over 50% of the material would need to be removed. Assembling the modules from machined plates reduces this waste, but it adds sealing complications between these plates. Casting the base geometry similarly reduces waste, but it adds machining complications because the casting process creates internal stresses within the material that cause the unit to warp during the machining process. Of these choices, casting is chosen because it reduces the material waste without creating additional sealing surfaces. Casting can also accommodate a range of materials, including aluminum which is used to make the machine chambers in this system.

6.2.1 Casting

The shared module geometry and relatively large number of modules needed for this system make casting a reasonable choice for creating the base chamber geometry. This cast body is then machined to meet the final feature tolerances required for successful assembly and operation.

There are four versions of the castings, shown in Figure 6.3, matching the four types of modules used in the final system. The shared geometry between all four versions is anticipated to make the molding process easier as elements of the casting core can be added or removed to change between modules, rather than requiring completely new cores to cast each version. The casting variety can be reduced, if necessary, down to two core types, modules with and without door features. The modules that transition between load-locks and standard tunnel units can be created from a casting that includes material for door features on both sides by removing the extra material from the door features that are not required. While this methodology results in higher material costs, it introduces additional flexibility on the manufacturing side by preventing an accumulation of transition units which may not need to be utilized. Additionally, the removed door feature material can be reused in the casting process.

Casting components for vacuum systems inherently carries some risk due to the porosity found in cast parts, such as the porosity seen in Figure 6.4. This porosity arises from impurities in the melted metal. The porosity creates a risk of internally linked pores that form leak paths through the chamber walls, compromising the ability of the chamber to hold a vacuum seal. Reducing the size of the pores in turn reduces the likelihood of linked pores forming a leak path through the walls. The risk of leaking through the walls can be further mitigated by using vacuum impregnation to plug the porosity holes. During vacuum impregnation, the surface level pores are filled with a liquid polymer. This polymer cures inside the pores, plugging them. While the polymer does not create a smooth profile flush with the metal surface, it does block any potential connected pores that would create leaks through the walls.

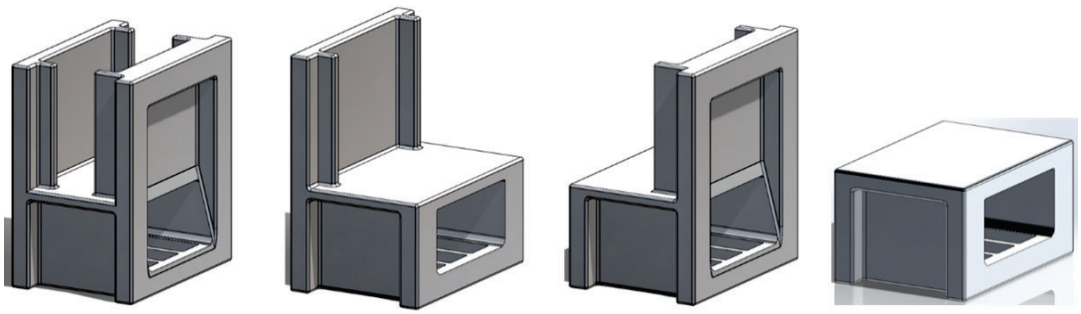

Figure 6.3: The castings have four different varieties. One includes features for doors on both faces, one includes features for supporting the door on the neighboring module, one includes features for a door on one side of the module, and one does not include any door features.

Another risk created by the casting porosity is that the pores could create a gap that spans another sealing interface, such as an O-ring seal. These pores could also create sharp edges that can damage an O-ring when it deforms during the preloading or sealing processes. The vacuum impregnation does not create a smooth surface when it seals the pores, so it does not address the porosity risk of creating leaks across sealing paths. This risk is mitigated by keeping the sealing path width much larger than the maximum pore diameter. The ratio between sealing path width and pore diameter can be increased by increasing the sealing path width and decreasing the maximum pore diameter.

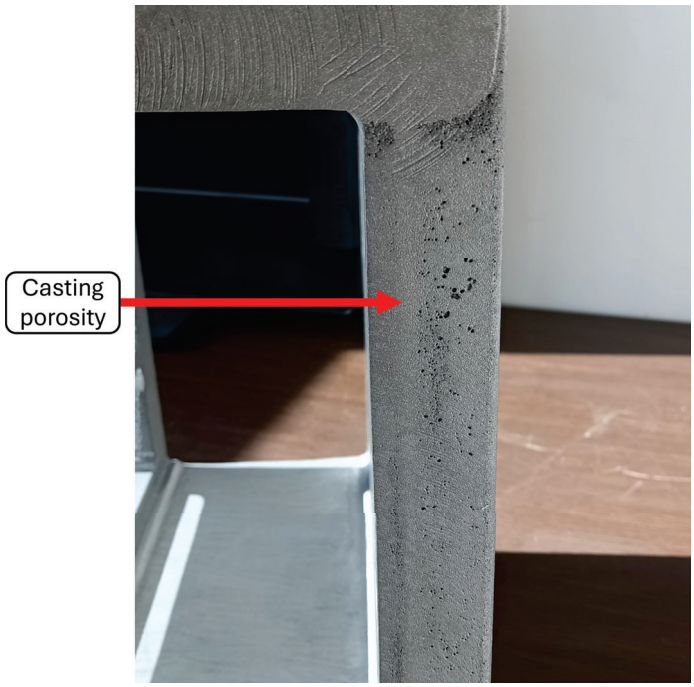

Figure 6.4: Cast parts often include porosity resulting from impurities in the liquid metal. This porosity creates a leak risk in parts made for vacuum sealed components, such as the machine chambers used in the drying section of the lyophilizer.

The machine chambers made for this system were cast by Yankee Casting in Connecticut. They were poured using a 50 mm x 59 mm x 22 mm 30 PPI size filter. This filter was used after initial testing using their default mesh demonstrated unacceptable porosity levels. The molds also underwent an argon gas purge before pouring to remove most of the oxygen in the mold, minimizing the oxides which are the primary cause of the porosity in the castings. Furthermore, the in-gates for the cast pour were reduced from a 1.5 in x 1.5 in gate to a .75 in x .75 in gate to reduce the velocity of the metal moving through the filters, further reducing the amount of oxides entering the castings. The units are cast from Aluminum 356-T6. A T6 heat treatment that hardens the aluminum was used to allow the aluminum to machine cleanly and so that the required dimensional tolerance can be achieved. Without the heat treatment, the cast aluminum remains soft and does not machine as cleanly, making it difficult to achieve the required geometric tolerances in the system during machining.

6.2.2 Chamber Machining

After the modules are cast, they must be machined to meet the tolerances required for assembly and function. This machining was done by the StartSomething machine shop in Somerville. The cast modules include an additional 10 mm of material thickness on their external surfaces to provide material to be machined away during this process. This thickness allows for warping in the module geometry during the casting and machining processes. Also, the material as cast does not have a proper datum for the external geometry relative to the internal tunnel, so this material must be machined in multiple iterations to measure and verify the creation of corners which can be used as an initial machining datum. A finished part compared to a cast module is shown in Figure 6.5.

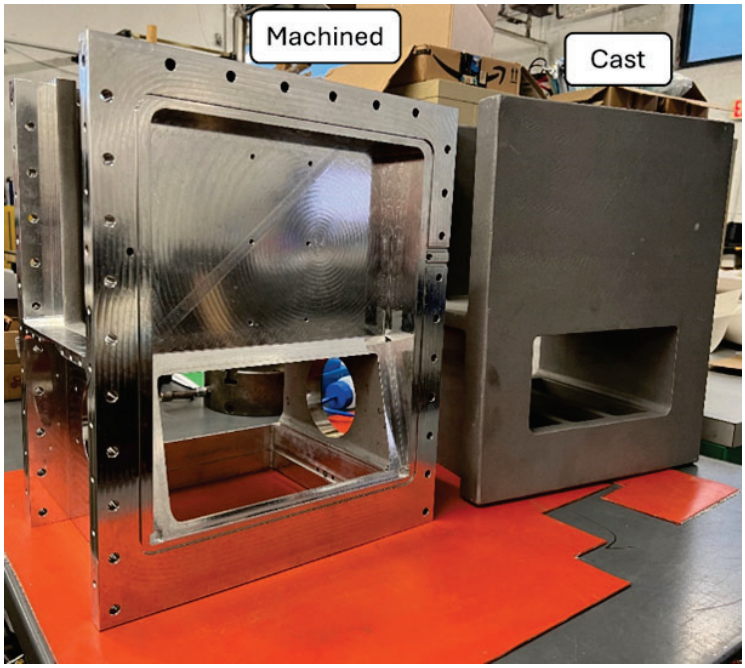

Figure 6.5: A final machined module is shown next to the cast module from which it is made.

6.2.3 Baseplate Machining

The baseplates contain mostly 2D geometry features. These features can be made with traditional machining, or the main features can be cut using a 2D cutting tool such as a waterjet. The baseplates also include dowel pin features which need to be machined in place. These dowel pin hole features provide the option to use pins for precisely locating the Planar Motors stator and the machine chamber relative to the baseplate. These holes must be machined to ensure that they meet the precise diameter requirements for the press fit of the pin into the hole, and to ensure that the pins are properly located on the baseplate. The baseplates are produced in bulk because they are the same for all modules other than the ends of the tunnel. At the ends of the tunnel, a special baseplate that is three stators wide is used to connect the three stators spanning the short dimension of the lyophilizer. This baseplate can be manufactured using similar methods, as it includes the same features as a single module baseplate.

6.2.4 Module Assembly

Once the chambers and baseplates are machined, the module can be assembled. This assembly includes putting the three module components together and sealing the interface between the stator and the machine chamber for vacuum. The most critical element of module assembly is the sealing of the stator to the machined chamber. This sealing is done with an injected thixoflex sealant. The seal created by this injected-in-place gasket performs significantly better when it adheres to both the chamber and the stator, blocking off potential leak paths between the two objects. Without this adhesion, the vacuum load can cause the sealant to deflect away from the walls, creating leak paths which compromise the vacuum seal [25]. To improve sealant adhesion, both the stator and the sealant groove on the machined chamber are sandblasted. This sandblasting increases the surface area at the interface, improving the sealant's ability to bond to the surfaces. A sandblasted chamber surface is shown in Figure 6.6.

After the machine chamber and the stator are sandblasted, they are ready for the assembly process. This process involves three steps: pre-sealing the machine chamber and the stator together, injecting the sealant into the sealant groove, and bolting the machine chamber and the stator to the baseplate. The machine chamber features designed for this process are described in detail in Chapter 2. During the pre-sealing step, Av-Dec Thixoflex Grey is applied to the chamber above the sealant groove, shown in Figure 6.7, and to the stator below the sealant groove position. This sealant is necessary due to the clearance between the stator wall and the stator pocket in the chamber, through which the injected sealant can leak out of its groove. This clearance is necessary to ensure that any stator will fit into any machine chamber. Once the pre-seal is applied to both components, the stator is inserted into the chamber. This insertion is done with both surfaces upside down so that the stator pocket is accessible. The top surface of the stator rests on the lip machined into the chamber, and a weight is used to ensure that the stator is fully seated on this lip. The Thixoflex Grey sealant is then left to cure, creating a barrier on the top and bottom edges of the vacuum sealant groove.

The second step involves injecting Thixoflex Self-Leveling Green into the sealant groove

Figure 6.6: The stator pocket in the machine chamber and the Planar Motors stator are sandblasted in preparation for vacuum sealing. This sandblasting increases the adhesion between the injected vacuum sealant and the two components, improving the vacuum seal fidelity.

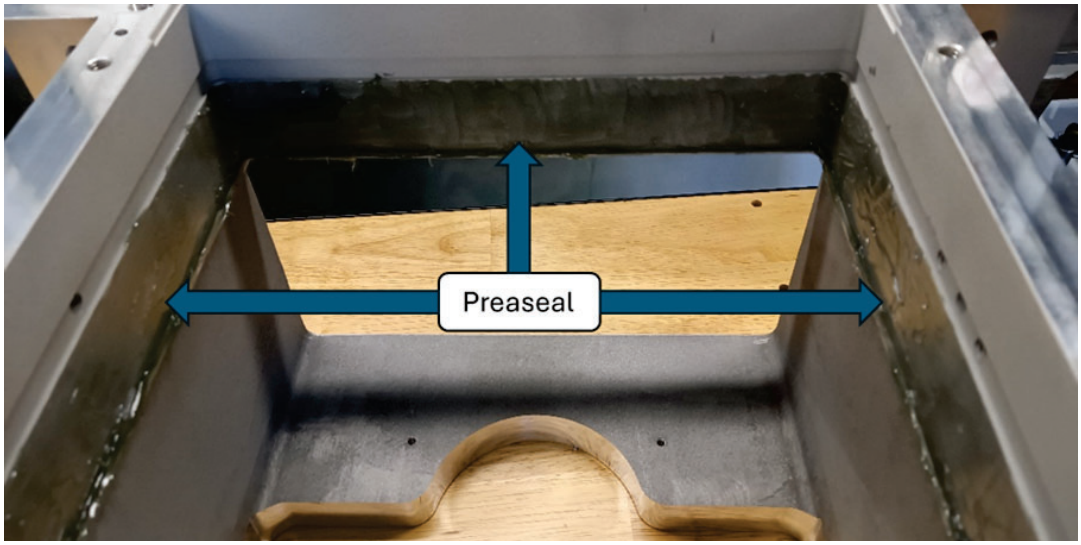

Figure 6.7: The presealant is applied to machine chamber and the stator to prevent the injected vacuum sealant from leaking out through the gaps between the machine chamber and the stator. These gaps are necessary to ensure any stator can be mounted into any machine chamber, but they must be sealed to ensure the vacuum seal can completely fill its sealing groove.

between the stator and the machine chamber. This sealant is injected from a cartridge using a pneumatic applicator to ensure consistent flow during the injection process. The Thixoflex Self-Leveling Green is a two-part epoxy that is mixed through a mixing nozzle during injection. The first injection step is to purge the injection line, which is done by ejecting the sealant into a waste receptacle through the mixing nozzle until air bubbles are no longer seen in the nozzle. Once the mixing nozzle is purged, it is connected to the injection line on the chamber, as shown in Figure 6.8. The sealant is then injected into the groove in the machine chamber while the outlet holes are monitored. Once the sealant begins flowing out of one of the two vent holes, that hole is plugged with a set screw. The injection continues until flow is seen through the second vent hole, at which point that second port is plugged with a set screw and the initial plug is loosened. The exit line is monitored for bubbles, and injection is considered complete when 30 seconds pass between visually observing bubbles in the exit line. After completing injection, the mixing nozzle is disconnected, and the sealant is left to cure for at least 12 hours before being subjected to vacuum loads.

The baseplate can be added to the module assembly while the injected sealant is still curing. This assembly is shown in Figure 6.9. This baseplate application should occur no less than four hours after the sealant is initially injected to ensure the injected sealant has set, even if it has not fully cured. An epoxy shim is added between the stator feet and the baseplate to ensure that the variation in the height of the stator and the depth of the stator pocket in the machine chamber do not prevent module assembly. Silicone grease is added to the stator and baseplate interfaces to prevent the epoxy shim from bonding the stator to the baseplate, so the parts can be separated if necessary for inspection or maintenance. A thin layer of this silicone grease is applied to the stator feet, the baseplate where it contacts the stator feet, and the machine screws which bolt the stator to the baseplate. The epoxy used for the shim is then applied to the stator feet before installing the baseplate. The epoxy layer is squeezed between the stator feet and the baseplate, creating a rigid layer matching the

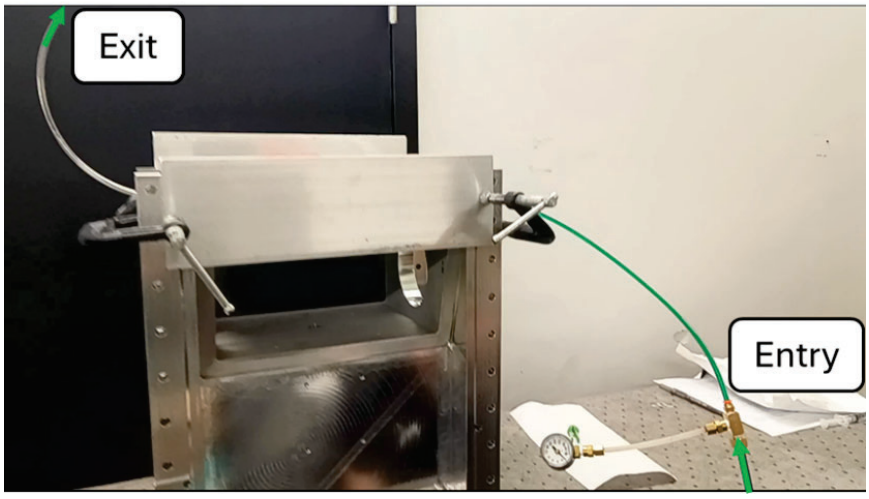

Figure 6.8: The Thixoflex Self-Leveling Green is injected through the sealant groove to create the vacuums seal between the stator and the machined chamber. The exit tube is oriented upwards and monitored for air bubbles to identify when the injection is complete.

height of the gap between them as it cures. Excess epoxy squeezes out through the overflow grooves so the epoxy layer reaches its needed thickness.

Once the sealants and epoxies have all fully cured, the module is vacuum tested to verify performance. This vacuum testing consists of closing all open ports and faces and cycling the chamber under vacuum, measuring the leak rate every ten cycles. This leak rate test verifies that the vacuum pumps will be able to overcome the chamber leakage to maintain the required vacuum conditions for sublimation in the drying section. This vacuum test is performed for the cooling chambers as well to test their sealing. While these chambers do not need to maintain vacuum conditions, their ability to seal impacts their efficiency in recirculating cold gas without absorbing excess energy from the environment.

6.3 Machine Table

The machine modules are assembled on top of the machine table. This table structure isolates the system from the floor on which it is built, ensuring that it can be successfully built on floors that do not require strict specifications. The structure also assists in the system assembly process and provides mounting regions for accessory hardware. The details of the table design can be found in Section 2.8. The following sections describe the process of assembling the table for use in this system.

6.3.1 Beam Assembly

Each beam for the machine table consists of two parts: the sliding rails and the structural tubing. A section of a fully assembled table beam is shown in Figure 6.10. The sliding rails are 1.5 in by 3 in double t-slotted tubing in 8 ft lengths. The structural tubes are 12 ft long,

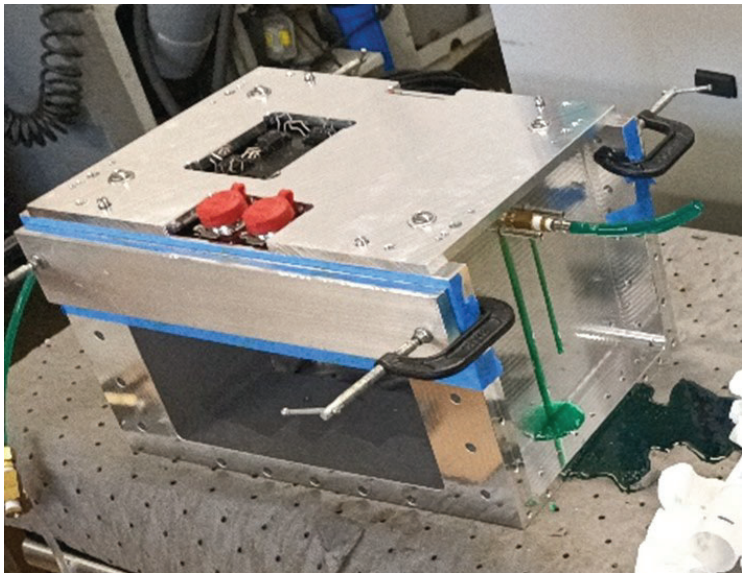

Figure 6.9: After the vacuum sealant injection is completed, the baseplate is added to complete the initial module assembly.

3 in by 6 in inch aluminum rectangular tubing with 0.25 in wall thickness. The 12 ft beams span the length of the structure. The 8 ft long sliding rails cover 2/3 of each beam, and additional rails are cut in half to complete the full 12 ft length. The difference in sizes is due to the availability of standard lengths for each part. The table needed to be 12 ft long, but the sliding rails only came in a maximum 8 ft length, necessitating the use of multiple parts. While this use of multiple parts creates a seam on the rails, the effect of this seam can be mitigated with careful alignment during the assembly process. The relative looseness of the sliding rail fit also allows it to accommodate small misalignment between consecutive rail sections.

The sliding rails are attached to the structural tubing with a structural epoxy. The structural epoxy used in this work is a two-part Loctite Hysol E-60HP Epoxy Adhesive. This epoxy was selected for its strength and its long set time. The beam deflection calculations shown in Chapter 2 show the moment in the overall beam at any given point. This moment can then be used to calculate the stress at the top of the beam. The epoxy must be capable of transmitting this stress between the structural tubing and the sliding rail to ensure that the sliding rail bends to match the profile of the table in deflection. If the stress at this interface exceeds the strength of the epoxy, then the epoxy will fail, causing the sliding rail to separate from the structural tubing. This separation reduces the overall stiffness of the table beams, increasing the total structure deflection. The epoxy must also fully coat the interface between the structural tubing and the sliding rail to create a rigid connection between the two components. Gaps in the epoxy layer create regions where the two components can deflect separately, increasing the stress on the epoxy layer and reducing the stiffness of the overall assembly.

The epoxy's long set time is used to reduce the risk of creating permanent positional errors during the assembly process. The epoxy set time describes how long it takes for the epoxy to begin permanently curing, after which motion can break its bonding and weaken its final

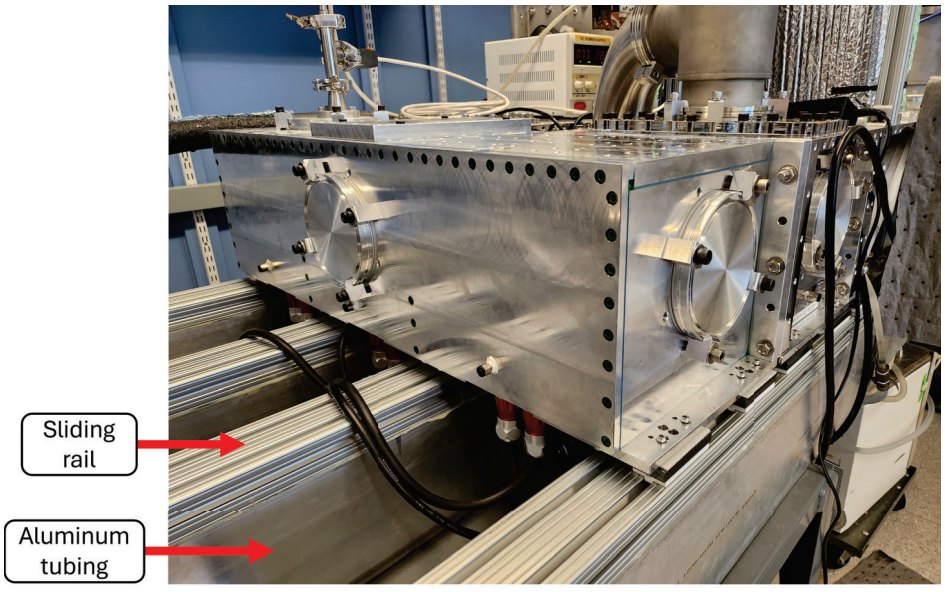

Figure 6.10: The table beams consist of sliding rails epoxied to extruded aluminum tubing.

state. Thus, the set time serves as a "working time" during which the parts being epoxied together can be moved and adjusted before clamping them down in their final positions for the full cure. Given the large bonding area and long length of the beams, it is expected to take substantial time to verify that all parts are in their proper positions before final clamping. Thus, the selected epoxy has a 2-hour set time.

The structural tubing and sliding rails are prepared for the epoxy process to improve the final seal strength. This preparation process improves the bonding between the epoxy and each component, which is particularly helpful given that aluminum can be a difficult material for forming strong eposy bonds. First, the mating surfaces are cleaned using a brush and a wet cloth. This cleaning is used to remove bulk dirt and dust that may have accumulated on these surfaces prior to the assembly process. Next, the mating surfaces are scratched using sandpaper. The sandpaper creates numerous cuts and small grooves in the mating surfaces, increasing their surface area for bonding. The preliminary cleaning helps reduce the amount of dirt and grime that may get pressed into these grooves during sanding. After sanding, the mating surfaces are cleaned again, this time more thoroughly. The sanding process creates numerous particulates which need to be cleared away to ensure strong bonding between the mating surfaces. This secondary cleaning process starts with wiping the surfaces using a wet cloth to remove the residue from sanding. Next, the surfaces are cleaned using ethanol or isopropyl alcohol (IPA) and standard paper towels until these towels no longer pick up residue. Finally, the surfaces are wiped using a higher quality cloth, such as a microfiber cloth, wet with ethanol or IPA, until the cloth picks up minimal dirt during the wiping process. At this point, the surfaces are ready for epoxying.

Epoxy is applied to both mating surfaces before bringing them together. This process ensures the epoxy layer will fully cover both mating surfaces. The epoxy is applied from its cartridge using a dedicated applicator and a mixing nozzle to properly mix the two-part epoxy during the injection process. Using the mixing nozzle on a cartridge replaces the need for mixing the two parts in a separate container, and the cartridge maintains the proper ratio for mixing during this application through the nozzle. Some of the epoxy needs to be purged through the nozzle to ensure the epoxy is fully mixed before it is applied to the mating surfaces, which can be done into a proper waste container. Once the epoxy extruded from the mixing nozzle is a uniform color, it is ready to be applied to the mating surfaces.

On the structural tubing, the epoxy is applied in a zig-zag pattern diagonal to the beam's primary axes. On the sliding rails, the epoxy is applied linearly along their length in pre-existing grooves on the sliding rails. Once the epoxy has been applied on both surfaces, they can be brought together. After the surfaces are connected, the sides of the mating surface are wiped down to remove excess epoxy which squeezes out. This initial wipe down helps reduce the risk of excess epoxy causing the beam to bond to the clamps used to hold the components in place during the full cure time. It is important to use enough epoxy such that some does squeeze out to ensure the mating interface is fully covered with epoxy; if the epoxy does not squeeze out the sides at the interface, then there may be areas which are not bonding the two components together.

The sliding rails and the structural tubing are then placed on clamps to prepare for alignment and clamping. The clamps are then brought together such that there are nominally a couple millimeters between the mating surfaces, ensuring the components can still be moved relative to each other. In this position, the beams are moved such that they are close to their

final positions, including doing the initial sliding rail alignment. The sliding rail's 8 ft and 4 ft segments need to be sufficiently aligned so that the modules can seamlessly slide between these sections without interference. This alignment process consists of moving a plastic slider in the slider groove back and forth across the gap between slider rails to ensure that it can move smoothly. If the slider cannot move smoothly across the gap in both directions, then the rail positions are adjusted so that this motion can occur unhindered. The slider groove used for this test is on the short side of the sliding rail, which corresponds to the long side of the structural tubing. While this groove is not one of the ones the modules will move through, aligning this groove ensures that the sliding grooves which will be used are not offset to the side of each other. Vertical misalignment is handled by the clamps squeezing the sliding rails onto the structural tubing. Once the test slider moves smoothly across the gap between sliding rails, the clamps are tightened to a sliding fit such that the beams can still be moved but they do not change position when released. The slider adjustment is performed again, this time on all four accessible slider grooves. Once the slider moves freely on all grooves, the rails are clamped down to the structural tubing. This clamping starts with the clamps closest to the interface between sliding rail segments, then moves outwards to the ends of the structural tubing. As the clamping progresses, the rails are aligned with the side face of the structural tubing to prevent large relative angles between different rail sections. Once all the clamps are secured, the sides of the assembly are wiped down to remove any excess epoxy that may have squeezed out of the sealing interface during the clamping process. The epoxy is then left to cure overnight. After the epoxy fully cures, the table beam is complete and ready for mounting on the table legs. The table beams were epoxied in pairs, where two beams were done each day over two days to create a total of four table beams for the structure.

6.3.2 Table Legs

The system table leg structure is a welded steel table, as shown in Figure 6.11. This welded steel table provides sufficient strength to support the table beams and the lyophilizer. Each welded steel table is sized to be long enough to support all four of the table beams so that they are on a shared reference surface. However, this surface can deflect under the weight of the table beams, causing a difference between their relative surface heights. This deflection can be reduced by adding a central support to the tables if necessary. The module orientation adjustability discussed in Chapter 2 can also help the modules account for small differences between the surface heights of the sliding rails after the full system table is assembled. These welded steel tables include a lower shelf which is used to store auxiliary components such as power supplies and controllers for the Planar Motors system.

The welded steel table is mounted on top of an aluminum plate and a rubber sheet. The aluminum plate is used to spread the weight load from the lyophilizer over a larger area, decreasing the pressure applied by the machine to the floor below. This load spreading reduces the risk of the welded steel table legs punching through the floor, reducing the strength requirement for the floor on which the lyophilizer is built. The aluminum plate used for this system measures 36 in x 18 in x 0.25 in.

The rubber sheet deforms to match the floor profile, ensuring the welded steel table can fully engage with the floor. The sheet also serves to damp vibrations that would be trans-

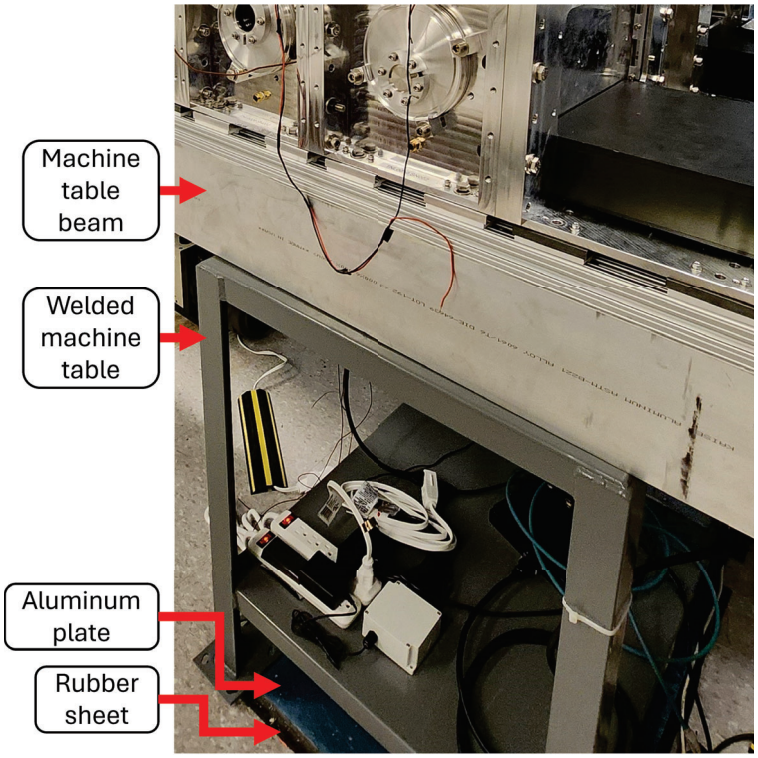

Figure 6.11: The machine table legs are welded steel tables. These tables provide sufficient strength to support the weight of the table structure and the freeze-drying system. The legs are placed on top of an aluminum plate to spread the lyophilizer weight load over a larger area on the floor than just the welded table feet. A rubber sheet is placed under the aluminum plate to absorb some of the floor surface profile and to damp vibrations that would be transmitted from the floor to the lyophilizer.

mitted from the floor to the system. This isolation helps keep the lyophilizer performance consistent regardless of outside disturbances such as people walking or other equipment in the room operating. The rubber sheet used for this system measures 36 in x 18 in x 0.25 in. The rubber sheet includes an adhesive backing that is used to attach it to the aluminum plate during assembly. The rubber sheet also has some areas cut out around its perimeter to create handholds under the aluminum plate. Once the rubber is attached to the aluminum plate and placed on the floor, it creates a suction effect that adheres the sheet to the floor. This adhesion can make it difficult to separate the rubber from the floor, which is necessary for positioning the legs during the assembly process. These cutouts create handholds that provide an ability to lift the aluminum plate and break the suction seal, allowing the plate and rubber to be moved.

The rubber sheet could be replaced with an air caster system. Setting up the system table legs with air caster capabilities would make it simple to move or reposition the lyophilizer, should the need arise. Air casters rely on creating a thin film of air between the table base and the floor. When the air flow is turned off, the system rests solidly in place. However, air caster systems can struggle with uneven flooring, so their use would limit the flooring on which the system can be built.

6.3.3 Table Assembly

The machine table assembly consists of mounting the table beams on top of the table legs. The table legs are spaced according to the calculations shown in Chap 2. The rubber sheets are attached to the aluminum load spreading plates, then these panels are placed at their designated locations on the floor. The welded steel tables are then placed on top of the aluminum plates. The tables are oriented such that they are positioned parallel to each other. The four table beams are then lifted and placed on top of the table legs. The beams are spread out to their nominal positions on top of the table to prepare for alignment. The table beams are self-aligned by sliding a module across the table while the table beams are free to move. As the module slides along the rail, it pushes the beams into a position which minimizes the frictional resistance to the sliding module. The alignment is first done on one pair of table beams corresponding to one side of the lyophilizer tunnel structure. Once these two beams are aligned, the three-stator loading plate is used to set the positions of all four beams on one end of the system. A single module is kept on each end of the aligned side of the table to ensure that those rails do not move out of alignment when the other half of the table rails is adjusted. With these three modules in place, another single module is used to align the table beams corresponding to the second side of the lyophilizer tunnel. Once the module can slide freely on the second table side, a second three-stator loading plate is placed on the opposite side of the table from the first one to ensure that all four table beams are in position on opposite ends of the table. If the table beams need to be moved to allow the three-stator loading plate to rest properly on all four table beams, the beams are moved to match the three-stator loading plate and the self-aligning is performed again on both pairs of table beams corresponding to each side of the freeze dryer tunnel. After this final alignment, the three-stator module can smoothly slide across the full table length. Upon reaching these final positions, the table beams could be attached to the welded steel tables. However, this attachment is not necessary because the module assemblies are stiff enough to maintain beam alignment once they are assembled on the rails.

6.4 System Assembly

6.4.1 Four Stator

Before moving to larger production of each module in this system, the modules are tested on the smallest potential fully operational system. This system would not need to be capable of continuous operation, but it would need to include freezing, drying, and sensing capabilities. The minimum system with these capabilities consists of four connected units: one loading, one cooling, one load-lock, and one drying/sensing chamber, shown in Figure 6.12. Because these modules were produced to validate the system before casting larger module quantities, they are manufactured from assembled machined plates.

The testing on this four-stator unit showed successful cooling and freezing could be achieved with the proposed module geometry. However, the cooling module used did not include an actuated door, instead requiring manual intervention to move vials into and out of the cooling chamber. Additionally, the drying chamber included a lowered ceiling for an old

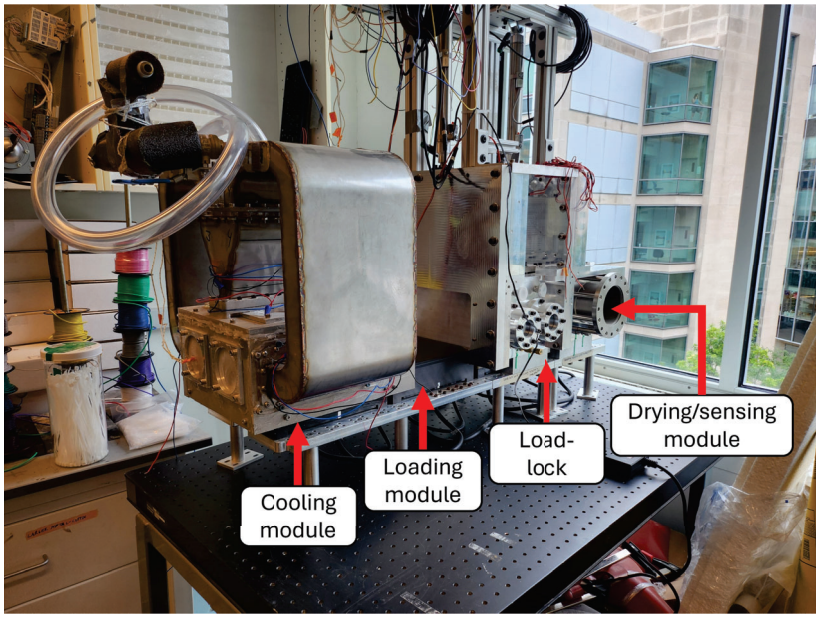

Figure 6.12: The first integration test for the freeze-drying system was done on a minimal unit system consisting of four modules: cooling, loading, load-lock, and drying/sensing. This system validated the general functionality of each module and their ability to interface with each other, but some geometric limitations caused it to see little lyophilization testing before moving to the larger system with updated chamber geometries. This four-module unit represents a minimum scale system with geometry matching that of the larger lyophilizer system. At this scale, the system can be used as a research and development tool whose results would translate directly to the larger scale system.

weight sensing design that preceded the system described in Chapter 4, preventing it from accommodating the final weight system. Thus, manufacturing moved forward on producing the castings for larger unit quantities rather than adapting this four-stator system for continued experiments. While the four-stator system was not used further in this work, it demonstrates a research-scale system which can be used for further evaluating the lyophilization of new formulations. This system is much smaller than the continuous lyophilizer, but it includes the same module geometry as the larger continuous system. Thus, lyophilization processes developed on the four-stator research unit can translate directly to the production system. This translatability resolves one of the existing problems in current lyophilization systems described in Chapter 1, where the processes developed in research must be re-determined on the production equipment due to the differences in their geometries.

6.4.2 Minimum Module Buildup

The full lyophilizer was tested in smaller subassemblies to verify integration performance. This subassembly testing makes it easier to debug any errors that arise during the integration process. The ability to maintain the required vacuum conditions is a growing risk as the number of units in the assembly increases. With each additional module, the number of sealing surfaces increases, which is expected to increase the overall system leak rate. Ad-

ditionally, the added modules increase the total volume that needs to be maintained under vacuum. Furthermore, the added tunnel length increases the head losses for the pump, in turn making it more difficult for the pump to successfully maintain the pressure conditions. These effects of these volume and sealing path length considerations on the vacuum system are modeled in [40]. During system assembly, it is easier to identify whether changes in the system pressure are due to expected losses due to the addition of more units to the system or leaks due to an error in the assembly. If the chamber pressure when pumped down after each subassembly addition is significantly higher than the expected pressure, that rise indicates that the assembly has an issue which is contributing to the increased leak rate.

Additionally, because the modules were delivered intermittently, the smaller subassemblies allowed immediate testing of parts as they were delivered. Accordingly, the subassembly build-up plan informed the priority for which modules needed to be machined first. The first subassembly level includes just two modules, a load-lock and its corresponding door module that mounts behind it. This subassembly is required to operate the load-lock chamber, as the load-lock module only includes one door so it cannot close both sides of its tunnel. This subassembly is done for each load-lock before it is added to the larger assembly. During this subassembly, the doors are shimmed into place to ensure that they move into the correct position for maintaining a vacuum seal when they close. The vacuum pump is connected through the side ISO100 port on the load-lock module with a pair of vacuum valves to allow it to alternate between vacuum and ambient pressure during operation. Further details about the load-lock module can be found in [25].

The second subassembly level adds a drying unit to the assembly. This assembly first enables testing of the top window surface on the drying unit for successful vacuum sealing. After the window performance is confirmed, the window can be replaced with an ISO100 window adapter piece which connects a condenser vacuum line to the drying unit. This addition allows the condenser line to be tested in situ with a drying unit. At this subassembly level, a pre-frozen test sample can be added to the system to validate that the vacuum successfully creates an environment that can sublimate water from a vial; however, without any sensing modality, this test does not provide significant information. Thus, this subassembly is still primarily useful for validating the assembly process and ensuring that units are successfully connected without creating additional leak paths.

The third subassembly level adds a second drying unit to the assembly. This second drying unit ensures that the system can include both a sensing window and a condenser connection. At this subassembly level, drying experiments can begin with pre-frozen vials. The drying process can be measured based on weight sensing units in the drying unit with the window. This subassembly enables testing of the drying measurement process, including verifying data collection stream and weight sensor capabilities in-situ. This chamber can also include a thermal camera to validate the thermal measurement system.

The fourth subassembly adds a cooling module which interfaces with the front side of the load-lock chamber. This five-unit subassembly is capable of the full lyophilization process, albeit in small batches and not continuously. This full process testing enables validation of each process chamber and evaluation of their interactions with each other.

The fifth subassembly adds an entry door to the cooling section. This subassembly simplifies the process for adding and removing pucks from the system, allowing fully automated cycling of vials through the system. However, because the pucks must still exit through the

cooling section, interrupting the cooling process, this subassembly does not operate continuously. The cooling section can keep its cooling state continuously while pucks exit, and because the drying process takes so much longer than the cooling process, this subassembly could be operated in a pseudo-continuous manner. However, it is more efficient to turn the cooling section off between sending in pucks for cooling to conserve resources.

The sixth subassembly adds the turn module to the system, along with a corresponding exit load-lock and door unit pair. Adding these modules separates the entry and exit of the system, enabling continuous operation. Because this subassembly has only one cooling section, the installed cooling section must perform both conditioning and freezing. Switching between these two states means that this machine section does not maintain a constant condition, as is planned for the final system.

The seventh subassembly adds a second cooling section to the system. This section must include both a door and a door support unit to fully integrate the section into the existing sixth subassembly level. Adding the second cooling section separates the conditioning and freezing processes. This addition can allow each section to maintain constant conditions during operation, as expected for the continuously operating system.

The remaining assembly consists of adding further drying modules to the system. These additional drying modules enable process section balancing so that the lyophilizer can be operated at higher capacity and the cooling section can be utilized more efficiently. The additional drying units also enable the testing of independent sensing stations to ensure that data can be collected from multiple regions within the machine simultaneously. Increasing the number of vials that can be held within the lyophilizer correspondingly requires an increase in the number of condensers used to ensure that the system can maintain the required vacuum conditions while sublimating larger volumes of water.

6.4.3 Auxiliary Components

6.4.3.1 FLIR

The vial temperature during drying is measured by a thermal camera. The thermal camera used in this system is a FLIR A35. This camera is mounted to the machine chamber wall through which it views the vials. An IR-transparent window is required for measuring temperatures of vials inside the vacuum with a thermal camera mounted outside of the vacuum. Because the camera's field of view is limited, it is not useful to try to replace an entire chamber wall with an IR-transparent material. IR-transparent material can also be costly, so a smaller window is significantly more economical. Given the camera lens size, a 50 mm diameter window is more than sufficient to provide an unrestricted view relative to the camera itself, if the camera is mounted close to the window. This window is mounted in a custom machined ISO100 blank, as described in Section 5.2.3. This lyophilizer uses zinc sulfide for its IR-transparent windows.

6.4.3.2 Weight Sensing Camera

The weight sensing camera used in this system is a Canon T7i. This camera is mounted to a simple t-slotted framing mount which bolts directly to the sensing chamber. The T-

slotted framing enables simple repositioning of the camera to evaluate the optimal position for weight sensing measurements.

Adding the camera to the system revealed several difficulties in implementing the weight sensing system described in Chapter 4. The final weight sensing camera setup, shown in Figure 6.13, includes the camera, an ambient light shield, and an LED strip used to illuminate the chamber below the camera. These additions are needed to ensure that the AprilTags on the ends of the sensing wires are sufficiently visible so the camera can successfully record their positions.

The AprilTag position measurement relies on the ability to identify the edge of the AprilTag in an image. This edge visibility depends on both the size of the border around the AprilTag and the color contrast between the AprilTag and its environment in the image. The AprilTags require a sufficiently wide white border outside of the black square border that defines the tag for the camera to consistently identify the AprilTag. Additionally, there must be sufficient illumination within the chamber for the camera to register the color

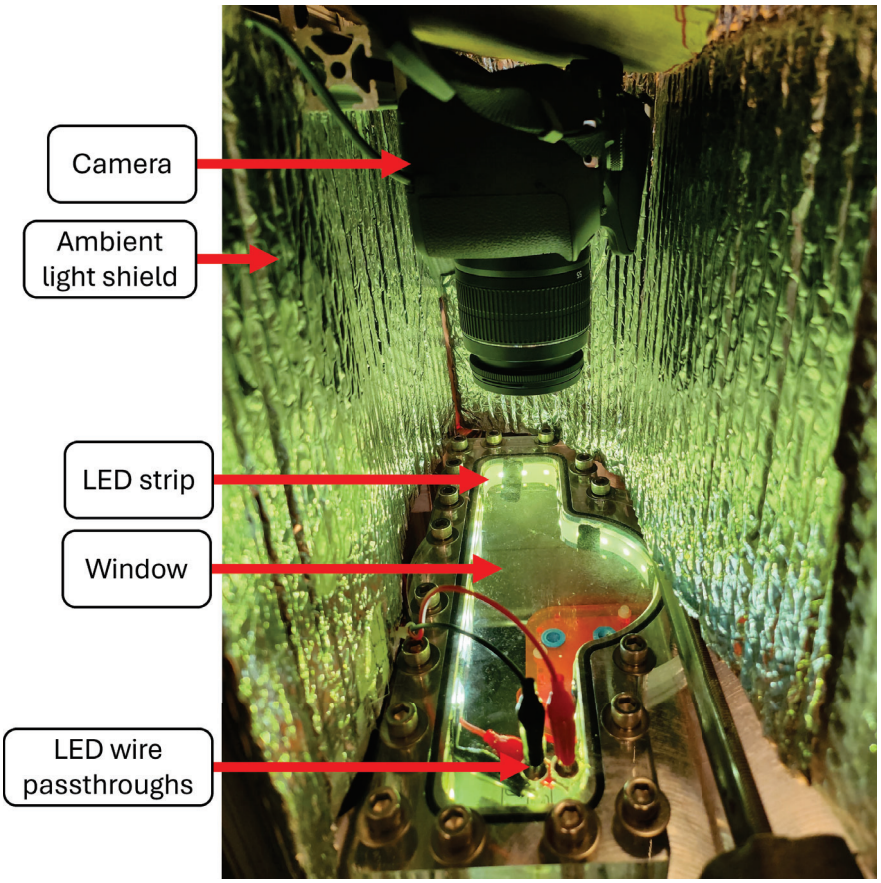

Figure 6.13: The camera used for weight sensing is mounted above the window on the drying module through which it images the weight sensing system. The top window cutout includes an LED strip inside the chamber which illuminates the April tags to help the camera locate their positions. A custom vacuum passthrough made from unused rivets provides power for these LEDs. The entire setup is surrounded by shielding to limit the reflectivity of the window itself.

contrast between the AprilTag and its environment. Further, the tags themselves must not be attached to the plastic nubs which hold them on the wire sensing arms with too much adhesive, as it can soak through the paper tag and make it reflective, reducing the contrast between the dark and light parts of the tag. Though the paper on which the AprilTag is printed has a base level of reflectivity cannot be completely removed, it can be mitigated by reducing the brightness of the light inside the chamber or by increasing the camera shutter speed. In practice, increasing the camera shutter speed while maintaining bright illumination was found to be more reliable than reducing brightness with a lower shutter speed.

The chamber is illuminated by an LED strip adhered to the inner wall of the machine chamber's window cutout. Two holes are drilled through the window, and a pair of stock rivets are inserted into these holes and sealed with silicone to create an electrical passthrough. The rivets include a straight section of material which serves as connection points on the inside and outside of the chamber and a flange which provides the sealing area for the passthrough hole. The LED strip has alligator clips soldered to its positive and negative leads that connect to the rivets inside the chamber. The LED strip is mounted to the chamber, rather than the window, to have the light directed orthogonally to the tags to further reduce the potential reflectivity of the tags. In the future, the lights could be integrated directly into the window to make it easier to move sensing stations between modules. The window itself also has some level of reflectivity. This window reflectivity led to ambient lighting sometimes obscuring the AprilTags in the camera videos. Additionally, the environmental lighting would change throughout the day, particularly around sunrise and sunset. These environmental lighting concerns are addressed by adding an ambient light shield around the sensing station so that light sources other than the internal LED strip illumination are blocked out.

6.4.3.3 PLC

The controls for this lyophilizer are managed by a PLC and MATLAB. The PLC interfaces directly with most of the hardware, such as the load-lock and vacuum system valves. MATLAB provides control instructions to the PLC based on the system state. The design of these control elements is outside of the scope of this thesis.

6.4.4 Scale Up

The production rate of the lyophilizer scales with the size of the machine, as described in Chapter 2. This scalability is verified on the manufactured equipment in two ways. First, the system can be operated with fewer pucks than the full capacity. This reduced capacity operation can simulate the limitations of how many pucks can fit in a smaller system. One downside of using this test method is that the conditions experienced by the pucks are not the same as they would be in a smaller system, as the vial density in the tunnel is decreased. This lower density is expected to lead to increased nonuniformity across pucks during drying, as the vials on the ends of each tray are not obscured from wall radiation by the vials on a neighboring puck. The sublimation rate could also be faster, as less vials in the system sublimating simultaneously would lead to lower vapor pressures in the vacuum

system, resulting in higher sublimation driving forces. These effects are expected to be small enough to prevent them from having significant impacts on these scaling test results.

The second scaling verification method involves adding modules to the system. The initial minimum module build discussed in Section 6.4.2 does not span the full length of the 12 ft table, as it only requires 9 ft at minimum to include both a conditioning and freezing section on one side of the table, as shown in Figure 6.2. This system does not include balanced freezing and drying times, as described in Section 6.1, where the drying section length limits the overall system production rate. Thus, adding drying modules to the assembly is expected to increase the lyophilizer's production rate without requiring the addition of more freezing modules. These drying modules are added to the system in pairs by removing the turn module and shifting it by 250 mm, making space for the new drying modules.

6.5 Vial Filling Integration

In a production setting, the continuous lyophilizer will interact with other pharmaceutical equipment rather than a manual operator. One key function of this additional equipment is moving vials from a production line to the lyophilizer. To validate this interface, a DispenseWorks Ring-Dex vial filling and capping system automatically fills vials, places them on the lyophilizer for freeze-drying, and removes them from the lyophilizer after they are dried. The DispenseWorks machine uses a robotic arm to transfer vials between the itself and the lyophilizer.

The robotic arm needs to reach past the sensing arms on the weight sensing system to be able to pick up or place vials on the pucks. The robotic arm head has a diameter of 40 mm that would contact the sensing arms when it tries to move the vials to and from the pucks. Thus, the sensing arms need to be moved out of the way of the robotic arm while it is interacting with the puck. This motion is achieved through two elements: the elephant trunk and the robotic arm cone.

The sensing arms move outwards from the vial when the spring wire deflecting arm is pulled downwards by the vial weight. These arms can be moved even further apart by pulling the basket down, as shown in Figure 6.14a. This pulling is done through an element called the elephant trunk, shown in Figure 6.14b. During puck loading and unloading, the puck first moves in front of the elephant trunk. The puck then moves towards the elephant trunk, causing it to engage with the vial basket. As the puck continues to move, the elephant trunk pulls the vial basket down to engage it with the corresponding puck alignment cone. Pulling the basket down causes the sensing wires to deflect outwards from the vial, moving them away from where the robotic arm travels.

The elephant trunk is an extended arm that pulls the basket down in preparation for vial loading and unloading. The trunk is mounted on the lyophilizer table and extends over the surface of the primary loading stator, in a position aligned with the vial filling machine's robotic arm track. The front edge of the trunk is positioned such that its bottom surface is above the neutral hanging height of the vial basket once the product in the vial is fully lyophilized. The surface of the root of the trunk is at a height such that when the basket is touching this surface, it is also pushed down into contact with the puck plate cone. The trunk profile follows an S-shaped curve to transition between these two surface heights. The

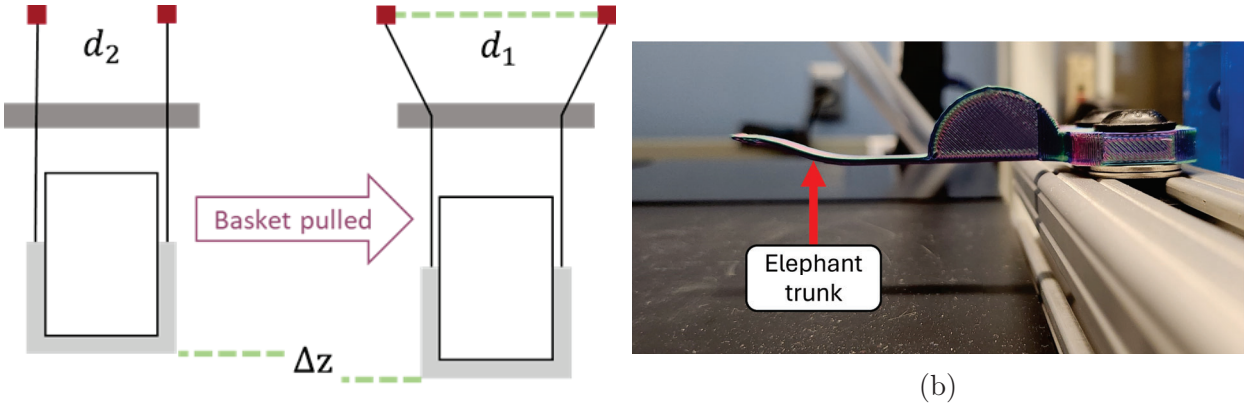

(a) (b)

Figure 6.14: (a) As the vial basket is pulled down by the elephant trunk (b), the sensing arms move outwards to create space for the robotic arm to grab the vials. This spacing is not sufficient to eliminate interference between the sensing arms and the robotic arm, so an additional element is added to the robotic arm to further push the sensing arms safely aside.

trunk thickness is set to 2 mm to ensure that it fits in the gap between the basket base and where the bottom of the vial sits when it is in the basket. The trunk width is set to 5 mm so that it can fit between the vial positioning posts where these posts are connected to the basket base. The trunk includes a semicircular area where it overhangs the stator to increase the overall stiffness of the elephant trunk system. Behind this semicircular area, the elephant trunk includes mounting holes to connect to a piece of 1.5 in t-slotted framing.

While the elephant trunk pulling motion helps to move the sensing arms away from the robotic arm, the motion is not sufficient to prevent interference between these two system elements. Thus, the robotic arm cone, shown in Figure 6.15, is added to the robotic arm. The robotic arm cone provides a smooth angled surface to redirect the nubs on the ends of the sensing arms away from the robotic arm gripper. This motion is necessary to ensure the sensing arms in the vial weight sensing system are pushed to the side instead of being pushed down. If the sensing arms are pushed downwards, they can be pushed out of the basket wire, dislodging the vial basket from the sensing arms. If this detachment occurs, then the sensing wires will move away from the basket wire, preventing it from reconnecting after the robotic arm leaves. If instead the sensing arms are pushed to the side, the hook at the top of the basket wire will remain engaged with the sensing wire's basket wire crook, ensuring that the vial basket remains connected to the sensing wires after the robotic arm retracts.

The robotic arm cone must function when the robotic arm gripper is in both its open and closed positions. The robotic arm gripper approaches the puck in its closed position when bringing a fresh vial to the tray, as it is holding the vial which it brings from the vial filling machine. The robotic arm approaches with the gripper in the open position when it is coming to the tray to pick up a vial which has completed the lyophilization process. The robotic arm cone system can be mounted directly on the gripper's moving arms, mounted to the robotic arm body, or mounted to the lyophilizer. Mounting to the lyophilizer reduces the potential error between the vial positioning and the cone because

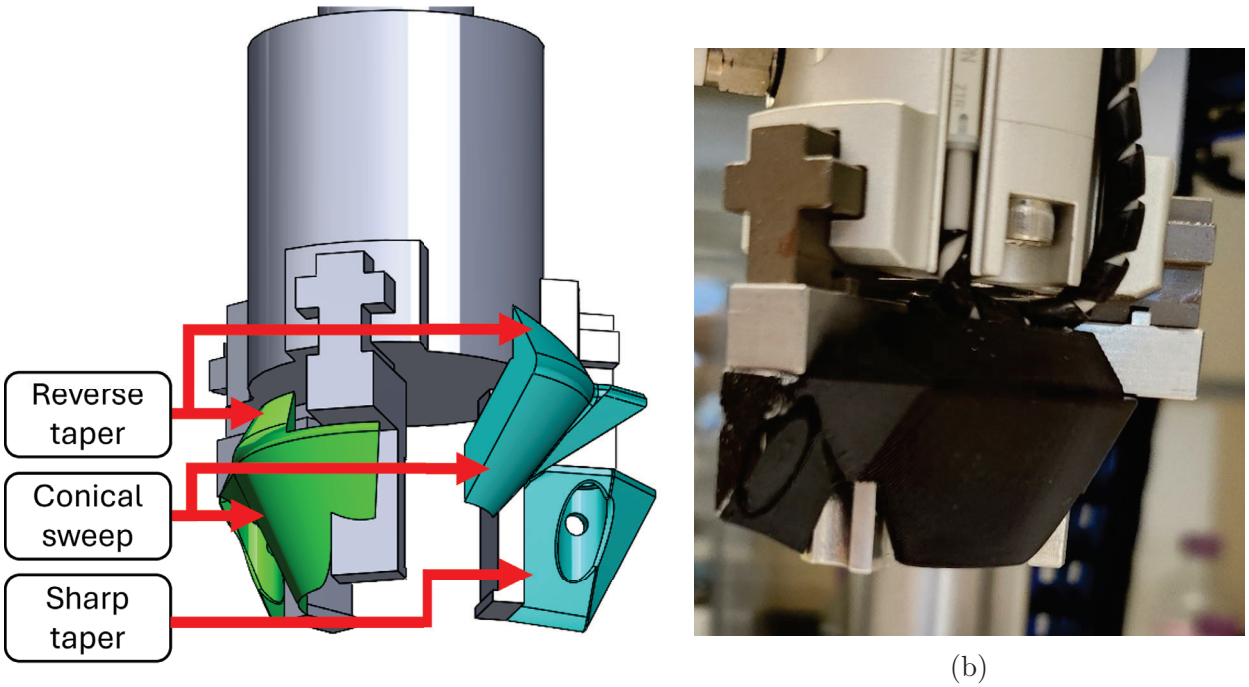

(a) (b)

Figure 6.15: The robotic arm cone consists of two parts mounted directly to the gripper fingers. These parts are shaped differently to account for the asymmetry in the interaction between the gripper fingers and the sensing wires on the puck. The angled surfaces gently move the sensing arms to the side without pushing them down so the spring wires do not dislodge from the vial basket. The reverse taper on the top part of the robotic arm cone components ensures these parts do not pull up on the underside of the AprilTag nubs on the sensing arms when the robotic arm pulls up and away from the puck.

they are both referenced to the stator. This mounting also keeps the cone independent of the vial filling system used. However, this mounting strategy requires features that the puck can avoid when entering and exiting the loading position, restricting its geometry. Mounting the cone features to the robotic arm ensures that the cone will not change positions between the gripper's open and closed positions, so if it works in one state then it necessarily works in the other state. However, this mounting requires the cone to be wider than the entire robotic arm, which would create too much deflection in the sensing wires. Mounting the cone directly to the gripper arms requires them to explicitly accommodate the gripper's open and closed positions, as well as avoid interference with the motion between these states. However, this mounting enables the smallest form factor for the robotic arm cone. Given the restricted geometry in the system, the mounting on the gripper arms is selected to minimize the overall robotic arm cone size.

The DispenseWorks robotic arm gripper includes three arms. One of these arms is perpendicular to the motion of the wires during loading and unloading, so it does not interfere with the sensing arms. The other two arms are symmetrically at an angle relative to the sensing arm motion. However, the spring wire sensing arms themselves are axisymmetric rather than planar symmetric. Thus, the requirements for the cone on each side of the

robotic arm are different. On one side of the gripper, the sensing arm is typically positioned directly below one of the gripper's fingers. On the other side of the gripper, the sensing arm is typically positioned between two of the gripper's fingers. In the first case, the cone specifically needs to shift the sensing arm from below the gripper's finger to the area between the fingers, then further to the side outside of the robotic arm diameter. In the second case, the cone simply needs to move the sensing arm from between the gripper fingers to outside of the robotic arm diameter. These considerations led to the robotic arm cone geometry shown in Figure 6.15.

The robotic arm cone geometry includes a sharp taper that is positioned within the gripper's finger area and a conical sweep that fills the volume between neighboring gripper fingers. The sharp taper shifts the sensing arm into the region between neighboring gripper fingers, where the conical sweep then guides the sensing arm out of the robotic arm's path. The robotic arm cone cannot infringe on the vial profile because that would block the robotic arm from picking up vials. The robotic arm cone also cannot extend too far outside of the robotic arm profile because it would then risk pushing down sensing arms on consecutive weight sensing systems. These limitations set the inner and outer diameters of the tapered elements.

The top outer edges of the robotic arm cone include a reverse taper. This reverse taper prevents the AprilTag nubs from catching on the robotic arm cone when the robotic arm retracts. The AprilTag nubs also include a conical body to help them engage with and be pushed aside by this reverse taper.

6.6 Future Work

The machine can be further extended with additional freezing and drying modules to demonstrate the production rate scaling with system length. Further system length extension requires lengthening the table structure which supports the system. This extension can be done by building another table and leveling its height to ensure that modules can seamlessly slide between the two table constructions. Newly building a structure for larger capacity could involve building a larger table from the beginning based on the table design described in Section 2.8.

The lyophilizer can also be built as a pair of machines that each move in a straight line, removing the need for the enclosed turn module. Instead, each end of the machine could be open and used for loading vials, and the lyophilizers would just operate in opposite directions. The machine pairing is needed so that the pucks travel back to the start of one machine when they exit the second machine. This layout could be less convenient for overall factory design, as the feed of fresh and lyophilized vials would interact with both sides of the machine, but it would simplify the components needed for the lyophilizer.

While the module casting is cost effective for producing many modules, the components had more geometric variation and warping during machining than initially expected. Future modules can include extra material in their internal tunnel geometry that can also be machined away. This additional material and machining would ensure that the tunnel is more accurately aligned with the module faces, which can reduce the difficulty found during manufacturing in trying to meet the required tolerances from cast surfaces that have relatively

large variability.

The lyophilizer in this work was integrated with an example vial filling machine from DispenseWorks, but other implementations might require different equipment to interface with this machine. Because the robotic arm cone described in Section 6.5 is specific to the robotic arm used in that DispenseWorks machine, integration with other hardware would potentially require additional steps for new equipment. However, the geometric insights provided through the development of this loading and unloading system can facilitate the integration of this system with other equipment.

Chapter 7

Conclusion

7.1 Conclusion

Lyophilization is a key step in the pharmaceutical production process where products are stabilized for storage and transportation. Commercial pharmaceutical production lines are increasingly moving towards continuous topologies, creating a need for continuous lyophilization. This work proposes a continuous, unit-dose lyophilizer designed to meet this need. This lyophilizer includes a novel thermal quench method for controlled nucleation during freezing and monitoring systems for individual vial temperature and residual moisture, allowing controlled product release.

The continuous, unit-dose lyophilizer built for this work was used to lyophilize solutions of sucrose, sucrose-mannitol, catalase, and lactate dehydrogenase (LDH). These products were each successfully lyophilized to a residual moisture content of less than 2 wt%. The activity of enzyme solutions was measured for the dried cakes, and no difference in activity was found before and after the lyophilization. These results demonstrate the ability of this lyophilizer to successfully freeze-dry products in a continuous system.

The primary contributions of this research are:

1. Design and experimental validation of a continuous, unit-dose lyophilizer for pharmaceutical products.

2. Scalable machine design tools for the deterministic design and analysis of this continuous lyophilization topology.

3. Design and experimental validation of a convective gas cooling chamber for the freezing process of lyophilization.

4. Implementation of a thermal quench method for controlled nucleation during freezing.

5. Design and implementation of a weight sensor used to measure the residual moisture of individual vials in-situ during lyophilization.

6. Scalable machine design tools for the deterministic design and analysis of this weight sensor for implementation on a variety of vial sizes and product formulations.

7.2 Future Work

This research into continuous lyophilization aims to bring the benefits of continuous manufacturing to the lyophilization process at the end of pharmaceutical production lines. This machine also provides opportunities for further research into the lyophilization process.

Future research can:

1. Investigate the effects of varying freezing profiles on the sublimation rate of vials.
2. Investigate the effects of different drying methods (such as using microwaves for power delivery) on the sublimation rate of vials.
3. Improve the weight sensor resolution to identify the end of secondary drying for each vial so that it can serve as a certification tool for product residual moisture. One such method for improvement is an optical lever system used to increase the effective sensing arm length, correspondingly increasing the sensor's deflection amplification.
4. Add thermal monitoring to the cooling chambers to directly measure vial temperatures in-situ during conditioning and freezing.
5. Compare cooling and sublimation rates to computational models used to predict lyophilization parameters for new formulations.
6. Reduce the vial tray size down to individual vial carriers to provide maximum flexibility in moving vials through the lyophilizer when they are finished with each process step.
7. Lyophilize a wider range of formulations on the continuous lyophilizer to demonstrate operational capabilities and determine lyophilization parameters for new formulations.
8. Determine lyophilization parameters for a new formulation on a minimum module research-scale machine and compare these parameters to those required on a larger production-scale machine to validate process consistency during scaling.
9. Compare the thermal quench nucleation method to other controlled nucleation methods such as vacuum induced surface freezing (VISF).
10. Investigate alternative methods for maintaining the cooling chamber conditions, such as using a refrigeration system for the conditioning chambers.
11. Investigate alternative injectable sealants from Thixoflex Self-leveling Green to improve injectability and vacuum sealing performance with chamber thermal cycling.
12. Add more weight sensing stations to the drying section, then down-sample from the data collected to determine the minimum required number of sensing stations needed to monitor the drying process.
13. Integrate cooling chamber exhaust lines into the feeds for upstream cooling chambers to improve nitrogen utilization.

14. Add heat shield fins to the drying chambers to maintain radiative power delivery uniformity while increasing the number of vial rows in each chamber.

15. Evaluate the effect of gaps in the cooling gas flow due to machine chamber interfaces or gaps between fans on the vial cooling rate.

16. Investigate the effects of increasing vial density on cooling and sublimation uniformity for the vials in the lyophilizer.

The research presented shows the development of the continuous lyophilizer to a production prototype level. Developing this continuous lyophilizer from production prototype to a factory ready machine includes the following next steps:

1. Add regeneration capabilities directly to the condensers so that they do not need to be manually removed and replaced when they reach their ice capacity.

2. Automate the assembly of the weight sensors.

3. Add QR codes to the vial caps so that the individual vials can be directly linked to their freezing and drying profiles for traceability.

4. Increase capping module capacity to handle higher product throughput.

5. Design the cooling chamber recycle ductwork for improved manufacturability, since maintaining alignment over the current geometry proved challenging during the welding process.

6. Add a purge chamber to the beginning of the lyophilizer to ensure a clean, dry environment.

7. Increase loading capacity from individual vials to multiple vials to improve throughput.

8. Add a stainless steel sheet to cover the magnetic levitation stators and fully enclose the lyophilization tunnel.

9. Test steam sterilization for cleaning the lyophilizer.

10. Investigate alternative or custom magnetic levitation systems built for this continuous lyophilizer to improve vacuum sealing and manufacturability.

11. Improve the load-lock design so that it can handle pressure differences in both directions rather than only one direction.

12. Add outer insulation or shielding to the machine chambers to prevent condensation buildup on the equipment, particularly in the cooling section.

13. Add cooling channels to the machine chambers to provide direct temperature control of the chambers.

14. Demonstrate continuous operation for multiple days at full capacity.

15. Choose a starting formulation and size a production prototype machine to meet that formulation's lyophilization process parameters.

16. Assemble the continuous lyophilizer in paired straight line machines to remove the need for the custom turn module in the drying section.

17. Add structural framing around the lyophilizer to accommodate vertically stacked lyophilizers to improve factory area utilization.

18. Integrate modules into larger combined chambers for permanent installations that are sized for specific formulations.

19. Build the machine chambers and corresponding process chamber components from stainless steel to match other pharmaceutical equipment materials.

References

[1] J. Snowman, "Lyophilization," in *Handbook of Downstream Processing*, Springer, 1997, pp. 203–234.

[2] D. Bisht and Z. Iqbal, "LYOPHILIZATION - PROCESS AND OPTIMIZATION FOR PHARMACEUTICALS," *International Journal of Drug Regulatory Affairs*, vol. 3, no. 1, 2018, ISSN: 2321-7162. DOI: 10.22270/ijdra.v3i1.156.

[3] *SC-BV25 Vaccine Transport Fridge 25L*.

[4] *-40C Upright Vaccine Freezer 110V | Ai DeepFreeze | 16 Cu Ft*.

[5] E. Shalaev, A. Soper, J. A. Zeitler, S. Ohtake, C. J. Roberts, M. J. Pikal, K. Wu, and E. Boldyreva, *Freezing of Aqueous Solutions and Chemical Stability of Amorphous Pharmaceuticals: Water Clusters Hypothesis*, 2019. DOI: 10.1016/j.xphs.2018.07.018.

[6] P. Hiwale, A. Amin, L. Kumar, and A. K. Bansal, *Variables affecting reconstitution time of dry powder for injection*, 2008.

[7] G. Reynolds, "The market need for reconstitution systems," *BioProcess Int*, vol. 4, no. 10, 2006.

[8] *Reconstitution of Solutions*.

[9] S. L. Nail and L. A. Gatlin, "Freeze-drying: Principles and practice," in *Parenteral Medications, Fourth Edition*, 2019. DOI: 10.1201/9780429201400-48.

[10] A. Sharma, D. Khamar, S. Cullen, A. Hayden, and H. Hughes, *Innovative Drying Technologies for Biopharmaceuticals*, 2021. DOI: 10.1016/j.ijpharm.2021.121115.

[11] L. Rey and J. C. May, *Freeze-drying/lyophilization of pharmaceutical and biological products: Third edition*. 2016.

[12] *The benefits of lyophilization in assay kit development*.

[13] D. M. Fetterolf, "Lyophilization," *Journal of GXP Compliance*, vol. 14, no. 4, pp. 52–60, 2010.

[14] S. U. Sane and C. C. Hsu, "Considerations for Successful Lyophilization Process Scale-Up, Technology Transfer, and Routine Production," in *Formulation and Process Development Strategies for Manufacturing Biopharmaceuticals*, 2010. DOI: 10.1002/9780470595886.ch31.

[15] *SP Scientific Advantage Plus EL-85 Console Lyophilizer with Stoppering*.

[16] *36 Sq Ft Hull 36fxs130-Ss30c Lyophilizer*.

[17] H. Kawasaki, T. Shimanouchi, and Y. Kimura, "Recent Development of Optimization of Lyophilization Process," *Journal of Chemistry*, vol. 2019, 2019, ISSN: 20909071. DOI: 10.1155/2019/9502856.

[18] A. C. Cartwright and B. R. Matthews, *International Pharmaceutical Product Registration*. 2016. DOI: 10.3109/9781420081831.

[19] R. Pisano, "Continuous manufacturing of lyophilized products: Why and how to make it happen," *American Pharmaceutical Review*, vol. 23, no. 3, 2020, ISSN: 10998012.

[20] *Loading and unloading systems*.

[21] trimleaf, *Harvest Right HRC100 Commercial Freeze Dryer*, 2024.

[22] P. J. Van Backstal, T. De Beer, and J. Corver, "A continuous and controlled pharmaceutical freeze-drying technology for unit doses," *European Pharmaceutical Review*, vol. 22, no. 6, 2017, ISSN: 13608606.

[23] United States Pharmacopeia, "<788> Particulate Matter in Injections," *Usp*, vol. 34, 2011, ISSN: 03634655.

[24] *PLANAR MOTOR INC. - Home*. URL: https://www.planarmotor.com/index.html.

[25] R. Flores, "Material Handling for Continuous Lyophilization Process," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Sep. 2022.

[26] Parker Hannifin Corporation, *Parker O-Ring handbook ORD 5700*, 2007.

[27] Wikipedia, *Airy points*, 2023.

[28] J. Verdirame, *Airy Points, Bessel Points, Minimum Gravity Sag, and Vibration Nodal Points of Uniform Beams*, Feb. 2016.

[29] R. Pisano, A. Arsiccio, K. Nakagawa, and A. A. Barresi, "Tuning, measurement and prediction of the impact of freezing on product morphology: A step toward improved design of freeze-drying cycles," *Drying Technology*, vol. 37, no. 5, 2019, ISSN: 15322300. DOI: 10.1080/07373937.2018.1528451.

[30] J. A. Searles, J. F. Carpenter, and T. W. Randolph, "Annealing to optimize the primary drying rate, reduce freezing-induced drying rate heterogeneity, and determine T'g pharmaceutical lyophilization," *Journal of Pharmaceutical Sciences*, vol. 90, no. 7, 2001, ISSN: 00223549. DOI: 10.1002/jps.1040.

[31] G. Assegehegn, E. Brito-de la Fuente, J. M. Franco, and C. Gallegos, *The Importance of Understanding the Freezing Step and Its Impact on Freeze-Drying Process Performance*, 2019. DOI: 10.1016/j.xphs.2018.11.039.

[32] J. C. Kasper and W. Friess, *The freezing step in lyophilization: Physico-chemical fundamentals, freezing methods and consequences on process performance and quality attributes of biopharmaceuticals*, 2011. DOI: 10.1016/j.ejpb.2011.03.010.

[33] L. De Meyer, P. J. Van Bockstal, J. Corver, C. Vervaet, J. P. Remon, and T. De Beer, "Evaluation of spin freezing versus conventional freezing as part of a continuous pharmaceutical freeze-drying concept for unit doses," *International Journal of Pharmaceutics*, vol. 496, no. 1, 2015, ISSN: 18733476. DOI: 10.1016/j.ijpharm.2015.05.025.

[34] J. Lammens, N. M. Goudarzi, L. Leys, G. Nuytten, P. J. Van Bockstal, C. Vervaet, M. N. Boone, and T. De Beer, "Spin freezing and its impact on pore size, tortuosity and solid state," *Pharmaceutics*, vol. 13, no. 12, 2021, ISSN: 19994923. DOI: 10.3390/pharmaceutics13122126.

[35] Y. Zhou, T. H. Gasteyer III, N. J. Grinter, A. T. Cheng, Y.-C. S. Ho, and R. R. Sever, *Method and system for nucleation control in a controlled rate freezer (CRF)*, 2014.

[36] R. Bursac, R. Sever, and B. Hunek, "A Practical Method for Resolving the Nucleation Problem in Lyophilization - BioProcess InternationalBioProcess International," *BioProcess International*, Oct. 2009. URL: https://bioprocessintl.com/manufacturing/formulation/a-practical-method-for-resolving-the-nucleation-problem-in-lyophilization-185124/?pageNum=2.

[37] K. E. Beech, J. G. Biddlecombe, C. F. Van Der Walle, L. A. Stevens, S. P. Rigby, J. C. Burley, and S. Allen, "Insights into the influence of the cooling profile on the reconstitution times of amorphous lyophilized protein formulations," *European Journal of Pharmaceutics and Biopharmaceutics*, vol. 96, 2015, ISSN: 18733441. DOI: 10.1016/j.ejpb.2015.07.029.

[38] J. A. Searles, J. F. Carpenter, and T. W. Randolph, "The ice nucleation temperature determines the primary drying rate of lyophilization for samples frozen on a temperature-controlled shelf," *Journal of Pharmaceutical Sciences*, vol. 90, no. 7, pp. 860–871, Jul. 2001, ISSN: 00223549. DOI: 10.1002/jps.1039.

[39] J. Luoma, E. Ingham, C. L. Martinez, and A. Allmendinger, "Comparison of techniques to control ice nucleation during lyophilization," *Processes*, vol. 8, no. 11, 2020, ISSN: 22279717. DOI: 10.3390/pr8111439.

[40] L. Stratta, "Continuous Freeze-Drying for Pharmaceutical Applications: Design, Prototyping and Process Development," Ph.D. dissertation, Politecnico di Torino, Turin, Jan. 2024.

[41] M. Moshinsky, "HVAC DUCT CONSTRUCTION STANDARDS METAL and FLEXIBLE," *SMACNA*, vol. 13, no. 1, 1995, ISSN: 0029-5582.

[42] D. Duncan, *Headspace moisture analysis for determination of residual moisture content in lyophilized pharmaceutical products*, 2016.

[43] S. Nail, S. Tchessalov, E. Shalaev, *et al.*, "Recommended Best Practices for Process Monitoring Instrumentation in Pharmaceutical Freeze Drying—2017," *AAPS PharmSciTech*, vol. 18, no. 7, 2017, ISSN: 15309932. DOI: 10.1208/s12249-017-0733-1.

[44] E. Joseph, "Residual moisture determination in lyophilized drug products," *Pharmaceutical Technology*, vol. 43, no. 11, 2019, ISSN: 21507376.

[45] S. Timoshenko, *Strength of Materials Part I: Elemental Theory and Problems*. Huntington: Robert E. Krieger Publishing Co., Inc, 1958, pp. 410–412.

[46] R. J. Roark, W. C. Young, and R. Plunkett, "Formulas for Stress and Strain," *Journal of Applied Mechanics*, vol. 43, no. 3, 1976, ISSN: 0021-8936. DOI: 10.1115/1.3423917.

[47] E. O'Connell, "Method for Continuous Inspection of Product Weight During Lyophilization," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Sep. 2021.

[48] M. J. Pikal and S. Shah, "The collapse temperature in freeze drying: Dependence on measurement methodology and rate of water removal from the glassy phase," *International Journal of Pharmaceutics*, vol. 62, no. 2-3, 1990, ISSN: 03785173. DOI: 10.1016/0378-5173(90)90231-R.

[49] A. Ganguly, S. L. Nail, and A. Alexeenko, "Experimental determination of the key heat transfer mechanisms in pharmaceutical freeze-drying," *Journal of Pharmaceutical Sciences*, vol. 102, no. 5, 2013, ISSN: 15206017. DOI: 10.1002/jps.23514.

[50] J. Park, J. H. Cho, and R. D. Braatz, "Mathematical modeling and analysis of microwave-assisted freeze-drying in biopharmaceutical applications," *Computers and Chemical Engineering*, vol. 153, 2021, ISSN: 00981354. DOI: 10.1016/j.compchemeng.2021.107412.

[51] P. Srisuma, G. Barbastathis, and R. D. Braatz, "Analytical solutions for the modeling, optimization, and control of microwave-assisted freeze drying," *Computers and Chemical Engineering*, vol. 177, 2023, ISSN: 00981354. DOI: 10.1016/j.compchemeng.2023.108318.