

Kepler Mission Operations Response to Wheel Anomalies

Kipp A. Larson¹, Katelynn M. McCalmont², Colin A. Peterson³, and Susan E. Ross⁴
Ball Aerospace and Technologies Corp., Boulder, Colorado, 80301

The Kepler mission completed its primary mission in November of 2012 and was approved for an extended mission two years beyond that, with the option of two more if the spacecraft remained healthy. One of the four reaction wheels failed during the primary mission but science collection continued successfully until six months into the extended mission when a second reaction wheel also failed. The steps taken to lengthen the life of the second wheel prior to its failure are outlined, as are the tests undertaken to attempt to recover the two failed wheels. Once the decision was made to abandon a three-wheel mission, a new method of attitude control using solar pressure to balance the spacecraft in roll was developed, and a test campaign was undertaken that resulted in an in-flight demonstration of planet detection less than five months after the initial concept was devised. The details of that test campaign are given, as well as a discussion of the changing mission operations philosophy that allowed the aggressive test schedule to succeed.

I. Introduction

THE Kepler mission was launched in April of 2009 into an Earth-trailing orbit with the goal of searching for Earth-sized planets in the habitable zone around sun-like stars. It accomplished this by staring at one field of view for long periods of time and using the transit method to detect planets. When a planet passes in front of a star in the line of sight of the telescope, the reduction of light from the star due to the planet can be measured by the telescope. The size and duration of the dip in light from the star can determine the size and orbital period of the planet. The Kepler field of view was ten square degrees, and the initial star field had 170,000 targets. To date, the Kepler program has identified over 3,500 planet candidates.

In order to detect Earth-sized planets the instrument must be able to detect changes in brightness on the order of 30 parts per million. This performance is a combination of the sensitivity of the detector Charge Coupled Devices (CCDs), the low instrument electronics noise, the natural variation of the star light, and the steady pointing of the spacecraft. For the entire primary mission of Kepler these factors were sufficient to allow the detection of planets Earth-sized and smaller. The steady pointing of the spacecraft was made possible by the use of its four reaction wheels and a set of Fine Guidance Sensors (FGS) co-located with the science CCDs.

The Kepler mission operations team, managed by NASA Ames Research Center, consisted primarily of four full time employees at Ball Aerospace and three more at the University of Colorado's Laboratory for Atmospheric and Space Physics (LASP). The team was sized to successfully run Kepler's operations through the end of the nominal mission and into its recently approved extended mission, and was based on the fact that after four years of looking for planets the operations were well understood and considered to be routine. While other Ball subsystem engineers were available to assist with anomaly investigations, the team size and budget were kept to a minimum. Any new products or operational approaches were closely scrutinized for need, and even for small changes development would take many weeks to go through the necessary reviews and testing. This minimalist approach would be greatly tested soon into Kepler's extended mission.

II. First Reaction Wheel Failure

In July of 2012 Kepler's reaction wheel number two failed unexpectedly. An Anomaly Review Board (ARB) was convened to investigate the cause of the wheel failure and look for possible mitigations that could be employed to prolong the life of the remaining wheels.

¹ Kepler Mission Operations Manager, 1600 Commerce Ave, Boulder, CO 80301TT-2, and AIAA Senior Member.

² Kepler Lead Flight Operations Engineer, 1600 Commerce Ave, Boulder, CO 80301TT-2, and AIAA Member.

³ Kepler Deputy Mission Operations Manager, 1600 Commerce Ave, Boulder, CO 80301TT-2, and AIAA Member.

⁴ Kepler Flight Software Lead, 1600 Commerce Ave, Boulder, CO 80301T-2, and AIAA Member.

By design, the spacecraft only required three reaction wheels to achieve the required pointing accuracy. Kepler's mission depended on collecting as much continuous data as possible. As a result, a decision was made to turn the wheel off and continue collecting science data, rather than take time away from the science collection by further testing wheel two. A diagram of the Kepler spacecraft is shown in Figure 1.




Figure 1. Kepler Spacecraft. *The spacecraft attitude is controlled by reaction wheels and thrusters on the spacecraft bus.*

III. Overview of Mitigations and Operations Changes

After the failure of the first reaction wheel, the spacecraft was successfully recovered to taking science on three functioning wheels. However, those three wheels provided no margin against another wheel failing, as one wheel is required for each dimension to properly maintain the vehicle's attitude. To reduce the risk of losing a second reaction wheel, several mitigations were put into place to preserve the life of the three remaining wheels. Since the Kepler mission relies heavily on extremely accurate and steady pointing, the nominal Kepler operations strategy was altered in the three wheel mission to accommodate the necessary prioritization of the reaction wheel health.

A. Raising Wheel Temperatures

All four reaction wheels are mounted on the exterior of the vehicle. Two wheels are mounted on each of the two baseplates that face the sun most during nominal Kepler science. In order to keep the solar arrays normal to the sun, the spacecraft is clocked 90 degrees four times a year to allow it to maintain the same field of view. Throughout a Kepler quarter, a pair of reaction wheels transitioned from sunlit to shaded and the other pair transition from shaded to sunlit. Due to this orientation, the wheels were on a relatively warm part of the spacecraft, but they also saw large variations in temperature, depending on the time of the quarter. After the failure of the first reaction wheel, it was a desirement to keep the wheels warmer at all times. Each reaction wheel pairing has a heater on the baseplate that it is mounted to that is controlled by flight software. The operations team updated the heater setpoints in the onboard EEPROM twice after the failure of the first reaction wheel to keep the wheels running warmer throughout the entire quarter.

B. Wheel Speed Constraints

Another artifact of the wheels' physical orientation on the vehicle was their momentum loading throughout a Kepler quarter. As the season progressed and the sun vector traversed the solar arrays, certain wheels would hold more momentum than others. As a result, at the edges of the quarter certain wheels would have minimal speeds and the others would be holding the majority of the momentum. This trend became more prevalent with only three functioning reaction wheels, which was an undesirable side effect. The vendor of the reaction wheels recommended keeping the wheels out of the sub-EHD (Elasto-Hydro-Dynamic) regime, which would keep the wheel speeds above 300 revolutions per minute (RPM). This constraint caused several operational changes, including how the vehicle's momentum was managed at the beginning and end of each quarter.

In order to offload some of the momentum from the wheel that would traditionally carry the brunt of the momentum, momentum biases were put into place at the reaction wheel desaturation events. Momentum biases shift the total angular momentum vector of the spacecraft such that the wheel speeds would be spread more evenly among the remaining three wheels, keeping the wheels out of the sub-EHD regime. One momentum bias was used during the first few desaturation events of the quarter, a second was used for the majority of the quarter, and a third was used for the final few desaturations. The operations team watched the wheel speed and momentum vector trending as the quarters progressed and made a call as to when a bias switch command was needed.

The second operational challenge that arose with the need to keep the wheel speeds high was performing maneuvers. Twice during the quarter and at the end of each quarter, the vehicle maneuvered to Earth-point from the science attitude to downlink the recorder on the high gain antenna. When the vehicle slewed, the reaction wheels provided the angular momentum necessary to perform the maneuver. In order to do that, the three wheels exchanged momentum, which often caused one or more wheels to lose speed or even traverse through zero RPM. Because of the sub-EHD constraint, it was desirable to keep the wheel speeds high before, during, and after the maneuver.

The maneuver times were known well in advance, so they could be combined with the scheduled reaction wheel desaturations to predict the momentum state of the vehicle at the start of the maneuver. Test results from the ADCS Matlab model and the System Test Bed (STB) simulation were combined with an analysis of known trends of the wheel speeds during the season to predict the speeds during and after the maneuver. If necessary, desaturations could be deleted, biased, or moved to after the maneuver to keep the wheel speeds above 300 RPM. The operations team worked with the Attitude Determination and Control Subsystem (ADCS) engineers to determine the best course of action.

C. Wheel Speed Reverse Biasing

The momentum loading profile was obviously changed when the mission went from four wheels to three. Traditionally, the wheels ran with repeatable tendencies each quarter with a pair of wheels always running in the positive spin direction and the other pair always running in the negative spin direction. One wheel pair would gain wheel speed magnitude as the quarter progressed and the other pair would lose wheel speed magnitude. In the three wheel mission, the momentum needed to be managed in a similar fashion, which resulted in one wheel always spinning positively, one wheel always spinning negatively and the last wheel changing directions mid-quarter. As part of the mitigation to possibly prevent another wheel failure, the two wheels that spun mono-directionally were reverse biased, meaning that they would now always spin in the opposite direction than they spun in the nominal mission. The reverse biasing strategy was achieved by tweaking the angular momentum vector and was planned out in combination with the sub-EHD regime avoidance.

D. Predicted Second Wheel Failure

The failure of the first reaction wheel was recorded in high resolution on the onboard recorder. As the vehicle was left without redundancy in its actuator set, the failure and the months leading up to it were studied to determine any warning signs, lessons learned and the failure profile. The wheel failure caused the operations team to examine the three remaining wheels with intense scrutiny. The figure of merit that was most used to evaluate wheel health was plotting the reaction wheel torque command versus speed to pick out any deviations from the normal trend. In healthy wheels, the amount of torque command needed to spin a wheel at a given speed is predictable.

In December of 2012 warning signs of elevated friction started to crop up in reaction wheel number four. Over a two month period the elevated but constant friction levels began to show occasional spikes to even higher levels, similar to those observed prior to the wheel two failure. These observations led to the predicted failure of the second reaction wheel.

IV. Wheel Four Failure

In anticipation of this potentially mission ending anomaly, the operations team at Ball had already been carrying out development testing of new modes of operating the vehicle using thrusters only. The first mode was similar to a wheel rest mode that was used as a mitigation step in January 2013. It was being investigated as an alternative to Safe Mode because the normal Safe Mode utilized wheels only, and positive attitude control would not be possible using only reaction wheels one and three. If a Safe Mode entry occurred on two wheels, there was a high probability that the spacecraft would quickly enter a lower-level Safe Mode controller known as Emergency Mode Control (EMC).

Without positive attitude control, EMC may be entered as a result of an uncontrolled drift of the spacecraft. EMC may be triggered by either an undervoltage condition, or by tripping a sun-avoidance threshold. An undervoltage condition would result if the solar panels drift away from direct sunlight resulting in a negative power balance condition. A sun-avoidance fault occurs if the angle between the sun and the telescope boresight enters a keep-out zone, to prevent direct or stray sunlight from entering the telescope, which would permanently damage the photometer. In either case EMC would take control.

However, EMC is an extremely undesirable mode with a very small set of low-rate telemetry and little commandability. Positive attitude control would be maintained using only thrusters, but this mode was only planned for short-term usage. In addition to being difficult to recover from, the burn rate in EMC would deplete the remaining fuel, leaving no options left for the spacecraft. The operations strategy was designed to avoid EMC at all costs. NASA Mission Management directed rapid closure of the thrusters-only Safe Mode development testing on March 27th, 2013. The new mode was called Thruster-Controlled Safe Mode (TCSM) and was enabled on the spacecraft on April 12th, 2013.

A. Thruster Controlled Safe Mode

TCSM is similar to normal Safe Mode. It maintains positive power and thermal balance by initiating a 90 minute rotation about the spacecraft solar panel normal (+Y-axis) and controls it to the solar vector. It controls using thrusters only, and uses half of the thrusters (5-8), which were always available during the nominal mission for EMC. The implementation of TCSM is relatively simple thanks to the ADCS control table software architecture, which allows for selecting pre-built attitude control actuator sets. By switching the “acquire sun” Safe Mode actuator set from the normal all reaction wheels set to an actuator set that includes only thrusters five through eight, TCSM can be activated automatically upon any Safe Mode fault trigger. Significant testing was required to ensure it would behave as expected considering it would most likely be initiated out of contact.

It was predicted based on Matlab modeling and STB simulations that the fuel usage was approximately seven thruster on-time seconds per hour, which while better than EMC, was still quite inefficient. At that time, the blow-down monopropellant hydrazine system had a tank pressure of 150 PSIA, and seven kg of fuel remaining. TCSM would burn fuel at a rate of about 0.8 kg every 10 days. Mission life following the loss of reaction wheel four would have been cut short to less than 90 days if TCSM was the only operating mode available. Therefore, the team needed to develop an entirely new spacecraft mode that did not require the maintenance of a Safe Mode spacecraft rotation. The new mode was called Point Rest State (PRS).

B. Point Rest State

PRS was an idea hatched by the Kepler ADCS team at Ball that was designed to minimize thruster usage by utilizing Solar Radiation Pressure (SRP) as a counter-balance torque against a one-sided thruster control deadband. By orienting the spacecraft in its known “Standby” attitude, with the solar panel normal (+Y axis) controlled to the

solar vector, and the instrument boresight (+X axis) controlled to solar ecliptic north, PRS only fires thrusters to maintain a detuned attitude error deadband (15° about the Z-axis, and 5° - 7° about the X and Y axes). Due to the spacecraft center of mass being offset toward the spacecraft aft, SRP causes the top end of the spacecraft to tip away from the sun. The attitude control system briefly pulses its thrusters to counteract the SRP, until the SRP again overcomes the one-sided thruster torque. The spacecraft therefore exhibits a bobbing motion. This mode also allows for motion about the low-gain antenna boresight (-Z axis) keeping the antenna pointed toward the Earth, making communications with the Deep Space Network (DSN) stations possible at all times.

PRS was more of an implementation challenge than TCSM. While the ADCS software control table architecture allows for many permutations of configuring parameters, PRS required utilizing spare entries in configuration tables to make it work. The foresight of the flight software designers at Ball to provide the spare parameter space for adding unanticipated data table entries was another requisite for achieving success going forward. The team began using the spacecraft in ways that were only meant for development and integration testing. Yet, the architecture made all this possible without a single costly flight software update.

PRS was a Point ADCS mode, which returned the spacecraft to a more nominal configuration than TCSM, a safe mode. Besides the detuned attitude control bandwidth, PRS utilizes all eight thrusters, with a smaller pulse-width of 100ms, as compared to 200ms for TCSM. Use of all eight thrusters also required warming up all catalyst bed heaters and activating both RCS latch valves. The star trackers were activated for precise attitude determination, rather than relying on the Coarse Sun Sensor (CSS) system and Inertial Measurement Unit (IMU) while in TCSM. The new mode was modeled on the System Test Bench (STB) giving a thruster usage of 0.4 thruster on-time seconds per hour. This vast improvement could provide almost five years of mission life in PRS.

PRS actually over-performed once deployed, at less than 0.2 thruster on-time seconds per hour. This was a new operating mode that the spacecraft could stay in for a long time, providing the entire operations team what they needed the most – time. Once the impending failure of reaction wheel four eventually happened, the team would need time to troubleshoot the failure, decide on an attempt at wheel recovery, and ultimately, figure out how to operate a meaningful science mission with the Kepler spacecraft while using only two functioning reaction wheels.

With TCSM in place, the team got to work on completing development of PRS. The original aggressive development schedule had PRS being uplinked and ready for activation on the spacecraft on June 7th, 2013. However, while an intentional on-orbit test of TCSM was never planned, on May 3rd, 2013, the team was surprised by finding the spacecraft in the new TCSM mode at the beginning of a scheduled DSN contact. The team was prepared to see that the anticipated failure of reaction wheel four had triggered the Safe Mode, but the cause was actually due to a star tracker fault, which is something that the team had become accustomed to managing throughout the life of the mission. TCSM had performed as advertised, as it maintained Safe Mode on thrusters at a fuel-burn rate of about seven thruster on-time seconds per hour. The unexpected on-orbit test of TCSM was a success, and there was still life left in reaction wheel four, as the spacecraft was even returned to collecting science using three-wheels.

This was, however, a wake-up call that PRS had better be ready soon. The PRS schedule was moved up a month, putting the Ball team in a race to beat the reaction wheel four failure. The decision to quickly deploy PRS marked a stark shift in risk posture for the team, and would serve as the template for how future operations would be carried out. This new operations philosophy required quick turn-around of operations products, less documentation, less oversight, and more autonomy given to the operations team to make the decisions and carry out operations as necessary to give the Kepler spacecraft a chance at a future.

After recovery from the first on-orbit TCSM, it was not long until it would be exercised again. On May 9th, 2013, the final PRS flight products were uplinked to the spacecraft. There was one other task that needed to be executed to make the transition to PRS possible, and this was to setup the new PRS parameter tables. This was planned to be performed at the next contact. At the beginning of the next DSN contact, on May 14th, the spacecraft was found to be in TCSM. This time it was the reaction wheel four failure. The team quickly decided it was best to complete the setup of PRS on the spacecraft after recovery and transition to it in order to save fuel. On May 15th, 2013, PRS was activated for the first time.

With the spacecraft in PRS as its new nominal mode, the team dusted off some of the previous development work for TCSM to become more fuel-efficient. Due to the expediency required for deploying the TCSM to the spacecraft, the ADCS control parameter trades were not entirely closed out, as there were some efficiencies left on the table. Some more STB testing confirmed that tweaks to the ADCS control gains would provide a more efficient TCSM that only used about three thruster on-time seconds per hour, or better than 50% fuel savings. This new mode, called Fuel-Efficient TCSM (FETCSM) was uplinked to the vehicle on June 27th, 2013.

V. Wheel Testing

After the second wheel failure another Anomaly Review Board was convened to decide what the next steps would be. The board decided that a series of tests would be conducted on both of the failed wheels to assess their health and viability for potential use. At this point Kepler had already achieved its primary mission lifespan and was into its two year extended mission. The program had to go back to NASA's Senior Review Panel in early 2014 to demonstrate it was healthy in order to pick up the other two years of its planned extended mission. As a result, there was a strong desire to make a decision quickly as to whether or not the wheels could be somehow repaired, or if a new mission concept might be necessary.

Both the wheel testing and later the attitude control development discussed below made use of the highly configurable Kepler ADCS. Most aspects of this control system are table driven, permitting significant reconfiguration without requiring flight software updates.

The wheel control algorithm provides the possibility of several different modes for wheel control. For flight testing of the failed wheels, all the wheels needed to be removed from use as attitude control actuators. This was already accomplished with the activation of TCSM and then PRS following the wheel failure. The next step was to put the wheels into a constant speed test mode which replaces the attitude control wheel speed request with a commanded test value. Wheel control gains and limits were adjusted to drive the wheels to the commanded test speed using the maximum commandable torque. Ground test data at the wheel component level and at the system integration and test level were reviewed to assist in evaluation of limits and expected results.

These special wheel control capabilities had not been used on the Kepler mission and had never been used in quite the same way they were for this special testing. Rapid prototyping of wheel control tunable features was implemented on the STB to develop the methodologies and configurations employed in the flight system for wheel testing.

Tuning of the approach included some reconfiguration of system fault protection and evaluation of the safety of fault protection during the planned wheel test progression. Particular attention was paid to the ability of the automated spacecraft safing mechanisms to retain attitude control authority if a safing event was triggered with wheel testing configurations in place. In addition, test abort contingencies were prepared. Wheel testing was performed with the motor temperatures held at the warmest possible levels. Relevant telemetry data was configured at the maximum rates for both real-time and stored data to facilitate real-time and post-pass trending analysis. The wheels torque command and measured speed data from the spin down periods were used to estimate the wheel friction and make the next decision.

A. Wheel Four Test

Wheel four was targeted for testing first because it appeared to totally "seize", and it was thought it could be easily eliminated as a candidate with which to move forward. A progression of testing was planned. First, determine if wheel four would spin in either direction. The commandable torque limit was raised from 0.2 NM to the maximum possible in the hardware, 0.26 NM. Wheel four was first driven at this new max torque to a commanded speed of +300 RPM for 30 seconds, then no torque commands for 30 seconds, and then driven at max torque to a commanded speed of -300 RPM for 30 seconds. The decision process after this test was that if the wheel did not move or moved a little and stopped during the 30 seconds, the team would then proceed to a "rocker" test that would alternate commanded direction in two second intervals for one minute to try and break it free. Alternatively, if the wheel moved and got to speed with something less than max torque, download the higher rate stored telemetry data and assess the next steps. The results of this first test were that the wheel did not move in the positive direction (the same direction it was turning when it failed) but did move to the commanded speed in the negative direction. The rocker test was therefore not executed.

Because wheel four moved freely in one direction, it was decided to retest it to see if moving it had freed it up in the other direction. A repeat of the first wheel four test was performed. Wheel four moved to the desired speed in both directions.

B. Wheel Two Test

The objective with wheel two was to assess whether it would perform consistently enough to be considered for reinserting into the controlling actuator set. This testing used constant speed commands with nominal torque parameters (0.2 NM max) to move +300 RPM for five minutes, no torque for three minutes to spin down, -300 RPM for five minutes, and no torque for three minutes again to spin down. Wheel two achieved the commanded speed in both directions.

C. Three Wheel Test

Analysis of the torque and speed data of the wheel two and wheel four testing revealed that both wheels were operating with wheel friction an order of magnitude larger than they had in the nominal mission. It was believed that the elevated friction levels of both wheels were too high for normal pointing control. However, the stability of this friction and the limitations it might place on pointing performance were still not known. A final test was proposed to return the spacecraft to three-wheel-controlled pointing to assess performance, and determine if a return to a wheel-controlled Fine Pointed mission was worth pursuing. Control parameters were adjusted to widen the control system integrator limits and the attitude error limits to allow the system to adapt to potentially varying levels of friction. This test was performed at the Standby attitude (with the telescope pointed to ecliptic north) using wheels one, two, and three.

The system was transitioned from PRS to Coarse Pointing using wheels. Pointing was maintained with all telemetry looking nominal for several hours. The torque levels on wheel two then began to elevate. Within 30 minutes, the maximum torque was being applied and wheel two could not keep up with the desired speed of the attitude control subsystem. After another 30 minutes of Coarse Pointing, the system Safed itself due to attitude error limit being exceeded. This was the final use of either wheel two or wheel four for the Kepler spacecraft and therefore the final attempt at three-wheel attitude control.

The three wheel system test on August 8th and 9th, 2013 was the transition point between the wheel testing efforts and the two wheel system development. The vehicle could not reliably maintain its attitude using three wheels because the friction in the failed wheel was greater than the commandable torque authority in the wheel. This proved obvious in a short duration test, showing that operating on three wheels was not a viable option in any capacity. Therefore, the team's mindset changed as a result of this test from trying to revive a three wheel mission to urgently developing a new precedent: a two wheel pointing mission.

VI. K2 – A New Mission

It was widely accepted by people on the program that the second wheel failure would mean an end to the Kepler mission. While there was little hope of continuing to find planets, the operations team did the one thing it could do, which was to gain the time necessary to take a shot at finding a new solution. Without the new fuel efficient thruster controlled Safe Mode and Point Rest State, the spacecraft would have run out of fuel in a matter of months. Now, Kepler could remain in Point Rest State for several years if necessary. The team then turned their attention to how to use a spacecraft with only two good reaction wheels.

The team had only five months to come up with a new concept and demonstrate it before the proposal was due for NASA's Senior Review. A call went out to the community in August of 2013 to solicit ideas for new missions for Kepler. At the same time, Doug Weimer at Ball Aerospace developed a novel approach to make use of Kepler's remaining wheels¹.

As shown in Figure 2, the concept involves aligning the spacecraft axis with the anti-velocity vector such that the solar arrays are normal to the sun vector. Due to the physical symmetry of the solar arrays the solar photon pressure balances the spacecraft in roll, allowing stability in the one axis that the two wheels cannot control. As the spacecraft moves in its orbit, it can remain pointed at one field of view as the yaw angle varies. The yaw angle is limited at one extreme by the need to prevent sunlight from entering the telescope, and at the other extreme by the minimum solar angle necessary for the array to remain power positive. These extremes allow for a dwell time of up to 82 days at a single field of view. The thrusters are used sparingly to make small adjustments to the roll angle due to the imperfect symmetry of the spacecraft. The wheels are also loaded with momentum every two days using a so-called "wheel resaturation" or "resat" operation to help keep the spacecraft more resistant to any imbalance in solar pressure along the solar arrays.

As a result, the concept allows for several more years of operations looking at brand new star fields along the ecliptic with pointing accuracy only somewhat degraded from the original mission. Before the new concept could be proven, however, the team first needed to learn how to balance on two wheels.

A. First Two Wheel Test

By August 29th, 2013, just three weeks since the decision to move forward with the development of the two wheel mission and just two weeks since the idea had been presented from the ADCS engineers, a two wheel pointing test was executed on the vehicle. The flight operations team worked closely with the ADCS engineers to take an abstract idea and translate it into the specific Kepler language of parameter tables and never before used test commands. The previous wheel testing efforts had not gone to waste; in fact, many of the same commands were used in the development and execution of the two wheel point testing. The STBs were used heavily in the two week

development time to work out commanding and fault protection strategies and the simulator was trusted to a high degree of fidelity. There were, however, many unknowns entering the two wheel test because not only had the vehicle never been flown in this manner, but it was never designed to use only two wheels and thrusters for attitude control.




Figure 2. K2 Operational Concept. The K2 mission uses the remaining two reaction wheels to control cross-boresight motion, and solar photon pressure to balance roll motion. (Image: NASA Ames)

For the two wheel test the vehicle was transitioned from thruster control to wheel control once the wheels were fully spun up using the same test commands that were used during previous wheel testing. The vehicle was monitored for several hours in two wheel mode before it transitioned back to thruster control and the wheels spun down. This test proved that the vehicle could be controlled using only two wheels and the mode was indeed stable. The test also clarified that the vehicle did not need to be run in a hybrid mode with thrusters controlling one axis and wheels controlling the other two. While it is true that three actuators are needed to control three dimensions, in this

design the sun plays the role of the third actuator. The about boresight roll axis of the spacecraft was not in the control loop at all during the two wheel test. This fact was a big step in the development of the new mission.

B. First Hybrid Control Test

With the momentum of the two wheel concept being vetted on orbit, the team worked towards a test on September 10th and 11th, 2013 to demonstrate the ability to adjust the about boresight roll axis attitude. Although the new concept relied on the sun to balance the third axis of motion, the reality is that perfect balance is impossible to achieve. The roll axis will drift over time and the magnitude of the drift is proportional to how far away the vehicle is from that perfect balance point. Therefore, the team devised a strategy to periodically use the thrusters to correct for any roll offsets or drift.

The thrusters had never been used in this capacity before. Their main use during the prime mission was to desaturate the reaction wheels after several days of momentum build up from solar pressure. They could also be used in Safe Mode, Emergency Mode and Detumble, but the system was never designed to be used for precise pointing adjustments.

The strategy centered on using pairs of thrusters to achieve the desired amount of correction. The vehicle has eight thrusters, each of which has a force vector that has some components in the body frame X, Y, and Z axes. By pairing thrusters together, one or more of these force vector components can cancel, creating a near pure adjustment in one of the axes. The X body axis of the vehicle is the about boresight roll axis, which is the uncontrolled axis in two wheel mode. Therefore, thruster pairs were devised that would be ideal to adjust the X-axis error, both in the positive and negative directions. These X-axis adjustments using thruster pairs soon became known as “tweaks”.

There were two flavors of tweaks entering the September 10th test: manual and automatic tweaks. Manual tweaks used test commands that fired thrusters in an open loop fashion. The manual tweaks specified which thrusters and for how long to burn, so the thruster pair and delta-V calculations had to be done beforehand with input from the ground. The automatic tweaks placed the thrusters in the control loop for a short amount of time, making the system hybrid controlled for that time. During that period, the flight ADCS control law calculated the control torque needed to fix the X-axis attitude error in order to command specific thruster on times. The thruster pairs used in the automatic tweaks were specified in an onboard parameter table.

Both the manual and automatic tweaks were exercised on orbit during the September 10th test, as it was unclear which direction the mission would head in without the empirical test data. These tweak tests gathered valuable data on using the chosen thruster pairs, as well as how the thrusters performed when only firing for very short periods of time (<100ms). The test also proved the repeatability of the two wheel entry, stability and exit.

To top it all off, the vehicle experienced an unrelated Safe Mode between the August 29th and September 10th tests that the flight operators had to recover from all while preparing for the unprecedented on orbit test. In the past, a Safe Mode entry would necessitate a series of Anomaly Review Board meetings and analysis before the team was given the green light to recover. This could take between a few days to several weeks depending on how quickly the anomaly was understood. In this rapid development approach however, once the cause was identified the team was able to recover and resume testing within one day.

C. First Hybrid Control Tests at the Velocity Vector

The first on-orbit test of the tweaks yielded results that helped formulate a two day test that took place from September 24th to September 26th, 2013. While the manual tweaks helped calibrate the thruster pair performance to the ground models, it was clear that the vehicle would need the capability to adjust the X-axis error without a person in the loop. Therefore, automatic tweaks were the only option moving forward.

The automatic tweaks that were performed using ground commands during the previous test showed promise, but they also demonstrated the need to have precisely timed onboard sequences to open and close the control window when the thrusters would be in the loop. This forced the operations team to converge on a set of parameters for the tweaks, as onboard sequences are more difficult to change quickly than commands in a procedure run from the ground.

The September 24th test implemented these new tweak sequences but it also broke new ground in that the vehicle was left in two wheel mode out of contact for the first time. The vehicle entered two wheel mode and tweaked on the 24th and was left until the contact on the 26th. This allowed the team to gain confidence in the longer term stability of two wheel operations. The test was successful and the vehicle was transitioned back to thruster control on the 26th while the team planned for the next test.

Thus far, all of the wheel testing and the two wheel tests had been performed at the Standby attitude: an attitude with a good communication link that was near to PRS. However, the plan for the proposed science mission was to operate with the telescope along the anti-velocity vector of the spacecraft’s orbit about the sun. This new science

attitude was roughly 90 degrees away from the Standby attitude and had no expected communication link with the Earth for a typical DSN contact with a 34 m station.

The next test, performed on October 10th-12th, 2013, would be the first time going to this new anti-velocity vector science attitude. Because the attitude faces the X-band communication antenna 90 degrees away from the Earth, all of the commanding necessary for two wheel operations, and anything else during those two days, needed to be sequenced. The two day test at the science attitude vetted these products and also continued to refine the parameters for the onboard automatic roll tweaks.

D. First Week Long Test

The next test was a weeklong test at the science attitude that took place between October 24th and 31st, 2013. The vehicle collected more science at the desired attitude, performed a resat for the first time in the blind and tested out new parameters for the attitude tweaks. The onboard products ran smoothly and, for the first time since the wheel failure, the vehicle ended the test at the Earth point, stored data downlink attitude.

In order to dump the solid state recorder, the high gain antenna must be precisely pointed at the Earth: a capability that was unproven in the new mission concept. STB studies had shown that an adequate communication link could be established with the high gain antenna by just using thrusters to hold the attitude. This theory was proven at the end of the first weeklong science test and the data was downlinked from the SSR.

For the first time since the wheel failure, the fine guidance sensors were used to try and glean more precise attitude data from the focal plane. However, FGS apertures did not see any starlight during the entirety of the weeklong test because the vehicle was not pointed precisely enough with the star trackers. This led to the team pursuing new options to fully understand the thermal environment of the anti-velocity vector science attitude to properly align the focal plane with the star trackers. This test also showed that the drift rates due to unbalanced solar torque were much different than those seen in the earlier October test. These two tests had the focal plane pointed at the same field, which caused the solar elevation angle with respect to the solar array normal to change. This created urgency to understand how the drift rate varied with respect to the entire range of solar elevation angles.

While this test revealed that there was still much work to be done, it did succeed in providing the first real indications of pointing performance in this new mode of operations. This included what became known as Kepler's "second light" image (Figure 3), a full-field image of the stars in the focal plane that looked nearly as sharp as the ones generated during Kepler's prime mission.

E. Mapping the Solar Ridgeline

At this point in the development, it was clear that the scientists needed to understand the magnitude of the uncontrolled axis' drift. The two October tests showed that different elevation angles had dramatically different drift rates due to the physical geometry of the spacecraft not being perfectly symmetric about any axis. The proposed science mission had campaigns that could cause the vehicle's orientation with respect to the sun vary by up to 82 degrees, meaning that there was a large span of angles that needed to be mapped for typical drift rates. Because the roll drift error of the vehicle increased on either side of the perfect balance point, and this balance point moved as it traversed the 82 degree swath needed for science campaigns, the line became known as the solar balance ridgeline.

A series of six tests were planned within one month to map key points along the ridgeline to better understand the magnitude and variation of the drift. The tests all had essentially the same concept of operations, but each one required approximately eight hours of continuous data to properly map the ridgeline at a specific solar elevation angle. Each test maneuvered the vehicle from approximately solar array normal (Point Rest State) to the necessary solar elevation angle about the Z axis of the vehicle. These maneuvers were performed on thrusters. Because science data was not being taken for these tests, it was not necessary to maneuver to the anti-velocity vector. This was desired due to the lack of communication link there. Once at the proper solar elevation angle, the wheels were spun up and given control. While on wheels, the X-axis roll drift was monitored and recorded to understand the drift at that elevation angle.




Figure 3. K2 Full Field Image (FFI). *This image shows all of the stars on the focal plane from the first set of tests at the anti-velocity vector. (Image: NASA Ames)*

One test was conducted on each scheduled contact for the entire month of November, gathering data at different solar elevation angles along the way. This test campaign ended in an optimized solar ridgeline “map” that could be used to inform the drift and desired attitude of the spacecraft throughout an entire campaign. In the end, while the ridgeline was found to be nonlinear, the results did not bring up any red flags in the plans to move forward with the new science mission.

While the rapid-fire testing of the solar ridgeline was ongoing in the month of November, the flight operations team was working in parallel to conduct another weeklong test at the science attitude. The October science tests were successful in proving the validity of theory and flight products, but there was increasing urgency to get the vehicle pointed precisely enough to utilize the Fine Guidance Sensors. Using the FGSs would reduce jitter, increase stability and eliminated the issue of thermal misalignment between the focal plane and the star trackers. Therefore, a test was performed from November 13th through the 19th to attempt to align the FGS apertures with stars on orbit and fold that attitude solution into the new two wheel attitude control scheme.

This test attempted not only to collect data with the FGSs, but also to transition the attitude determination of the vehicle from the star trackers to the FGSs, what is known as transitioning from Coarse Point to Fine Point. While the initial alignment of the vehicle was good enough to get it into Fine Point, a configuration error caused the vehicle to enter Safe Mode on the last day of the November weeklong test. The flight operations team managed to recover the vehicle from Safe Mode, maneuver to stored data downlink and dump the contents of the solid state recorder without causing any slip in the testing schedule for the ongoing solar ridgeline calibration tests. Even though the weeklong

test did result in Safe Mode, there was added light at the end of the tunnel in that the vehicle had successfully transitioned to Fine Point: a necessary step in the evolution of the new science mission.

VII. K2 – Commissioning Tests

With all the major elements of performing the new mission concept demonstrated, it was time to put them all together in a single test. The development schedule was leading up to a 30-day test in January 2014, and the 82 day FOV window for that attitude opened in mid-December. This afforded the opportunity to get a sneak peak of the 30-day test FOV. Each campaign looks at an entirely new star field, and fine pointing adjustments must be made to center the stars in the science apertures and also in the FGS apertures. This test would allow the science team to exercise the analysis of the science data to generate an attitude update for fine-tuning of the spacecraft pointing.

A. Preview of the New Field of View

The plan was to maneuver to the new FOV attitude and execute a six-day mission from December 12th-19th, 2014. The spacecraft would be in the two-wheel mode demonstrating a regular resat interval every two days, and closed-loop thruster X-axis roll tweak every 12 hours. There were also planned attempts at transitioning to Fine Point guidance, having solved the issue from the previous test. But this all relied on the guide stars falling on the FGS star apertures.

The spacecraft is in coarse-point prior to attempting fine-point. Coarse pointing utilized the star trackers, which are mounted on the aft deck. Fine pointing utilized the FGS CCDs, which are located on the focal plane along the telescope boresight. During the early part of the 82-day campaign, the sun is more directly pointed on the aft portion of the spacecraft. Throughout the 82-day campaign, as the sun vector moves forward along the solar balance ridge, the thermal characteristics along the telescope change, and the coordinate rotation between the star trackers and the FGS sensors on the focal plane is variable with temperature. This makes the transition from the star trackers to the FGS dependent on loading the correct star tracker to telescope boresight alignment quaternion as a function of solar angle throughout the campaign.

The closest thermal conditions from the Kepler prime mission to the early campaign of the new mission was during the roll to the summer Kepler prime season. On the Kepler prime mission, the Star Tracker to boresight quaternion was changed for each season. So the team loaded the summer season quaternion to have the best chance at successfully transitioning to Fine Point. The solar angle difference between Kepler prime summer season and the beginning of the December test, however, was still greater than 10 degrees. Even so, it was a good test of all the tuning parameters and their sensitivities hopefully working in concert to acquire and maintain Fine Point guidance.

The test planned two five minute and two one hour periods of Fine Pointing. Then it was to transition and stay in Fine Point for the final several hours of the test before returning to Earth point attitude. The first four attempts were successful and Fine Pointing was successfully demonstrated. The final attempt, however, was not successful at maintaining Fine Point Lock. The initial discouragement felt by the team was quickly replaced by elation when the cause was determined. It turned out the pointing was so good that it was within the threshold of the Kepler prime mission on three or more reaction wheels which triggered a Fine Point Lock flag. Once the flag was set, if the attitude error grew beyond that limit fault protection would automatically command the spacecraft back to Coarse Point. It turns out this was an important piece of data, as the team now knew they needed to simply change the Fine Point Lock threshold to be appropriate for the two-wheel pointing mode.

B. Making K2 More Fuel Efficient

The holiday season for the Kepler mission was never too kind for allowing holiday time off, due to historically untimely anomalies, and this year was not too different. While the team had a break from actual testing, and no spacecraft anomalies to attend to, they needed to prepare for the 30-day test in January, and another problem became apparent. There was a fuel resource issue. All the testing up to this point was necessary, but it was costly in fuel. Rough analysis showed there was less than two years of fuel remaining to maintain a rolling 82-day campaign schedule. This would not be looked upon favorably by the NASA review board that would be tasked with approving the two-year extended two-wheel mission.

The engineers needed to sharpen their pencils and figure out the actual fuel budget remaining. There was also a need to look across all the elements of the new mission to find fuel efficiencies. A trade study was carried out to determine the optimal X-axis roll tweak interval to maintain pointing requirements and minimize fuel cost. This was dependent, however, on how far the solar vector was from the solar balance ridge. The balance ridge is not completely linear so the roll drift rate varies throughout the campaign. The trade study not only provided guidance

for what the tweak interval should be, but also how often thrusters will actually fire, and thus, what the fuel cost would be.

In addition, looking closely at both the resat and tweak performance over the previous tests revealed some techniques that could be utilized to again save fuel. These new methods were tested on the STB, and planned to be executed at the next spacecraft test. The newly informed fuel budget, including the efficiencies gained, predicted there was over three years of mission life left, which translated to more than 12 total campaigns. The new mission, which had recently been dubbed K2, had new life, again.

C. First Pre-Campaign Checkout

The K2 baseline concept of operations included a checkout period at the beginning of each campaign. This would involve maneuvering to the new campaign attitude, thermally settling, transitioning to Fine Point, and taking a few hours of science data. Then it would maneuver to Earth point and downlink the data. This was so the new FOV pointing could be analyzed to determine the optimum pointing for science. The attitude update would be made, and the spacecraft would be maneuvered back to the campaign science attitude for the remainder of the campaign.

While the planned 30-day test started in mid-January, the team decided to demonstrate this checkout concept in early January by doing a three-day test from January 6th-9th, 2014 at the target FOV. This would allow them to demonstrate the new fuel-efficient resats and now optimal three-hour tweaks. It also allowed them to baseline all the mission products as flight certified onboard the spacecraft, which were up to now still development versions. The team was brimming with confidence that they were out ahead of the 30-day test. But the spacecraft still had other ideas. The checkout test concluded, but Fine Point was never achieved. The guide stars were never in the FGS POI apertures. This proved how important knowing the star tracker to boresight alignment quaternion was to achieving Fine Point.

The 30-day test was scheduled from January 16th to February 13th, but it became apparent that the alignment could change drastically enough, that even if a new attitude solution was derived based on the 10-day old data, it may not put the guide stars in the FGS POI apertures. The need for knowledge of the Star Tracker to boresight alignment quaternion throughout the whole campaign was made apparent.

D. The 30-Day Test – Collecting Real Science Data

The team concluded the best option for starting the 30-day test was to not implement an attitude update, and instead return to using the fall season quaternion, which would be closest to the solar angle at the start of the 30-day test. This experience had a beneficial side-effect, however, in that it forced the team to come up with contingency scenarios. They devised a plan to store an abort sequence onboard. A 70 m DSN antenna actually provided enough signal margin that signal acquisition with the spacecraft was possible. The plan was to maneuver to the science attitude, thermally settle, and attempt transition to Fine Point. Then a 70 m DSN station was scheduled to provide telemetry about whether Fine Point was achieved or not. If not, the contingency sequence could be executed, and a maneuver to Earth point would occur and the data could be downlinked to determine the next move.

This contingency plan would turn out to not be necessary, however, as a different problem manifested itself when the spacecraft was found to be in Safe Mode at the beginning of the 70 m contact on January 21st, 2014. Upon recovery, the data showed the cause of the Safe Mode was due to a star tracker fault that the team quickly determined could be prevented by changing strategies of the fault management. There was another unrelated anomaly that was more concerning. The good news was that Fine Point was achieved and maintained for almost four days. The bad news was that it ended due to an anomalous shutdown of the photometer electronics as a result of an apparent glitch of its power supply. The FGS system was powered down as a result of the anomaly, and so the spacecraft transitioned to Coarse Point.

The power supply circuit that shut itself off was evaluated by Ball engineers and deemed safe to turn back on. Despite the hardware issue and the unrelated Safe Mode entry, the team was able to recover the spacecraft and start collecting science again within ten days.

The test also validated a new activity that might further reduce the fuel use. As part of the pointing checkout a set of reference pixels, which are a subset of the science pixel target set, were downlinked on X-band via a 70 m DSN station and used to synthesize an attitude adjustment to better line the stars up in the center of their apertures. The command for the adjustment was successfully sent in the blind via a 34 m DSN station. This method does away with the previously planned maneuver to Earth point attitude to downlink pointing data via Ka band. Without that maneuver to and from Earth point, the mission may be able to increase the number of campaigns by 10 percent.

The rest of the test went off without a hitch, and resulted in the spacecraft remaining in Fine Point for nearly ten days. The pointing that resulted was a factor of ten better than what had been seen at the beginning of the two-wheel tests. The new fuel efficiencies that had been built in to the tweaks, resats, and maneuvers were fully validated and

demonstrated a potential K2 mission lifetime of close to three years. In addition, one of the science targets for this test was WASP-28b that was a previously detected exo-planet. The K2 data, shown in Figure 4, clearly showed that the spacecraft was capable of detecting Jupiter sized planets with the potential for detecting even smaller ones². Despite the many issues encountered with this test, in the end it demonstrated that the K2 mission concept was proven, and that the team could prepare for the first full 82 day science campaign.




Figure 4. K2 Light Curve. Light curve detected in January 2014 for WASP-28b. (Image: NASA Ames/T. Barclay)

VIII. The Changing Approach to Mission Operations

The first signs of Kepler's second wheel starting to fail showed up in December 2012. As shown in Fig. 5, by May it was clear that failure was imminent, and a fix was put in place in a matter of weeks to ensure that the failure would not prove fatal. In June there was still lingering hope that one of the failed wheels could be resurrected, but the clock was ticking towards the December 2013 deadline for NASA's Senior Review. A rapid series of tests with a set of commands never intended for use on orbit proved that the three-wheel mission was indeed over. Not giving up, the team devised a new approach to pointing, and was able to take it from PowerPoint concept to on-the-ground proof of planet detection in less than five months. To understand how this remarkable timeline of success was able to be realized, it is helpful to consider the evolution of the team's approach to mission operations during this time.

During the commissioning of Kepler many anomalies were encountered. Each one was approached, appropriately, with a high degree of caution and conservatism in terms of the response. The team had very little operational experience with the spacecraft, and one wrong move could have ended the mission before it really began. With that mindset in place, any change to the baseline plan, which had been developed over several years, required a significant process of testing and review prior to approval.

This process would begin with filling out a multi-page change request form. This would define the activity and identify any operational prerequisites, relevant flight rules, affected subsystems, downlink requirements, and hazardous commands. The systems engineer and each subsystem lead would have to review the activity and procedures and sign off on it. The command script would be run several times on the STB, and any command products would be checked in to an internal tracking system to be configuration managed and reviewed by multiple

people besides the developer. This entire package would then be brought forward to Ball management, the Ames customer including the Kepler Mission Director, Project Systems Engineer and Deputy Program Manager, as well as the Jet Propulsion Laboratory (JPL) Project Manager and several other people from JPL program management. This entire group would then need to debate the finer details of the plan before approving it.




Figure 5. K2 Timeline. Following the reaction wheel four failure, progressive on-orbit tests were performed to characterize the ability to operate on two wheels and to demonstrate the viability of the K2 Mission concept.

It was hard to argue with that level of rigor considering what was at stake, and it was a credit to the entire program that people from top to bottom felt a very personal connection to the mission and a strong sense of responsibility to get things right. However, this methodical, conservative, and very expensive approach to running mission operations was entirely inappropriate to the task at hand following the second Kepler wheel failure.

The mitigation steps used after the wheel two failure followed a similar process of development. They were investigated and reviewed through an Anomaly Review Board, made up of Ball engineers, Ames management and engineering, the wheel vendor, and various consultants from JPL and industry, for several months prior to approval. The products required to implement them took weeks to develop after that, and then a few weeks more to fully deploy and evaluate.

Once wheel four started exhibiting symptoms similar to what had been seen in the review of telemetry after the wheel two failure, it became clear that the traditional process would no longer work. If wheel four followed the same failure timeline, then it was only a matter of weeks before the wheel would fail. Considerable time had already been used just to create a schedule to develop a strategy to allow the spacecraft to safe itself without losing control and then go to Point Rest State. That schedule had a timeline of two months to develop the new PRS mode, test the changes on the STB, review the products, get them all approved and put them on board. Because of the warning signs in the spacecraft wheel telemetry, the decision was made to condense this schedule to one week.

From that point on the test program proceeded on a roughly two-week cadence where the test concept was conceived, debated, tested on the ground, and implemented all within that two weeks. Long term goals were outlined, but the nature of the program meant that each step was something that had never been done before. The time to demonstrate a new mission concept was only a few months, so any significant setback might derail the entire effort. The program was therefore based on being as aggressive as the team could be without risking that a test would fail completely and have to be re-evaluated and repeated.

Each test was built on the success of the previous one, and the next test could not fully be planned until the results from the previous test were understood. This meant that the goals for the next test were continuously being re-evaluated up until a day or two before the test itself. While a few long-term high level goals were identified, the

program was essentially created as it went along. It was only because of this willingness and ability to be extremely flexible that success was achieved.

The situation that drove the progress towards a K2 mission was very unique. First, the necessity of generating operational products in time for the second wheel failure drove the team to use the minimal amount of process deemed safe to get the job done in time. Second, the delivery of a proposal to the NASA senior review meant that there was a finite amount of time to develop and test the new mission concept. Third, the expectations were very low from the community. When the second wheel failed, it was widely expected that the mission would be over and no new use for the spacecraft would be found that could justify the money necessary to keep it going. All of this led to the acceptance of a much greater risk profile than anyone had previously been comfortable with.

What this risk profile meant was not an abandonment of the types of checks and balances that had previously been employed, but rather the creation of an extremely slimmed-down version of them. For instance, rather than taking the time to produce PowerPoint slides to support frequent meetings with several layers of customer management, the Ball Mission Operations Manager held brief phone tag-ups with the Ames Mission Director on an almost daily basis. This was for status only – no in-depth questions, no slides, and no follow-on questions that the Ball team would have to take time to hunt down answers for. This way the customer always knew what was going on, but was not taking away team resources in order to dive into details unnecessarily.

In the more traditional view of mission operations, a great deal of time and effort is put into developing a plan, and any changes to that plan are not looked upon favorably unless they are absolutely necessary. In the development of K2, a completely new test was being devised, tested, reviewed, and implemented roughly every two weeks. This necessitated a great deal of flexibility. Test details could change rapidly over a couple of days, and occasionally at the last minute. Rather than having a set process for changes or hard deadlines, the team would collectively decide if a change was necessary, if there was enough time to test and review it adequately, and if the risk of a last-minute change was worth the potential benefit. The change could therefore be proposed, debated, and approved all in an afternoon within the Ball mission operations team.

The only rule was that at any time any team member at Ball or LASP could throw a “red flag” and halt the process if they felt the safe execution of the test was at risk. Each team member was empowered to propose a change, challenge a change or halt the test at any time regardless of the rank of the team member. Not only did this approach make certain that each team member felt that their voice was heard, it also proved to be tremendously effective. Every member of the Ball and LASP team managed at some point to make at least one good “catch” that would have otherwise resulted in a significant test problem.

The team, although small, contained a very helpful level of diversity. One member had been on the team for six years and had a good knowledge of not just how the spacecraft worked but what types of things had gone wrong in the past. One member had been on the program for three years with a background in flight software. She brought a history of work on other programs with similar architectures as well as an appreciation for the more rigorous configuration management of software. The two other members were new to the team at the time of the second wheel failure. One had a background in manned spaceflight, with a strong attention to detail and an appreciation of what can go wrong when decisions are not questioned sufficiently. The final team member was fresh out of graduate school with a strong ADCS background. She had worked at LASP as a student operator on Kepler and understood a lot about working on console and generating command scripts.

Having a diverse team proved crucial, especially since the team was so small. The variety of backgrounds and skill sets was a great asset, but perhaps the biggest benefit was that each person had different comfort levels with different aspects of the test development. The standard practice within the team for creating new operations products was to ensure that the person reviewing or approving the product was not also the developer of the product. This check was maintained, but without the accompanying documentation that had previously gone with it.

The team relied on each person’s level of comfort to decide which parts of new command products or activity steps needed the greatest scrutiny. If even one person was uncomfortable, the team would stop and discuss the concern, and do more work as necessary until the concern was alleviated. There were no real instances of “we’ve always done it, so therefore it must be okay”, or “we’ve never done that, so it must be a bad idea.” If each person on the team, with their own set of biases and varying levels of knowledge about the products, was comfortable, then it could be assumed that the products were ready to send to the spacecraft. Even with rapid turn-around and often little overlap between individual team member’s tasks, this proved to be very effective in eliminating “gotchas”.

Despite the extreme pace of product development, this sensitivity to the group’s level of comfort also led to periodically taking a step back to look at how reviews, documentation and configuration management were being handled. The goal was to ensure that the success of the tests was due to an effective, slimmed down process rather than simply “getting lucky” with no self-inflicted mistakes. Without the same level of oversight by the customer and

upper levels of management, this willingness to self-examine and question how things were being done was important to keep the team from letting success create blind spots that could create problems later.

Considering the rapid testing schedule and the small team executing it, it is remarkable that no significant in-flight mistakes were made. The few stumbles that were encountered were firmly in the category of learning how to do something that had never been done before. This slimmed-down process of mission operations product development may not be appropriate to the front-end of programs when time and money are more plentiful and risks are tolerated less, but it is a lesson in what can be done when creativity, flexibility, and hard work are allowed an environment in which to flourish.

IX. Conclusion

Flying a spacecraft with a small team is not easy even under nominal circumstances. That difficulty was exacerbated by the failure of one of the prime subsystems on the spacecraft that put the team in a potential mission-ending situation. Rather than feel rightfully satisfied that the spacecraft had completed its primary mission and move on, the team did the one thing it could do. They quickly gave the spacecraft the ability to survive the wheel failure and maintain a fuel-efficient attitude, allowing for the possibility, however remote, of finding a way to continue to collect science data of some kind.

With the spacecraft safely in Point Rest State, it only took one clever idea to give the team something to work towards. Given free rein to look for creative solutions and accept more risk, the team methodically built a test program that was aggressive, safe, and ultimately extremely successful. They were able to take the concept from back of the envelope to full flight demonstration in only four months. This was accomplished without flight software upgrades, without bringing on extra staff, and within the budget constraints of the nominal extended mission. The program that everyone had given up on was reborn with capabilities far better than anyone had hoped for, allowing the search for planets around other stars to continue.

Acknowledgments

The authors would like to thank Jeff Pugliano and Kevin Moore for keeping the System Test Benches working, and Mike Packard, Lee Reedy, Jason Gabbert, Crystal Salcido, and the students at LASP for many long days creating, reviewing and executing flight products.

References

¹Putnam, D., and Weimer, D., "Hybrid Control Architecture for the Kepler Spacecraft," 37th Annual AAS Guidance and Control Conference, Breckenridge, CO, Feb 2014, paper 14-102

²Howell, Steve B., et al, "The K2 Mission: Characterization and Early results," Submitted for publication to PASP, March 2014

This article has been cited by:

1. Katelynn M. McCalmont, Kipp A. Larson, Colin A. Peterson, Susan RossK2 Pointing Enhancements and Performance . [[Citation](#)] [[PDF](#)] [[PDF Plus](#)]
2. Colin A. Peterson, Kipp A. Larson, Katelynn M. McCalmont, Susan RossK2 Mission Operations: Finding Balance in Year One . [[Citation](#)] [[PDF](#)] [[PDF Plus](#)]