aPPEL academy of program/project & engineering leadership

OCO-2: A Second Chance to Fly

Case Study

Artist's impression of Orbiting Carbon Observatory-2 (OCO-2), which makes precise global measurements of atmospheric carbon dioxide (CO2) to help scientists better understand its sources and "sinks."

appel.nasa.gov

Photo Credit: NASA/JPL-Caltech

Introduction

The early morning of February 24, 2009 was cold, wet, and beautiful. At 1:55 a.m. Pacific Standard Time, a Taurus XL rocket rumbled away from the ground at Vandenberg Air Force Base in California carrying the Orbiting Carbon Observatory (OCO). OCO was on its way to orbit—but not for long.

The payload fairing failed to separate from the launch vehicle, preventing the final stage from boosting OCO into its injection orbit. The surviving pieces of OCO splashed down in the ocean near Antarctica. Twenty-four hours later, the project closeout process for the mission began, which included capturing the team's lessons learned.

The OCO mission was an Earth System Science Pathfinder project run by the Jet Propulsion Laboratory (JPL). OCO was intended to join the A Train, a constellation of Earth-observing satellites, where it would make precise, time-dependent global measurements of atmospheric carbon dioxide (CO2), helping scientists better understand the sources and "sinks" of CO2. This information was considered so critical that, following the loss of OCO, scientists and world leaders were invested in the discussion about how to proceed without it.

Three days after the launch mishap, a proposal for a replacement mission—Orbiting Carbon Observatory-2 (OCO-2)—was sent to NASA Headquarters for review. In September 2009, NASA presented a plan for a new mission that was nearly identical to OCO to the Office of Science and Technology Policy and the Office of Management and Budget. By early 2010, the OCO-2 team received Authority to Proceed (ATP).

This rare second chance to fly presented the OCO-2 team with unique opportunities and challenges. The first was to actively apply lessons learned and knowledge from their original mission to a second flight. While they benefitted from having gone through the project life cycle once, the OCO-2 "build to print" philosophy—which meant the new satellite would be an exact copy, to the extent possible, of the first OCO—directly opposed the notion of improving upon or deviating from the original mission plan. The second was an unexpected chance to work with a live datastream from another project. This gave the team an opportunity to improve raw data interpretation, fine-tune their scientific methods, and rework their software.

To Rebuild: From OCO to OCO-2

OCO went down in February of 2009. With it went the hope of obtaining critical and urgent information about CO2 that could help the world's scientific and political communities better understand how to address this key contributor to climate change. The three-spectrometer instrument on OCO, with its unprecedented sensitivity, had been poised to deliver the most detailed picture yet available of the behavior of this greenhouse gas.

CO2

CO2 traps reflected sunlight and is difficult to eradicate once in the atmosphere. Scientists know that CO2 is absorbed into "sinks," such as oceans and growing plants, where it is removed from the atmosphere. But they need to discover more about where it comes from, where else it goes, and why.

But by early 2010, there was hope again. The President's Budget, released on February 1, included funding for a second OCO flight. On March 8 the news became official: OCO-2 received authorization to rebuild.

A year had passed between the loss of OCO and the official approval of OCO-2. But the project team had not been idle. Many team members were focused on early parts procurement: locating the materials they would need for a rebuild of both spacecraft and instrument. Tasked by NASA Headquarters to present a rationale for a second mission, the group set out to determine how long it would take to reach the launch pad again. "The original goal we were handed down was two years," said Randy Pollock, OCO-2 Project Architect. "[NASA] wanted a 24-month rebuild. But we couldn't get there. Finally, we came back with 28 months after talking to all the vendors and working out long-lead parts issues. In those first few months [after OCO failed], we made a very intense effort to figure out how we could redo this."

Meanwhile, others members of the team were busy closing out OCO and capturing lessons learned.

The OCO-2 science instrument is prepped at JPL.

Photo Credit: NASA/JPL-Caltech

The lessons learned process

OCO-2 Mission Operations System Manager Patrick Guske led the lessons learned initiative, which began shortly after the failed launch. “We were looking ahead with great anticipation and great hope that the [reflight] project would be picked up. And in fact it was picked up almost immediately, but we still had to close down OCO. So we looked at this as a perfect opportunity to make our lessons learned document really, truly usable.”

The process was comprehensive. From engineers to contractors to secretaries, every member of the OCO team was encouraged to submit lessons. Four overarching criteria were identified for each lesson:

- It should be phrased positively.
- It should be achievable.
- It should not assign blame.
- It should offer a solution.

In addition, Guske said, “Any lesson that recommended a change needed to be about making things work, not making things better.”

Senior members of the OCO team reviewed the collected lessons learned, excluding duplicates as well as any that presented a complaint rather than a focused lesson. The resulting lessons fell into three broad camps:

- Those that identified problems impacting the schedule or causing extra work for employees.
- Those that escalated cost for the project.
- Those that contributed to a problem in process resulting in confusion.

The lessons were laid out in a simple framework so that everyone on the team could understand each lesson and proposed solution. Some lessons noted a problem on OCO, with accompanying solution, while others identified a positive experience on OCO with encouragement to repeat it on OCO-2.

Overall, 78 lessons learned were identified for implementation. Guske then assigned an individual to each lesson. That person was responsible for implementing the lesson (or lessons) assigned to them. The lesson did not have to be implemented according to the solution provided if that suggestion proved unworkable, or if a different, better, solution presented itself. The important thing was that the intent of the lesson learned was realized to the benefit of OCO-2.

To ensure that the lessons were implemented, Guske entered them into an Excel spreadsheet that noted the subject of the lesson, a summary of the issue, the person responsible for implementing the lesson, lesson status, and relevant comments. He then checked in periodically with the individuals responsible and updated the spreadsheet with the status. By April 2013, 55 of the 78 lessons had been implemented successfully.

OCO Lessons Learned Implementation for OCO-2
Status as of April 2013

- **Implementation successful: 55**
- **Intent of lesson met: 2**
- **Identified as a best practice: 1**
- **Business as usual: 9** (no specific lesson to implement)
- **Partial implementation: 1**
- **Attempted but unsuccessful: 1**
- **Lesson declined: 3**
- **Awaiting status determination: 1**
- **Awaiting implementation: 5**

The four of the five lessons still "awaiting implementation" in early 2013 had been assigned to Guske himself and were eventually addressed. In the case of the partially implemented lesson, one aspect of the lesson was completed but the others could not be finalized. The three lessons declined were done so by the project itself due to cost or other factors.

The OCO-2 philosophy: “Build to print”

When the official green light for OCO-2 was received, the team was elated. But the Formulation Authorization Document came with very specific direction. “The basic instruction from our NASA sponsors was to take full advantage of our existing designs, drawings, schematics, and documentation to the extent possible so that we could mitigate technical, schedule, and cost risk. We made sure that we were true to that philosophy. We didn’t make a change unless there was a compelling reason to make that change,” explained OCO-2 Project Manager Ralph Basilio. To accomplish this, the team intended to follow a build-to-print approach in order to make a “carbon” (as they joked) copy of the original project. Essentially, that meant they should not do any engineering. “Take the drawings out, go get the parts built, take the assembly instructions out, follow them the same way they were done before,” clarified Pollock. “There were a lot of things that we knew we could do much better. But we didn’t. Because if we started down that path, pretty soon it would be a whole new design.”

To accomplish this directive, the team followed the dictum that “better can be the enemy of good enough.” Under a tight timeline of 28 months and faced with an enormous task, they knew that any alteration to the product as it was being built—even if the change improved the output—increased the risk of schedule, cost, or technical slips. So the goal for each step of the rebuild was to meet requirements without creating an inadvertent problem elsewhere by “improving” anything.

As a result, the team lived with certain things that were less than perfect but met requirements. “With OCO-2, we stopped at good,” stated OCO-2 Spare Instrument Manager Tom Glavich. One example he cited involved the project’s focal planes. The instrument had three, each with slightly different technology. When the focal planes returned from the manufacturer, testing revealed they weren’t ideal. “We could have spent another million dollars and waited a couple of months and gotten some better focal planes. Or at least had another focal plane run. But we looked at what we had and said: these really do meet specifications so they’re good enough.”

Another area that was defined as “good enough” was the flight instrument software. By the time OCO-2 became active, the person who had created the instrument software for the contracting company had retired. Faced with an information gap, the team was unable to make the minor tweaks they would have preferred to do. Instead, they left things exactly as they were: good enough.

Required redesigns: When “build to print” wasn’t possible

Despite the project’s “no engineering” approach, the team made room for necessary changes based on the 78 formal lessons learned. As OCO-2 Project Scientist Mike Gunson pointed out, “There was the opportunity to learn from the experiences and challenges of the first OCO to overcome certain technical issues.”

Eventually, as the project progressed, a need for additional alterations emerged as well.

Issues Driving Redesign Efforts on OCO-2

- **Lessons learned**
- **Parts obsolescence**
- **Inconsistent results**
- **Emerging concerns**

The lessons learned provided a clear blueprint for change that was considered critical. The team determined that any changes resulting from lessons learned or problems that arose during the rebuild had to accomplish one of two things: address a core problem or solve something as-yet unsolved.

Parts obsolescence presented a pressing problem. During the early parts procurement process, the team discovered that a number of critical items that had gone into the construction of OCO—ranging from minute electronics to larger pieces such as the cryocooler—were unexpectedly unavailable. Often, while the change itself appeared insignificant, it necessitated some redesign to resolve incompatibilities that emerged as new parts were incorporated into the existing blueprint.

One of the most baffling issues the OCO-2 team faced was inconsistent results. Despite a commitment to doing things on OCO-2 exactly as they had been done on OCO, the end result was occasionally different—and unsatisfactory. When this occurred, the team was required to spend time and money resolving an issue that was never expected to be a problem in the first place.

While the “build-to-print” approach worked well overall, it could not account for concerns that emerged over the course of the rebuild. Such concerns included things over which the team had no control, such as the emerging need to switch launch vehicles, and problems that became apparent over time, such as issues associated with seasonal variation.

Redesign based on lessons learned

“From a lessons learned perspective, there were some higher-than-desired technical risks that we attempted to fly with on OCO. [For OCO-2], we had an opportunity to provide a product that was more resilient and gave us a better chance of meeting our Level 1 requirements,” said Basilio.

Three of the biggest risks OCO flew with related to the instrument. One lesson learned pointed to a stray light issue on two of the three detectors. The fix was simple: adding an extra piece of metal provided enough shadow to eliminate the problem.

A second lesson corrected for an issue with the instrument’s detectors. Known as “residual image,” the problem was similar to what happens when a person looks at a camera when the flashbulb goes off. In the darkness that follows that sudden burst of light, the person sees an after-effect for several seconds. The OCO instrument was expected to experience a similar “ghost image” problem as it transitioned from light to dark or dark to light. The following seconds of data would be unusable. The team could have worked around the problem, but the potential impact of lost data was considered important enough to warrant changes. To resolve the issue, the team carefully cherry picked the detectors they used and ran a residual image check before installing them permanently. According to Gunson, resolving this problem did necessitate some engineering work. “Ensuring that there were no residual image properties on the focal plane arrays led to a redesign to operate them at a lower temperature.”

The third change addressed a slit-misalignment on the OCO instrument. One of the spectrometers was “wall-eyed”: it looked slightly to the side. The lesson learned from OCO suggested revisiting the cantilevered system, which appeared to cause the problem. In response, the team installed a stabilizer and removed the cantilever entirely. This provided support for each spectrometer on the front and back ends, resolving the issue. Preliminary examination of in-flight data from OCO-2 indicates that the spectrometers now meet alignment specifications.

“In theory, we could have followed the philosophy blindly and just decided to fly the same instrument, warts and all. Or we could do what we did: go through [the lessons learned] in a very systematic process and implement those changes that would certainly improve our odds of meeting our Level 1 requirements,” said Basilio. “We had to be careful about not biting off more than we could chew and then finding out that better is the enemy of good enough, so it was a bit of a balancing act. But I would say that, in general, the team did a very good job in implementing lessons learned that truly put us in a better position to be successful when we operated OCO-2 in space.”

In many cases, employing the hard-won lessons from OCO saved time and effort for the OCO-2 team. On OCO, the team dealt with multiple problems involving the room temperature vulcanization (RTV) silicone that held the instrument’s optical lenses in place. Application was challenging, and in some cases the RTV flowed onto the lenses themselves. Because the lessons learned process captured the OCO team's solutions to the problem, it was not an issue on OCO-2. "We knew we had the right procedures in place for getting the RTV in the right place. We knew how to get it there, we knew how to check that it was actually there, and we knew how to make sure that the lenses stayed pristine and clean," said Glavich.

Another solution identified during the lessons learned process helped simplify calibration. On OCO, the team decided to use the moon, not stars, to calibrate the system and instituted it as a manual procedure. "It was a much brighter target and much easier to use in the calibration process," said Ray Welch, Flight System Manager on OCO and part of OCO-2. For OCO-2, the team incorporated the change as an automated computer procedure in the Mission Operations Software with commanded inputs from the flight team, increasing the efficiency of the process.

Lessons learned concerned document management processes as well as hardware development. On OCO, the team faced a problem with paperwork storage and access. They used the DocuShare system to store daily communications, customer-contractor communications, review packages, and other material. "It got to be called DocuHide," said Welch. The problem was so pervasive that DocuShare was cited in four distinct lessons. All were implemented successfully.

Some lessons learned focused on the importance of very specific, one-time behaviors. A lesson titled "Observatory Shipping" called for one individual, known as the "observatory transportation lead," to assume responsibility for the observatory while in transit from Orbital Sciences Corporation's headquarters in Dulles, VA to their Gilbert, AZ plant. An individual was identified, their name was put into the lessons learned chart, and they traveled with the shipment, making sure the observatory reached its destination along with all relevant ground equipment. A similar process was instituted for the shipment from Gilbert to Vandenberg Air Force Base.

Not all lessons learned pointed to obvious problems or required solutions beyond ensuring attention was paid to doing things in a way that supported the long-range goals of the project. One example was a lesson called "Allocation of Resources." The summary stated: "The impact of reassigning staff from their primary responsibility to support a task that is 'critical' must be evaluated (and quantified) for future impact to the primary responsibility, with resources/margin being assigned for the lost/delayed effort." The individual responsible for implementing the lesson and the status of the lesson were listed as the same: "Business as Usual." In other words, said Guske, "This needed to be looked at from a system point of view for the entire life cycle of the project. We drew attention to it by writing a lesson learned. Essentially, we said, 'You need to be aware that if you start pulling people off [their main jobs], you have to figure out where down the line that could be a problem.'" This kind of detail and focus helped the project remain on track.

Redesign required by obsolescence

One of the biggest obstacles to the rebuild was parts obsolescence. “What happened in the time between OCO and OCO-2 is the economy crashed and a number of companies that provided hardware for OCO had gone out of business. Or were still in business but had lost capabilities,” said Glavich.

Members of the team spent months out on the road meeting with vendors to determine whether the items they needed were still available and what the timing would be to produce long-lead parts. One example was the glass used for the instrument’s lenses. “It was very special glass,” said Pollock. He raced to ensure the vendor still had the particular “melts”—the specific glass mixture—needed. Fortunately, they did.

Other materials proved harder to find. The company that manufactured the S-band transmitter on OCO was bought out, and the new company discontinued the needed model. So the team was forced to use a different model, which necessitated some redesign work.

Another challenge involved finding silicon cells for the core sun sensors. Ultimately, the team was forced to use Gallium Arsenide solar cells instead. “That caused a redesign, not only of the bracket but also the software and the electronic hardware because the gains were all different. The sensors operated more efficiently so we got too much signal, which created a problem,” said Welch. “Electronics are also a big issue. All the components that you buy have a lot of electronic piece parts in them. And they’re continually becoming obsolete so you have to get something else to replace them. And that sometimes requires a redesign of a circuit. So we try to follow ‘identical,’ but unless you build everything at once, you run a risk of not being able to achieve that.”

The cryocooler also presented a problem: OCO-2 had to use a different model because the one used on OCO was a flight spare from a different project. “Well, there isn’t a spare to the spare,” said Guske. “So we had to find another. Then the interface for that cryocooler was similar [to the one on OCO] but not identical.” This necessitated minor design alterations to render the crycooler compatible with OCO-2.

Redesign due to inconsistent results

The project had to grapple with several situations in which they did something exactly as it was done on OCO yet the results were not the same. These unhappy surprises mandated changes for which the team could not plan in advance.

In some cases, the results in question had to do with testing. “I don’t think that’s unusual. I think when you move through a project, you get some results that may be on the edge. But you don’t notice they’re on the edge. Then, when you test the next [spacecraft], which has slight variations, the results are over the edge. The first one would probably have worked. The second we made sure worked," said Welch.

In other cases, the inconsistencies involved the hands-on work by technicians. One such issue concerned the paint on some of the spacecraft panels, which were painted black on OCO to improve thermal performance. For OCO-2, they hired the same company to use the same procedure and the same materials. What they couldn't account for was the personnel: a different technician arrived to do the job. When he finished the job, the paint appeared perfectly applied. But it peeled off during the first vibration test. The problem, which may have been due to poor surface preparation before painting, could not be fixed. Ultimately, the team was forced to remove the paint entirely and fly without it. Some of the spacecraft's thermal characteristics were altered as a result, but the team was able to demonstrate that the change would not impact the mission.

A similar problem affected the special black coatings that were applied to precision machine parts on the instrument. The company that had done the coatings for OCO was still in business, but the technician who had applied the material was no longer employed there. Glavich said, "They were as much an art as a science to put on. [The company] had somebody else trying to do it on OCO-2, but they could not get the coating on properly."

Considering the original technician for OCO, Basilio said, "I think he took with him the secret of success, if you will. Maybe he just had the right touch that was developed over 40 years of working with these very fine parts that allowed him to produce the highest-quality product. Maybe the new person only had a few years of experience. With these really precise machine parts, sometimes it does take somebody, or a set of individuals, who have nurtured and developed skills that are only learned through decades of experience to get the quality you need." The team eventually had to invest in hiring a different company to take care of the coatings so the parts would meet specification.

Redesign due to emerging concerns

Despite the team's clear mandate to do things exactly as on OCO, some things had to change as problems became clear over time.

One such problem involved the launch vehicle. OCO had launched on a Taurus XL provided by Orbital Sciences Corporation, which was also responsible for the OCO spacecraft. Keeping in mind its build-to-print philosophy, OCO-2 intended to use the same rocket. It made sense: the rocket was an ideal size for the needs of the mission, the interface between hardware was established, and the spacecraft was already qualified for the Taurus XL. Altering anything would impact the launch schedule. But on March 4, 2011, a change became inevitable.

The failure of OCO was related to its launch vehicle: the payload fairing did not separate, preventing the final stage from propelling OCO into its injection orbit. Following careful investigation, the failure was deemed a one-time problem and OCO-2 was cleared to use the Taurus XL as well. But that changed when the payload fairing on a Taurus XL carrying the satellite Glory failed to separate. The OCO-2 mission could no longer accept the risk associated with a Taurus XL.

“We were part way through when the Taurus failed on Glory,” said Pollock. “At that point, we were still working toward the original plan to launch on a Taurus in 2013. But everyone knew that wasn’t going to happen.” Instead, another year passed as the Taurus XL contract was canceled and the project searched for a new launch services provider.

They ended up selecting a proven launch vehicle: the United Launch Alliance Delta II 7320 rocket. In addition to its reliability, the powerful Delta II simplified the spacecraft’s ability to reach orbit. According to Welch, “The launch vehicle that OCO used had a larger dispersion of where the injection orbit could be. So we launched into an orbit that was much lower than the A Train, which meant we needed about 20 days of thruster firings to move up to orbit. With OCO-2, I think we only had six, maybe eight, burns in total. Because the Delta is a much more accurate launch vehicle, we got up a lot closer to start with and had a lot less effort to transition to the mission orbit.”

The change in launch vehicle required some engineering efforts as the interface was different, but overall the redesign was surprisingly limited. The OCO-2 spacecraft was generally compatible with the Delta II and the launch services provider used a SoftRide adapter to further minimize differences. According to Pollock, “This shock absorber that sits between the rocket and the spacecraft gave us the flexibility to sort of move the peaks of the vibration curves around a little bit.” The bigger effort required by the shift in rockets was in qualifying the spacecraft to work with a Delta rather than a Taurus.

During the delay dictated by the new launch vehicle, a more significant redesign became necessary. The reaction wheels used on OCO, which control the motion of the satellite in space, were failing on spacecraft already on orbit. The OCO-2 team began the process of vetting the wheel design. Part way through the process, they realized that the problem was not being resolved. The failures on orbit were continuing yet the manufacturer could not identify the problem. OCO-2 was subsequently redesigned to accommodate the new reaction wheels. This necessitated extensive modifications to the spacecraft flight software, spacecraft hardware, and test program. It was challenging for the team to make the decision to change the reaction wheels—and make all of the associated alterations—with the launch date looming, but the risk associated with the original wheels was too great to ignore.

In addition to emerging problems involving hardware, during tests for OCO-2 the team became aware of a potential risk pertaining to seasonal variation. Everyone knew the changing seasons could impact operations for an Earth orbiter like OCO-2. But they had not pinpointed a concern until testing revealed that a serious problem might involve interactions with their main ground station in Alaska. OCO-2 was designed to send data to the station through an X-band antenna that was hard-mounted to the spacecraft. During the winter months, when the OCO-2 antenna pointed toward the ground station, it would do so in darkness. But in summer, when OCO-2 pointed toward the tracking station, it would run the risk of pointing directly into the sun as well because Alaska is illuminated nearly 24 hours a day during that season. This was a significant concern, as the antenna beam direction and the instrument boresight are both fixed with respect to the observatory. As a result, pointing the antenna to the ground station could have inadvertently pointed the instrument to the sun. The problem might never have occurred: there were fail-safes in place, such as fault protection, that should have prevented the antenna from pointing right at the sun. But OCO-2 decided not to take any chances. To mitigate risk, the team performed tests to determine the extent of the issue and implemented changes to ensure a problem could not occur.

Ultimately, the proposed 28-month period from the failure of OCO to the launch of OCO-2 took much longer. “It’s taken over five years, and that’s for two reasons. One, it took NASA about a year to officially green light the new project,” said Pollock. The other, he explained, was the change in launch vehicles.

Over the course of those five years, engineers, technicians, contractors, and many other dedicated members of the OCO-2 team did everything they could to ensure a successful rebuild. Meanwhile, the OCO-2 science team was doing everything possible to ensure the post-launch mission would deliver meaningful data quickly and accurately.

International collaboration: From OCO to GOSAT to OCO-2

Unlike the team members working to rebuild and retest the hardware, the science team did not have to reconstruct their work after the failed launch. “For validation, algorithm, and Scientific Data Operations System (SDOS), we didn’t lose anything when the instrument crashed. We just hit a giant pause button,” said Annmarie Eldering, OCO-2 Deputy Project Scientist. But the team didn’t pause for long.

On January 23, 2009, a month before the launch of OCO, the Japanese sent GOSAT (Greenhouse gases Observing SATellite) into orbit. Like OCO, GOSAT intended to study greenhouse gas emissions. But its focus was broader: the satellite measures concentrations of methane in addition to CO2. GOSAT is backed by the Japanese Aerospace Exploration Agency (JAXA), the Japanese Ministry of the Environment (MoE), and Japan’s National Institute for Environmental Sciences (NIES).

Immediately after the loss of OCO, GOSAT’s Project Manager contacted OCO’s science team and invited them to contribute to the analysis of data provided by GOSAT’s instrument, the Thermal and Near Infrared Sensor for carbon Observations (TANSO) Fourier Transform Spectrometer (FTS). This offer enabled the OCO team to begin working with their own algorithms and software despite the failure of their initial mission.

The OCO-2 and GOSAT teams were already acquainted, having formed working relationships in the time leading up to the launch of both missions. Originally, the two teams had intended to compare data sets between the two missions with a goal of combining the information to provide a more informed understanding of atmospheric CO2. Instead, the OCO-2 team was renamed Atmospheric CO2 Observations from Space (ACOS) and both teams focused exclusively on GOSAT data. The teams’ data retrieval algorithms, however, would differ. This difference in method would contribute to a greater understanding about the global distribution of CO2, including how its sources and sinks vary across time, location, and season.

The collaboration benefited both OCO-2 and GOSAT, providing insight into the efficacy of each team’s data collection processes. For GOSAT, the collaboration helped identify inaccuracies in their pre-launch calibration parameters, enabling the team to maximize the accuracy of their data. The benefits to the ACOS/OCO-2 team were profound as well. “The level of collaboration and connection was unprecedented in my earth science experience,” said Eldering. The work allowed the team to identify errors—and resolve them—well in advance of the launch of OCO-2. “We had orders of magnitude improvement in our code because of the time and the effort that was invested in preparing to work and working with the GOSAT data.” The team was able to speed up their algorithm, markedly streamlining the way their system processed data.

“For the launch of OCO, I would judge that the algorithm was the least mature of the many, many parts of the overall system,” said Gunson. “I suspect we would not have advanced or matured the system that we have today, both in an understanding of what the appropriate algorithms are and the actual code that is the implementation of those algorithms, as well as the pipeline system that has to manage the process, without some exposure to a live datastream, such as the GOSAT data.” Working with the data, he maintained, “enabled us to overcome 90% of the kind of things that could have delayed our ability to supply a high-quality science product to the community.” The work is expected to enable OCO-2 to meet its Level 1 requirements far more quickly than was expected for OCO.

The work also trained the scientists to identify the optimal data for analysis. By familiarizing themselves with the process of working with a live datastream, they became adept at determining which data were valuable and which were not worth their time. This focus on the quality of the data resulted in an unexpected boon when one of the scientists noticed an extra signal: the presence of light in places it was not expected. Assuming it was due to a calibration error, the team was ready to attempt to fix the problem. But it turned out not to be a problem at all. Eldering said, “That was the beginning of a realization that you can see the impact of plants doing photosynthesis and sending out light as a fluorescent signal. You see that in our data. And if you don’t account for that, you actually get an error of one and a half parts per million in the measurement. We never would have known that except for looking at real data and pouring over it.”

This discovery, known as solar induced fluorescence, had several advantages. First, it helped the team better account for a known error source in retrieving estimates of atmospheric carbon dioxide. Computer codes were subsequently revised to account for the signal when processing OCO-2 data. Second, it created a new science product for distribution to the public, enhancing the value of the OCO-2 mission. “It will answer some very fundamental questions about what I’ll call ‘food security,’” said Basilio. The new data record increases in value when examined in conjunction with information from other satellites. “For example, SMAP, the Soil Moisture Active Passive measurement mission, will provide soil moisture measurements. And we can provide a sampling of the plant fluorescence, which is a direct measurement of plant health. A whole new set of scientists is going to be looking for this data.”

Despite the similar overarching purpose of the two projects, working on GOSAT data did not happen overnight. While upper-level processing of data was similar between the two projects, lower-level processing was completely different. As a result, the OCO-2 software had to be adapted so that it was flexible enough to handle GOSAT as well as future OCO-2 data.

For a science team that worked with GOSAT data for years, transitioning back to their original mission could be challenging. Because the GOSAT instrument is an interferometer and OCO-2 is a spectrometer, they make fundamentally the same measurement but do so in different ways. To interpret the data correctly, the team must remain cognizant of the way parts of the calibration are different. Pollock noted the importance of getting the team used to the idea that “while the raw data may look similar, it cannot be interpreted the same way.”

Overall, the team agreed that the opportunity to work with GOSAT greatly benefited OCO-2. “Our experience with GOSAT not only helped [the science team] validate their algorithms and their science results, but it’s definitely provided a great advantage for us in OCO-2,” said Basilio. “We’ve been able to minimize or mitigate our technical risks with these science retrievals. We even signed up to a shorter turnaround time: being able to produce data sooner than the OCO mission. For example, on OCO we signed up to produce retrieved estimates of CO2 within nine months of completing the in-orbit checkout period. In OCO-2, we signed up to six months because of this collaborative effort with GOSAT.”

A United Launch Alliance Delta II rocket launches with the Orbiting Carbon Observatory-2 (OCO-2) satellite onboard from Space Launch Complex 2 at Vandenberg Air Force Base, Calif. on Wednesday, July 2. OCO-2 will measure the global distribution of carbon dioxide.

Photo Credit: NASA/Bill Ingalls

The future: OCO-2 and beyond

OCO-2 reached its final operating orbit in early August 2014 and has begun returning data on global atmospheric CO2 measurements. With a successful launch behind them, the team can put the specter of starting from scratch once more—this time as OCO-3—behind them.

Some, however, are intrigued by the concept of an OCO-3. During the lead-up to the OCO-2 launch, the team constructed a flight-spare instrument to mitigate risk. “OCO was built with very few spares,” said Pollock. “So when we started OCO-2, we only had a handful of components we could reuse. On OCO-2, we decided to buy a full set of spares. Then, to take advantage of the time [during the wait for a decision about a new launch vehicle], we went ahead and put the spare instrument together.”

Hopefully, it won’t remain a spare instrument forever. “We were authorized to develop another project with this hardware,” said Basilio. The intention is to fly the spare instrument on the International Space Station (ISS) as OCO-3. OCO-3 will provide complementary information to OCO-2. According to Pollock, there’s a strong rationale for basing another observatory in low Earth orbit. “[OCO-2 is] in a sun-synchronous orbit, which makes the data processing much easier to analyze but also means you can’t see how the CO2 levels change over the course of the day as plants pull in carbon dioxide and let it back out. The ISS is not in a sun-synchronous orbit. It will allow us to start to understand a little bit about the diurnal cycle.” The potential mission has been identified as a project in the President’s Strategic Initiative Fund and has potential for eventual funding.

Whatever the future holds for the project, the team achieved what its original members set out to do nearly 15 years earlier: launch an on-orbit CO2 observatory that would provide critical scientific information to help answer questions about climate change. They were not able to stick to the original 28-month timeframe: launch vehicle difficulties derailed that. And they couldn’t adhere entirely to the “build to print” philosophy, as intentional alterations due to lessons learned and unanticipated changes resulting from emerging concerns had to be addressed. But they used their time effectively from both a technical and scientific standpoint. Ultimately, the team stayed true to their intentions while maximizing the value of their mission.

Follow OCO-2 on its mission.