Sensitivity Analysis of Heat Rejection and Propellant Management Technologies for Nuclear Thermal Propulsion Architectures

Robert J. Hetterich[1] and Mitchell A. Rodriguez[2] and Stephen J. Edwards[3]
NASA Marshall Space Flight Center, Huntsville, Alabama, 35808, United States

Cryogenic fluid management (CFM) technologies are very important for enabling a wider range of missions that utilize advanced cryogenic propulsion concepts such as nuclear thermal propulsion (NTP). Technologies for thermal and cryogenic propellant management allow for vehicles to take full advantage of the higher efficiency NTP systems for longer duration human interplanetary and deep space robotic missions. When developing these technologies, it is important to understand the sensitivities of key performance parameters (KPPs) at the system and overall mission level due to the ways the technologies interact with each other and other subsystems to influence the overall vehicle. The Advanced Concepts Office (ACO) at NASA's Marshall Space Flight Center recently developed an integrated system model of a human Mars NTP mission to evaluate the impacts and sensitivities of CFM technologies on the overall vehicle and mission. This paper will cover the buildup of the model and highlight major sensitivities and breakpoints encountered, as well as future work in improving the existing models and sensitivities being evaluated.

I. Introduction

High efficiency propulsion is an important technology for enabling high energy, roundtrip human Mars missions. Total mission duration is understood to be a major driver for health and safety of the crew, so reducing the time crew are exposed to microgravity and radiation environments is a long-standing objective. However, reducing roundtrip time comes at the price of significant increases in energy requirements, placing a premium on propulsive efficiency. Nuclear Thermal Propulsion (NTP) presents a propulsion option with a much larger efficiency than chemical propulsion, but with operational and system complexity advantages over concepts employing electric propulsion. The combination of efficiency and thrust has led to NTP often being considered for human missions to improve mission schedule and benefit crew health and logistics [1, 2].

NTP concepts typically employ cryogenic propellants to achieve the highest efficiencies, and so cryogenic fluid management (CFM) technologies are important for enabling NTP concepts. Many of these CFM technologies are in active development so it is important to understand how they impact the performance of the overall system. To gain a better understanding of these CFM technology impacts, the CFM Portfolio Project from NASA's Science, Technology, and Mission Directorate (STMD) tasked Marshall Space Flight Center's (MSFC) Advanced Concepts Office (ACO) with performing a sensitivity study. This iteration of the study is primarily interested in sensitivities of system leakage, cryocooler performance, intermediate stage cooling temperature, tank residuals, and ullage.

For this work, a conceptual integrated roundtrip human Mars architecture model was leveraged. The Mars mission durations and environments make for a demanding operational profile for the CFM technology portfolio. The integrated nature of the architecture model allowed for sensitivities and interactions to be captured and propagated to performance at the system of systems level. The overall goal for this study was to use the integrated model of the human Mars NTP architecture to generate top level sensitivity data which can help to inform CFM Project about their technology investments.

---

[1] Aerospace Engineer, Advanced Concepts Office, NASA/MSFC ED04.
[2] Aerospace Engineer, Advanced Concepts Office, NASA/MSFC ED04, AIAA Member.
[3] Mission Architect, Advanced Concepts Office, NASA/MSFC ED04, AIAA Member.

II. Representative Architecture Overview

The representative NTP architecture being used to evaluate the CFM technology sensitivities is an 850-day round trip, crewed Mars mission. In this architecture the crew would perform a 30-day surface stay at Mars before returning to Earth. The vehicle transporting the crew to and from Mars would use hydrogen propellant for the NTP engines, thus requiring the use of cryogenic fluid management technologies. It should be noted that the architecture concept used in this sensitivity analysis is one option from a large trade space of possible human Mars architectures, and no decisions have been made on a final human Mars transportation architecture.

A. Vehicle Configuration Overview

The NTP vehicle being modeled with the CFM technologies in this representative architecture consists of four primary element types, as shown in Fig. 1. The core stage primarily consists of the NTP engines and provides the primary propulsion for the vehicle. The hydrogen tanks would all be commonly designed and provide propellant to the core stage. This notional vehicle configuration consists of 13 hydrogen tanks which are mostly dropped off the vehicle as they are used up during the mission. These tank elements are where most of the CFM technologies modeling takes place since they are where the cryogenic propellant is stored. The power and propulsion bus element is a structure to which the drop tanks mount. This element would provide power to the rest of the vehicle after it has been fully assembled and route propellant between the tanks and core stage. Finally, there is the transit habitat element. The transit habitat is the element which would carry the crew and logistics during the roundtrip to and from Mars. All the different elements would be launched individually and aggregate into the full vehicle configuration.

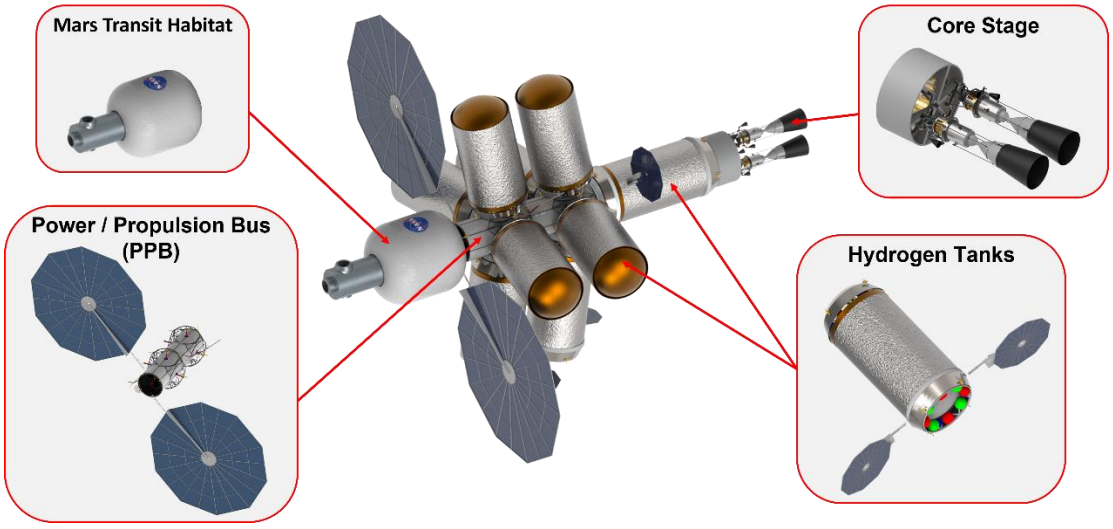

Fig. 1 Notional NTP vehicle configuration.

B. Mission ConOps Overview

In this sample mission's concept of operations (ConOps), the crew would depart for Mars in the 2039 window, perform a 30-day surface stay, and return to Earth within 850 days. The vehicle would be launched as separate elements which aggregate in an optimized medium Earth orbit (MEO). The assembled vehicle would then perform multiple burns with the NTP engines to raise its orbit apogee. After raising the orbit and removing the hydrogen tanks used up during the burn, the vehicle would transfer to a near-rectilinear halo orbit (NRHO) where it could pick up the Mars transit habitat and logistics, which had been waiting at Gateway. The vehicle would take the habitat from NRHO to a lunar distance high Earth orbit (LDHEO), where it does a rendezvous with the crew launched in Orion and prepares for the roundtrip Mars mission.

The crewed NTP vehicle would carry out the trans Mars injection (TMI) burn in October 2039 and drop any additional used up hydrogen tanks during the burn. The crew would coast to Mars for about 261 days and then the NTP vehicle would perform the Mars orbit insertion (MOI) burn. In Mars orbit, the crew would rendezvous with the crew rover and lander which would be sent to Mars in an earlier cargo mission. The crew would descend in the crew rover / lander, carry out the 30-day surface stay inhabiting the rover, then meet up with the Mars ascent vehicle (MAV), also sent by an earlier cargo mission, and ascend back into orbit. The MAV would rendezvous with the transport vehicle so the crew can move back to the transit habitat. The return trip would begin with the trans Earth injection

(TEI), have a deep space maneuver (DSM), then end with the Earth orbit insertion (EOI) burn. The return trip was optimized to about 538 days. Once back at Earth, what is left of the Mars transit vehicle would rendezvous with Orion, which the crew take back to Earth. The empty transit habitat would then be returned to NRHO where it could dock with Gateway and prepare for the next mission. Throughout the crewed mission, trash would be dumped from the vehicle before major burns and empty hydrogen tanks would be dropped off the vehicle to save mass. The full crew portion of the mission can be seen in the bat chart in Fig. 2.

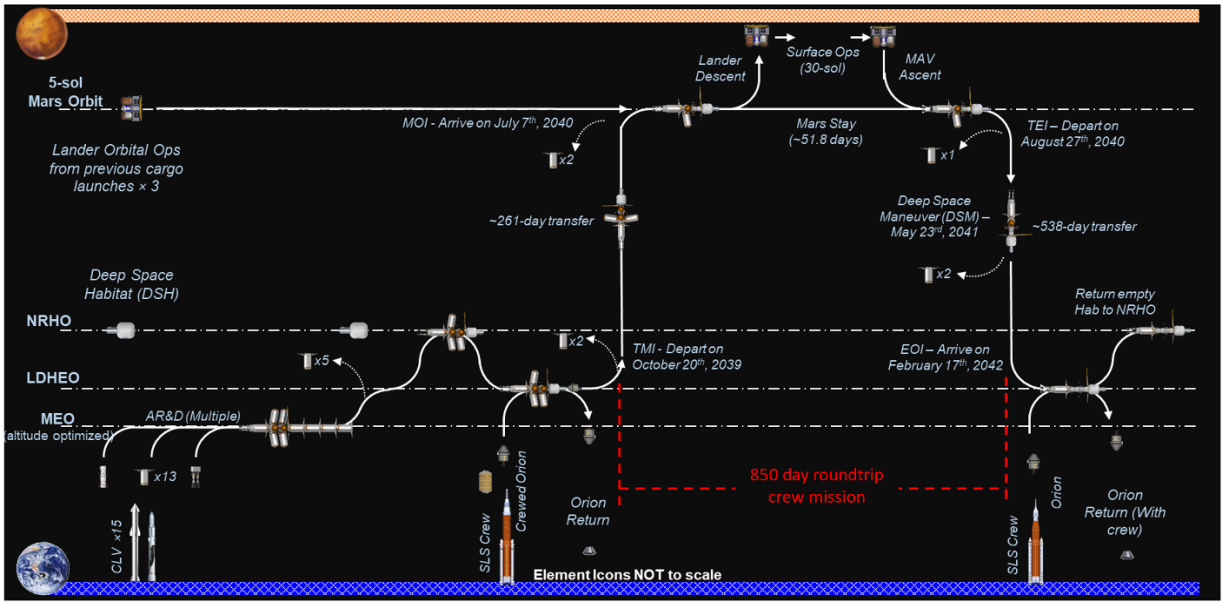

Fig. 2 Notional NTP round trip, crewed Mars mission bat chart.

III. Integrated Systems Modeling Approach

A. Vehicle Element Modeling

For this work, the vehicle elements and mission ConOps were all modeled and integrated using Dyreqt. Dyreqt is a Python module developed in ACO for multi-disciplinary analysis and optimization (MDAO). The framework is built to allow for generalized synthesis and execution of space transportation architectures [3, 4]. Dyreqt is built on top of another NASA developed tool, OpenMDAO, which is a generalized framework for efficient solving, integration, and optimization of multidisciplinary problems [5]. In Dyreqt, each element in the ConOps and the ConOps itself can be modeled as groups of interacting component models. Each of these component models can be at varying levels of fidelity and can be connected to other components to exchange data using "resources" such as propellant, heat, power, etc.

In the Dyreqt models for this study, the vehicle was broken down into the four primary elements that were outlined in Fig. 1. This way the Elements could model separate launches and attach / detach as needed in the Dyreqt model. Fig. 3 shows a breakdown of the component models making up each of the main conceptual vehicle elements in Dyreqt and how they interact with each other by exchanging resources. The black arrows in this figure represent other miscellaneous data connections outside of resources, such as tank surface area. The habitat, PPB element, and core stage were mostly generalized to fixed masses based on a more detailed, conceptual NTP vehicle point design executed by ACO's design team. However, in addition to the fixed mass, these elements also include a power system scaling model since power is the main way the CFM technology trades interact with the rest of the vehicle. The power management and distribution models were primarily built based on equations outlined by Metcalf [6]. The solar array and battery scaling in these models were based on MegaFlex solar arrays and lithium-ion batteries.

The hydrogen tanks, on the other hand, were modeled with a more detailed breakdown of component models since most of the CFM technologies of interest reside in the tank elements. The tank elements were modeled as zero-boil-off systems with active cooling and a two-stage cooling approach. The tank model was broken down into component models for multi-layer insulation (MLI), cryocoolers, power system, radiator, etc. to capture how the sensitivities of CFM technology trades would propagate through the rest of the system. These models, which size component masses

and power / heat / propellant interactions, were a mix of level-0 physics-based models and regression-based models using data internal to ACO or from the CFM Portfolio Project. It should also be noted that the leaky valves model currently assumes that all of the hydrogen tanks' valves are leaking externally, which represents a worst case scenario for leaking propellant.

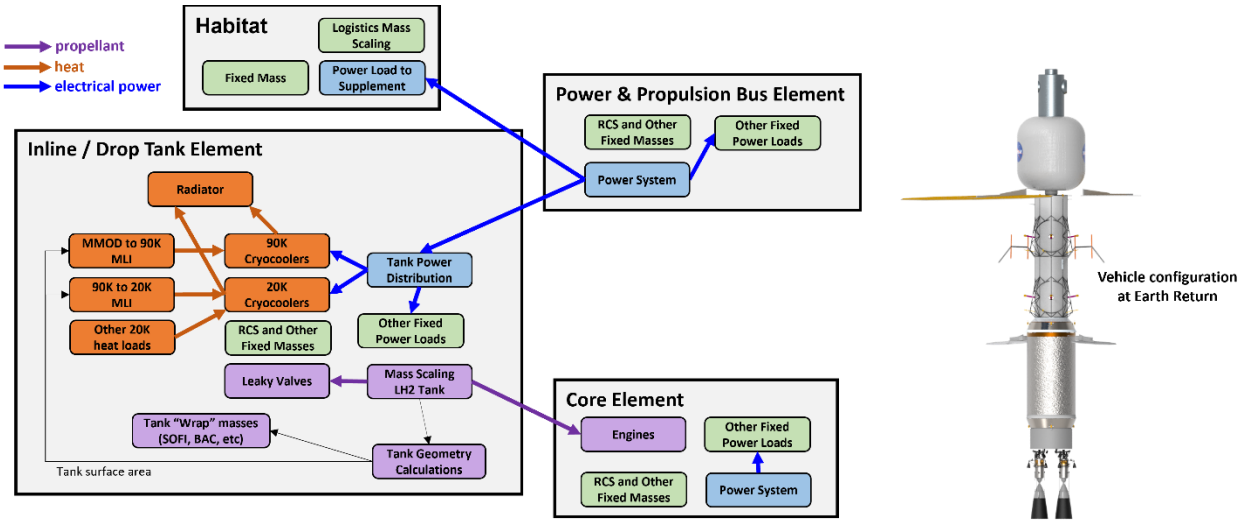

Fig. 3 Vehicle elements modeled as a network of interacting component models with Dyreqt, shown for the vehicle state at Earth return.

A more detailed view of the data being passed along with the heat between the CFM related component models can be seen in Fig. 4. An environment temperature of 176 K was assumed at the outside of the vehicle based on a separate, more detailed thermal desktop analysis. The only computed heat loads to the cryocooler models were the heat leaks through the MLI. Other heat loads were not explicitly sized but were treated as fixed heat loads based on test data.

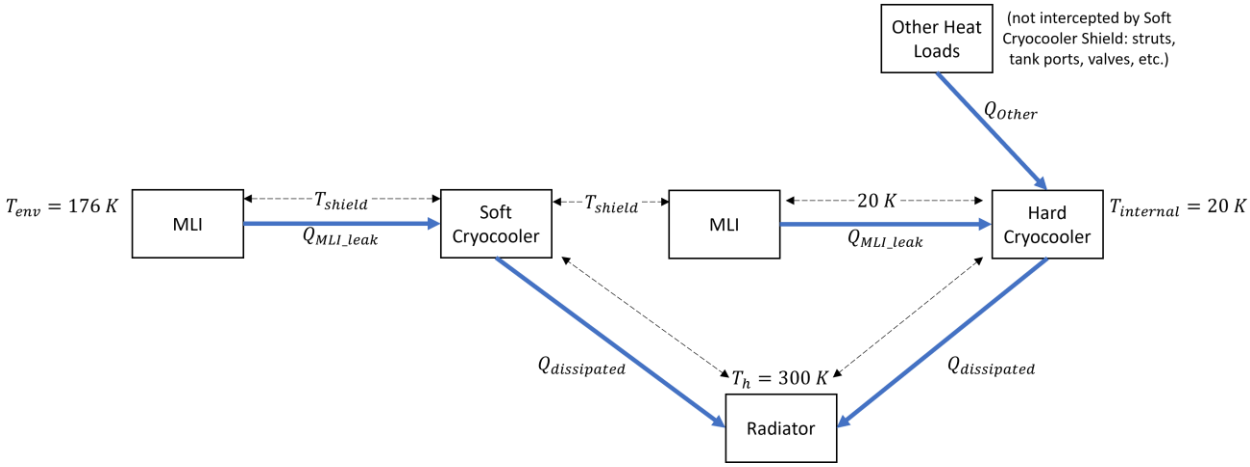

Fig. 4 More detailed heat data connections between CFM components in the LH2 elements.

B. Mission ConOps Modeling

The ConOps side of modeling in Dyreqt is like the element modeling in that events occurring during the mission ConOps can be broken down into a series of connected component models. The element models are flown through these mission event models and can interact with other elements during the mission. As the vehicle configuration changes, such as drop tanks being removed, the connections between the elements can change with it. The ConOps

shown in Fig. 2 can be broken down into a more detailed sequence of events where different elements in the mission will connect, disconnect, dump trash, carry out burns while connected, do things in parallel, etc. The mission modeling in Dyreqt is very flexible in the ConOps that it can represent. In Fig. 5, for example, a zoomed in snippet of the bat chart ConOps can be broken down into a series of discrete events. Each of the event blocks contain a model at varying levels of fidelity to simulate how the vehicle elements and their components interact with each other or themselves. Some events may just contain a rocket equation calculation, some compute the total logistics load, some may move a mass between elements, some may add launch vehicle mass constraints for optimization, etc. With the Dyreqt framework, the system model can simulate interactions between components, between elements, and how those interactions change with the addition / removal of other vehicle elements across the mission ConOps. Note that in this study, the mission ConOps was not varied and was left relatively static to keep the focus on the CFM technologies. The mission roundtrip times and major NTP burn ΔVs were provided by a Copernicus deck and were left fixed in this analysis.

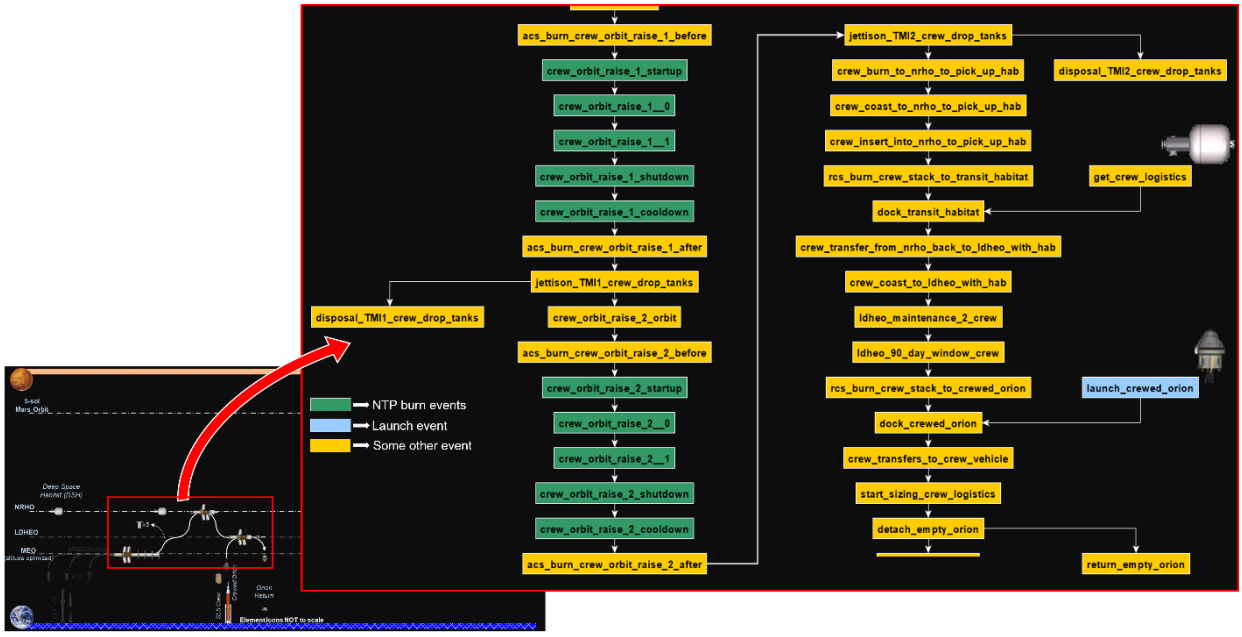

Fig. 5 Mission modeled as a series of connected event models with elements in Dyreqt.

Using this integrated model, parametrics of interest can be run by changing the value of the inputs to any of the event components or vehicle subelement components. Through the component model connections, these parametrics will be propagated to impacts at the system and system of systems levels. This will allow for the evaluation of sensitivities on overall mission closure for the different CFM technology key performance parameters of interest.

IV. Model Execution

For this study CFM Project identified a handful of technology trades in which they had the most interest, as outlined in Table 1. These trades in particular focus primarily on cryocooler performance, valve leak rates, and tank conditions. The ranges for these parameters were also provided by CFM Project and were based on both the current state of the art and their KPP targets. However, some of the ranges were iterated on and tightened over the course of the study as the sensitivities were run, as some of the original ranges were found to produce too many infeasible vehicles. The final set of executed ranges is what is shown in Table 1. Many outputs of interest were recorded and looked at, but the primary output of interest for these sensitivities ended up being the sized propellant load and volume of the liquid hydrogen (LH2) tanks. The tank volume output is the major constraint determining when the vehicle is no longer feasible due to launch vehicle volume limits and a new tank must be added to the vehicle.

Table 1 Parametrics to run through the integrated system model.

Trade Category	Parameter	Units	Lower	Upper
Valves / Leakage	Valve leak rate	SCIM	0	300
Cryocoolers	Heat Lift, Soft Cryo	W	20	180
Cryocoolers	Integration Losses, Soft Cryo	unitless	0.1	0.2
Cryocoolers	Heat Lift, Hard Cryo	W	10	100
Cryocoolers	Integration Losses, Hard Cryo	unitless	0.1	0.2
Cryocoolers	Cold Side Temperature, Soft Cryo	K	70	120
Residual Propellant	LH2 Tank Residuals	unitless	0.03	0.05
Pressurization	LH2 Tank Ullage	unitless	0.03	0.05
Dry Mass	Cryo Tank Element Basic Dry Mass Scaler	unitless	0.95	1.05

To properly generate the sensitivities of interest, it is not sufficient to independently sweep each parameter one at a time. The sensitivity of the outputs to each parameter can potentially shift and change depending on the settings of the other parameters due to interactions and dependencies in the model. If the valve leak rate is very high, for example, then that means the propellant load will be larger, which will increase the surface area of the tanks, the heat leak through the MLI, and therefore the impact of the cryocooler parameters. To capture any dependencies or interactions, a Sobol sequence design of experiments (DoE) was created where the nine parameters of interest were varied simultaneously within their ranges. The DoE with about 2000 data points was run through the integrated Dyreqt model and results of interest were collected for each data point. Of the points run through the simulation environment, 99% ran successfully and returned data.

Surrogate models were then fit to the resultant data, with the goal of quickly and accurately predicting the outputs across the parameter space of interest outlined in Table 1. The surrogate models allow for sensitivities to be quickly shown anywhere in the provided parameter space as well as how those sensitivities vary across the space, without having to rerun the original simulation environment every time. Due to the nonlinear nature of parts of the response space, neural networks were used to train the surrogate models for the outputs of interest. The surrogate model fit to the hydrogen tank volume data was able to predict the training data with a mean error of 0.002% and standard deviation in the error of 0.28%. About 99.5% of the training data was predicted with an error of less than 1%, with a few outlier points at the extremes of the space with errors closer to 5 or 6%. To check for overfitting, a separate 650 point testing data set was also run and showed similar goodness of fit. Histogram distributions of the model fit error and model testing error for the hydrogen tank volume surrogate, as an example, can be seen in Fig. 6.

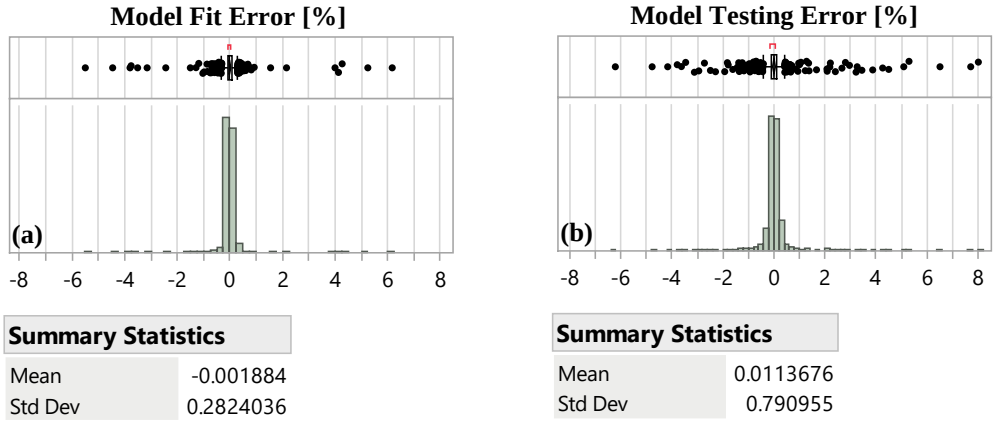

Fig. 6 (a) Model fit error and (b) model testing error for the hydrogen tank volume surrogate models.

V. Sensitivity Results

A. Creating Views of the Data

Using the surrogate models built from the DoE results, partial derivatives can be extracted for each of the nine parameters of interest and used to visualize sensitivities within the parameter space. Since the partial derivatives are extracted from a neural network equation, the values of the input parameters can be changed, and the equation will quickly update the plotted sensitivities. The plots shown in Fig. 7 show the partial derivatives of the volume for the LH2 drop tank elements with respect to a few of the technology parameters as examples. The red number below each plot indicates the current value of that input parameter and the red number on the left shows the resultant LH2 drop tank volume for those parameter settings. Additionally, the black line in each plot is the partial derivative for drop tank volume with respect to each parameter while the red cross-hair shows the current location on the curve of the parameter / response value. The top row of plots shows the tank volume sensitivities for one slice of the parameter space and the bottom row of plots shows how these sensitivities change when looking at a different slice of the parameter space. If there are interactions between any of the parameters, the slopes will also change as the curves shift.

As an example, the leak rate per valve parameter was changed from 1 standard cubic inch per minute (SCIM) to 250 SCIM while all the other parameters were left at their current values in Fig. 7. In addition to the partial derivative curves shifting upwards due to the increase in the volume response when going from 1 SCIM to 250 SCIM, some of the curves also changed shape due to interactions. The LH2 tank residuals partial derivative curve became steeper, for example, going from an average slope of 929 m^3 (9.29 m^3 / %) to about 1283 m^3 (12.83 m^3 / %). This is the expected interaction because as the leak rate goes up, the sized propellant load also goes up, which means that the residual percentage will be of a larger amount of propellant. Like this example, any of the other parameters can also be varied within the ranges from Table 1 to explore different slices of the parameter space.

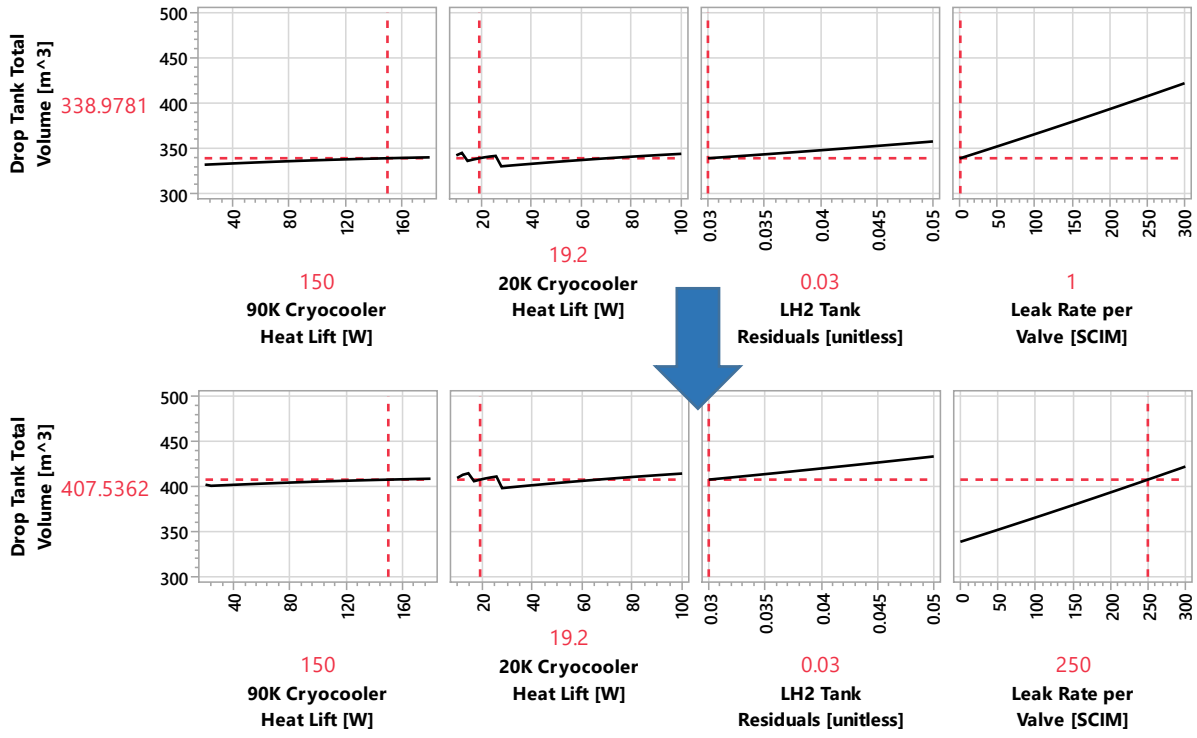

Fig. 7 The sensitivity of the response with respect to the shown parameters changes as the leak rate parameter setting changes from 1 SCIM per valve to 250 SCIM per valve.

With the data and surrogate models that were generated, constraints of interest can be overlaid onto the plots to see where in the parameter space the architecture becomes infeasible. For different slices of the space, parameter limits can be identified that put the architecture just at the limit of feasibility. For the drop tank total volume response, for example, a constraint can be overlaid based on the max volume of a given launch vehicle shroud. For this analysis, it

was assumed that the hydrogen tank elements would be launched by a Starship-like vehicle. Some CAD analysis with the Starship fairing and a more detailed design of the hydrogen tank element showed that the tanks could have a max volume of about 345 m^3 before they no longer fit in the fairing [7]. The plots shown in Fig. 8 are the same as the examples from Fig. 7, but with the additional element of a tank volume constraint overlaid. The constraint was overlaid at a volume of 345 m^3. The shaded area above this number represents the infeasible part of the parameter space where the propellant load of the tanks is too large to fit in the launch vehicle shroud or where an additional tank needs to be launched and added to the architecture. The white area below represents the feasible portion where the tank volume to close the mission is under 345 m^3 for the given number of tanks. In Fig. 8, the leak rate per valve parameter was varied to pinpoint, for this slice of the parameter space, the highest leak rate that could be supported before the tanks become too large. In this plot, the leak rate per valve could be pushed up to about 24 SCIM, assuming the other parameters were left fixed at their "baseline" values.

It is important to emphasize that when looking at breakpoints and evaluating potential limits for technology parameters, it matters where one is looking in the parameter space. The plots highlighted in the previous paragraph show a leak rate limit of 24 SCIM for a specific slice of the parameter space. However, if looking at a different slice of the space where all the other parameters are at their "best" values (right sized heat lift, lowest residuals, lowest integration losses, etc.), the sensitivities move, and the leak rate limit goes up to about 120 SCIM. So, when evaluating sensitivities in an integrated, multidimensional space like this one it can be helpful to have an understanding of how the parameters trade off against each other.

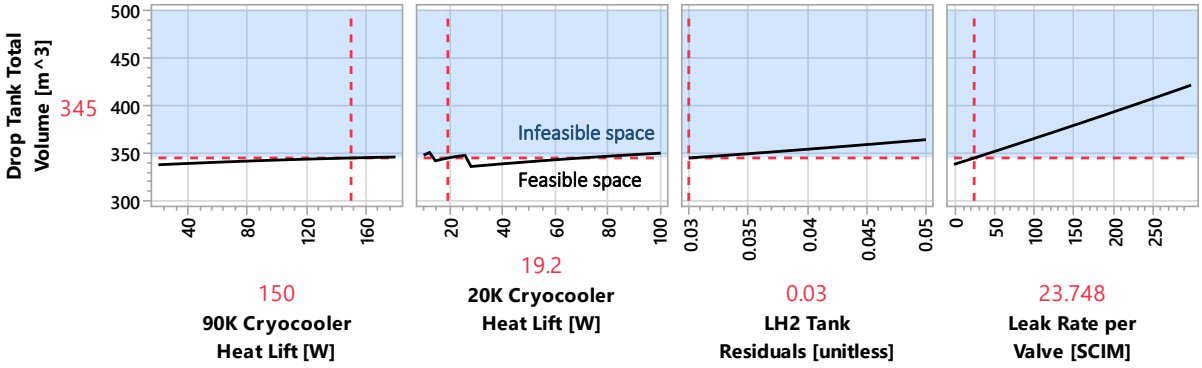

Fig. 8 Volume constraint overlaid onto the sensitivity plots can show for a given slice of the space where the architecture becomes infeasible.

To get a better view of how the parameters trade off against each other with respect to constraints of interest, sensitivity plots can be combined into a contour plot. In Fig. 9, two 1D sensitivity plots with a constraint overlay are combined into a 3D plot and then the top-down view is taken to generate a contour plot. This is a very useful view of the data because it shows how the leak rate per valve and the tank residuals trade with each other along the tank volume constraint boundary. In the contour plot from Fig. 9, moving along the 345 m^3 boundary says, all else constant, for each 1% increase in the LH2 tank residuals, the max allowable leak rate per valve is reduced by about 34 SCIM. So when making decisions on technology investments and the most efficient allocation of money, this view of the data allows a decision maker to ask themselves, "is it easier to decrease the tank residuals by 1% or decrease the leak rate per valve by 34 SCIM?"

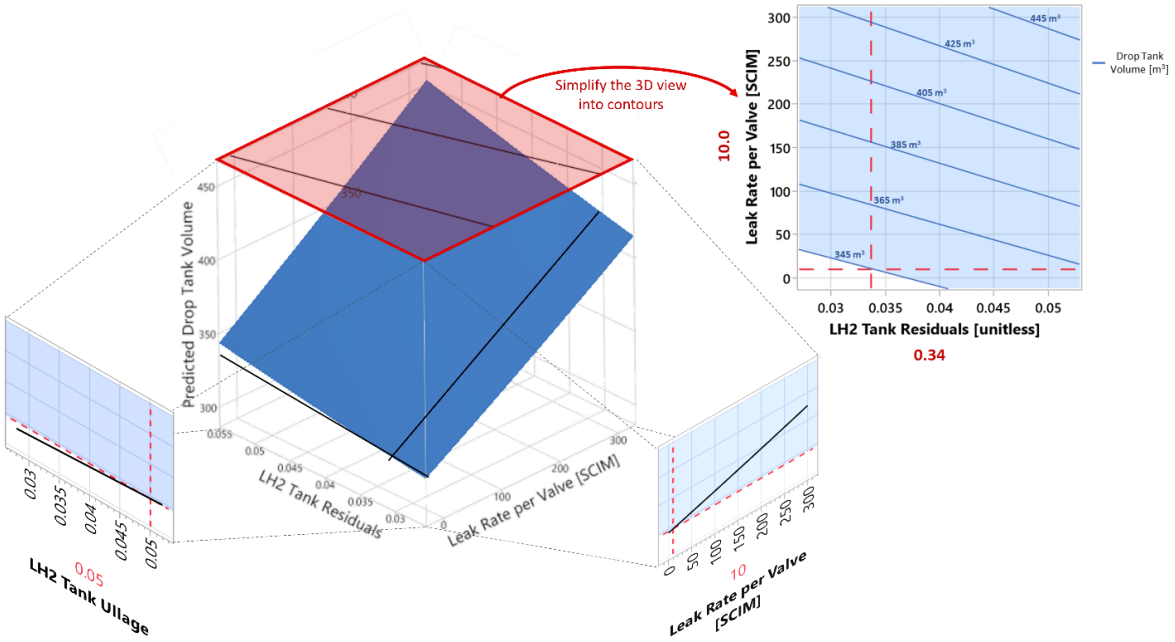

Fig. 9 Can combine 1D sensitivity constraint plots into a 3D view, then simplify into a contour constraint plot.

B. Analyzing the Full Data Set

Using views of the data described in the previous section, the full range of technology parameters can be further analyzed. Fig. 10 shows the full set of partial derivative plots for all nine parameters with the tank volume constraint overlaid. From this slice of the parameter space, the cryocooler integration losses, shield temperature, and tank dry mass scaler noticeably do not have a large impact on the sizing of the vehicle. The integration loss parameters are decrements applied to the lift capacity of the cryocoolers, which primarily impacts the vehicle by increasing the required power load due to the reduced lift. Evidently, the increase in power load is not enough to cause a significant increase in mass. Although the integration loss does have interactions with other terms as the surface area and heat load to the tanks increase, the impacts are relatively flat throughout the entire space. The flatness of the shield temperature parameter is primarily due to offsetting impacts. As the temperature goes up, the power load of the soft cryocooler goes down primarily due to decreased heat leak through the outer MLI. On the other side, the power load of the hard cryocooler goes up primarily due to increased heat leak through its MLI as the shield temperature goes up. The offsetting of these impacts results in a relatively flat plot for this model. Although it looks flat compared to the other parameters, the shield temperature plot is parabolic with a minimum that moves around depending on the location in the parameter space.

While most of the curves called out in Fig. 10 are approximately linear or have some basic curvature, the 20K cryocooler heat lift sensitivity is of particular interest due to its nonlinear shape. This sensitivity resembles a piecewise curve where each piece steps down with a decreasing slope. For each section of the curve, the required tank volume increases as the heat lift increases due to the cryocoolers being oversized for the heat load that needs to be lifted. The transition points where it steps down are the points where the heat lift per cryocooler gets large enough that a cryocooler can be removed. The cryocooler component model in Dyreqt was built so that it right sizes the number of cryocoolers based on the heat load, nominal lift per cryocooler, and any adjustments to the lift based on temperature. In that view of the curve the number of active cryocoolers steps down from three, to two, and then to one at the long part of the plot (always with one spare). Once the lift is high enough that only one cryocooler is needed, then any increase in lift will just make the cryocooler more and more oversized. From this plot the most efficient 20K heat lift at the system level would be at the point where there is one operating cryocooler which is right sized for exactly the needed heat lift. This is the trend that should be expected for this plot as a single, larger cryocooler can be more mass / power efficient than many very small cryocoolers. The 90K cryocooler has some similar behaviors, but the number of required 90K cryocoolers does not vary in this slice of the parameter space. The 90K impacts are also generally smaller than the 20K parameter impacts due to the much lower power load that is required by the 90K cryocooler.

One other trend to call out from these plots is that the leak rate per valve by far has the dominant impact on the sizing of the vehicle. The vehicle used as a reference for this analysis has twelve drop tanks due to the large amount of hydrogen needed for a roundtrip Mars mission and hydrogen's poor density. For the round trip mission and long assembly time, this equates to a large amount of propellant being lost from the vehicle as the leak rate increases.

Another overlay constraint of interest is the burn time of the NTP engines. However, it was found that the constraint limit for engine burn time, which is on the order of several hours, was never really an active constraint compared to the tank volume constraint. With two engines on the crew vehicle and 25 klbf of thrust each, across the parameter space the burn time per engine never became an issue before the hydrogen tank size. So, the analysis and visualizations will primarily focus on the hydrogen tank volume response.

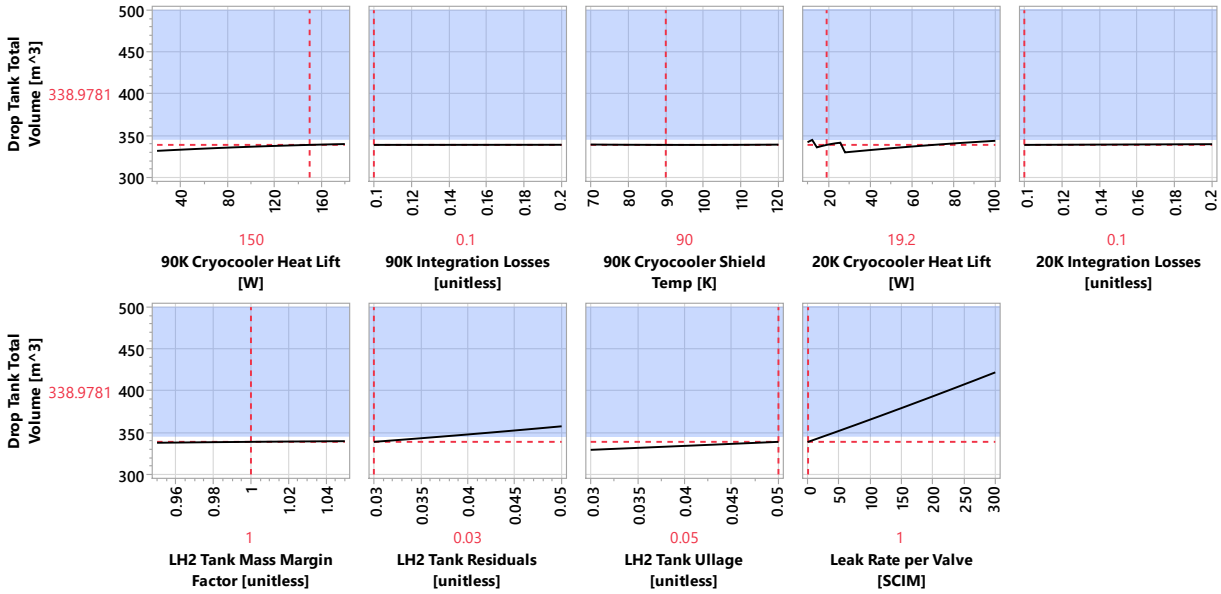

Fig. 10 Sensitivity plots for LH2 drop tank volume with respect to all parameters of interest. Plots at baseline settings with a 345 m^3 constraint overlaid.

For the baseline slice of the parameter space shown in Fig. 10, each technology parameter can be pushed to their feasibility limit. The results for this can be seen in Table 2. The two parameters with the largest impact across their ranges, leak rate and propellant residuals, had finite limits. The other parameters were able to be pushed all the way to their "worst" value (from a systems performance perspective) while the architecture remained feasible from a tank volume perspective. For the valve leak rate parameter, the limit values represent a conservative, worst-case scenario since all the LH2 tank valves are assumed to be leaking externally and losing propellant in this model.

Table 2 Technology parameter limits subject to a max drop tank volume of 345 m^3. Each limit assumes all other parameters are at their "baseline" values.

Trade Category	Parameter	Units	Limit Value
Valves / Leakage	Valve leak rate	SCIM	23.75
Cryocoolers	Heat Lift, Soft Cryo	W	Feasible within full range
Cryocoolers	Integration Losses, Soft Cryo	unitless	Feasible within full range
Cryocoolers	Heat Lift, Hard Cryo	W	Feasible within full range
Cryocoolers	Integration Losses, Hard Cryo	unitless	Feasible within full range
Cryocoolers	Cold Side Temperature, Soft Cryo	K	Feasible within full range
Residual Propellant	Residuals	unitless	0.0369
Pressurization	Ullage	unitless	Feasible within full range
Dry Mass	Cryo Tank Element Basic Dry Mass Scaler	unitless	Feasible within full range

Next, the "absolute" limit for each technology parameter can be found by pushing the parameter's value as high as it can go while setting all the other variables to their "best" value (from a systems performance perspective). The results of doing this for each parameter can be seen in Table 3. From this analysis it can be seen that in the best-case scenario, all of the parameters result in a feasible architecture within their full range except for the leak rate per valve. The valve leak rate can be pushed up to an absolute maximum of 120 SCIM before the architecture is no longer feasible from a tank volume perspective. This further re-enforces the fact that the leak rate technology parameter has the largest impact on performance with the given ranges.

Table 3 Technology parameter max limits subject to a max drop tank volume of 345 m^3. Each limit assumes all other parameters are at their "best" values.

Trade Category	Parameter	Units	Limit Value
Valves / Leakage	Valve leak rate	SCIM	120.12
Cryocoolers	Heat Lift, Soft Cryo	W	Feasible within full range
Cryocoolers	Integration Losses, Soft Cryo	unitless	Feasible within full range
Cryocoolers	Heat Lift, Hard Cryo	W	Feasible within full range
Cryocoolers	Integration Losses, Hard Cryo	unitless	Feasible within full range
Cryocoolers	Cold Side Temperature, Soft Cryo	K	Feasible within full range
Residual Propellant	Residuals	unitless	Feasible within full range
Pressurization	Ullage	unitless	Feasible within full range
Dry Mass	Cryo Tank Element Basic Dry Mass Scaler	unitless	Feasible within full range

Because leak rate was shown to be the most impactful parameter, contour plots were used to show the tradeoff between the other technology parameters and leak rate in Fig. 11. For the flat contour plots like the integration losses, increasing or decreasing the integration losses doesn't buy any additional leak rate margin. The integration parameters don't have a large overall impact and don't have much interaction with the leak rate so, whether the integration losses are 10% or 20%, the max leak rate for feasibility doesn't really change. The heat lift parameters, as expected, allow for the most leak rate when they are closest to being right sized for the max heat load. However, a cryocooler will not generally be sized exactly right for every use case's heat load. In practice, one design will likely be used for multiple different scenarios. For this architecture model and vehicle, the contour plot for the 90K cryocooler heat lift shows that each one Watt the cryocooler is over-sized, the max leak rate that can be supported goes down by about 0.2 SCIM. In the case of the 20K cryocooler, for a large single cryocooler, each one Watt the cryocooler is over-sized decreases the max feasible leak rate by about 0.7 SCIM. With two smaller 20K cryocoolers this goes up losing 1.9 SCIM of max feasible leak rate per Watt each active cryocooler is oversized. This type of information can be used to inform how much valve leak rate would need to be improved for a more generalized cryocooler design to close with this architecture.

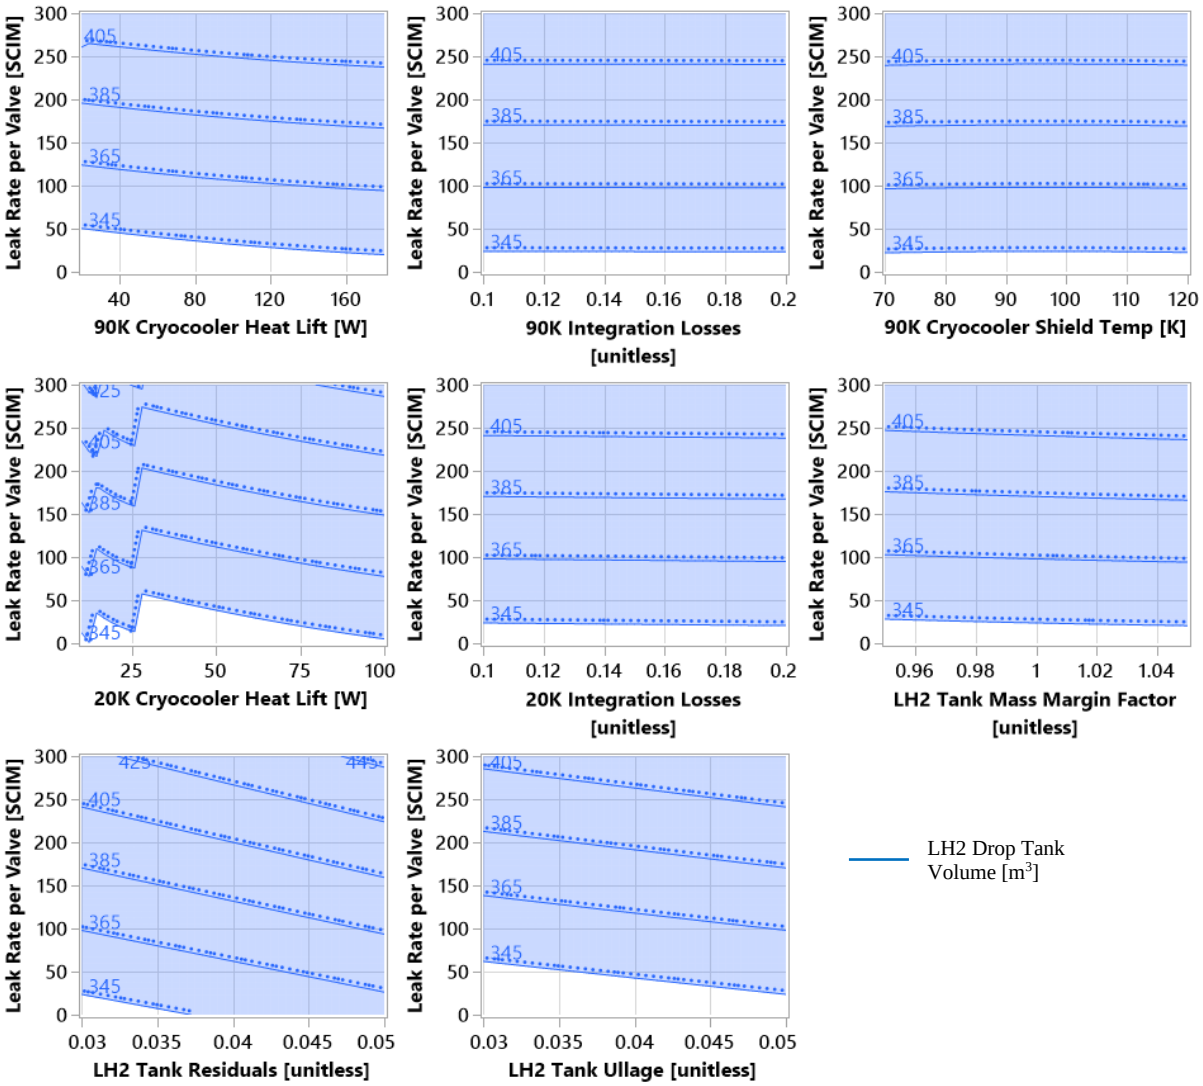

Fig. 11 Contour plot of LH2 drop tank volume for each technology parameter with respect to leak rate per valve.

Due to the multivariate nature of the problem, it's hard to draw many definitive conclusions about any one parameter without knowing the settings of the other parameters. Some of the analyses and data views shown here can help to better inform how variables trade off and interact throughout the parameter space. However, there will always be some uncertainty across the technology portfolio and where each parameter will end up. The surrogates and data set developed in this work can be used in future iterations of this study for additional multivariate analyses, using uncertainty ranges for the technology parameters. Analyses such as probabilistic risk assessment, for example, can leverage uncertainty distributions and the speed of the surrogate models to run Monte Carlos and develop an understanding of where to reduce uncertainty.

VI. Conclusion

The ACO's analysis was able to show sensitivities of CFM technology parameters of interest at the integrated system level for a crewed NTP vehicle. The integrated architecture model was able to generate a data set which captured individual parameter sensitivities and interactions between parameters. This data and corresponding surrogate models showed that leak rate per valve was the most influential parameter on the overall system closure, for

the given ranges. They were also able to show break points for the parameters, how those break points moved in the parameter space, and tradeoffs between technology performance parameters. Additional refinements will be made to further improve the component models to enable analysis of additional parameter sensitivities and relationships.

References

[1] Mars Architecture Steering Group, "Human Exploration of Mars Design Reference Architecture 5.0", Drake, Bret G., ed., National Aeronautics and Space Administration, NASA-SP2009-566, Washington, DC, July 2009.

[2] Borowski, Stanley K., McCurdy, David R., and Packard, Thomas W., "7-Launch NTR Space Transportation System for NASA's Mars Design Reference Architecture (DRA) 5.0", AIAA-2009-5308, August 2009.

[3] Trent, D. J., "Integrated Architecture Analysis and Technology Evaluation for Systems of Systems Modeled at the Subsystem Level," Ph.D. thesis, Georgia Institute of Technology, 2017. URL https://smartech.gatech.edu/handle/1853/59264.

[4] Edwards, S. J., Diaz, M. J., Mavris, D. N., and Trent, D., "A Model-Based Framework for Synthesis of Space Transportation Architectures," *AIAA SPACE and Astronautics Forum and Exposition*, AIAA, Orlando, FL, 2018. doi: 10.2514/6.2018-5133.

[5] Gray, J. S., Hwang, J. T., Martins, J. R. R. A, et al, "OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization," Structural and Multidisciplinary Optimization, Vol. 59, 2019, pp. 1075-1104.

[6] Metcalf, Kenneth J, "Power Management and Distribution (PMAD) Model Development," NASA /CR-2011-217268

[7] SpaceX, "Starship Users Guide," 2020. URL https://www.spacex.com/media/starship_users_guide_v1.pdf.