

Testing and Modeling of the Mars Atmospheric Processing Module

Anthony C. Muscatello,¹ Paul E. Hintze,² Anne J. Meier,³ Elspeth M. Petersen⁴, Jon A. Bayliss,⁵ Ricardo M. Gomez Cano⁶, Rene Formoso,⁷ Malay G. Shah,⁸ Jared J. Berg,⁹ Bruce T. Vu,¹⁰ Alexander R. Walts,¹¹ and Rupert U. Lee,¹²
National Aeronautics and Space Administration, Kennedy Space Center, FL 32899

and

James G. Captain¹³
Vencore, Kennedy Space Center, FL 32899

Here we report further progress in the development of the MARCO POLO/Mars ISRU Pathfinder Atmospheric Processing Module (APM). The APM is designed to demonstrate in situ resource utilization (ISRU) of the Martian atmosphere, which primarily consists of carbon dioxide (CO₂). The APM is part of a larger project with the overall goal of collecting and utilizing CO₂ found in the atmosphere and water in the regolith of Mars to produce methane and oxygen to be used as rocket propellant, eliminating the need to import those to Mars for human missions, thus significantly reducing costs. The initial focus of NASA's new ISRU Project is modeling of key ISRU components, such as the CO₂ Freezers and the Sabatier reactor of the APM. We have designed models of those components and verified the models with the APM by gathering additional data for the CO₂ Freezer and the Sabatier reactor. Future efforts will be focused on simultaneous operations of the APM and other MARCO POLO/Mars ISRU Pathfinder modules.

Nomenclature

at. %	=	atomic percent
°C	=	degrees Celsius
CAD	=	Computer Aided Design
ISRU	=	In Situ Resource Utilization
J	=	Joules
K	=	degrees Kelvin
MFC	=	Mass Flow Controller
SLPM	=	Standard Liters Per Minute
W	=	Watts

¹ Chemist, Applied Science Branch, UB-R3-A, and AIAA Non-member.

² Chemist, Applied Science Branch, UB-R3-A, and AIAA Non-member.

³ Chemical Engineer, Applied Science Branch, UB-R3-A, and AIAA Member.

⁴ NASA Intern, Applied Science Branch, UB-R3-A, and AIAA Non-member.

⁵ Electrical Engineer, Testing and Design Branch, NE-L1, and AIAA Non-member.

⁶ NASA Intern, Applied Science Branch, UB-R3-A, and AIAA Non-member.

⁷ Software Engineer, Avionics Branch, NE-EA, and AIAA Non-member.

⁸ Engineer, Engineering Analysis Branch, NE-M6, and AIAA Non-member.

⁹ Engineer, Engineering Analysis Branch, NE-M6, and AIAA Member.

¹⁰ Engineer, Engineering Analysis Branch, NE-M6, and AIAA Senior Member.

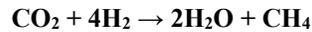
¹¹ NASA Intern, Applied Science Branch, UB-R3-A, and AIAA Non-member.

¹² Engineer, Analytical Laboratories Branch, NE-L3, and AIAA Non-member.

¹³ Scientist IV, Chemical and Biological Sciences, ESC-854, and AIAA Non-member.

I. Introduction

The Mars APM is designed to demonstrate collection of CO₂ from a simulated Mars atmosphere and combine it with hydrogen gas (H₂) through the use of the Sabatier reaction to produce CH₄ and H₂O. This reaction is shown below:



The MARCO POLO (Mars Atmosphere and Regolith Collector/Processor for Lander Operations) project was initially designed for an analog field demonstration of an integrated Mars ISRU propellant production system at a scale sufficient for a Mars Sample Return mission. The field demonstration was not funded, but most of the modules have been built and tested to varying degrees using alternative funding sources, as was described in earlier publications.¹⁻³ The project continues as the Mars ISRU Pathfinder to continue to develop the technologies needed for human Mars missions in the 2030s. Designed and constructed at the NASA Kennedy Space Center (KSC), the APM (Fig. 1) can be divided into two subsystems, the CO₂ Freezer subsystem and the Sabatier subsystem. The CO₂ Freezer subsystem (Fig. 2) extracts the CO₂ from a simulated Mars atmosphere through freezing with a pair of cryocoolers (one freezes while the other supplies CO₂). The Sabatier subsystem (Fig. 3) combines CO₂ and hydrogen (H₂) to make methane (CH₄) and consists of the Sabatier reactor, product separation and recycling of excess reactants (H₂ and CO₂).

H₂ is obtained by first collecting H₂O from simulated Martian soil using the KSC Regolith Advanced Surface Systems Operations Robot (RASSOR 2.0) to deliver regolith to the KSC Hopper on the KSC Lander where it would be sent to the Soil Processing Module (SPM) at Johnson Space Center (JSC). The KSC Lander incorporates a Mock Oven built by KSC to demonstrate regolith delivery and unloading. The KSC Water Cleanup Module (WCM) removes the impurities from the extracted water, followed by electrolysis to yield H₂ and O₂ in the JSC Water Processing Module (WPM). In a flight-like system, the O₂ would be dried, liquefied, and stored, while the H₂ is sent to the Sabatier Subsystem in the APM where it is reacted with pressurized CO₂ feed from the CO₂ Freezer Subsystem. The Sabatier reactor is started with an excess of H₂ and then scaled back to the optimal molar ratio of 4:1; the excess H₂ is recycled through the use of a membrane separation module that removes H₂ and CO₂ from the CH₄ product after water condensation and a recycle pump that allows the reaction to always run with some excess of H₂, preventing CO₂ from contaminating the CH₄ product.

A virtually integrated test with the Lander, the regolith Hopper, APM, WCM, RASSOR 2.0 excavator rover and simulations of the SPM and WPM was conducted in 2016 to demonstrate the compatibility of the system. A set of tests were run on the APM to determine the temperature profile of the Sabatier reactor so that an accurate heat transfer model could be developed; data from previous runs was also compiled so that it could be used as a validation method for the model. The second set of tests were focused on finding the optimal ratio of H₂ to CO₂ that would keep the reaction stable and at the same time keep the CH₄ close to 100% pure. Efforts were also directed towards generating and updating schematic and computer models to match the system.

Prior work^{2,3} has demonstrated that the CO₂ Freezer Subsystem operated well and exceeded the required 88 g/h freezing and supply rate. It freezes ~70% of the incoming CO₂. From this test data, estimates for required power to

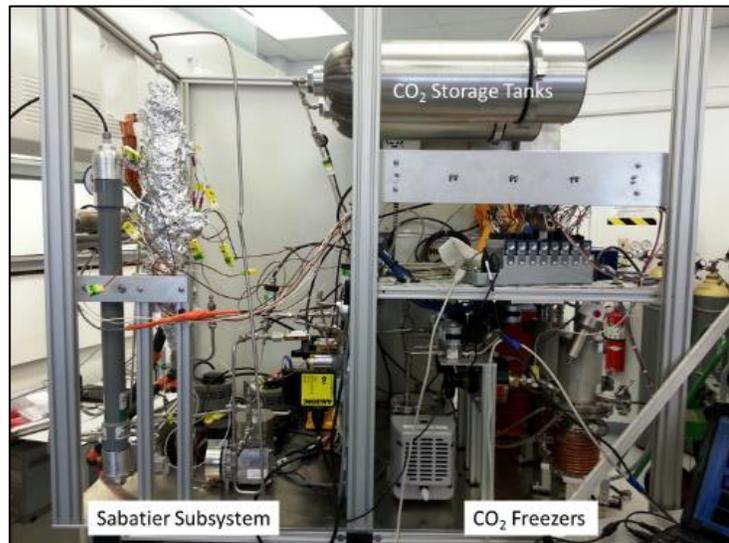


Figure 1. Photo of the APM.

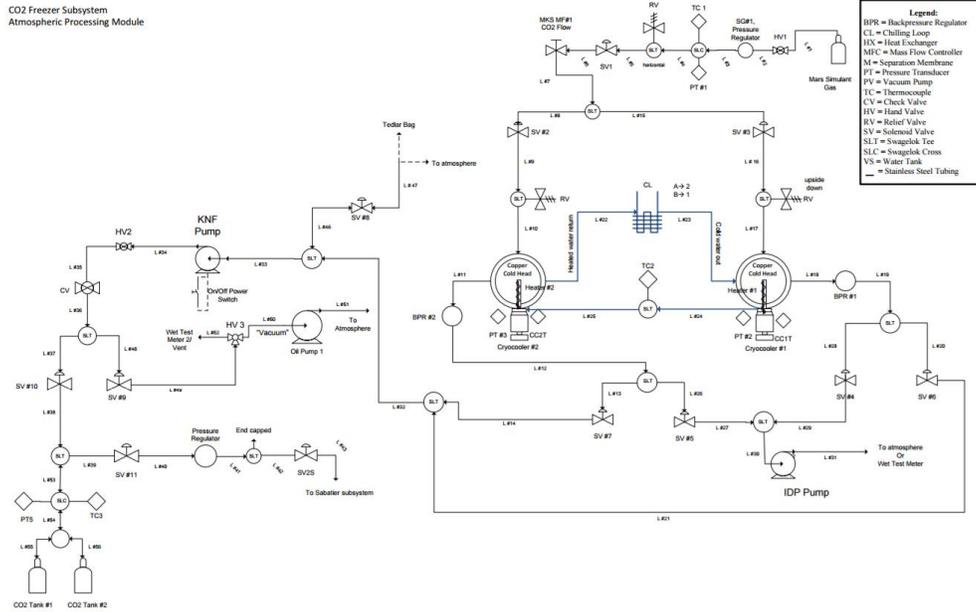


Figure 2. CO₂ Freezer subsystem schematic.

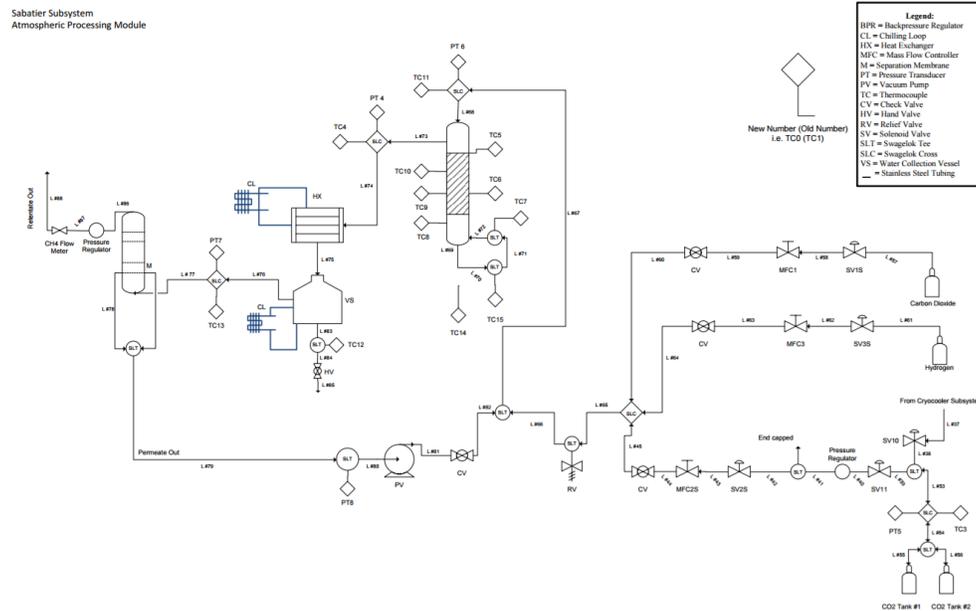


Figure 3. Sabatier subsystem schematic.

freeze CO₂ at Mars pressure were calculated, averaging 0.22 W/g CO₂ frozen during 1.4 h cycles. This was only 108% of the theoretical amount needed to cool the gaseous CO₂ to 150 K and change it to the solid phase without consideration of heat leaks from the freezing chamber walls. This work is contributing to full-scale human Mars mission ISRU system designs, e.g. 680 W of cryocooler lift is required to freeze 3.1 kg CO₂/h for further processing. The Sabatier Subsystem also operates well using CO₂ from the Freezer Subsystem. The reactor is efficient when combined with the recycling system, producing pure (99.9+%) CH₄ obtained at the required 32 g/h rate. There is an average of ~6% shortage in the amount of water produced (64-70 g/h vs. 72 g/h expected), which is not accounted for by the <<1% of loss in the methane product. We continue to seek the cause of this shortage.

Recently, additional integrated tests of the CO₂ Freezers and the Sabatier Subsystem have been performed and faster and slower production rates were tested: 1.0-1.6 (nominal 1.2) SLPM feed to CO₂ Freezers resulted in 87-71%

of the incoming CO_2 being frozen and requiring 4800-5400 J/g of power from the cryocoolers. The Sabatier Subsystem works at a CO_2 feed rate of 0.3 to 1.2 SLPM (0.75 SLPM nominal, 550°C max T observed), but some carbon monoxide was observed in the CH_4 after higher flow rates, which is an indication that the ruthenium on alumina catalyst was damaged. Better LabVIEW automation was implemented, with automated sequences being initiated by the operator instead of operation of individual parts. We also performed “virtual” integrated tests with other KSC systems at KSC in May and September 2016 with physical integration of the WCM on the Lander and virtual integration of the APM planned for October 2017. Also, the Sabatier Subsystem was operated with unseparated Mars Gas Simulant with good results, but argon and nitrogen contaminate the CH_4 , which is expected. Downstream processing, such as cryodistillation to remove and potentially use the Ar and N_2 , would be needed. Our long term goal is to continue to refine ISRU technologies for potential robotic Mars missions such as the proposed Mars ISRU Pathfinder in the mid- to late-2020s.

II. CAD Modeling and Visio Schematic

Depicted below are some of the parts that were modeled or acquired from vendors to incorporate into the APM assembly CAD model. The CAD models for commercial off the shelf (COTS) products are shown in Fig. 4, A-F, while the NASA KSC designed custom parts are shown in Fig. 4, G-I. The full APM assembly is displayed on the right of Fig. 4.

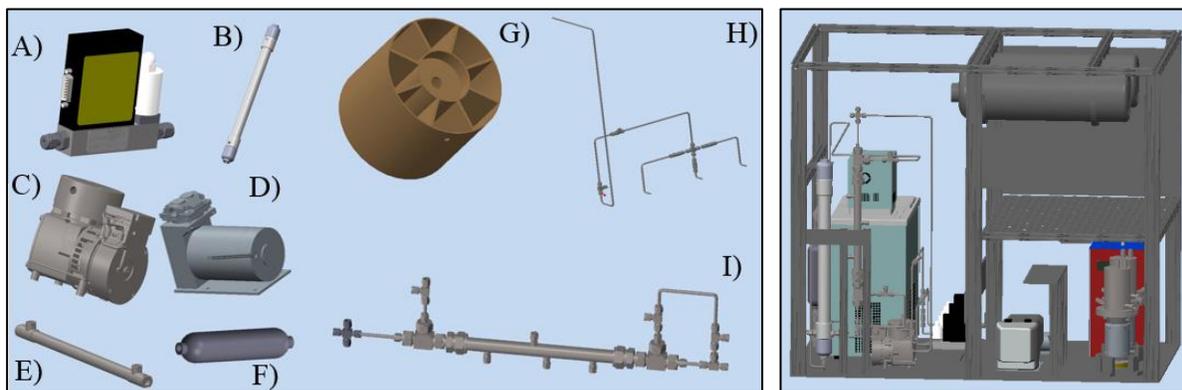


Figure 4. CAD model parts (not to scale in this image): A) Aalborg GFC17 Mass Flow Controller (MFC), B) Air Products nitrogen membrane Separator, C) KNF N022 pump, D) KNF NPK09 pump, E) Exergy LLC heat exchanger, F) Swagelok water collection vessel, G) Cryocooler cold head, H) Tubing from Aalborg MFCs to Sabatier inlet and from Sabatier outlet to membrane separator I) Updated Sabatier reactor. Right: full assembly CAD model.

III. Results

A. CO_2 Freezer Testing and Modeling

1. Model Comparison: Ferris Wheel Vs. Branching Design

Long duration tests were performed with the original “Ferris Wheel” cold head design that had been optimized to collect and supply at least 88 g CO_2 /h in a 1.4 h freezing/1.4 h sublimation cycle in our prior work. Figure 5 shows the dry ice collected on the Ferris Wheel in a 1.4 h cycle after removal of the collection chamber as quickly as possible to visualize the actual performance of the cold head. This was done to assist with calibration of Computational Fluid Dynamics (CFD) modeling of the freezing process. Prior work with a clear polyvinylchloride (PVC) test stand showed that the dry ice is perfectly transparent during the accumulation process. Thus, the white frost in Fig. 5 is probably water ice freezing out from the laboratory air during the removal of the chamber. As during previous developmental testing with the PVC chamber, cracks in the dry ice were visible, a result of the



Figure 5. Dry ice and water ice frost accumulation on the Ferris Wheel cold head.

warming of the cold head after turning off the cryocooler. The dry ice appeared to be uniformly distributed on the copper cold head, although photos looking down through the channels between the fins show thicker accumulation at the bottom, which could be explained by the CO_2 getting colder as it flows from the top to the bottom of the chamber.

Several approaches were tried to obtain reasonable results for the CFD models, which have to deal with the direct conversion of gaseous CO_2 to solid dry ice, an unusual situation in modeling. An approach called “Volume of Fluid” (VOF) proved to be successful and was used as described below. The VOF method uses volume fractions to track the relative amount of solid-phase CO_2 to gaseous CO_2 .

The results of the VOF modeling and other investigations of the design of heat exchangers led to the design of a potentially improved cold head, shown in Fig. 6. Designated the “Branching” cold head, the design was too complex

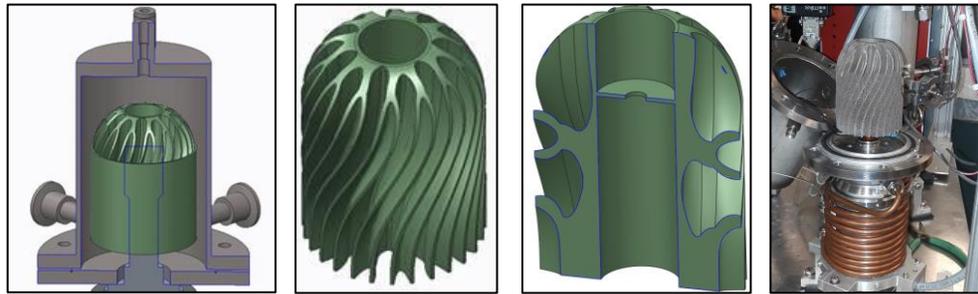


Figure 6. Branching cold head design drawings and photo of the final version mounted in the CO_2 Freezer.

for standard machining, such as that used to make the Ferris Wheel design, which was cut from a copper rod by drilling the center sections and using wire electrical discharge machining (EDM) to cut the

volume between the fins. We had contacted the Marshall Space Flight Center (MSFC) Additive Manufacturing (AM) group in anticipation of needing their assistance. They kindly supplied both an AM version of the Ferris Wheel and the Branching cold head made from GRCop-84, an 88 at.% copper, 8 at.% chromium, and 4 at.% niobium alloy developed by the Glenn Research Center (GRC) for rocket nozzle applications.⁴ Pure copper is difficult to use in AM due to its tendency to oxidize in an air atmosphere during laser sintering. The thermal conductivity of GRCop-84 at 150 K is ~ 280 W/m-K,⁴ which is 70% of that for pure copper at the same temperature (400 W/m-K). The AM version of the Ferris Wheel was planned to be used to determine the effect of the reduced thermal conductivity. Modeling indicated that the effect would be negligible compared to the low thermal conductivity of dry ice at 150 K, so those tests were not performed. The Branching cold head weighed 843 g after final standard machining of the central cylinder to obtain a tight fit on the cryocooler cold tip whereas the Ferris Wheel mass was 254 g. The exposed area of the Ferris Wheel is 415 cm^2 and that of the Branching cold head is 739 cm^2 . The Branching design also featured an outer shell printed with a latticed structure which added additional surface area for CO_2 to grow on, but this area is not included in the 739 cm^2 . It was expected that the Branching design would collect at least twice as much dry ice as the Ferris Wheel.

Simulations were conducted to compare the predictions of dry ice growth between the Ferris Wheel and the larger-area Branching design with possibly better thermal conductivity. CAD drawings of the Ferris Wheel and Branching designs are shown in Fig. 7. A solid outer shell was used in place of the lattice structure for the Branching design to

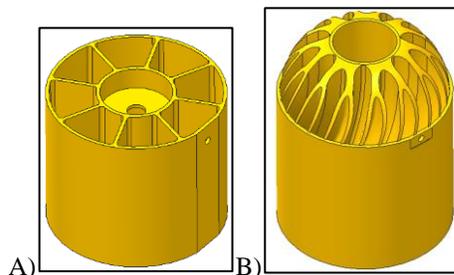


Figure 7. CAD drawings of the Ferris Wheel cold head (A) and the Branching design (B).

reduce the complexity of the simulation. Two computational methodologies were used: one that included the gas flow and a simplified one that treated the domain with no flow. The treatment of the model with no flow was done based on observed growth patterns from an earlier experimental run where the cryocooler chamber was disassembled after growing dry ice on the Ferris Wheel (i.e. Fig. 5). It appeared the dry ice grew fairly uniformly on the outer surfaces especially at the top where the incoming gas impinged directly onto the mounting screw.

In the models, the copper was set to a fixed temperature of 150 K. The outer chamber walls had an environmental convective boundary condition (lab temperature of 20°C with a low convective heat transfer coefficient).

2. Dry Ice Growth: Flow Considered

For the Ferris Wheel cold head, the flow pattern within the chamber caused the dry ice to grow in such a way that the inner volumes were not filled in completely, whereas the Branching design did fill in. This can be seen in Fig. 8. The majority of the ice in the cold heads was present within the inner volume, and it was predicted that only a thin layer existed on the outer shells.

With the incoming warm gas impinging directly onto the top of the Ferris Wheel, the model predicted very little ice growth at the top where the mounting screw is located. This contradicted the experimental results where there was visible ice growth in the same location. Therefore, we can assume the model tended to over-predict the effects of flow on ice growth. These results may indicate a minimum expected dry ice growth when compared to the no-flow assumption in the next section.

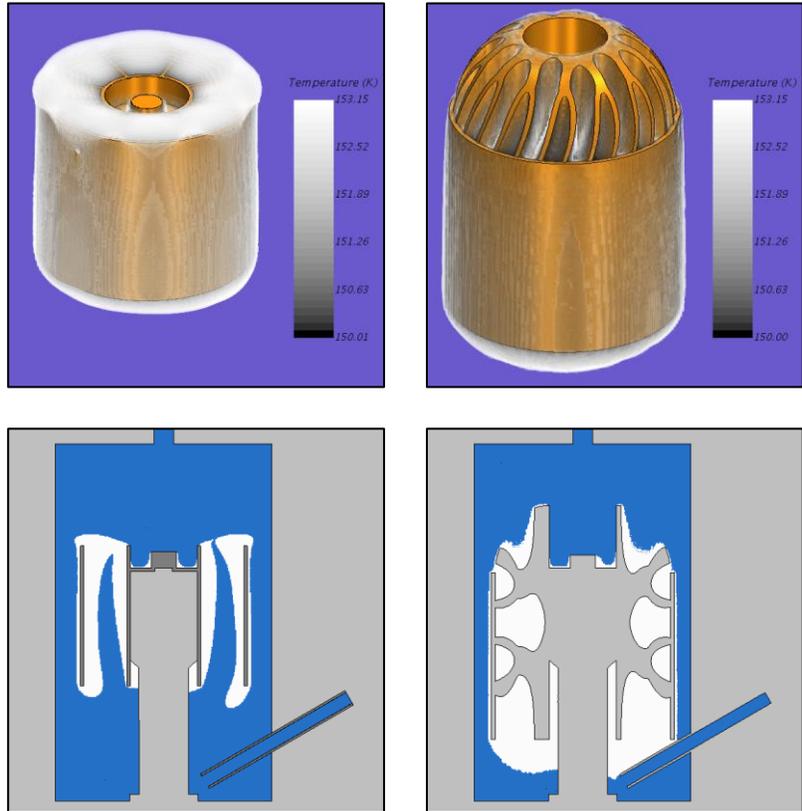


Figure 8. Drawings of the Ferris Wheel cold head (left) and the Branching design (right) with predicted dry ice accumulations at steady-state using the Flow Considered model.

3. Dry Ice Growth: No Flow Considered

In these simulations, the entire domain was treated as a solid region with the properties of gaseous CO_2 . When the temperature was equal to or below 153.15 K, the thermal conductivity was updated to reflect that of dry ice. The thermal conductivity of gaseous and solid CO_2 used was 0.02 W/m-K and 0.622 W/m-K, respectively.

Figure 9 shows that much more ice was predicted to grow on the copper surface for both cold head designs. Using this method, the entire inner volume of the Ferris Wheel was filled in with dry ice, which seemed like a more reasonable prediction. We can assume these simulation results are closer to the maximum ice one can expect from a given design.

4. Dry Ice Predicted Mass

The predicted dry ice mass in grams is shown below in Table 1 for the two simulation approaches described above, assuming a dry ice density of 1562 kg/m^3 . The entry for “Branching (Lattice)” is only an estimate that includes the small amount of ice that would be present assuming the solid shell is replaced with a lattice structure.

The experimental results of the Branching design currently available are from shorter-duration tests which have not fully reached steady-state. A long-duration test (7 h) of the Ferris Wheel gave an estimated dry ice mass of 406 g, which is more than either of the Ferris Wheel simulations predict. One possible factor affecting this could be that the density of the dry ice varies as a function of the deposition rate, which is currently not well known. So while we cannot explicitly quantify the expected dry ice mass for a given cold head geometry, the models may be used to predict an estimated improvement between two designs. Experimental results for the Branching cold head for a 6.33 h test were 492 g of dry ice, which is also greater than the steady state prediction. The experimental results do not include the cooling time to reach 150 K, which was significantly higher for the Branching (Lattice) design (45 min) than the Ferris Wheel (8 min).

Table 1. Comparison of Predicted Dry Ice Mass at Steady State for Two Cold Head Designs

Simulation Type	Ferris Wheel, g	Branching, g	Branching (Lattice), g
Flow Considered	207	296	312
No Flow Considered	339	388	404
Experimental Results	406 (7.0 h)	NA	502 (6.33 h)

The dry ice prediction grew by 64% and 31% for the Ferris Wheel and Branching designs, respectively, when comparing between the No-Flow and Flow simulations (a 29% increase in the Branching (Lattice) design). This indicated that the Ferris Wheel design may be much more susceptible to flow effects than the Branching design. With respect to just the “Flow Considered” simulations, the predicted dry ice for the Branching design is 51% more than the Ferris Wheel design. However, when considering the “No-Flow” case, the Branching design appears to only have a 14% improvement over the Ferris Wheel (19% improvement if the Lattice is considered). The conclusions we draw from these simulations are:

- The model seemed to be better suited to provide cold head *comparisons* as opposed to explicit predictions for an individual design.
- The improvements in the predicted *total* dry ice accumulation between the Branching and the Ferris Wheel designs are on the order of 15 – 50% based on the no-flow and flow considered cases. It is expected the actual improvement would lie somewhere in between these two extremes.
- The model assumed the cryocooler would have enough lift to maintain the entire cold head geometry at 150 K. From experimental results it is clear that the cryocooler had difficulty in maintaining this set point temperature; only near the end of the experimental run did the temperature approach the correct value.
- The relative *rate* of accumulation for each design may be estimated by normalizing the results when compared to completed steady-state experimental runs.

5. *Optimization*

As a complement to the CFD simulations of cold head performance and to address questions of operational cycle optimization, a spreadsheet-based analysis was also developed. We collected time-series data during multiple freezer runs with different coldheads and observed how their CO₂ collection performance varied over the duration. Given a goal of maximizing the total mass collected over an indeterminate operational period, a straightforward analysis can determine the optimal cycle time for a particular coldhead.

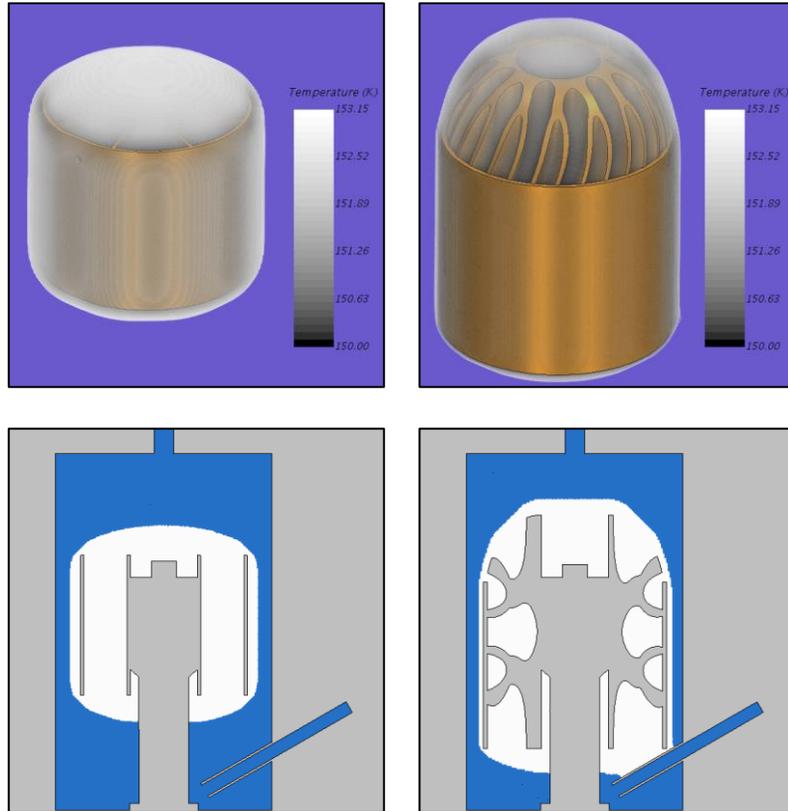


Figure 9. Drawings of the Ferris Wheel cold head (left) and the Branching design (right) with predicted dry ice accumulations at steady-state using the No Flow Considered model.

The collection performance of a coldhead degrades as its surfaces and volumes accrete the relatively insulating solid CO₂ and the total thermal resistance of the system increases. A comparison of three actual coldheads and a dummy “Ideal” design are shown in Fig. 10. (The Starburst cold head was a predecessor to the Ferris Wheel design with less surface area (192 cm²) and slightly higher mass (260 g).²) While the individual curves differ in important ways, they all demonstrate similar macro behavior, with an early period of maximum performance followed by gradual reduction in collection rate.

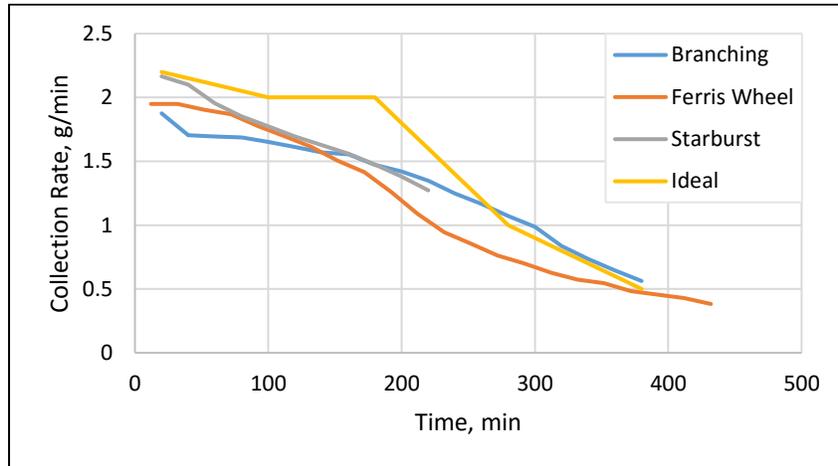


Figure 10. Cold head performance comparison.

It is assumed that in any operational scenario, the cold head assembly will start at a warm initial temperature and require a finite cooling period until it is sufficiently cold to initiate collection of solid phase CO₂. The length of this period mostly depends on the thermal mass of the coldhead, with secondary effects such as parasitic heat leaks to the chamber walls also present. This unavoidable phase with no collection of CO₂, combined with the behavior of the coldhead performance curves, immediately suggests the existence of an optimized process duration that maximizes the average CO₂ collection rate over a cycle.

Figure 11 shows a simplified representation of a notional production cycle, with the associated coldhead temperature change. The durations of each phase of the process (cooling, freezing, warming and sublimation) are not reproduced to scale. However, we do assume that the freezing and sublimation periods are of equal time so that pairs of cryocoolers could operate simultaneously, each on different parts of the cycle, yielding “continuous” production of CO₂.

During the cooling phase, a fixed quantity of time and energy are expended before any product collection can begin. Once freezing commences, the rate of collection inevitably decays as the cold head saturates itself. At some point, it becomes more efficient to end the current freezing phase, begin sublimating, and repeat the cycle than to “ride” the performance curve into the unfavorable region of the plot.

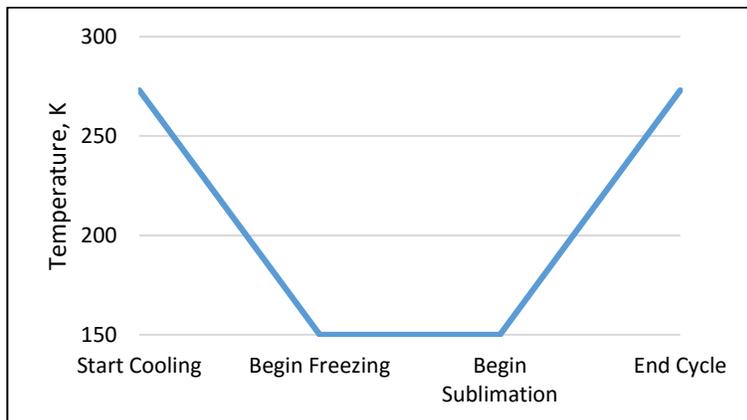


Figure 11. CO₂ collection cycle overview.

Applying this reasoning, it is possible to plot the average collection rate of each coldhead versus the cycle time. That is, the mass of CO₂ collected during the freezing phase divided by the total time elapsed over all phases of the cycle. Figure 12 compares these characteristic optimization curves. Each has a steep early region where the average rate increases rapidly, followed by a broad plateau and eventually a gradual slump back towards the x-axis.

The Ferris wheel and Starburst designs exhibit nearly identical optimal cycle times around 134 minutes, with

freezing phase durations of 70 and 61 minutes respectively. The Branching design is most efficient just beyond the longest duration tested, around 315 minutes, with 135 minutes of active freezing. The differing chilldown times are also evident in the starting locations of each curve. Referring back to Fig. 10, note that this optimization logic dictates that we only utilize the earliest portion of a coldhead's performance profile. Said

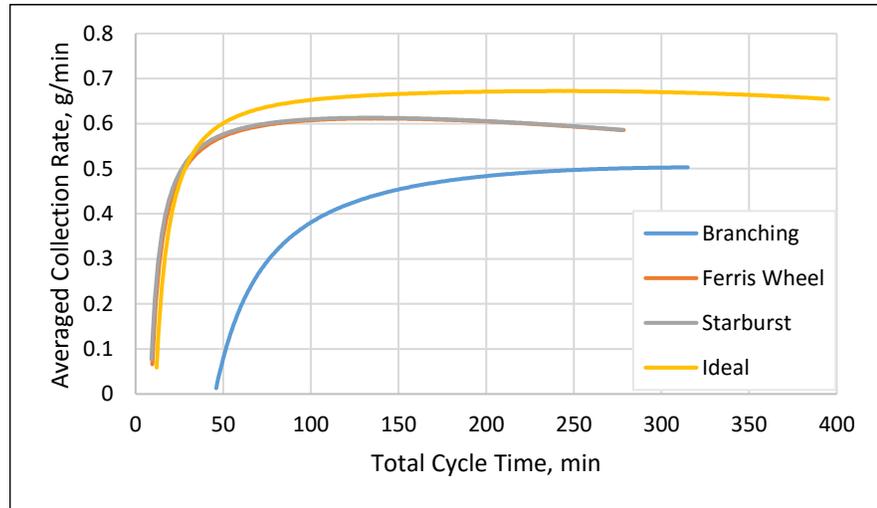


Figure 12. Averaged performance of cold heads at different cycle durations.

another way, once collection performance falls below 5-10% of the initial peak, it becomes more efficient to restart the cycle and eventually begin freezing again at the maximum rate.

Consequently, the initial freezing rate of a coldhead and how long it can maintain this rate are the most critical performance parameters. Also key is a low thermal mass, which minimizes the chilldown time. While both of these factors were deemed important from the beginning, we underestimated their relative value in comparison to other goals like maximizing surface area and maintaining collection performance late into the freezing cycle. The Branching design was created with the idea that relatively long freezing cycles would be typical, and therefore a much larger thermal mass was justified by superior CO₂ collection results in the steady state as indicated by CFD analysis.

Indeed, Fig. 10 shows that after approximately 180 minutes, the Branching design surpasses the Ferris Wheel and Starburst coldheads and continues to dominate them until all approach zero after ~400 minutes. However, the initial time “debt” spent cooling this massive cold head is so large that it can never be “repaid” by later performance. The optimal freezing phase duration for the Branching design barely enters the region where its performance exceeds the Ferris Wheel. The initially lower performance of the Branching design is also interesting, likely caused by greater radiation and convective heat transfer losses to the environment due to the larger geometry and proximity to the chamber walls.

Here, the purpose of the “Ideal” coldhead collection rate profile becomes clear. This curve serves as a strawman target for future cold head designs and envelopes the performance that can be reasonably expected from the current experimental configuration. Therefore, even a design that starts at the maximum collection rate and shows minor performance degradation over the first 180 minutes only yields a ~10% increase in average collection rate as seen in Fig. 12. The optimal freezing duration for the Ideal design is 116 minutes, before large reductions in collection rate occur. This notional performance profile is helpful in restraining expectations of the magnitude of performance gains possible from coldhead design and should guide engineering efforts towards other elements of the system that could provide larger benefit.

The importance of early freezing performance has guided a further iteration of the cold head design, the “Tuning Fork” discussed below. The behavior of the optimization curve also grants a few more insights into a final system concept of operations. The wide performance plateau that surrounds the true maximum collection rate allows making small compromises in CO₂ collection in exchange for other benefits. For example, the average energy efficiency of a continuous process increases significantly as the cycle time is extended due to fewer chilldown phases. Likewise, lengthening the cycle time also reduces the number of power cycles the cryocooler must endure.

Initially, the main factor in determining an appropriate cold head design was based on maximizing the final steady-state accumulation of dry ice. However, it has become clear that ideal cycle times does not require running to steady-state. A more suitable cold head design may feature rapid cool-down times and significant early-stage buildup as opposed to a design that may give higher steady-state accumulation, but much longer cool-down time. The Tuning Fork design was introduced as a concept based on this approach. The overall dimensions (height (5.33 cm) and diameter (6.35 cm)) are the same as the Ferris Wheel design to simplify comparison. Figure 13 shows drawings of the Tuning Fork design with either no or 10 radial cuts to increase surface area and reduce mass.

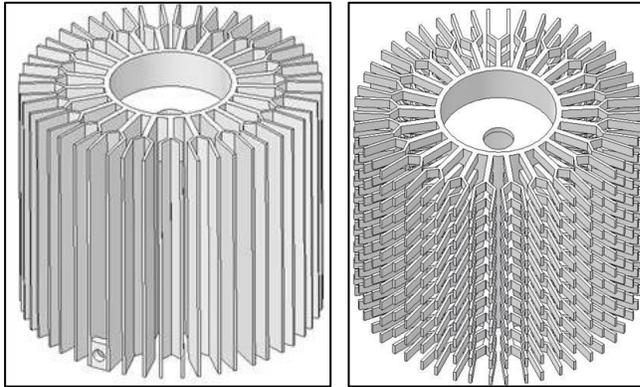


Figure 13. CAD drawings of the “Tuning Fork” cold head with zero (left) and ten radial cuts (right).

The idea behind cutting the Tuning Fork was to reduce the mass of the cold head and to make the design more efficient. As more and more cuts are introduced, the volume rapidly approaches an asymptotic limit of 50% of the original (with no cuts). However, the available surface area of the cold head is reduced until the number of cuts exceeds 32. Despite this initial reduction in surface area compared to the base with no cuts, the total available surface area is still greater than the Ferris Wheel design for any number of cuts. The relative area and volume ratios compared to the Ferris Wheel can be seen in Fig. 14. Construction of the Tuning Fork is underway and the first set of follow-on experimental runs will be conducted using the basic Tuning Fork design with no cuts.

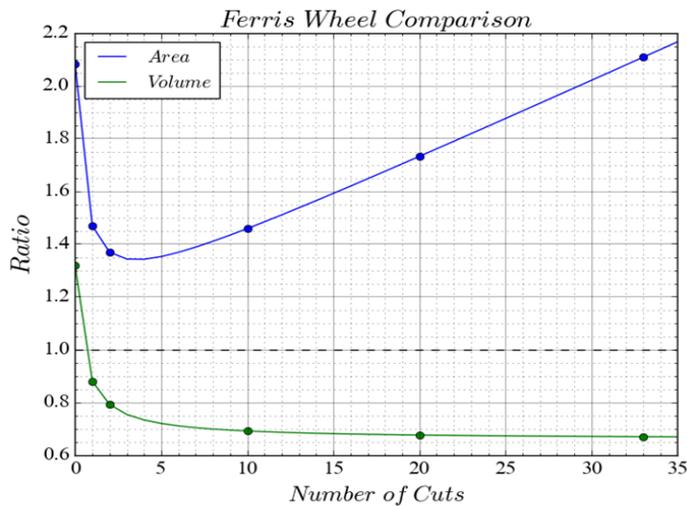


Figure 14. Plot of the effect of the number of radial cuts in the Tuning Fork cold head design on the ratios of its area and volume relative to those of the Ferris Wheel cold head.

B. Sabatier Reactor Temperature Profile Testing

Modeling of the Sabatier reactor requires detailed thermal information. Figure 15 shows the placement of all the thermocouples (TCs) in the Sabatier reactor. The gas stream first runs through the reactor to be preheated (L #68 and L #69) then it is directed into the catalyst bed (L#70-72) inside the reactor where the reaction takes place. L#70-72 allows some of the heat in the feed gas to dissipate to the environment before entering the catalyst bed. The gas stream exits the reactor through L #73. TC14 is inserted into the preheat tube that is inside the reactor, while TC15 and TC7 are at the inlet and outlet of the preheat loop respectively. TC6, 8, 9 and 10 are inside the catalyst bed and measure the temperature of the gases at different positions throughout the bed during the reaction. TC5 measures the surface temperature of the reactor and was not included in this analysis.

In order to develop an accurate temperature profile for the reactor, additional TCs were installed at strategic locations including the inlet and outlet of the reactor and preheat loop as well as multiple locations inside the reactor near the catalyst bed and one TC (TC14) in the preheat tube inside the reactor that had variable lengths into the reactor. The different positions of TC14 are shown in Fig. 16. The naming convention for these positions is shown in Table 2.

A total of four tests were conducted for this purpose, one for each configuration of TC14 (one test had to be repeated due to an early termination of the run). Figure 17 shows the results for the 41.9 cm (“Long”) position of the movable thermocouple and Tables 3 and 4 summarize the results of the test series. Application of these results to the Sabatier model is discussed in a later section.

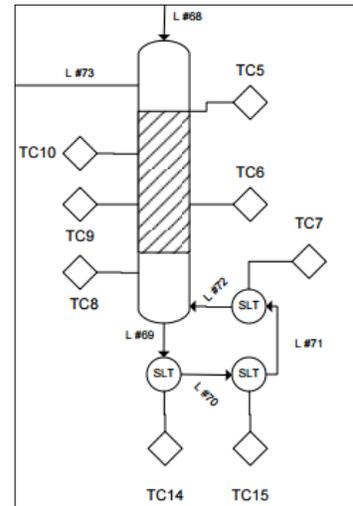


Figure 15. Schematic drawing of the Sabatier reactor.

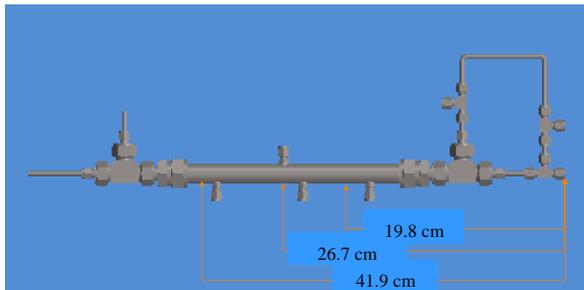


Figure 16. TC14 Locations in the Sabatier Reactor.

Table 2. Naming Convention for TC14

Configuration Name	Length (cm)
TC14-A	41.9
TC14-B	26.7
TC14-C	19.8

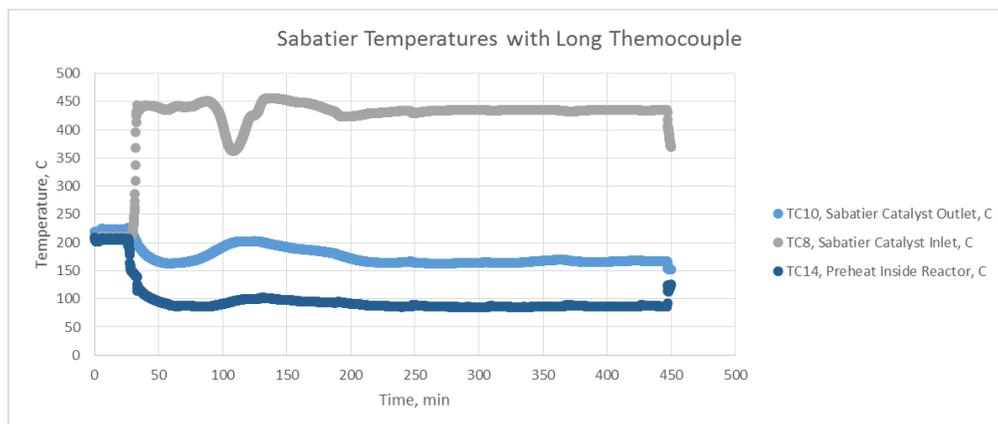


Figure 17. Sabatier Reactor Temperatures for the 9/29/2016 Run

Table 3. Temperature Averages at Steady State for Each Run

Averages at Steady State								
TC14-A	TC14-B	TC14-C	TC15	TC7	TC8	TC6	TC9	TC10
86.83			148.86	159.65	433.41	273.83	259.92	165.44
	192.17		132.20	115.84	418.75	256.24	247.91	167.72
		326.64	143.86	119.59	422.42	260.78	252.95	161.53

Table 4. Temperature Averages at Steady State for All Three Runs

Averages at Steady State								
1	2	3	4	5	6	7	8	9
TC14-A	TC14-B	TC14-C	TC15	TC7	TC8	TC6	TC9	TC10
86.83	192.17	326.64	141.64	131.69	424.86	263.62	253.59	164.90

C. Sabatier Catalyst Damage Study

A 0.5% ruthenium on alumina catalyst in the form of pellets was used for the Sabatier reactor in the APM. The catalyst pellets are commercial off the shelf (COTS) and purchased from Aldrich. Testing has shown that the catalyst provides a high conversion rate in this setup (up to 90% in a single pass), but the operational lifetime may be shorter than desired. The longevity of the catalyst appears to be affected by anticipated, off-nominal operating conditions.

Off-nominal conditions were tested to determine a range of suitable operation parameters for the reactor. During a test with flow rates higher than nominal, the Sabatier catalyst reached 586°C at feed gas flow rates of 1.25 SLPM for CO₂ and 5.0 SLPM for H₂. Following that test, the Sabatier reactor was unable to maintain a steady reaction resulting in temperature swings throughout the reactor. Gas Chromatograph (GC) analysis of the product gas revealed carbon monoxide production after this high flow rate test. The catalyst was removed from the reactor in December 2015 and replaced with unused catalyst. Figure 18 shows that the appearance of the spent catalyst differs from the unused catalyst. The spent catalyst is much lighter in color than the black, unused catalyst, and the spent catalyst contains many broken pellets.

During nominal operation, the reactor undergoes a slow preheat up to 210°C, then the temperature inside the reactor rises to temperatures between 453°C and 467°C within approximately 7 minutes before stabilizing at approximately 450°C at the inlet. The temperature of the reactor typically decreases as the distance from the inlet increases. It is believed that this sudden change in temperature is thermally shocking the catalyst and causing sintering and physical damage to the pellets. The pellets expand in the elevated temperatures, but are restricted by the packed bed reactor which may also be causing physical damage to the catalyst. Sintering can occur in ruthenium at temperatures greater than 500°C, irreversibly damaging the catalyst. Sintering can cause a reduction in catalytic surface area due to crystal growth of catalytic phase, the support area may decrease due to support collapse and pore collapse, or sintering could cause a transformation of catalyst from a catalytic phase to a non-catalytic phase. Okal and Kepinski⁵ showed that at temperatures of 600-700°C, the smallest ruthenium particles migrate and coalesce to form clusters, but a majority of the ruthenium particles remain on the surface similar to the as-prepared catalyst. Ruthenium oxide (RuO₄) can be formed in the presence of oxygen and will cause oxidation of hydrocarbons.



Figure 18. The spent catalyst (left) was removed from the Sabatier reactor after it no longer performed as expected. It is much lighter in appearance and has many broken pellets compared to unused catalyst (right).

After damaging the catalyst in high flow rate testing, a study was started to determine the cause of catalyst failure that occurred after approximately one year of intermittent, successful Sabatier reactor operation.

D. Catalyst Testing

The catalyst was tested in conditions similar to those that occur in the APM reactor chamber to determine the cause of the catalyst failure. To mimic the conditions in the reactor, the catalyst went through thermal shock cycles packed in a small stainless steel tube. Unconstrained catalyst pellets in an open crucible were also exposed to the same thermal shock cycles to act as a control. The diagram in Fig. 19 shows the catalyst packing configuration for testing. The stainless steel tube had a 1/2" OD and was three inches long. The ends were packed with glass wool in the first five cycles; however, the high temperatures caused degradation of the glass wool. The wool was replaced with glass tape for the remaining five cycles. Each catalyst tube underwent a total of ten thermal cycles.

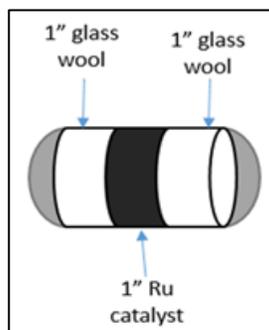


Figure 19. The catalyst was packed into a 1/2" OD, 3" long stainless steel tube in order to mimic the constricted conditions in the Sabatier reactor.

For the thermal shock testing, a Whip Mix Pro 100 ceramic furnace was used because it had the capability raising a platform into a preheated furnace providing a thermal shock to the catalyst placed on the platform. The furnace used a vacuum and nitrogen gas line to provide an inert environment for the catalyst during testing. The catalyst arrived in an oxidized state, but to prevent further oxidation and to more closely match the conditions in the reactor, the inert environment was used during each run.

Two packed tubes like those shown in Fig. 19 were initially made, one for each thermal shock temperature. One tube was cycled at 450°C, the nominal operating temperature of the reactor, as a control. The other tube was cycled at 600°C, the maximum temperature reached before the catalyst in the reactor failed. Each heating cycle consisted of a furnace preheat to the desired temperature, the catalyst tube and free pellets, placed in a stainless steel crucible, were set on the furnace platform and mechanically lifted into the heated furnace. The catalyst remained at the elevated temperature for 20 minutes, then the furnace reduced temperature over approximately 2 hours until the platform lowered at 200°C. The tube and loose pellets were weighed before and after every cycle.

After the first five cycles, each tube was unpacked, and all the pellets were analyzed using a scanning electron microscope/energy dispersive X-ray (SEM/EDS) and optical microscope. The same catalyst was repacked in the tubes, and after five more cycles, they were analyzed using a field emission scanning electron microscope (FESEM), SEM/EDS and X-ray photoelectron spectroscopy (XPS). A third trial was conducted with fresh catalyst packed in a tube in the same manner as the first two tubes and run through ten thermal shock cycles at 450°C. The pellets inside the third tube were analyzed visually.

E. Catalyst Testing Results

After the first five thermal shock cycles, the catalyst pellets from the 600°C and 450°C test groups were compared to unused catalyst and the spent catalyst that had been removed from the reactor in December 2015. First, the pellets were photographed using an optical microscope at an apparent magnification of 55x (Fig. 20) to examine the condition and color of the pellets. The catalyst that had been removed from the reactor in December was much lighter in color than the unused catalyst. The 600°C cycle pellets appear to be slightly lighter in color, and the 450°C cycle pellets are similar in hue to the unused pellet. The SEM/EDS analysis did not reveal any trends.

The pellets were analyzed again after 10 thermal shock cycles. The pellets were observed to be physically damaged. One pellet split in half at the lower thermal shock temperature.

The catalyst pellets were examined under SEM/EDS using the EDS mapping. The results can be seen below in Fig. 21. Unused catalyst was compared to the catalyst removed from the reactor in December. The ruthenium element maps do not show a trend at 300x or 1000x magnification.

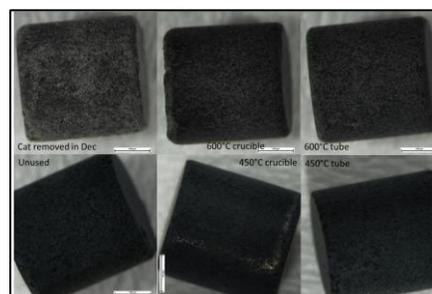


Figure 20. The catalyst pellets were photographed after five heating cycles using an optical microscope at an apparent magnification of 55x.

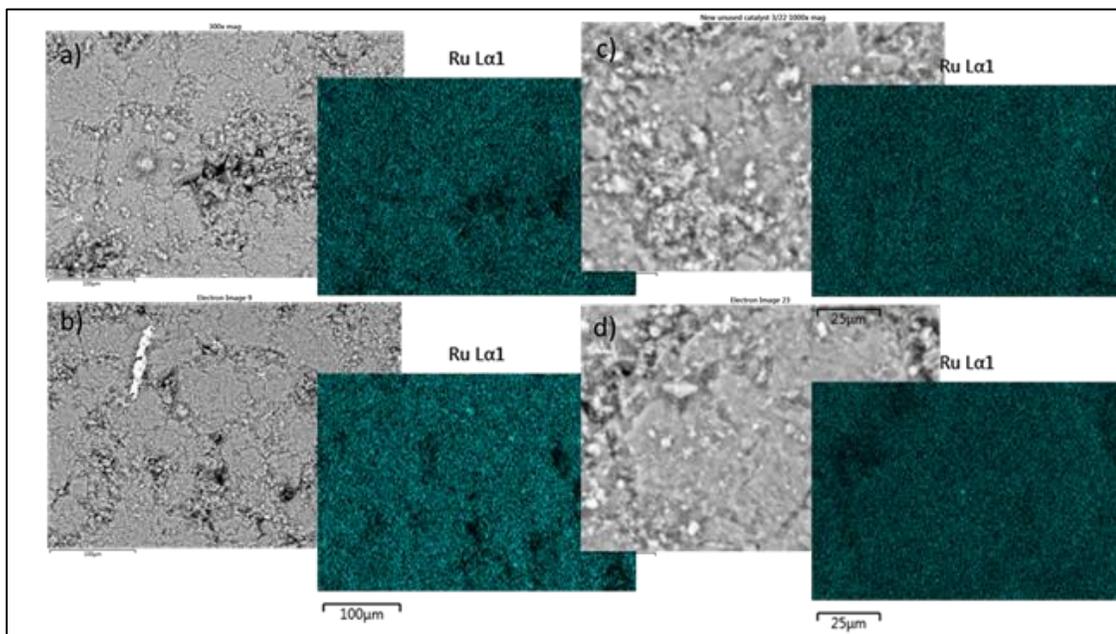


Figure 21. The ruthenium content on the catalyst pellets was mapped using SEM/EDS. The SEM image for each magnification is overlapped by the EDS ruthenium map of that location. a) Used December 300x magnification b) Unused Catalyst 300x magnification c) Unused Catalyst 1000x magnification d) Used December 1000x magnification.

The catalyst pellets from the thermal shock trials, unused catalyst, and spent catalyst were imaged with the SEM/EDS at a very low magnification. The images in Fig. 22 show that the condition of the edges of the catalyst pellets and the surface condition of the pellets do not exhibit a trend in physical appearance changes or damage related to pellet use conditions.

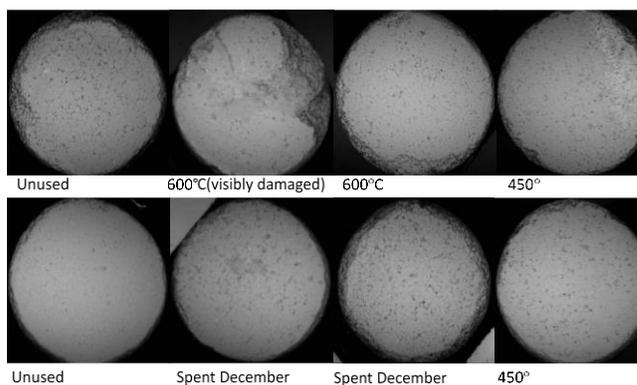


Figure 22. The catalyst pellets were imaged using the SEM/EDS on low magnification. There does not appear to be a trend in the amount of physical damage to the pellet compared to the use conditions of the pellet.

A FESEM was used to collect images of the pellet surfaces at 15,000x magnification of the unused pellets and spent reactor catalyst pellets. The images of spent reactor catalyst have disperse, white spots covering the surface. The unused catalyst images have fewer or no white specks. Representative images of this trend can be seen in Fig. 23.

The last test carried out on the ten-cycle pellets was XPS analysis. Sample pellets from each group (unused, spent catalyst from the reactor, 600°C thermal shock tube, and 450°C thermal shock tube) were analyzed. The results show that, in general, for the surface scans, the new catalyst has two peaks in the aluminum and one in the ruthenium binding energies, whereas the used catalyst show one peak in the aluminum and two or more peaks in the ruthenium binding energy. Representative images of the trends in bond energies of surface peaks are displayed in Table 5.

A third tube of fresh catalyst was prepared and tested at 450°C for 10 thermal shock cycles. This tube also exhibited pellet damage (Fig. 24).

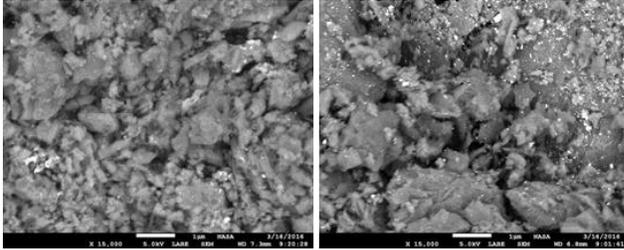


Figure 23. The unused catalyst (left) shows fewer bright specks than the spent catalyst (right) when imaged at 15,000x magnification on an FESEM.

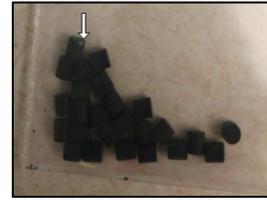


Figure 24. The 450°C thermal shock experiment was repeated with fresh catalyst. The second experiment also resulted in physical damage to the pellets.

Table 5. The surface peaks from XPS analysis of pellets representing the trend of one peak in aluminum and two or more peaks in ruthenium for the heated pellets and two peaks in aluminum and one peak in the ruthenium for the unused catalyst.)

Type of Catalyst ↓ Peak →	Aluminum 2p	Ruthenium 3d
Unused Catalyst		
450°C Heat Shock		
Spent Catalyst Removed December 2015		

F. Catalyst Testing Discussion

During the thermal shock trials, the pellets packed in the stainless steel tube showed signs of physical damage, but the unconstricted pellets that were heated in the crucible did not show signs of physical damage. The packing procedure may have played a role in the pellet damage, but the majority of the observed physical damage is likely due to the constriction caused by the pellets expanding when heated and pressing into each other and the walls of the tube. Prior to the thermal shock trials, it was believed that the damage to the catalyst pellets in the Sabatier reactor was caused by the 600°C temperature anomaly, however, cycling the pellets at 450°C, the nominal Sabatier reactor operating temperature, caused substantial physical damage.

The observed change in pellet color, turning from black to gray, occurred more rapidly in the catalyst cycled at 600°C. It is unknown if the gray color is indicative of catalyst degradation or failure, but the catalyst removed from the reactor were starkly lighter in hue than unused catalyst as seen in Fig. 20.

The lower magnification SEM mapping does not show an apparent difference between the unused and spent catalyst; however, the high magnification, FESEM images show a trend of an increase in the number of white specks in the spent catalyst. The white specks are believed to be ruthenium. The brightness of the spots indicates that it is an element much heavier than the oxygen and aluminum that make up the alumina support pellet. The brightness of these spots rules out silica from the glass wool packing material. Some EDS analysis and the test configuration indicate the specks are ruthenium and not copper from the cold heads. The visible white specks on the spent reactor catalyst could be caused by ruthenium sintering and coalescing into clusters that are large enough for the FESEM to detect. Fewer visible ruthenium specks on the unused catalyst could indicate that the ruthenium is still very dispersed across the surface of the pellet.

One of the potential sources of catalyst poisoning is the thermal paste used to mate the cryocooler cold finger with the copper cryocooler cap. The compound used, Céramique 2, is likely not a poison due to its chemical composition of synthetic oil, aluminum oxide, boron nitride, and zinc oxide, which are not known to be ruthenium catalyst poisons. The thermal paste had to be reapplied when the heater was repaired in cryocooler #2 meaning that the previous application had dispersed from the cryocooler cold finger and passed through other components of the APM after it left the cryocooler. It was a small amount of product, but its effects on the other components are unknown.

The XPS results show that aluminum and ruthenium bond energy spectra at the surface of the pellet are consistently different on the unused pellet compared to the pellets that had undergone thermal shock. The spent pellets from the reactor and the 450°C and 600°C thermal shock pellets showed very similar peak trends to each other. This indicates that the high temperatures experienced by the pellets or the rapid increase in temperature was the driving factor in the change and not other factors such as poisoning. Further analysis is underway.

The results from the XPS analysis indicate a clear trend between the used and unused catalyst. Analysis is complicated by the fact that the carbon peak overlaps the ruthenium peaks. Some literature is available on typical ruthenium on alumina binding energy peaks, but more work is needed to interpret the results.

If the alumina pellets that currently support the ruthenium catalyst in the packed bed cannot withstand the operating conditions, other options will need to be explored. A miniature reactor could be assembled for testing various catalysts at the upper and lower bounds of operating conditions. Various sized pellets and other support structures with heavier catalyst loads could be tested. A heavier catalyst load with a robust support structure has been successfully implemented by Hamilton Sundstrand for the Sabatier reactor currently in flight on the ISS. If the pellet damage and heat damage is caused by the thermal shock experienced by the pellets, the gas flows rates to the reactor could be increased more slowly so that the temperature of the pellets ramps up more slowly and prevents the formation of hot spots.

IV. Sabatier Reactor Modeling

Extensive work was performed to model the Sabatier reactor in the APM. This work has already been published and was reported at the 47th International Conference on Environmental Systems (ICES 2017).⁵ The conclusions of this study follow:

A thermal and kinetic model of the APM Sabatier reactor and small scale Sabatier reactor was created using Python. The model was developed after a catalyst performance failure occurred in the APM Sabatier reactor in 2015 during higher than nominal flow rate testing. Since high flow rates will damage the catalyst, the model will help optimize the throughput of the catalyst bed, and be validated experimentally to build trust in the model. The model will continue to evolve including considering transport by convection in the fluid, diffusion inside the catalyst pellets, reaction kinetics, and thermal effects of the gas-phase and solid-phase aspects of the reactor system. The reaction diffusion investigation of the catalyst will be coupled to the mass and energy balances of the fluid, as well as the

pressure drop considerations across the reactor. The continued development of this modeling will allow us to optimize systems that can predict and determine acceptable flow rates and temperature profiles in order to avoid damaging the catalyst and troubleshoot system tolerances. It will also help with scaling larger systems or systems in series for future design of a Mars ISRU propellant production system.

Ru/Al₂O₃ catalyst characterization analysis resulted in subtle changes that were observed in both XRD [X-Ray Diffraction] and XPS [X-ray photoelectron spectroscopy]. The XRD crystal structure on the pellets that encountered higher reaction temperature (>500 °C) were missing RuO₂ characteristic reflections, and the cooler portions of the catalyst bed (reactor middle and outlet) had reflections corresponding to RuO₂, while the reactor inlet did not. The XPS observations revealed that a stronger oxide layer was present on 'healthy' or unused catalyst as well, but not present on the damaged catalyst. If the phase change from γ -Al₂O₃ to α caused pores to collapse and obstruct the surface layer of Ru, this may have been why performance for Sabatier selectivity and reaction temperatures declined. Overall, catalyst exposure to reaction temperatures greater than 450 °C caused the thermal shock on the catalyst surface, losing the ability to selectively form CH₄.

Work on modeling smaller and larger Sabatier reactors is underway to better understand their dynamics and will be reported at future conferences.

V. Conclusions

The CAD modeling of the APM has provided the necessary information for the modeling of the CO₂ Freezer and Sabatier components. Modeling of the CO₂ freezing process has provided great insight into ways to optimize the rate and the mass of dry ice collected by precise design of the copper cold head used. Testing of alternate cold heads has generated the data needed to validate the models to ensure their utility for designing full-scale CO₂ freezers for robotic precursor and human missions to Mars that will be built and tested in the next few years.

The measurements of the steady state temperatures in the Sabatier reactor were crucial to verifying the validity of the reactor models. Results of the modeling were reported in a separate report and the models have been successful in reproducing the experimental data for small Sabatier reactors. The models will now be used to design full-scale reactors for Mars missions as well.

The Sabatier catalyst pellet damage appears to be caused by elevated temperatures or thermal shock and not poisoning or interference by the alumina support structure. Stainless steel tubes and the furnace could be used to test the effectiveness of ramping the temperature more slowly to the reaction temperature, and simple changes to the procedure could be made to insure the reactor temperature does not increase too rapidly. A miniature reactor could be constructed for testing other catalysts, catalyst loading weights, or catalyst supports if the current conditions will not meet flight requirements.

Work is underway to simultaneously operate the APM and other Mars ISRU Pathfinder modules at KSC in the near future, including data sharing between the modules to continue the development of the system, determination of any significant interactions of the modules, and guidance for the designs of future systems.

Acknowledgments

We thank the many KSC interns who have contributed their expertise and enthusiasm to this project. Dr. Muscatello also greatly appreciates the civil service labor funding provided by the Space Technology Mission Directorate (STMD) Game Changing Development (GCD) Next Generation Life Support (NGLS) project as well as contractor labor and materials funding by several KSC Continuing Technical Competence (CTC) and Independent Research and Technology Development (IR&TD) projects.

We also especially thank Omar Mireles and Zachary Jones of MSFC for using their Additive Manufacturing capabilities to produce alternate cold head designs, which are essential to the success of the CO₂ Freezer modeling efforts.

References

¹Interbartolo III, M.A., et al. "Prototype development of an integrated mars atmosphere and soil-processing system." *Journal of Aerospace Engineering* 26.1 (2012): 57-66.

²Muscatello, A., Devor, R., and Captain, J., "Atmospheric Processing Module for Mars Propellant Production," American Society of Civil Engineering Earth and Space 2014 Conference, 2014, pp. 1-46.

³Muscatello, A. C., Hintze, P. E., Meier, A. J., Bayliss, J. A., Karr, L. J., Paley, M. S., Marone, M. J., Gibson, T. L., Surma, J. M., Mansell, J. M., Lunn, G. M., Devor, R. W., Captain, J. G., and Berggren, M., "Mars Atmospheric

In-Situ Resource Utilization Projects at the Kennedy Space Center,” American Society of Civil Engineering Earth and Space 2016 Conference, 2016.

⁴Ellis, D. L., Keller, D. J., & Nathal, M. (2000). Thermophysical properties of GRCop-84, NASA/CR-2000-210055.

⁵Okal, J., & Kępiński, L. (2009). Sintering of Colloidal Ru/ γ -Al₂O₃ Catalyst in Hydrogen. *Catalysis letters*, 128(3-4), 331-336.

⁶Meier, A.J., Shah, M.G., Hintze, P.E., Muscatello, A.C., and Petersen, E., “Mars Atmospheric Conversion to Methane and Water: An Engineering Model of the Sabatier Reactor with Characterization of Ru/Al₂O₃ for Long Duration Use on Mars,” 47th International Conference on Environmental Systems, ICES-2017-161, 2017.