A new way to look at the performance of passive cold-chain shipping containers

A 'performance curve' chart for containers could simplify the design process

By Bernard McGarvey, Eli Lilly and Co.; Geoff Kaiser and Shreyas Panse, Cold Chain Technologies; and Paul Harber, Modality Solutions

ABOUT THE AUTHORS (L TO R)
Paul Harber is principal at Modality Solutions, LLC. Dr. Bernard McGarvey is Senior Engineering Advisor at Eli Lilly & Co. At Cold Chain Technologies, Geoffrey Kaiser is Director of Engineering & Technical Services, and Shreyas Panse is Thermal Design Engineer.

Specifying a suitable container for shipping temperature-controlled life sciences products usually involves matching the design elements of the container (insulation quality, amount of gel refrigerants, etc.) with the trade lane (the transportation route) that the package will follow. Many pharma companies verify this performance by doing field tests (shipping the container along the specified route) and monitoring overall performance. From this comes the traditional practice of specifying "summer" (warm season) and "winter" (cold season) configurations, or, when possible, an "all seasons" configuration.

What follows are some fairly technical approaches to generalizing this design practice. The model we are following is something commonplace to engineers: the "pump curve" that measures how a given pump will perform under varying conditions of flow and pressure. The intent here is to show that some current practices in handling temperature-controlled shipments—such as the temperature at which packages are maintained prior to shipping—need to be examined more closely; and that a "performance curve" for container designs could simplify the design process.

Engineers routinely select equipment such as pumps, valves and agitators. This article expands this routine selection process to passive shippers intended to maintain a specified range of temperatures from the point of origin to the destination. Once the engineer has identified the requirements, the next step is to choose a specific piece of equipment that has the capability to meet these requirements. For any given type of equipment (or passive shipper), there are many vendor possibilities. Furthermore, each vendor will have a wide range of options. In order to make a decision, the engineer will initially use some of the requirements to eliminate many of the potential candidates—for example, material of construction, cost or ease of maintenance. Once the list has been narrowed in this way, the engineer seeks a performance curve for the equipment. For the pump, the vendor presents a performance curve consisting of a graph of pressure developed vs. the flow rate through the pump. An example of such a performance curve is shown in Fig. 1.

Performance curves share a number of common attributes:

(1) The curve plots the equipment capability vs. system requirements. For example, in the case of the pump curve, the system pressure represents the capability, and the requirement is represented by the flow rate.

(2) The performance curve differentiates design space within which the equipment can perform and where the equipment cannot perform. In Fig. 1, region "A" represents the possible ranges in which the pump can satisfy the design requirements, whereas region "B" represents the conditions in which the pump is improperly sized for its intended use.

(3) No engineer will specify equipment without a performance curve. Without a performance curve, the selection process breaks down to an expensive, time-consuming trial-and-error activity. Lack of a performance curve indicates that the vendor may not understand the capability of the equipment well enough for the engineer to use the equipment with confidence.

(4) Vendors have many customers with different demands, and it is prohibitive for the vendor to create individual performance curves for all of these different demands. For the pump example, one customer application might use water and another customer application might use ethylene glycol. These two liquids have different densities and viscosities. It would be impractical for the vendor to have a different performance curve for each density and viscosity that could be used by the customer. Over the years, methods have been developed that effectively normalize the performance curves so that a single curve can be applied in many different applications of the equipment, making the curves applicable over a wide range of demands.

(5) Pumps, valves and agitators are complex pieces of equipment when viewed from the perspective of the physical relationships that are occurring when the equipment is being used. The utility of a performance curve is that the critical relationships within the operating equipment are captured in a single performance curve that can be applied in many seemingly different situations.

With respect to (passive) shipper selection, the ambient lane temperature profile represents the requirement, demand or thermal challenge that the container must meet, and the expected duration that the container can operate for a given ambient temperature profile represents the capability or performance of the container. The table below summarizes the comparison between a pump performance curve and a container performance curve.

A performance curve for the container moves the container design process closer to the design processes used for other equipment. A container performance curve could be represented generically as shown

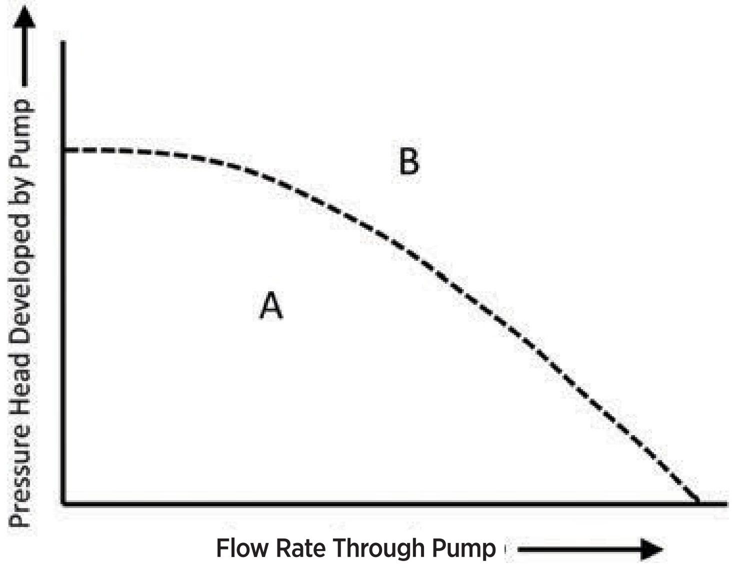

Fig. 1. Typical Pump Performance Curve

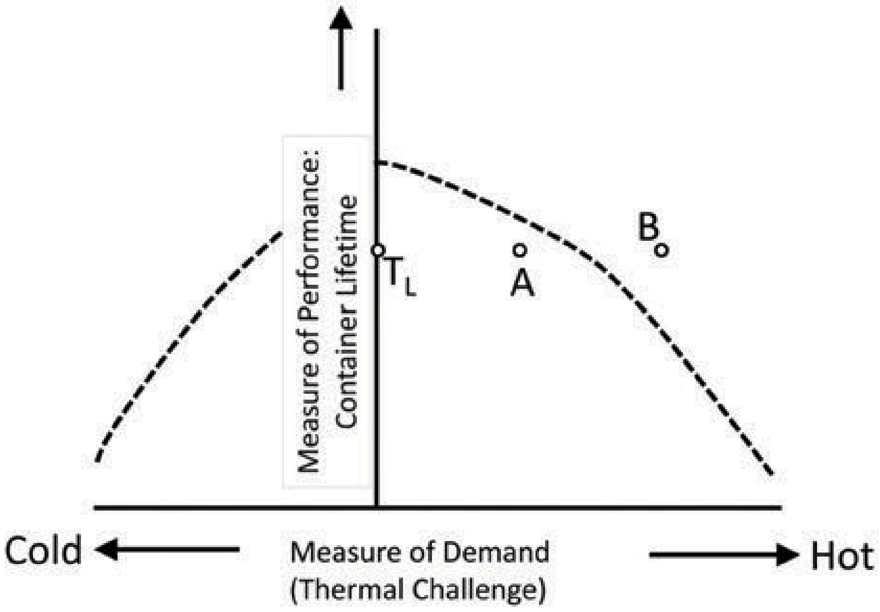

Fig. 2. Passive Container Performance Curve

Table 1. Equipment selection criteria

Item	Pump Selection	Passive Shipper Selection
Elementary Considerations	Purchase cost, maintenance	Purchase cost, complexity of assembly
Performance Capability	Pressure delivered	Container lifetime
Operating Demand	Liquid flow rate required	Ambient temperature profile container must withstand
Selection Criteria	Based on performance curve	Currently based on trial-and-error testing, performance curves not available

in Fig. 2 below. Note that in Fig. 2, the thermal demand measure may be a positive (hot) or negative (cold) ambient condition (this "thermal challenge" metric is defined below, with Fig. 4).

If such a curve were constructed for a given container, it could be used in exactly the same way as performance curves are currently used for other engineered systems. In order to decide if a container is suitable, the customer requests the performance curve from the vendor. Suppose the container is required to maintain 2–8°C for a time TL as shown in Fig. 2. For a particular real ambient temperature profile, the thermal challenge metric is calculated and the point is plotted along a horizontal line through the lifetime value TL. If the point is within the region defined by the curve (A), the container can maintain the product temperature within 2–8°C up to a time TL. However, if the point is outside the performance curve (B), the container cannot maintain the product temperature within 2–8°C up to a time TL. A performance curve that defines shipper capability allows the customer to quickly and easily determine if the container can satisfy the thermal requirements for the bulk of the ambient temperature profiles it is likely to see in real shipments. Given two vendors (one that has a performance curve for its containers and one that does not), a customer would favor using the former rather than the latter—all other things being the same.

The above discussion shows that the performance curve approach discussed here satisfies the first three attributes (1) to (3) indicated earlier. It is not obvious, however, that this approach will have attributes (4) and (5). Given that this type of analysis has not been applied to passive containers before, the existence of these attributes must be verified. In order to test if attributes (4) and (5) hold for passive containers, a set of runs were performed using a validated model for a passive container where the ambient temperature was kept at a constant value until the container failed (product temperature went outside 2–8°C). This container is complex because it contains two layers of phase change material (PCM). The inner layer of PCM phase changes at 5°C, and the outer ice layer phase changes close to 0°C. The lifetime of the container was plotted against the (constant) ambient temperature as shown in Fig. 3.

Of course, in real shipments the profiles are not constant. However, if the performance curve is robust, plotting the results for real shipment profiles should follow the same curve. Fig. 4 shows the results when 30 real shipments were simulated in the same model.

The data points follow the curve very closely. Such results show that the performance curve is applicable to a wide range of behavior in the ambient temperature profiles. This means that the container performance curve also has the attributes (4) and (5) discussed earlier for general performance curves. Note that the performance metric used in the performance curve is the mean profile temperature up to the point of failure of the container (between t = 0 and t = container lifetime).

The passive container performance curve shown in Figs. 3 and 4 has a number of very interesting properties that are a direct indication of the complexity of the thermal behavior of the container, and have significant implications for the performance of this type of passive container:

(**a**) There is a constant ambient temperature that maximizes the lifetime of the container. Even though the desired product temperature is 5°C, the optimal constant ambient temperature is 16°C—very different from 5°C. The constant ambient temperature that maximizes the container lifetime is called the "sweet spot" temperature, denoted TSS.

(**b**) The reason why TSS is much higher than 5°C is due to the use of the outer ice PCM. This ice provides protection against higher ambient temperatures. If TSS was 5°C, many of the ambient profiles that the container will be exposed to would be above TSS since in real shipments most of the time the ambient temperature is above 5°C. However, given that TSS is actually 16°C, the container will be exposed much more to temperatures below TSS. This means that in the container, the aspect of the design that protects against high temperatures also exposes the container on the cold side. In effect, TSS defines what hot and cold really means from the perspective of the container.

(**c**) If and when shipping delays are encountered, the forwarder may offer to hold the shipper in a 5°C cold room. The performance curve shows that this will result in lower safe storage duration, as the ice will cool the product below 2°C. It would be much better to hold the shipper as close to its sweet spot temperature as possible to ensure a maximum excursion-free storage time. Unless and until all gels have been exhausted, the remaining energy in the shipper will drive the payload temperature below 2°C. Holding the shipper at the sweet spot is generally the best option. This practice conserves the energy in the shipper and it will then continue to protect the payload for the maximum time after the hold period is over.

(**d**) The performance curve is highly nonsymmetric around the SST. For example, the container lifetime at 26°C (10°C above TSS) is about 50% greater than the container lifetime at 6°C (10°C below TSS), which is expected from a theoretical perspective. This has a significant implication from a design perspective since it has a major impact on what the optimal TSS should be for a certain set of real ambient shipment profiles. For example, suppose the demand on the hot side is represented as a constant ambient temperature of 29°C and the demand on the cold side is represented as a constant ambient temperature of 3°C. If the performance curve was symmetrical about TSS, it would make sense to design a container with TSS = (3+29)/2 = 16°C. The lack of symmetry, however, means that the optimum value of TSS should be less than 16°C.

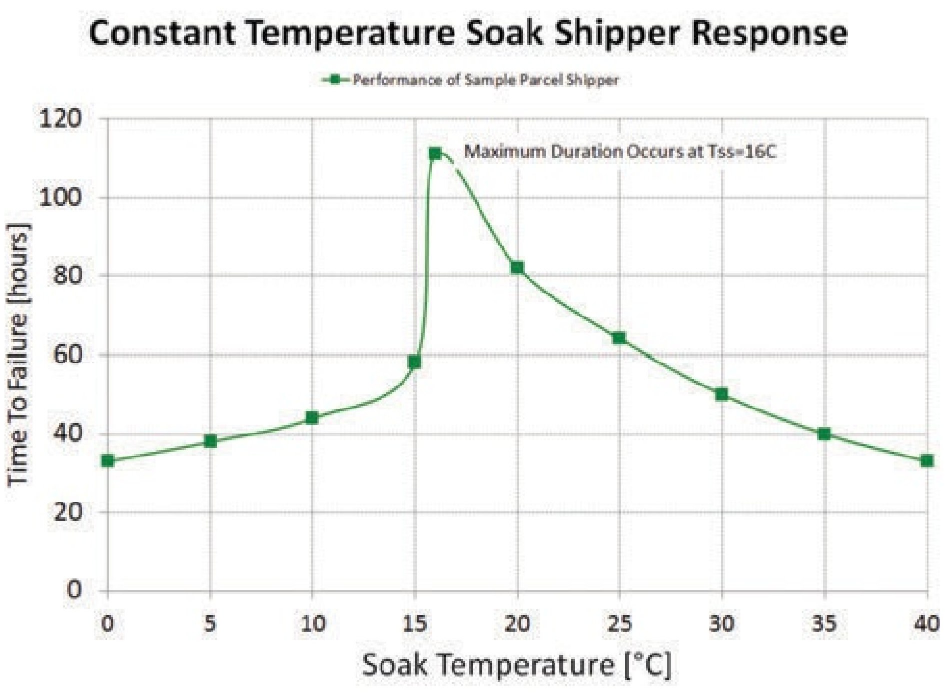

Fig. 3. Passive Container Performance Curve

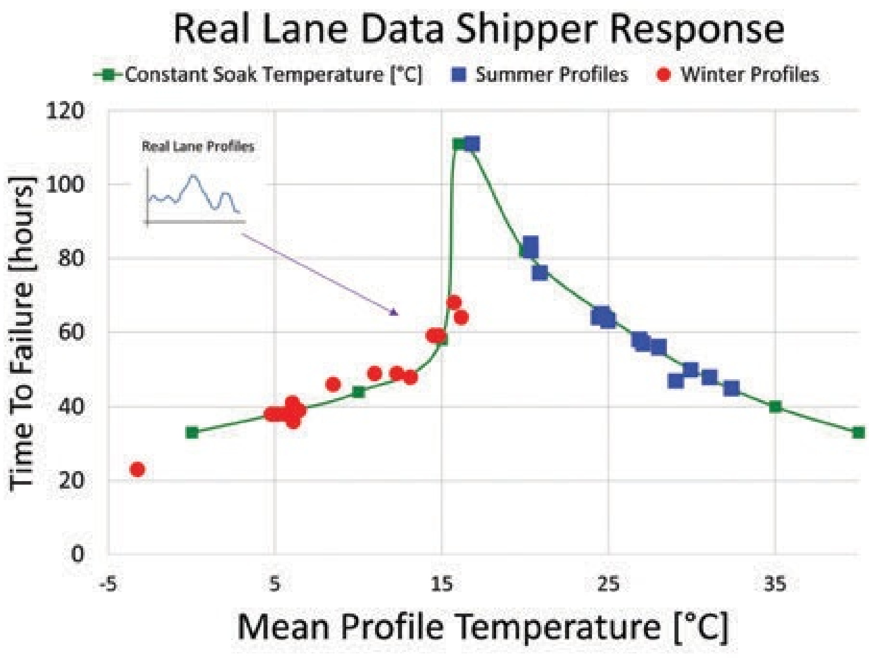

Fig. 4. Passive Container Performance Curve Robustness Results

(**e**) Even if the customer only knows the value of TSS and the customer does not know the complete performance curve, it is still possible for the customer to use the theory of the performance curve. The metric can be evaluated for each real ambient profile and ranked based on this metric. With this information, an Operational Qualification (OQ) ambient temperature profile can be chosen so that the metric for the OQ curve shows that the OQ curve is more thermally challenging than the bulk of the real ambient temperature profiles. Since the performance metric is the mean profile temperature from the beginning of the profile to the time the container failed and since without the performance curve it is not known how long it will take for the container to fail, the metric cannot be used directly. In this case, the maximum (or minimum on the cold side) of the metric (from the beginning to the maximum time that the container is required to protect the product) is used to rank the profiles.

Note that the form of the performance curve will change when the container changes from one design to another. It is important to recognize that the curve can be modified by changes which do not change the physical components. The conditioning of the PCMs is a prime example of how the container performance curve is changed simply by changing the initial conditions of the container before the product is loaded and the container is shipped.

The use of performance curves to design passive containers as described above is not perfect, but it is an improvement over current approaches. In particular, if a real ambient profile has a very high temperature spike (or a very low dip), the performance curve may not reflect the impact correctly. However, this situation can be handled by a separate analysis on the location, size and length of temperature spikes (and dips). PCm