Formulating the CO_2 to Glucose Challenge

Ademir-Paolo Vrolijk, Ph.D.

The George Washington University

avrolijk@gwu.edu

April 2022

This document summarizes the formulation process of National Aeronautics and Space Administration (NASA) CO_2-to-Glucose Challenge. It highlights the important decisions that shaped the problem that participants would solve when they competed. These decisions were primarily made by NASA's subject matter experts in related fields of CO_2-based manufacturing. This document also summarizes the Challenge's outcomes viewed through the formulation lens.

The Challenge aimed to find and demonstrate an efficient pathway of converting CO_2 to glucose, a conversion that would be highly valuable during long-duration stays on Mars. The challenge launched in 2018 and ended in 2021.

Contents

1 NASA's Technology Goals

1.1 Using CO_2 as an In-Situ Resource

NASA plans to land astronauts on Mars in the 2030s. This is an expensive endeavor especially considering the infrastructure and consumables needed to keep the crew alive. To address this issue, subject matter experts (SMEs) across the agency are investigating how resources on Mars could be used to create the needed products instead of transporting them from Earth. These systems would reduce the launch costs and provide the crew with a degree of self-sufficiency [CO1].

CO_2 could be a critical Martian resource for this endeavor. During their stay, the crew will need organic consumables like pharmaceuticals, nutrients, adhesives, and fuels [NASA Centennial Challenges Program (CCP)149, CO3, CO1]. Here, the carbon atoms in CO_2 can form the building blocks for these products. SMEs at NASA Ames—drawing on their expertise in synthetic biology, regenerative life support, and CO_2-based manufacturing in the space context [CO3, CO23]—propose engineered bacteria to convert the CO_2 into those products [CO1, CO3]. These bacteria will need a source of energy, and glucose—itself an organic compound—is one of the best candidates for their "food" [CO1, CO3, CO4]. As such, the ability to convert CO_2 into glucose—and other valuable sugars—is a useful capability: it can enable the crew to create various complex products locally.

1.2 Developing Efficient Pathways to Convert CO_2 into Glucose

At a high level, the design of a CO_2-to-glucose conversion system involves two questions. First, how will this system fit with NASA's (planned) Mars infrastructure [CO23]? The space context imposes constraints on any technology, and this would be no exception. CO_2 manufacturing SMEs describe two important parameters at this early stage of development [CO23]. First, volume. If the conversion system were the size of a room, it would not be feasible [CO22, CO23]. Second, power. If the system required megawatts of power to operate, it would not be feasible either [CO23]. So, the system's implementation on Mars imposes limits that need to be taken into account.

Second, what conversion method will the system use? A bioreactor *could* perform the function of turning CO_2 into glucose. In fact, a biological approach would be easier: producing glucose via plants, microorganisms, or enzymes is common in various industries [CO2]. And organisms can, likely, be engineered to perform this task by leveraging other technologies like CRISPR and gene editing [CO23]. However, any biological system has drawbacks that make it a less desirable option for Mars[1]. First, they are large: the system needs large tanks with liquid for the organisms in the reactors [CO23]. Second, they lag: starting and stopping the organisms from producing their products can take a lot of time, which means that controlling the process is difficult [CO3, CO23]. Lastly, they are fragile: the reactor's conditions need to be closely controlled to keep the organisms alive [CO3]. Faced with these drawbacks, the SMEs wanted to try a different approach for producing glucose.

Specifically, the NASA SMEs wanted to know whether a physiochemical system would be feasible. Such a system would be very valuable: compact, fast, efficient, responsive, and robust [CO3, CO4, CO23]. One subject matter expert (SME) succinctly described how he thought a hypothetical system would look like:

> It would be a single or two-step process with little to no waste. Incredibly efficient in terms of bond breaking— The energy needed to break and make bonds. And [it would be] highly reliable, doesn't ruin your catalysts, and doesn't get really dirty and you need to wash the whole thing with an acid. [CO22]

But this approach had its own risks. Previous work has shown the conversion of CO_2 into other carbon molecules, but only into products with less than six carbon atoms—which glucose has [CO12]. Additionally, any conversion from one molecule to another requires more and more energy. So, while the SMEs acknowledged that a "college student" could "hopscotch" their way from CO_2 through the different intermediate products to glucose

[1] Note that the crew would still need bioreactors, and associated bacteria, to manufacture the more complex products. But reducing this system's footprint by reducing its dependence on biological approaches was considered valuable.

[CO10], finding a pathway that's efficient and not wasteful was the big issue. Additionally, physiochemical conversion was many people had researched or implemented [CO24, CO2]. For Earth applications, one does not need bacteria food from CO_2—there are many cheap, biological sources for this [CO2, CO3]. Per one SME: "very few people have done anything of strong significance. It's all very new" [CO3]. As such, the approach would be "extremely hard" [CO23, see also CO26] but also "not economically favorable" in the SMEs' eyes [CO3].

2 Opening the Conversion Problem

2.1 Betting on Different Approaches and Outsiders

There were less risky options to feed bacteria for NASA's aims—both in approach and product. Converting CO_2 into bacteria food is a complex problem [CO3], and SMEs considered solving it through their regular innovation or problem-solving funding channels. But they would only be able to make "some incremental changes" through these channels with known players [CO23]. For example, while glucose is "the gold standard" for energy for bacteria, other products could perform that role as well [CO10]. Among these, acetate is "one of the better products" [CO10]. But this is a tradeoff between performance and uncertainty. Acetate contains less energy than glucose, which means it is not as good a food source. However, it is easier to convert from CO_2: there is less uncertainty in pursuing the acetate route. Thus, acetate would be a "very likely outcome that we can use, versus [glucose,] an unlikely one that would be great if it did work" per one of the SMEs [CO10].

Nevertheless, SMEs decided to bet on *both* acetate's safe bet *and* the glucose long shot—forming a "suite of approaches" to address this problem [CO23]. The same SMEs started a "collaborative agreement" with a lab at Stanford for the former [CO3]. The external partner chased the acetate conversion, and has since broken records for efficiency and yield [CO10]. Separately, a collaboration of several universities and institutes began working on interrelated biomanufacturing projects, including the CO_2 to acetate pathway. Led by UC Berkeley, they won a multi-year, multi-million-dollar NASA grant to do this work [CO10].

For the latter, the SMEs wanted to try something different. Here, SMEs decided that a challenge would be an appropriate avenue for this problem for several reasons [CO3]. First, while the function of converting CO_2 into usable products was not novel, no one had yet developed a physiochemical pathway of doing this efficiently. SMEs knew they needed to push the field in the direction of converting CO_2 to glucose, encouraging or incentivizing the right people and their institution. They hoped to jumpstart widespread commercial activity on these kinds of conversions [CO3]. Second, SMEs did not know who may have had a potential solution to this problem. CO_2 conversion is a nascent field; people with relevant expertise could have been in a "business, or within the academic realm, or wherever" [CO3]. A challenge would reach more people (and organizations), especially those "not traditionally part of the NASA stakeholder base" [CO23]. Third, SMEs wanted to encourage momentum behind the problem that would sustain it "financially, legally, [and] politically" [CO23]. A challenge would connect to the public, shine a light on the issue, and get interested parties to form a long-term ecosystem around it [CO3, CO23]. SMEs expected solvers to "become part of [the] journey [and] come along with [us]" [CO23]. So, high-level discussions at NASA headquarters decided to go the open innovation route [CO2], and "throw the challenge out to the world and see who's been thinking about this" [CO3].

While the challenge would reach people people "from everywhere" [CO23], SMEs did not expect whomever to solve this problem [CO2, CO23]. They had a set of people in mind that had a better chance of solving it. These were outside the aerospace industry and, predominantly, in green chemistry:

> People in the green chemistry arena, we think, will be the most interested in this. We've reached out to several companies. If you look, there's kind of an XPRIZE challenge right now, using CO_2 as a resource. It's a much larger and less pointed challenge than what ours is. People within that realm of expertise. There are [the] Green Chemistry societies— It's chemical engineers, particularly people who are looking at CO_2 conversion technologies [CO2].

The SMEs' betting strategy was about spreading the risk and maximizing the chances of success. Funding the partnership on acetate was very likely to yield good results. But there would be a chance that the challenge would make some progress as well. They could incorporate this pathway into NASA's regular innovation channels if it did. One SME described how that would occur:

> Now that being said, if anyone in the challenge starts to make this in a decent way and– And it's not a guarantee, I'm just saying that it's possible that we could look at the winner or winners and we could say, "gee, we would like to help you keep moving forward on that." Find a collaborative way to keep moving forward on that as well. [CO10]

2.2 Scoping the Challenge's problem

After deciding to pursue a challenge, SMEs' next hurdle was scoping the challenge problem. Deciding on the challenge problem and how to measure its solutions was not easy [CO33]. This process included several teams: HEO Advanced Exploration Space (AES), NASA Ames, the CCP, and ad-hoc input from SMEs at DoE and National Academy of Sciences [CO4, CCP149]. Here, NASA's senior SMEs on the topic held the most sway in these decisions. They "[would] drown out other voices" in this discussion [CO25], ensuring that the challenge would be in line with the technical direction they believed was most promising.

To make sure the problem was possible at all, the SMEs relied on first-principle calculations and previous work done by NASA and others. They started their scoping by calculating what manufacturing rates of glucose were theoretically possible under the conditions imposed by the hypothetical settlement on Mars. These would dictate the order of magnitude for the system's footprint: how large the system would be and how much power it would need [CO23]. They also compared these estimates to conversions to intermediate products to gauge whether they were in the right ballpark [CO23]. These were all to ensure that the problem was not physically impossible from the outset. If it exceeded the size and power upper bounds by a lot, then it would not be a good path to pursue.

With the feasibility of the problem established, formulating a feasible challenge was the next task. There were three aims that the SMEs were trying to balance: addressing NASA's need, encouraging non-traditional activity on this topic, and judging solutions both fairly and accurately [CO2, CO3, CO23]. The need to balance these three had a profound impact on the challenge problem.

2.3 Focus Areas for the Formulation Process

2.3.1 The Challenge's Deliverables

With these aims in mind, SMEs considered what solvers would deliver in response to the challenge. Their initial idea was to ask solvers for a "plan" to create the conversion system [CO3]: requiring a description of their pathway from CO_2 to glucose, with the appropriate analysis to back that up.

Initially, this deliverable was suggested *instead* of going directly to a (prototype) production system [CO3]—what SMEs wanted in the first place. SMEs thought people or organizations with little background might want to give the challenge a try. And since there would not be a way to transport the reactors to one NASA site, the judges would potentially have to spend resources to "[go] to places or [deal] with products that are just not ready for primetime" [CO3]. "That can create[d] an administrative burden" that they did not want to bear, per one SME [CO3].

However, this deliverable would only resolve so much of the uncertainty of the solution; it did not demonstrate that the solver's plan could actually work [CO10]. As with the hopscotch example, people could describe pathways that make sense on paper. But a paper solution alone would leave much of the implementation uncertainty unaddressed [CO22]. SMEs mentioned several issues that could differ between plan and demonstration: uncertainties, and limits, in the workings of different catalysts; micro-interactions of compounds creating unwanted products; changes in system behavior with temperature changes and or the presence of oxygen; and uncertainties in behavior depending on *how* materials are introduced in the reactor [CO3, CO22]. In short, describing the pathway and its system would show promise and build confidence in the approach, but a demonstration would resolve much more of its uncertainty.

SMEs decided on two competitions: the first to plan the conversion system (Phase 1) and the second to build it (Phase 2). In addition to the uncertainty around the quality of the solutions, having to self-fund the whole problem would be expensive per the SMEs [CO3]. Having two separate competitions gave SMEs an opportunity to award a (small) prize to the plans that looked the most promising [CO3]. This would give a leg up to solvers who might not be as well funded as other teams but still might have solid ideas for solving this problem [CO3]. A two-phased challenge also would provide a gate to screen for solution quality and teams' technical abilities in the first round [CO12]. Thus, both sides were primed for a more complex second phase: SMEs, through a better picture of the participants and their solutions; and solvers themselves, through their work on their plans. With the difficulty of the conversion problem in mind, SMEs wondered who, and how many, would show up to solve the problem: "if they're heavily funded large industries or if they're academicians or small start-ups" [CO3]. These discussions also involved the challenge requirements, which could have influenced participation: e.g., the mass and phase (solid or liquid) of the sample and the footprint of their system [CO9]. CCP would release the challenge phases in a staggered manner: Phase 1 on its own and Phase 2 at a later date.

To better understand the operation of the solvers' systems, judges would conduct site visits in Phase 2. SMEs expected the challenge systems to be "large and complicated" [CO12]. Unlike other Centennial Challenges, transporting these to a central location might not be feasible [CO12]. Instead, the challenge judges would fly out for a site visit, seeing the solvers' operation and output in person [CO12, CO4]: During this visit, solvers would have up to seven hours—the estimated maximum length of the judges' stay on-site [CO9]—to create their sample [CO12]. Having the judges verify the operation of the system and the contents of the samples would reduce the uncertainty in the solutions [CO4]: providing "proof that we're seeing CO_2 go into it, and we're seeing product come out, and we're going to know what that product is" [CO3].

Given these uncertainties, the formulation team decided that they would use the performance of the solvers in Phase 1 to set "realistic performance criteria" for Phase 2 [CO9]. Getting external input—before solvers actually solved the problem—would allow them to make better decisions on what solvers could do [CO3, CO9, CO2]; this would allow them to shape Phase 2 to be successful. And while the requests for information provided NASA with some feedback here, having them provide all that information in a formal challenge deliverable would be "the best way to obtain the necessary information" [CO9]. The delay between Phase 1 and 2 would also allow changes to be made without the paperwork—and potential embarrassment—of making big changes to the rules after they had been released [CO9]. The prize purse for this challenge was a total of $1M [CO9]. Prizes for Phase 1 and Phase 2 would be $50k and $750k, respectively (up from $500k initially) [CO4, CO5, CO9, CO19]. The bonus round in Phase 2 set aside $100k of the $750k prize for a separate award. Winning (or even participating) in Phase 1 was not required to participate in Phase 2.

2.3.2 The limits on the Footprint of the System

Initially, the SMEs imposed numerical limits on the system's footprint to fit the Mars implementation. Early drafts of the rules gave specific volume, power, and mass upper limits: if solutions exceeded these, they would not be valid [CO1]. Subsequent drafts removed or relaxed these numbers almost completely: stating that the system should fit within 25 ft^2 [CO6], and later 100 ft^2 [CO9]. In the end, the Executive Program Management Council (EPMC)—the challenge's final request for authority from NASA senior management to proceed—removed the limits on the system's footprint altogether [CO9].

Two reasons contributed to the decision not to specify any limits. First, judging the footprint fairly across different kinds of systems proved to be a difficult problem. Even excluding biological systems, there were many pathways—with as many processes—that solvers could create to perform the conversion [CO22]. These processes involved specific infrastructure, with their own space and power requirements, resulting in a wide range of potential systems. SMEs were concerned that consistently measuring the footprint of the varying solutions would be hard [CO22], and might not result in a fair judging process. At worst, this could have biased solvers towards certain processes—and thus, certain solutions—at too early a stage instead of focusing on glucose production [CO22].

Second, SMEs were afraid of imposing limits that would disqualify teams from getting close—but not quite achieving—the conversion goals. While size and power requirements would be front of mind for NASA's applications, SMEs did not want to dismiss potential solutions. In particular, they were concerned that these would hamper solvers so much that they would not complete the challenge, or not participate at all [CO22]. So by removing the explicit footprint limits, they could avoid these issues of fairness in judging and restricting solutions. Here's one SME describing their concern:

> [I]n doing the Centennial Challenge, you don't want to push it in such a way that it becomes impossible. That people say, "I could have done it, but it became so impossible that I couldn't do it. I could have gone 50% of what they're asking for." I don't want to eliminate that in the first step. [CO23]

Instead, reducing (or optimizing) the system's footprint would occur later. In deciding to remove the limits entirely, the attendees at the EPMC felt that Phase 1 should be easier for participants [CO2, CO22]—allowing them to focus on the conversion itself. Phase 1 would then "be more of a casting stage—spreading the net, seeing how many fish you can get" [CO2]. Simply having (non-optimized) estimates of what the footprint of such a system could be was valuable. Per one SME, "the main focus is: can [a participant] even do it. And if you do it, tell me what is the footprint" [CO23]. In that vein, the EPMC also suggested that the footprint limitations could be accommodated later, as "NASA will have time to work on the footprint limitation in the future when missions are defined" [CO9]. Here, the SMEs agreed with this change, adding "that creating glucose is the first priority of the Challenge, scaling a technology that can accomplish this can come second" [CO9].

Despite not having explicit limits, the SMEs stressed and incentivized a smaller footprint. While SMEs deferred the requirements to make the solutions fit into NASA aims, they felt that they needed to reinforce the space application of this technology [CO25]. As such, likely in line with the sentiment expressed in [CO23] above, the rules emphasized and incentivized small footprints. In both Phases, the rules informed potential solvers that "to increase the potential for use in space missions, scalable, low mass/power/volume systems are sought" [CO7, CO19]. In Phase 1, the system's footprint was an important factor under the scoring criteria of Applicability of Proposed System for Space Missions (itself 25% of the total score for Phase 1) [CO7]. And in Phase 2, a bonus competition related to the system's footprint was added [CO7, CO25]. Here, $100k of the prize purse would reward submissions' "effectiveness for future application in space missions" [CO19]. It would specifically grade solutions on making the solutions efficient in terms of power and conversion, the ease of scaling the operations, and the difficulties of operation [CO19]. Notably, the criteria did not capture the mass of the system—the formulation team could not come up with a way to fairly measure the equipment needed to produce the samples [CO25].

2.3.3 The Purity of the Sample to be Produced

The presence of contaminants in the sample A sample's glucose mass is not the only thing that determines its success as bacteria food. SMEs knew that certain compounds and intermediate products of the CO_2 to glucose conversion would be detrimental if they were present in the bacteria's food source. For example, even a sample with 90% glucose and 10% of certain other products would be "no good" [CO23]. Solutions that did not account for this would, ultimately, not be able to fulfill NASA's goal [CO2]. However, like scaling this technology to fit the Mars implementation, SMEs also deferred any requirements on the purity of the output.

Rules that would limit contaminants in solvers' solutions were hard to implement for a few reasons. First, measuring the composition of a sample was hard [CO3]. Quantifying all (potentially detrimental) products in the solvers' samples would be too complex a task for judges [CO22]. Another pathway that SMEs explored was the samples' compatibility as a food source [CO3]. This would be a (biological) test to ensure that the solvers' output would be compatible with the kinds of bacteria that SMEs would hope to feed [CO4]. However, this test would not be "even-handed" to all conversion approaches or candidate bacteria [CO23, CO9, CO3]. Second, it was possible to refine outputs to be compatible with different bacteria that SMEs would want to feed [CO2]. In this view, purification would be an extra step to the conversion, not an essential part of it [CO2]. Requiring that solvers tack this on would make the challenge too "large" [CO2] and add "complexity" [CO9]. Lastly, SMEs estimated that additional requirements to make solutions more compatible with its space application would make the challenge too expensive to solve [CO22].

In the end, SMEs decided to address the contaminant problem outside of the challenge context and did not add quantitative rules on contamination. [CO9]. The importance of demonstrating the conversion from CO_2 to glucose took precedent over how pure its sample could be [CO11, CO9]. Per one SME on the CCP team: "if we were able to just get the glucose molecule, regardless of whether it was compatible with [the test we had proposed], it would be considered a tremendous success" [CO2].

But the importance of contaminants was not forgotten. Instead of explicit rules governing the solutions, SMEs stressed their preference. They inserted statements in the rules that would—hopefully—focus the solvers on making something that would (eventually) be compatible with the Mars settlement goal [CO9]. For example, this line appeared in the rules for both Phase 1 and 2: "Likewise, the ability to make target compounds at high efficiency and specificity, and *with minimal contaminants and/or toxic by-products, is preferred*" [CO7, CO19, emphasis mine]. Additionally, SMEs felt that the solutions would still provide enough information to (somewhat) address the contaminant criterion at this stage. In judging the solutions, they would apply their expertise in the NASA context to identify which (kinds of) solutions would be most promising. One SME said it best:

> And sort of innately, we'll know– If we get a product stream and it contains a lot of toxic heavy metals in it, or it's highly acidic, or very, very salty, or fill in the blank. We'll go, "ok, in accordance to the rules, it's fine," but me as a NASA person will go, "we'll never be able to use this thing." So it's possible it still could be a winner in the challenge but not a viable candidate for our uses. And we're ok with that because it will still be progressing the field. [CO22]

Glucose versus other Carbon Products in the Sample The output of the conversion process was another area of uncertainty. Solvers' systems could create several intermediate products that, like acetate, could be food for bacteria [CO3, CO11]. These formed a ladder of increasingly higher carbon molecules (i.e., 2, 3, 4, 5, and 6-carbon products). Each were progressively more difficult to create from the CO_2 molecule, but also progressively more efficient as a food source: "not only is [a higher rung] more difficult, but it's also a better product" [CO2]. But this risked diluting the challenge's focus: solvers could chase high production rates of these intermediate products without actually producing any glucose. The SMEs had internal discussions about what it would mean for solvers to produce different combinations of these kinds of products. In these discussions, questions like the following would arise: "what if you got the glycerin [a 2-carbon compound], the lowest on there, if you have a pure amount of that, versus a very unpure amount of glucose? ... So, which one do you like better?" [CO2].

In the end, the SMEs skewed the challenge's scoring towards the outcome they favored. The challenge as a whole would be pushing to "improve technology that is able to convert CO_2 into other molecules" [CO2], which the SMEs acknowledged would be an important technology to have in the future [CO3, CO2]. But to keep the focus on feeding bacteria efficiently, the SMEs decided that the intermediate products were valuable enough to score them too. SMEs weighted the scoring of the sample along those lines. Here, each carbon product would be weighted higher per their utility in feeding bacteria, with glucose receiving the highest weighting (the most important one) [CO7, CO19, CO5]. Thus, the rules would incentivize—but not mandate—the kinds of solutions that would best fulfill this function.

This presented a scenario where solvers could win without accomplishing the main goal.

To counter this, SMEs tried to emphasize and incentivize solutions that addressed that goal. SMEs emphasized the glucose goal by positioning it at the top of the weighting factors and emphasizing it very strongly in the text of the rules [CO5, CO7, CO9] and other public descriptions of the challenge [CO12]. For example, using statements like "D-glucose being the most preferred" substance [CO9, CO7], or stating they were "looking to push the envelope, so to speak, to move [CO_2 conversion] towards [these] sugars" [CO12].

2.3.4 The Production Rate and Sample Size

Both the sample size and the time allotted to produce it fluctuated. At first, they limited the system's operation time to 4 hours [CO1, CO5] and asked for specific amounts of product to analyze in grams. One of the last drafts of the rules went up to 4 grams in 4 hours. This number was based on amounts needed to perform three analyses in a lab accurately—smaller amounts were at risk of running into physical detection limits [CO9]. But at the same time, there was a concern about requiring a certain sample size and having solvers undershoot that amount. The EPMC specifically stated their "concern that 4 grams in 4 hours might be too much" [CO9]. Other concerns raised also included shifting the focus of the solver from producing some glucose to producing a lot of lesser product: "What if they produced just 1 gram of glucose (which is ultimately what we want)? They would not meet success requirements the way the rules are written now" [CO9].

Ultimately, SMEs decided to remove production rate/sample size limits. Initially, the uncertainty of estimating how these solvers would perform pushed the decision on what the minimum production rate would be until after Phase 1 [CO9]. Contrary to earlier versions of the rules, Phase 1 did not include how much product the challenge required [CO7]. Instead, they gave solvers a production window and stated that their samples would need to be "enough" to be analyzed by specific tests at NASA Ames [CO12, see also CO9, CO19]. Thus, the rules did not explicitly require solvers to scale their systems to a determined output. Instead, solvers would be allowed to interpret that requirement based on their expertise and available infrastructure. Their samples would be scored by the mass fractions of the desired products in the sample [CO19]. Note here that even with the Phase 1 results received, SMEs decided not to require a specific sample size for Phase 2 [CO19].

Despite removing these limits, SMEs communicated the importance of scaling through the points. They incentivized the design of the solutions to be scalable by scoring these kinds of solutions more heavily [CO7, CO19]. In phase 1, this was part of the Applicability of Proposed System for Space Missions category, with 25% of the score [CO7]. In Phase 2, this criteria formed 30% of the bonus round's score [CO19].

2.3.5 Excluding Biological Solutions

In contrast to decisions on the other focus areas, NASA SMEs made a definitive choice on what conversion approaches solvers should (not) use in their solutions. Specifically, solvers were to avoid biological approaches to convert CO_2 to glucose. There were two reasons why: the relevance in addressing NASA's need and difficulties judging the solutions. The former was about the technology family that the SMEs wanted to develop: SMEs wanted to "push past what biology can do and set bar high enough so it's a stretch for people" [CO11]. Because "biomass [was] not going to solve the problem" [CO11] of providing NASA with an efficient method of feeding bacteria on Mars [CO12], the goal of the challenge was to find a better pathway, not just any pathway [CO3, CO24]. The latter was about the judging of the biological solutions. SMEs were concerned that this was an area where solutions might seem operational, but the underlying problem had not been solved. Here's how one SME described their concern: "there are ways to skirt the system to make it seem like you're doing it and it's not going to work" [CO3]. So, the rules [CO5, CO7] and presentations about the challenge [CO12] made it very clear—from the start—what the challenge was hoping to accomplish.

This exclusion would (mostly) extend to products that could be used in solvers' reactors derived from plants or bacteria. SMEs clarified the intent of the rules after the challenge was posted. As with other CCP challenges, the solvers would submit questions to the formulation team. While some solvers asked for—neutral—clarifications of the rules, some asked for the SMEs' blessing on a proposed approach. This was the case here too. Per one of the SMEs, some solvers were trying to be "cute" [CO23]. While they knew that biological approaches were not allowed, they asked whether compounds derived from organisms would be allowed instead. The problem here is the supply chain to the Mars settlement: they would have no way—other than biological processes—to replace these compounds once they ran out. So, they were not considered valid entries for the challenge [CO13]. However, the SMEs decided not to clarify this as an update to the Phase 1 or Phase 2 rules [CO19, CO9]. One SME stated that he "did not want to throw away" solutions that had some biological processes in them, despite the problems above [CO24].

3 The Outcomes of the Challenge

3.1 Feedback from (Potential) Solvers before the Solution Submission

The SMEs received strong pushback from solvers interested in pursuing the biology approach [CO23]. SMEs received this feedback in the requests for information that preceded the challenge's launch and as informal questions submitted by (potential) solvers [CO4, CO13]. Pursuing the conversion using only non-biological means was the exception, not the norm; several teams (with expertise in biological systems) tried to push the rules towards these systems anyway [CO4]. Imagining themselves as a solver, the lead SME acknowledged that that constraint was counterproductive and a "thorn in their side" [CO24]. However, SMEs did not want this challenge to be "swamped" with bio-solutions that could likely outperform others just because they are more mature or developed [CO24]. So, they stuck to their exclusion. But this explanation did not seem to appease all interested teams. One SME said that they told solvers, "'sorry that's not going to work,' and they weren't happy about it" [CO10].

This pushback worried the SMEs. It signaled to the SMEs that those following the challenge also realized that it would be difficult, especially since the bio-pathway had been closed off. Thus, they were worried that the challenge would not see healthy participation. Per a SME: "we're taking a bit of a chance. We're going to restrict our participant base a bit to focus on it more diligently. And we're crossing fingers that we get the response we need to host a healthy challenge" [CO3]

3.2 SMEs' Reflections on the Solutions

Despite their initial apprehension going into the challenge, the SMEs were generally pleased with the solutions' quality and quantity. I explain both below.

3.2.1 Quality

Teams delivered successful solutions in both Phase 1 and Phase 2, taking home the prizes in both. In Phase 1, the SMEs estimated that at least the majority of the submissions were good, with some even being excellent: a "third were excellent, a third were good to ok, and a third were not very applicable" [CO22]. Here, the top five teams took home $50k [CO20]. In Phase 2, three teams successfully demonstrated their conversion system, taking home equal shares in the $650k prize purse [CO21]. Additionally, these same teams won the Phase 2 bonus round [CO21]: one industry team won the $50k top prize, and the other two took home $25k. As a testament to the quality of the highest performing solutions (in Phase 1), one SME remarked that the judges were happy to reward teams for their performance, not just because they followed the rules [CO22].

How did the judges judge their confidence in the solutions? While the Phase 2 rules did contain some quantitative ways of calculating the solvers' final score, a lot of the scoring throughout this challenge was more qualitative. Moreover, Phase 1's criteria did not include a quantitative output at all. As such, success—particularly in Phase 1—meant speaking the same language and showing that you knew the material well [CO3]. SMEs wanted solvers to show that they had been thinking about this problem for a long time, not "like someone had an idea in the shower and they [wrote] it down" [CO22]. They expected solvers to show that they had gone through the appropriate literature out there and ensure that they did not propose something that had (partly) already been tried and failed [CO22]. It even came down to the references that the solvers would use. Per an SME, they were looking for:

> "... markings of a clear understanding of the problem, providing background to their solution. Where it's being derived from. If they can provide a logical argument as to why this would be a good way to go after this problem. And then there is understanding their chemistry. Or whatever field it is that they are using to solve the problem, including references." [CO22]

SMEs thought "the solutions were hopeful and promising of getting somewhere" [CO22]. They acknowledged that these solutions made some headway into the CO_2 to glucose conversion problem [CO20, CO21, CO22]. And they found this work valuable even with the limited ability to match the rules to NASA's requirements [CO22]. The solutions presented some new conversion pathways: some combined known chemical processes in new ways, others performed known conversions in new ways (with the latter being more valuable) [CO22]. At the same time, however, SMEs remarked that they did not see anything surprisingly novel in the solutions. They did not expect to see a "miracle cure" and did not see one either [CO22]. Instead, solutions delivered "middle of the road, good solid progress" on non-biological conversion systems [CO22]. Here's how one SME summarized his view on the advancements made by the solutions:

> [There] wasn't anything that we went "Wow! Oh my goodness, this is out of this world." It's chemical engineering. The solution space is fairly well-defined whether or not you're innovative in that area. A lot of the advances right now are just iterative improvements on old systems. Finding a slightly better catalyst, or one that lasts longer, or a lower temperature to operate it, or finding ways to make less waste products. ... There's mild to medium innovation here [in this challenge]. And it's good. It will advance– It will push this field forward. Everybody typically hopes for a miracle solution to things, right? And physics usually doesn't allow it. [CO22]

The SMEs' strategies to balance their aims influenced solvers and their solutions. Most solvers explored conversion methods that relied fully on non-biological approaches, and some even pursued versions of their systems that could work in NASA's context. In Phase 1, for example, some solvers explicitly described how their solutions addressed NASA's aims by, e.g., being able to regenerate their catalysts [CO14] or limiting their consumables to exclusively Mars in situ resources [CO14, CO18]. Additionally, three teams won prize money for accommodating the bonus objectives in Phase 2 of "efficiency, scalability, and reliability" [CO19].

While the challenge generally made progress on conversion systems, the biological exclusion rule was still an issue for some solvers. According to the SMEs, only one solver pivoted from the initial path of something that would not work to something that would [CO10]. The ones that could not—or did not want to–pivot submitted non-compliant solutions anyway [CO24]. In the same vein, SMEs remarked how it was the academic teams that were less likely to venture outside of their wheelhouse of knowledge; in contrast, industry teams were "scrappy and [tried] to pull in what they need when they need it" [CO22].

3.2.2 Quantity

Overall, SMEs "got a lot more applications than [they] thought [they] were going to get, which was excellent" [CO10]. Over a thousand teams—1415 in total—showed an interest in participating; hundreds of teams—210 in total—completed the challenge's preregistration form [CO34]. In the end, 24 different teams participate across both phases [CO34]. In Phase 1, 20 teams submitted solutions. Of those, CCP classified five teams as academic teams (e.g., PI-led research groups), seven as industry teams (e.g., start-ups), and eight as other (e.g., unaffiliated research teams or hobbyists) [CO34]. In Phase 2, the number of teams dropped to eight. Of those, CCP classified two as academic teams, three as industry teams, and three as other [CO34]. Only four teams, all of whom were winners in Phase 1, participated in both phases [CO34].

The SMEs were also surprised about the number of compliant solutions they received. The pushback from teams working on biological systems painted a different picture than the challenge's outcomes [CO10]. One member of the formulation team even remarked that the challenge might not have been as hard as the SMEs envisioned it would be for external solvers [CO26]. Here is how one SME described his view of the number of "good" solutions (in Phase 1):

> Initially, we were quite worried that enough viable ideas were actually going to be submitted, and we got plenty, and we were quite happy with that. The five that were selected [in Phase 1]— There were some that were better than others, but they were all above the bar of what we were thinking, and we were pleasantly surprised ... I think you can say [I'm] surprised from the perspective of "we didn't expect so many people to apply and so many good applications to be submitted as well. [CO22]

Lastly, the challenge attracted teams that were both known and unknown to the SMEs. Because of their knowledge of the field, SMEs expected certain individuals and their institutions to participate in the challenge, even reaching out to several companies who would potentially be interested [CO2]. This had a lot to do with their capabilities [CO23, CO2]. SMEs expected some teams to participate, and they did, but others did not [CO22]. SMEs also expected that the challenge could attract "an entirely new cadre" of individuals and teams [CO23, see also CO3], in particular, those "who typically don't participate in NASA calls" [CO22]. And they did as well. SMEs described several teams—mainly start-ups—who "came out of the blue" and participated in the challenge [CO22]. While most teams had strong backgrounds in CO_2 conversion technology [CO10], several solvers did not have previous experience in the aerospace industry [CO28, CO29, CO30, CO31]. One of these non-aerospace solvers expressed that the challenge helped them realize that "NASA will need serious chemistry research for Mars exploration tech" [CO30]. This team even went so far as to explicitly say that they wanted to pivot their business to become a supplier to the aerospace industry (including NASA) based on the work in the CO_2 to glucose challenge [CO30].

References

Table 1: References used in the "Formulating the CO_2 to Glucose Challenge" Case Narrative

Reference	Date created	Description
CO1	Sep 11 2017	First draft of CO_2-to-Glucose RFI
CO_2	Apr 18 2018	Interview with SME CC8 about the start of CO_2
CO3	May 3 2018	Interview with SME CC11 about the start of CO_2
CO4	Oct 27 2017	Centennial Challenges Program presentation to STMD
CO5	Apr 4 2018	Second draft of CO_2-to-Glucose RFI (with comments)
CO6	Apr 4 2018	Released CO_2-to-Glucose RFI
CO7	Aug 16 2018	Final Phase 1 rules
CO8	Aug 16 2018	CO_2-to-Glucose FAQ V1
CO9	Aug 16 2018	Internal discussion and resolutions of issues brought up by the EPMC
CO10	Oct 10 2019	Interview with SME CC11 after phase 1
CO11	Mar 14 2018	Kickoff meeting CCP and Common Pool
CO12	Feb 28 2020	Phase 2 webinar
CO13	Oct 2 2018	CO_2-to-Glucose FAQ V2
CO14	May 16 2019	Summary of winner's solution (Phase 1)
CO15	May 16 2019	Summary of winner's solution (Phase 1)
CO16	May 16 2019	Summary of winner's solution (Phase 1)
CO17	May 16 2019	Summary of winner's solution (Phase 1)
CO18	May 16 2019	Summary of winner's solution (Phase 1)
CO19	Sep 16 2019	Final Phase 2 rules
CO20	May 16 2019	Press release detailing Phase 1 winners
CO21	Aug 24 2021	Press release detailing Phase 2 winners
CO22	Oct 18 2019	Interview with CC11 on the results of Phase 1 and expectations for Phase 2
CO23	Apr 1 2020	Interview with SME CC28 on the start and formulation of the challenge
CO24	Feb 26 2021	Informal conversation with CC26 and CC11 on Phase 2 formulation
CO25	Aug 26 2020	Informal conversation with CC26 on CO_2 challenge formulation
CO26	Jun 5 2019	Informal conversation with CC26 on Phase 1
CO27	Oct 8 2020	Interview with solver CO1U1 on participation in Phase 1
CO28	Oct 2 2020	Interview with solver CO2SB1 on participation in Phase 1
CO29	Oct 2 2020	Interview with solver CO3U1 on participation in Phase 1
CO30	Oct 1 2020	Interview with solver CO4SB1 on participation in Phase 1
CO31	Oct 24 2020	Interview with solver CO5SB1 on participation in Phase 1 (written questions)
CO32	Jul 3 2019	Feedback from solver CO5SB1 based on CCP questionnaire
CO33	Jul 29 2020	Meeting with CC26 on challenge formulation
CO34	n/a	CCP list of all CO_2-to-Glucose participants and winners