

Dellingr: Reliability lessons learned from on-orbit

Larry Kepko, Luis Santos, Chuck Clagett, Behnam Azimi, Dean Chai, Alan Cudmore, Scott Starin, James Marshall,
John Lucas
NASA Goddard Space Flight Center
Space Weather Laboratory, Greenbelt, MD, 20771; 301-286-2728
Larry.kepko@nasa.gov

ABSTRACT

Dellingr, NASA Goddard Space Flight Center's (GSFC) first 6U CubeSat, was deployed from the International Space Station (ISS) on November 20, 2017. The primary objective of the mission was to apply and appropriately tailor GSFC knowledge and capability to design and build a CubeSat that increased resiliency and capability, while containing costs. The Dellingr spacecraft is a mixture of COTS and in-house components and includes two science instruments – an advanced gated time-of-flight ion-neutral mass spectrometer (INMS) and a boom mounted fluxgate. While a traditional GSFC spacecraft approach includes detailed analysis, design, testing, and extensive reviews, the Dellingr team adopted a “build, test, fix” approach to identify and correct potential mission ending issues. Yet, despite extensive testing, Dellingr immediately experienced unexpected major anomalies once on orbit. Using a flatsat and the insight gained from extensive on-orbit engineering data, the team was able to alleviate some of these anomalies and recover some of the lost functionality. The extensive set of lessons-learned is driving changes to our systems architecture, flight software, and testing approaches, and has provided valuable insight into what is required to produce a NASA CubeSat science mission with a moderate assurance of mission success, while containing resource requirements.

INTRODUCTION

Dellingr is a 6U CubeSat developed internally at NASA's Goddard Space Flight Center (GSFC). The objective was to leverage commercial CubeSat development and apply and appropriately tailor GSFC knowledge and capability to design and build a CubeSat that increased resiliency and capability, while minimizing mission lifecycle costs. The spacecraft, described in a previous paper¹, is a mixture of commercial off the shelf (COTS) and in-house subsystems. This paper focuses on issues that have arisen while on-orbit, our approach to identifying and solving or working around these issues, and lessons learned from those experiences.

On-orbit events, anomalies, and the team's responses are presented in sequential order, as they occurred during the mission, rather than organized by subsystem. We feel this provides insight into how our team approached determining the source of abnormal behavior and recovered from potentially mission-ending anomalies.

OVERVIEW

Dellingr was primarily an engineering demonstration mission, but also contained two science instruments: a boom-mounted fluxgate magnetometer, and a novel gated time-of-flight ion-neutral mass spectrometer (INMS)¹. The INMS is a key instrument for future ionospheric Heliophysics missions and validating the instrument and characterizing its performance was a key

post-launch mission objective. Dellingr was designed to keep the INMS entrance aperture, located on the +Y 2U face, in the ram-pointed direction – i.e. pointed in the direction of spacecraft motion – so that ionospheric particles could enter the aperture to be measured. This orientation required an active attitude-control subsystems with several key components in order to achieve the desired control. The instrument required several weeks of outgassing before we could turn on the high-voltage, a delay that impacted our ability to turn on the instrument given spacecraft anomalies that appeared soon after launch.

A primary feature of the Failure Detection and Correction (FDC) scheme for Dellingr consisted of a 25-hour reset that power cycled the spacecraft. The spacecraft is reset as part of daily operations, so that it does not interfere with planned activities. There is also a 25-hour reset that will power-cycle the spacecraft if the daily reset fails to reset the spacecraft or the spacecraft exits the daily operations cycle. The 25-hour counter restarts every time the spacecraft reboots. This power cycle was intended to automatically clear any potential issues and single event upsets, while also reducing the FDC development effort. While the 25-hour reset rescued the satellite on more than 1 occasion, it sometimes made troubleshooting and scheduling more difficult in other situations.

An important aspect of lessons learned from Dellinger operations is the difference between making a spacecraft mission reliable versus making it resilient. From the standpoint of achieving mission success with an appropriate and tolerable level of risk and resource expenditure, it is potentially more cost-effective to focus on resiliency rather than exclusively on reliability. We discuss this further at the end of the paper.

Table 1 provides a summary of significant spacecraft events that occurred after Dellinger made it to space. In the sections that follow we step through these events.

Date	Event
Pre-deployment	Spacecraft turned on inside the NanoRacks deployer
20-Nov-2017	Deployment from ISS
20-Nov-2017	Contact on first pass
30-Nov-2017	Anomalous gyro and sun pointing accuracy
19-Dec-2017	GPS unresponsive
21-Dec-2017	Determined FSS unusually noisy
25-Jan-2018	Spacecraft data deleted – unknown cause
26-Jan – 6-Feb 2018	Unable to talk to flight computer. continuous reset state
6-Feb-2018	Recovery
6-Feb-2018	INMS turned on; ion data collected
14-Feb – 21-Feb-2018	Startup RTS001 V1 uploaded
22-Feb-2018	INMS filament burn in
28-Feb-2018	Startup RTS001 V2 uploaded
26-Mar-2018	Spacecraft data deleted – unknown cause
29-Mar – 5-Apr-2018	B-dot V1 upload
12-Apr – 17-Apr-2018	B-dot V2 upload
1-May – 7-May-2018	B-dot V3 upload (attempt 1)
8-May – 15-May-2018	B-dot V3 upload
1-Jun-2018 - present	INMS neutral commissioning

Table 1. Dellinger on-orbit event list

DEPLOYMENT AND INITIAL CHECKOUT

Dellinger was launched to the International Space Station on August 14, 2017, via the SpaceX Falcon-9 launch of CRS-12. Ejection from the NanoRacks 6U deployer was scheduled for the morning (EST) of November 20, 2017. The Dellinger team watched a live feed of deployment, but as the countdown reached zero on the first attempt there was no indication of ejection. The NanoRacks team reset the deployer and made a few more attempts to eject Dellinger, all of which failed. The ISS then entered eclipse, and the NanoRacks team, in conjunction with ISS personnel, discovered that connectors on the

deployers were swapped with an empty slot and the release commands were not being received by the Dellinger deployer. The ISS operations team proceeded to command the correct connector after exiting eclipse, with a successful deployment just after noon EST.

Even with the low resolution live feed it was quickly apparent that Dellinger had experienced the first anomaly of on-orbit operations, just seconds into the mission: The magnetometer boom and UHF antenna deployed immediately upon ejection, despite a 30-minute built-in delay timer. Figure 1 shows a high-resolution image of Dellinger deployment. The magnetometer boom is seen clearly pointed to the lower part of the picture, as are the UHF antenna.



Figure 1. Photographs of Dellinger’s deployment confirmed that both the magnetometer boom and the antenna deployed immediately upon ejection from the NanoRacks deployer.

The first downlink pass was at 8:00 PM (EST) at the NASA Wallops Flight Facility (WFF) UHF ground station. The satellite responded on the first contact attempt. Telemetry showed that the satellite had automatically entered sun-pointing mode, with the 6U face pointed towards the sun, and was maintaining a healthy power profile.

Analysis of the onboard data collected prior to deployment is consistent with the satellite having turned on while still inside the ISS. Figure 2 shows the battery voltage and temperature for the first hours of Dellinger’s mission and includes pre- and post-deployment data. The system clock booted into a default epoch, as it was intended to sync with GPS once deployed, so the initial time of the pre-deployment data is unknown. But we can infer the most likely scenario for when this initial deployment occurred.

As expected, the spacecraft telemetry at the first turn-on of the spacecraft indicated a fully charged battery

system. The battery voltage depleted slowly over the course of 17 hours until the safety circuits, triggered off the low voltage, shut down the power system. At this point, the satellite was still inside the deployer with uncharged batteries. The satellite turned on again after deployment from the ISS (approximately hour 24 in Figure 2), once the batteries had recharged sufficiently via power from the solar panels. The time between the EPS shutoff due to low battery voltage and the turn-on once batteries were recharged is unknown, since the internal clock was reset. The 7-hour gap in Figure 2 is for illustration purposes only.

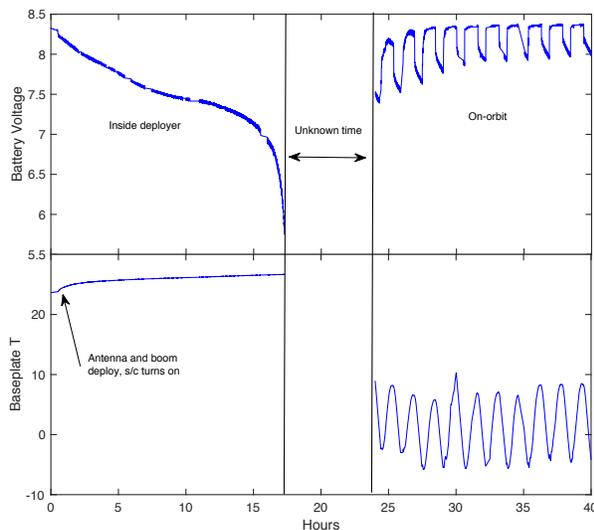


Figure 2. On-orbit battery voltage (top) and temperature (bottom) telemetry indicates that Delligr had turned on inside the deployer while onboard the ISS. After deployment the spacecraft tumbled until the batteries recharged sufficiently to turn the spacecraft back on, at which point it acquired the sun, and entered stable sun-point mode.

The subsequent telemetry shows that the satellite recharged and maintained a healthy battery voltage while sun pointing, depleting the batteries slightly during periods of eclipse, as expected. The 17-hour profile of pre-deployment temperature shows that Delligr was in a stable, controlled temperature environment of 23.6°C (74.5°F), very close to the 23.9°C (75°C) temperature of the ISS. It is therefore unlikely that deployment occurred while the satellite was in the NanoRacks deployer while attached to the ISS robot arm, since some type of temperature cycling should have been observed. We also eliminate ground handling, launch, and ISS storage prior to pre-deployment activities, since we cannot develop a viable theory that would cause a 17-hour disengagement of the safety inhibits while the deployer is under axial preload. In addition, a mechanical failure of the separation switch is unlikely due to its triple independent

redundancy. The most likely timeframe is between removal of the deployer preload aboard ISS and robot arm operations. Our leading theory suggests that the clustering of the sensors on a small surface area of the spacecraft combined with the NanoRacks deployer geometry lead to an inadvertent release of the switches.

Figure 3 shows the location of the three Delligr separation switches, which are concentrated in a corner of the 2U face (bottom left). This side of the satellite rested against the deployer door. From delivery to pre-deployment, a static preload was applied to keep the satellite from moving during launch and handling. During ISS operations, the preload was removed by an astronaut, thereby removing pressure at the door, and causing the door to move 0.025". In principle, the pusher plate should have had enough force to depress the switches against the door and maintain the inhibit. It is believed that interactions between pusher plate, clearances on the rails, ganged switches, door movement, and friction caused a distribution of pusher plate force that was not sufficient to keep switches pressed down in the off position. We confirmed that the satellite waited 30 minutes after disengagement of the switches to deploy the boom and antenna (Figure 2).



Figure 3. Delligr's separation switches were located at the bottom left corner of a 2U face and depressed by pressure against the NanoRacks deployer door.

Inadvertent power on of the spacecraft while onboard the ISS did not pose a risk to the spacecraft or the ISS, due to the system design and testing performed before launch. The team tested what would happen if the boom and antenna release mechanisms failed (or were unexpectedly deployed) while inside the canisterized dispenser. This test showed no hang-ups and smooth ejection, as we were able to verify in real-time watching the video feed. By design, all Delligr transmissions are initiated by the ground, and the satellite did not include a beacon. In retrospect, if Delligr had included an automatic beacon it would have radiated while inside the deployer, with potential unforeseen consequences. Power system and battery protection circuits monitored the battery levels and kept them from entering a dangerously low voltage level. Although unexpected, the

system design and on-ground testing mitigated a potentially mission-ending anomaly.

Once deployed (into space), the recorded telemetry (after hour 24 in Figure 2) confirmed that once the satellite's batteries had recharged, the satellite automatically entered sun-point mode, as designed, with the 6U face pointed toward the sun and reaction wheels providing 3-axis control. Figure 2 shows that the satellite recovered nicely from the low voltage state, and fully charged the batteries after several orbits.

Evaluation of the satellite state and the deployment anomaly was the focus of the first 2 days of operations. With the power-positive state of the spacecraft the team went into the Thanksgiving holiday break, and successfully contacted the spacecraft 5 days later to begin spacecraft checkout in earnest.

SUN POINTING PERFORMANCE

Dellingr uses a combination of 6 coarse sun sensors (1 on each face) and 2 fine sun sensors (FSS) on the +Y 6U face to obtain knowledge of spacecraft orientation, 3 reaction wheels to point the spacecraft, and magnetotorquers embedded in each solar panel to control momentum. Initial checkout of the ACS system indicated several issues of concern.

When in sun-point mode, the FSS provides the angles of the spacecraft coordinate system with respect to the sun, with 1 towards the sun, 0 perpendicular, and -1 anti-sun. While performing checkout of the attitude control performance, the team noticed the sun vector was very noisy, particularly in the x component. The satellite +Y (6U) face should be pointed to the sun during sun pointing mode. The high frequency, large amplitude, noise observed in Figure 4 would indicate the spacecraft was rapidly moving its orientation with respect to the sun, but the reaction wheels were incapable of moving the spacecraft that quickly over such a large angle. We concluded that the spacecraft attitude was not moving as rapidly as indicated by the sun vector solution, but that the data from the FSS was noisy, for unknown reasons.

The effects of this noisy sun sensor data were twofold. First, the reaction wheel torque commands (Figure 4 middle panel) shows that the spacecraft was heavily torquing the wheels to respond to what it believed to be pointing errors, often to maximum torques and with rapid switches in direction. Second, the sensor noise made it impossible for the Kalman filter to converge on a good attitude solution to achieve eventual ram pointing for mass spectrometer science. Since this impacted INMS science operations, this issue halted the progression of the commissioning phase until the team could determine the source of the noise.

Another anomaly observed, shown in the bottom of Figure 4, is the occasional increase of reaction wheel speed to maximum values during eclipse. The speeding up of the wheels, which was later attributed to a software bug, by itself was not a major concern, as the small CubeSat wheels could not destabilize the spacecraft, and after exiting eclipse the spacecraft appeared to reacquire sun-point mode fairly quickly. Since it was not related to the pointing accuracy issue observed during sun pointing, and as it did not seem to adversely affect the system, we deferred further investigation of the wheel speed-up issue.

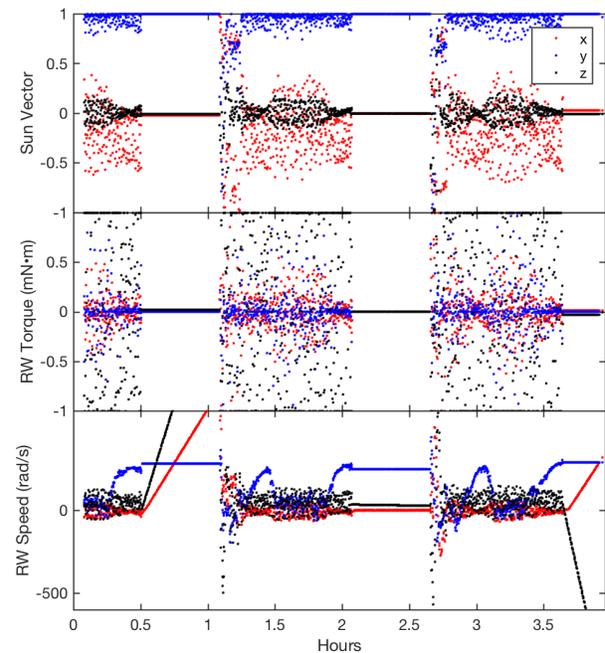


Figure 4. The ACS system exhibited noisy sun vector solutions (top panel), particularly in the x component, which resulted in heavy commanding of the reaction wheels (middle panel). Reaction wheel speed also exhibited unexpected behavior during eclipse. (bottom panel), sometimes accelerating to maximum value.

Dellingr carried two fine sun sensors. The first was an in-house developed FSS (WFSS) that had been demonstrated through ground testing to have a high level of accuracy. The second was a GomSpace fine sun sensor (GFSS), as backup, with a lower performance accuracy but with flight heritage. The default was to use the WFSS, and the satellite included a failure detection and correction action that would automatically switch to GomSpace sensor data if no or stale data from the in-house sensor was obtained.

Nominal spacecraft telemetry includes sun vector from the fine sun sensors but it does not indicate which sun

sensor was used, in an effort to reduce spacecraft bus telemetry size. The team commanded the spacecraft to retain individual sun sensor data to investigate sun sensor performance. Figure 5 shows simultaneously collected Y-axis (6U face, nominally pointed towards the sun) data from the in-house and GomSpace sun sensors. The figure shows the WFSS produced noisy and erroneous sun vector data, while the GomSpace accurately reported stable sun-pointing. Stable pointing was also confirmed by the energy entering the system from the solar panels. Erratic sun pointing such as the one reported by the in-house sun sensor would show non-constant power harvesting, but this was not observed. Despite a noisy FSS, stable sun-pointing was achieved since the control algorithm was designed such that random noise errors canceled. The resulting torque commands on the reaction wheels were a major concern, however, because they were driving the reaction wheels very hard and changing direction.

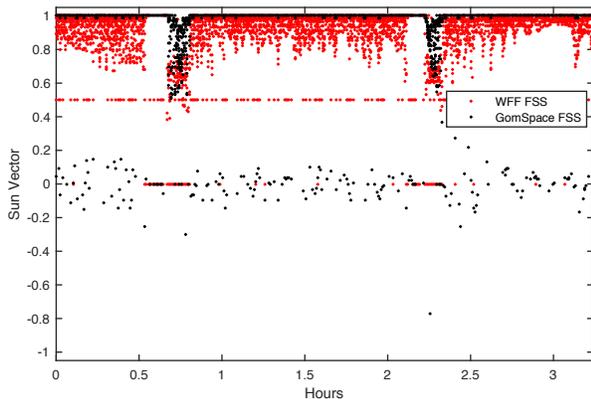


Figure 5. Delligr was commanded to send data from both sun sensors in an attempt to pinpoint the source of noisy sun vector data. Sun-pointing is indicated by a value of 1.

Diagnosing and solving the problem without a second sun sensor would have been difficult. The information in-hand indicated noisy in-house sun sensor data and the failure detection and correction on-board did not switch to the secondary sun sensor because noisy data was not considered a credible failure mode. The team switched off the in-house WFSS, to verify that the FDC would switch automatically to the backup GFSS. As expected, FDC performed as designed, and the sun vector data became very clean once the in-house WFSS was switched off, then noisy again immediately when the in-house WFSS was turned back on (Figure 6). Figure 6 also shows the reaction wheel torque, demonstrating a more stable control while utilizing the GomSpace sun sensor.

While we could manually turn off the noisy WFSS, the FDC 25-hour daily reset would automatically revert back

to using the noisy WFSS. A permanent fix required a software change to prioritize the GomSpace sun sensor over the in-house sun sensor. While working on such patch other, higher priority, issues appeared and temporarily halted the sun sensor patch work.

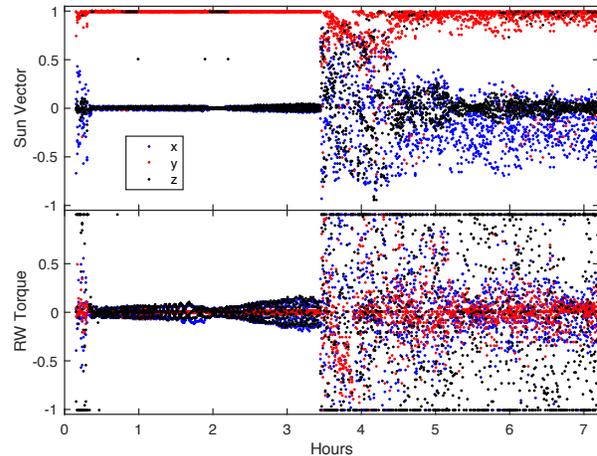


Figure 6. Delligr carried two fine sun sensors. From ~0.5-3.5 hours, the GomSpace FSS was active and indicates accurate sun pointing with Y-axis measuring 1 (pointed to sun), X and Z near 0 (perpendicular to sun) when sunlit. At hour ~3.5, we switched to the WFF FSS, and saw an immediate increase in the noise level. WFSS sun sensor data were the default for the first few months of the mission

THE IMU ISSUE

Delligr carries a Sensoror STIM210 Inertial Measurement Unit (IMU) to provide body rotation rates. However, it has not been providing accurate data since deployment. Recreating and debugging the issue has been complicated by the accommodations made to integrate the IMU into the spacecraft. The interface to the IMU is serial, but all serial interfaces to the flight computer were already allocated. We added an I²C to Serial converter, which required a TTL-RS422 converter. This added two additional points of failure for the IMU. It also added another device to the I²C bus, which already had several devices on it, including the reaction wheels. The I²C bus is not ideal, since a failure in one I²C bus component risks taking down the entire bus and thus all other devices on it, unless I²C isolators are utilized.

The IMU itself was designed to be used at a high data rates (up to 2000 Hz), which is significantly faster than our ACS cadence of 1 Hz. Reading the IMU at 1 Hz using the IMU's external trigger mode proved to be too noisy, so the flight software was modified to read the IMU at 10 Hz and provide ACS with a sliding average. This further burdened the flight computer and the I²C

bus. The fact that the IMU is used for attitude determination and control complicates the testing of the device. We are using the “42” simulator that allows us to simulate attitude and trajectory dynamics for testing of the flight ACS software on the flight computer but simulate the IMU for these tests.

These accommodations have complicated debugging the IMU issue, as we have no access to the operation of the converters in flight. Some raw data from the IMU is stored in telemetry, beyond just the sliding average, in the form of housekeeping data, but only a fraction of the raw messages from the IMU are stored to save downlink and radio resources. Unfortunately, the Cyclic Redundancy Check (CRC) values from the IMU are not stored; only a flag that notes if the CRC check passed. It would be preferable to have the CRC with the rest of the



Figure 7. Dellinger flatsat

raw data to confirm on the ground that the data from the IMU is not being corrupted by the converters.

Our investigation into the issue has thus far been inconclusive. The data as read on the flight computer appears to be valid from the IMU, as the CRC check is passing. However, the data itself appears incorrect. The data repeats the same values over hundreds of consecutive readings, a behavior that we did not observe during testing.

We have attempted to recreate the issue on the ground through use of a flatsat (Figure 7). The flatsat has been assembled from engineering units and spare parts. It is based around a custom backplane for PC-104 connectors which allows components to be added and removed with ease. One thing we are lacking is a spare Cadet radio. Instead, we are using a Cadet simulator program developed by NASA's Independent Verification and Validation (IV&V) Facility. This allows us to complete the loop between the flight software and ground control software to view the telemetry in real-time. Despite

having a spare IMU and a reasonable approximation of the other hardware on board, we have not been able to recreate the issue using the flatsat.

If we are not able to recreate the issue on the ground, we have the option to upload new applications to the spacecraft. For information gathering purposes it would be preferable to already have some diagnostic functionality built in, rather than rely on a software patch upload. For example, being able to send commands from the ground and receive data back from the IMU would be beneficial. Unfortunately, most of our byte-level interactions with the hardware are contained in a hardware library that is not accessible by our software applications.

The fact that the IMU shares the I²C bus with other devices introduces more uncertainties. The reaction wheels have experienced issues on the I²C bus during flight, which may cause the IMU problem, be a symptom of it, or due to interactions with another device on the bus. Unfortunately, the original flight software did not provide a flexible interface to isolate the power on the devices, as power state was tied to the spacecraft's mode of operation. This was altered later in the mission.

LOSS OF GPS

Dellinger carried a Novatel GPS (OEMV-1G) that failed on December 16, 2017, less than 1 month into the mission. The GPS telemetry shown in Figure 8 appears to show the moment the GPS failed. For the first 2 hours of the interval, the GPS reliably calculated the location of the spacecraft (top panel). During this time, it alternated between pulling ~327 mA and 440 mA; the majority of the telemetry was at 327 mA. Near 2:15, the telemetry showed a sudden change in the current draw, down to 280 mA, and a slight drop in temperature, suggesting a component failure on the GPS card. The next set of telemetry, just before hour 5, showed no current draw and zeros for the GPS-derived location.

The loss of the GPS was a major blow to the mission, as it eliminated the real-time ephemeris that was necessary for flying the spacecraft in a 3-axis, ram-pointed orientation. The best orientation for INMS science would have been to establish a fixed attitude relative to the local-vertical, local-horizontal (LVLH) frame with the INMS entrance aperture in the ram direction, and the Y-axis of the spacecraft (6U face) in a direction to optimize power production over the orbit (dependent on beta angle). Establishing an LVLH frame with no onboard capability to directly sense the Earth or the direction of nadir requires two computations. First, the FSW was designed to calculate the LVLH frame relative to inertial space using output from the GPS receiver. This would provide a target quaternion in the inertial frame.

The second computation is to determine the inertial attitude well enough to align the body with the LVLH target. A computationally efficient extended Kalman filter (EKF) was designed for the mission. Because the rotational rates in an LVLH-fixed frame are very low, an IMU with good rate resolution at low rates was needed for the attitude propagation step of the EKF. To make corrections to the attitude estimate, the EKF would compare the sun sensor measurements to the expected values from the onboard solar ephemeris, and the magnetometer measurements would be compared with an onboard magnetic field model. This magnetic field model was driven by the GPS location measurements. All four of these information sources – GPS, magnetometer, FSS, and onboard sun ephemeris plus the magnetic field model – were necessary for Dellingr to fly in an LVLH-fixed attitude. The loss of GPS meant that the spacecraft could not maintain an LVLH attitude. An onboard ephemeris calculation was not possible because it would have been too computationally intensive for the flight computer.

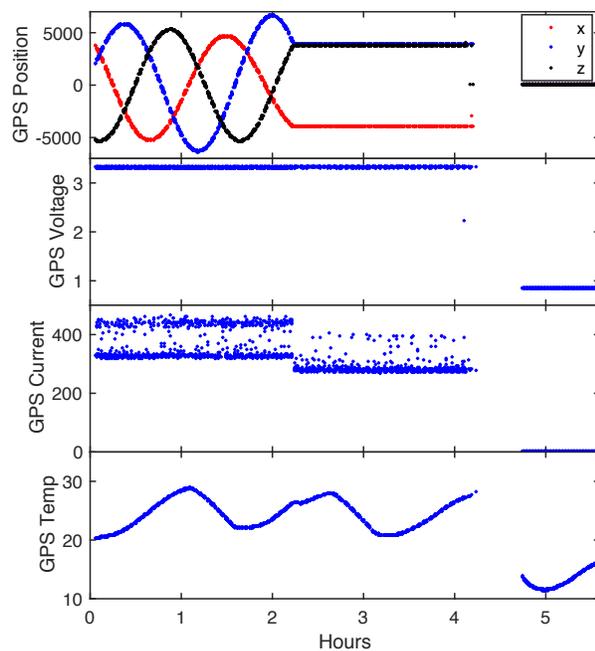


Figure 8. Interval when the GPS failed.

THE RESET PROBLEM

In the middle of January, the spacecraft started to show resets outside the expected daily reset. At first, the data showed a potentially thermal-related issue that caused multiple system resets a few minutes after entering eclipse, and then again a few minutes after exiting eclipse (Figure 9). Resets are indicated by an increase in the reset counter (top panel). The frequency of the resets rapidly increased to the point that by January 27th the spacecraft had entered a state of constant resets, once

every 63 seconds, rendering ground communication with the flight computer impossible.

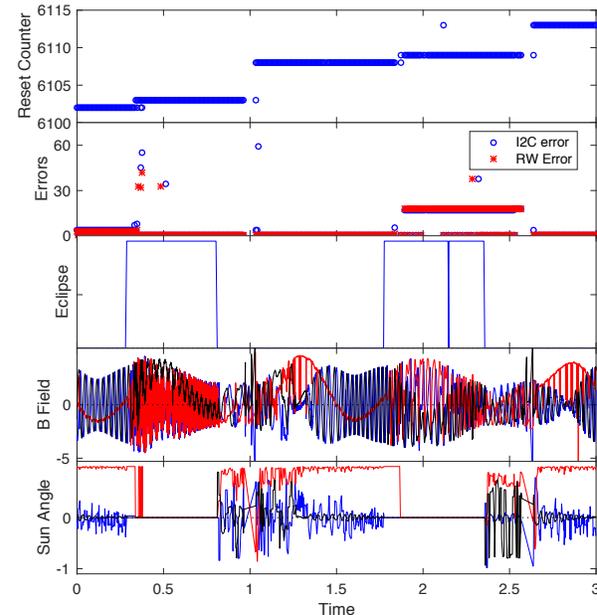


Figure 9. Data from January 26, 2018, showing an example of resets and I2C errors seemingly associated with entering and exiting eclipse. Resets are indicated by increases by the reset counter.

In normal operation the flight computer sends data to the radio for later download. Furthermore, the radio is connected directly to the voltage busses without switches. Since we could not communicate through the flight computer, on January 31st the team sent commands directly to the radio to download data, to test if the reset was localized to the on-board computer or if the power system was resetting the entire satellite. During these ground passes, we were able to communicate to the radio without interruption, thus isolating the reset issue to the on-board computer, rather than the EPS. If the EPS had been triggering the resets, the spacecraft would have rebooted during the typically 5-7 minute duration of the passes. As the system booted up we noted that the FSW was sending limited telemetry to the radio – telemetry that included the boot-up logs. Review of this telemetry indicated that the flight computer crashed at the same point of the boot cycle, 63 seconds in, when it attempted to communicate with the I²C devices. We were able to reproduce the crash on the ground with the flatsat and isolated the issue to a bug in the I²C driver. The bug is in the code that is supposed to prevent simultaneous access to the I²C driver by different tasks in the FSW. If two or more tasks were trying access the I²C bus at the same time, one would gain access to the bus, and the others would wait with a 10 second timeout. If the task that had access was delayed due to I²C device errors, the waiting tasks would eventually timeout and access the bus

anyway, causing the crash. Correction of the device driver was found to be a simple one-line modification.

Although we had isolated the problem, we could not alter the driver onboard the spacecraft to avoid the crash. It was not in an area of the FSW that we could modify through ground uploads. The EPS had a 4-minute watchdog timer on the I²C line. If it did not see activity for 4 minutes it would power cycle the spacecraft. But during the 63 second reboot sequence, the flight computer utilized the I²C line just enough so that the watchdog never tripped.

From January 27th-February 5th, 2018, the spacecraft rebooted continuously, every 63 seconds, reaching over 13,000 resets over this period. During this time, the satellite was inoperable, and it appeared to be an end-of-mission failure mode. The team had begun to plan a check-in schedule, perhaps once/week, to see if the issue had resolved itself.

The Cadet radio has a ground-commanded reset line that can be used to reset the on-board computer. By itself, that was not helpful, as the flight computer was already resetting. However, the team theorized that if during a ground pass we could constantly command the radio to reset the flight computer so that it never booted far enough into the sequence to service the I²C line, the EPS watchdog would, after 4 minutes of silence, power cycle the spacecraft. It was a longshot to be sure.

On February 6th during a pass at WFF, the team attempted this “back-door” reset but had to wait until the next pass to see the results. After the second pass, approximately 90 minutes later, the team received an email from the ground operator:

“We just confirm Dellinger back to business”

It worked. By constantly commanding the radio to reset the flight computer we had effectively jammed it and triggered an EPS watchdog reset. During the interval of constant resets, the team had theorized that the issue was high traffic on the I²C line. The current theory revolves around degraded components on the I²C line, making the line unstable. After recovering the spacecraft, the team immediately disabled the reaction wheels to decrease I²C traffic, as the reaction wheels are by far the worst offender on I²C traffic, with each of the 3 wheels communicating on the bus at 10 Hz.

Later that same day, the team received another email, from our project manager:

“Victory. [We were] able to turn on INMS in the ion mode. We had no resets between the last pass and this one.”

By turning off the wheels, we had removed the resets completely. But without wheel control the spacecraft could no longer maintain a sun pointing orientation. Still, we were finally able to turn on the INMS, and a few days later, on February 9th, we had spectra (Figure 10). Validating the INMS instrument was a key goal of the mission, and the team rightly celebrated the recovery.

But just a few weeks later, Dellinger again found itself in a state that threatened to end the mission.

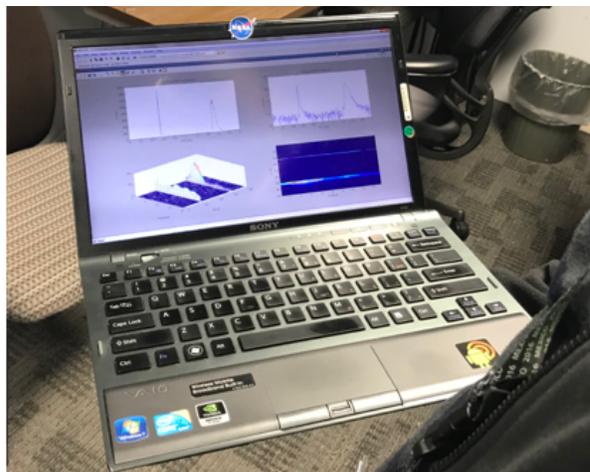


Figure 10. The first ion spectra from Dellinger.

THE SPIN PROBLEM

The revised operations plan was to keep the wheels and ACS system off during most of the week, and let the spacecraft slowly spin about the sun-pointing axis while INMS was turned on. At various times in the orbit the 2U face would spin through the ram, and the INMS instrument would collect data. It appeared that the reaction wheels could remain operational for a little more than 24 hours without inducing the recurrent 63-second reset loop. Therefore, at the beginning of the week we would turn on the reaction wheels and ACS system to ensure the spacecraft remained nominally pointed at the sun, and to unload any accumulated momentum.

Yet it soon became clear that the ACS system was not keeping the 6U face pointed towards the sun and, in fact, the ACS system seemed to be having no effect on the attitude at all. On March 6th, 2018, by fitting a sinusoid to the science magnetometer data, we saw the first indication that Dellinger was in an uncontrolled, fast-spinning tumble (Figure 11). The maximum observed spin rate was about 105°/sec (17.5 RPM) primarily about the body-Y axis (6U face), and all the reaction wheels were saturated at maximum speed. The ACS system was designed to work at spin rates below about 5.0°/sec/axis, so Dellinger was operating well outside the design limits. The ACS system appeared unable to unload the spin

momentum. The tumbling rate faced some initial skepticism because it wasn't clear to the ACS engineers how Dellinger could have reached that high spin rate (about 20x Dellinger momentum capacity). Also, 17.5 RPM is just under the Nyquist frequency where one can deterministically assess rate based on the 1 Hz sampling rate. Between that and limited and noisy data, it seemed serendipitous for Dellinger to be in an uncontrolled yet observable state. Nonetheless, the de-tumbling effort soon commenced. Determining how the spacecraft entered the uncontrolled, rapid spin rate would have to wait.

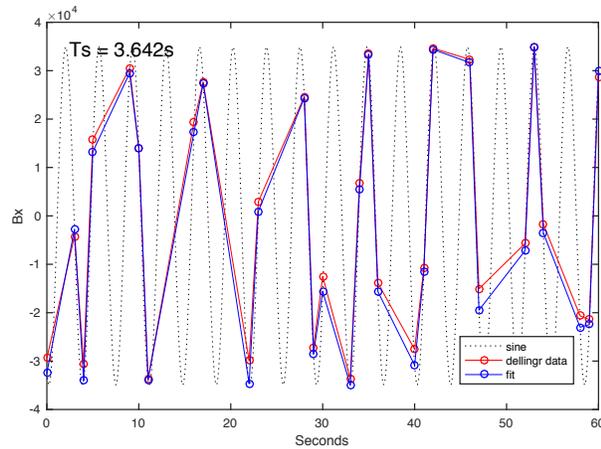


Figure 11. Sine wave fits to the x component of the magnetic field data from March 6th provided accurate assessment of the spin rate, even with unevenly sampled data.

RECOVERY FROM UNCONTROLLED SPIN

It was clear from the beginning that options to de-tumble Dellinger were limited to magnetic sensing and control. The onboard ACS software included a coded but untested “B-dot” control algorithm. A B-dot algorithm is a reliable, time-proven, simple magnetic detumbling algorithm originally proposed in *Stickler and Alfriend* (1976)². The B-dot control law is given by

$$m = -K\dot{b}$$

where m is the induced dipole, K is a positive gain matrix, and \dot{b} is the rate of change of the measured magnetic field in the body frame (hence “B-dot”). However, the default 1 Hz sampling rate of the magnetometer and the 0.5 Hz magnetic-control actuation cycle introduced a significant amount of time/phase lag into the system, making the original onboard B-dot de-tumble implementation ineffective at the spin rate Dellinger was in. The team would have to develop a new B-dot software patch and upload it to the spacecraft.

The first version of the patched B-dot control software eliminated the off-pulse cycle of the magnetorquers, effectively increasing the actuation cycle to 1 Hz. While 1 Hz sensing and actuation is still suboptimal, this controller was a relatively straightforward implementation such that, even if it couldn't completely de-tumble Dellinger, it should restrict the spacecraft to a lower rate. At the time, before the root cause for tumbling was identified and before sufficient characterization was done on the magnetic control loop, it was deemed the appropriate action. The first B-dot patch was uploaded to the spacecraft over the course of a week but was deemed ineffective in affecting the rate one way or another.

The second version of the patched B-dot control software took advantage of the parallel processing nature of the cFS architecture by placing the B-dot on an asynchronous schedule relative to other processes thereby giving B-dot the capability to off-pulse the magnetorquers and, more importantly, sample the magnetometers 250 milliseconds apart, reducing the measurement delay necessary for effective B-dot control. On a given B-dot control cycle, B-dot first turned off all three torquers, waited for some parameter defined duration of time, sampled the magnetometers, waited for 250 milliseconds (also a commandable parameter), sampled the magnetometers a second time, and then commanded the torquers for the rest of the roughly 1 second actuation cycle period. The second version of the patched B-dot software also turned out to be ineffective. However, the additional parameterization gave us much-improved insights to our magnetometer, which became the key to the final, successful implementation.

After the second failed attempt to de-tumble Dellinger, the team experimented with several timing parameters within the B-dot software and found that the magnetometer measurements were sometimes nulled or had unexplained biases. It turned out the B-dot algorithm was competing with the original magnetometer data acquisition and processing application; in a given 1 second period, the third magnetometer query (outside of B-dot), depending on its relative timing with respect to B-dot, could corrupt each other's results (a so called “race-condition”). While we were crosschecking the precise timing of these various processes against our working theory with observed flight data, the team got another valuable insight from the magnetometer instrument engineer about the serial nature of how magnetometer data acquisition is handled. Up until this point, the B-dot implementation team had an incorrect understanding that upon a magnetometer data query, the software will return the latest measurement after about 150 milliseconds. That turned out to be partially true.

The Dellinger magnetometer, upon query, immediately returns data collected from the time of the *last* query and collects data *for the next* query in the following 150 milliseconds.

The third and the final B-dot software patch had features to avoid race-condition with the science data acquisition application while also sampling the magnetometers three times within a 1-second period. It utilized the data from the last two queries that corresponds to time from the first two queries collected within the 1-second period. This final implantation is given by

$$m_{b\dot{d}ot} = -K * (b_{meas,3rd}) - b_{meas,2nd}$$

following the timing diagram in Figure 12.

Dellinger was successfully despun over the weekend of May 19-20, 2018.

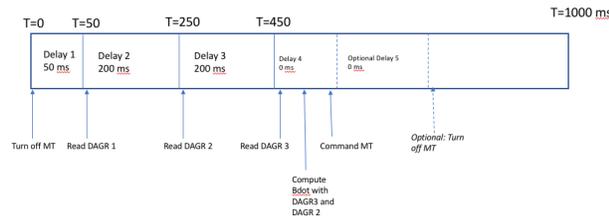


Figure 12. Timing diagram of the final B-dot algorithm.

The leading theory to explain the root cause that brought Dellinger to the high-rate tumbling state appears to be a magnetic control software implementation error that occurs during eclipse. The Dellinger ACS was designed to send zeroed torque commands to the reaction wheel during eclipse while momentum unloading of the wheels continued. It was discovered during in-flight operations that the last reaction wheel command before eclipse entry was *effectively stuck* in the command queue and caused the wheel speeds to linearly increase till saturation (see Figure 4). At the same time, the magnetic control loop continuously attempted to unload the wheel momentum for a prolonged period. This continuous magnetotorquing spun up the spacecraft to a high spin rate.

Although we do not have data from the interval that spun up the spacecraft to the uncontrollable state, we have seen the effect throughout the mission. An example is shown below in Figure 13. Prior to eclipse the spacecraft was spinning about the y-axis, as indicated by the magnetic field measurements, with the x and z component seeing a quadrature magnetic field. Upon entering eclipse near 0.2 hours, reaction wheel Z (2U face) accelerated to maximum speed, and the spin orientation of the spacecraft changed. The spin period

(bottom panel) calculated via fits to the magnetic field observations tells the story. The z axis spin rate remained relatively constant at ~80 seconds, but the x and y axes saw an initial decrease to ~35 seconds (i.e., sped up), followed by a slow decrease that was due to the magnetotorquers attempting to unload wheel momentum. By the end of the eclipse, the spacecraft was spinning with a 20s spin period about the z axis, then recovered quickly once leaving eclipse.

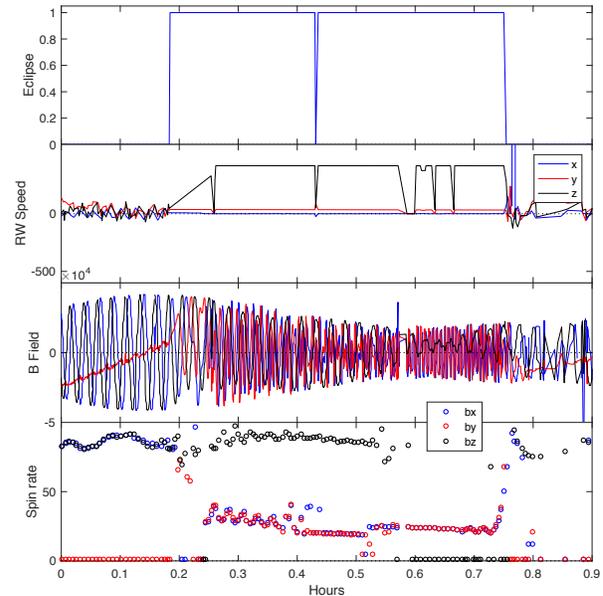


Figure 13. Example of Dellinger body rate changes during eclipse. Spin fits of zero indicate a bad fit and should be ignored.

The aforementioned ACS FSW implementation error would have been uncovered during the more typical Goddard FSW development process. However, such a process was not in place for the low-cost Dellinger development effort. The limited amount of system-level (ACS end-to-end) testing did not uncover the error because it did not include the real-life scenario where we had a failed GPS and IMU, which together caused the poor attitude control performance upon entering eclipse, and a rogue magnetic control loop unchecked by an IMU. The de-tumbling recovery process also could have been smoother had we had a more robust implementation of a B-dot controller. The failure to have a working robust algorithm combined with limited people resources and documentation made the final effort take much longer.

FLIGHT SOFTWARE PATCHES

During the course of on-orbit operations, several updates have been made to the Dellinger flight software including command sequence tables to enable more efficient passes, a new B-dot control application to de-spin the

spacecraft, and a new INMS instrument application to maximize the INMS data collection given the current state of the spacecraft hardware. The flexibility and features provided by the Core Flight System (cFS) enabled these updates to be made.

Dellinger uses NASA/GSFC's core Flight System flight software which is being used on class B missions such as the Lunar Reconnaissance Orbiter (LRO), the Magnetospheric Multiscale (MMS) mission and the Global Precipitation Measurement (GPM) mission. The cFS provides the maintenance features that enabled the changes that have been made on Dellinger. The cFS features that enabled the maintenance included the ability to:

- Upload files to the on-board file system;
- Verify the size and CRC of a file on the file system;
- Manage on-board files;
- Individually replace cFS applications;
- Compress cFS applications to reduce uplink bandwidth;

The first software changes involved updating the on-board command sequence tables known as Relative Time Sequence (RTS) tables. The RTS tables are sequences of commands that allow routine on-board operations to be automated. The RTS table changes included updates to account for malfunctioning hardware, and updates to allow more efficient day to day operations.

The first major patch to the Dellinger flight software was a replacement cFS application that implemented the B-dot algorithm in order to de-spin the spacecraft. The existing Attitude Control System (ACS) cFS application was designed with a built-in B-dot mode, but it was determined the mode would not work without changes. While the most logical path would have been to fix the ACS B-dot code and upload a new version of the ACS application, the ACS is the largest loadable application, and would have taken many weeks of passes to fully upload the new code. Rather than replace the ACS application, the team decided to implement the B-dot control algorithm in the GPS application, which happened to be the smallest flight software to upload. The team was able to code the B-dot algorithm, simulate its ability to de-spin the spacecraft, and upload it in eight parts. The new B-dot application was designed to stop the current ACS and reaction wheel applications when B-dot is active. The B-dot control application sends commands to the cFS Executive Services to stop the ACS and Reaction Wheel cFS applications. The B-dot control can be commanded to idle mode in case we want to run the normal ACS and Reaction Wheel apps.

After the B-dot patch was able to successfully de-spin the spacecraft, the team focused on improving the data collection from the INMS instrument. The INMS instrument is also controlled by an uploadable cFS application, so a new update was designed to improve the ability to collect data from the instrument given the state of Dellinger. The INMS update consisted of two changes, one minor change to allow INMS data collection in any spacecraft mode, and one major addition to enhance our ability to operate the spacecraft. The major addition included a new feature called the "Command Sequence Engine" (CSE) that allows a sequence of 16-bit sequence commands to be packed into a single ground command to the spacecraft. The sequence commands include: select telemetry filter table, start RTS, turn on Power Supply Electronics (PSE) switch, turn off PSE switch, enable INMS data collection, disable INMS data collection, delay, enable B-DOT mode, and reset the spacecraft. With the addition of the CSE, the operators can send a single command to operate the spacecraft for up to a day, primarily collecting data from the INMS instrument.

Future flight software updates will include modifications to the Reaction Wheel (RW) cFS application, and potentially updates to the ACS cFS application to account for hardware failures.

OTHER ON-ORBIT ISSUES

Operations

The Dellinger operational plan was to use a traditional mission operation center (MOC) located at Goddard in Greenbelt, MD and the Wallops UHF Ground Station (GS) located in Wallops Island, VA. Dellinger, classified as a "Do No Harm" mission, operated under a slightly higher risk posture than a more traditional Goddard mission. Therefore, the team did not perform as much extensive MOC testing in regards to contingency planning and how operations would be conducted if we experienced significant software or hardware problems on orbit, compared to missions with a lower risk posture. The project plan had assumed a 1-month engineering checkout period of the Dellinger satellite followed by normal science operations. Normal operations were envisioned and designed to support ground contacts twice a day Monday through Friday, with most of the downloading automated to reduce both staffing and cost.

The initial checkout after deployment worked mostly as planned. Sun-pointing worked and Dellinger was power safe which allowed the team to spend time methodically going through subsystem checkout. As described above, once we began checkout we discovered a number of different hardware and software problems. It became apparent that the initial checkout was turning into a

fulltime engineering exercise in finding solutions and workarounds to the issues. The MOC team did not have the engineering depth and knowledge specific to Dellinger to address these challenges.

Due to the on-orbit challenges facing the missions, we established an Alternate Mission Operation Center (AMOC) for the engineering staff to operate, perform mission diagnostics, and send corrective software patches to restore as much functionality as possible. This approach allowed the engineering team to operate directly with the Dellinger system and not be encumbered with the MOC operating processes. The engineering team operating the AMOC was able to interact directly with Dellinger, begin the process to diagnose the issues, and work to find solutions to restoring science operations. The MOC staff was used for just standard data retrieval and archiving which they were setup to perform for the normal mission operations. The recovery effort was coordinated by the engineering staff and the MOC team was called when the engineering team needed their normal support role. This change allowed us the streamline recovery efforts and reduce operating costs by not needing both teams operating concurrently.

The AMOC was setup in the Dellinger lab alongside the flatsat hardware to support on-orbit diagnostics as necessary (Figure 7). From the AMOC we were able to interface with the Wallops ground system directly and perform ground passes to collect diagnostic data, monitor Dellinger health, and upload patches to fix or work around anomalies. A tremendous amount of engineering knowledge has been gained by the staff in how to operate and optimize ground passes based on the Cadet radio and the Wallops ground system. For example, the team learned to look ahead to ground track effects on data quality (primarily whether the pass was over land or water), and their impacts to daily operating plans, and schedule critical activities accordingly. In general, this AMOC continues to be a tremendous learning experience for the engineering team and will have a major impact on future designs of GSFC's CubeSat missions.

Data deletion

Dellinger's radio memory bank was deleted twice during the mission without an identifiable reason. In both cases the satellite was neither in the range of the WFF ground station nor in the same geographic location. The first incident occurred on January 24, 2018, around 10:23 am EST, losing approximately 650,000 packets of data. The location was estimated to be over the south Atlantic Ocean near South Africa and the South Atlantic Anomaly (SAA), as shown in Figure 14 (top). The second incident occurred on March 24, 2018 at 6:33 pm EST, losing around 700,000 packets. This incident

occurred in the middle of the Europe as shown in Figure 14 (bottom).

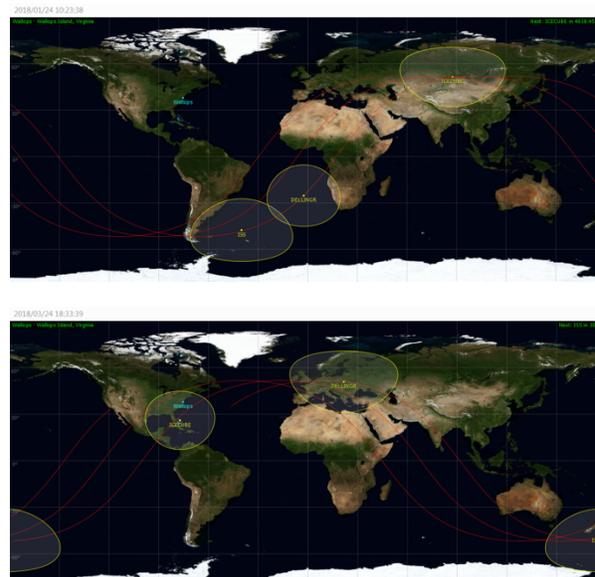


Figure 14. Approximate locations of the first (top) and second (bottom) data deletion events.

At this point in time we do not have a good explanation for the data deletion events. Direct command of the satellite by a UHF ground station seems unlikely since knowledge of Cadet radio operations is required since it will reject any commands not formatted correctly. While a single event upset is possible, particularly since the first event occurred near the SAA, we do not have an explanation for why an SEU would cause data deletion. An SEU does not seem likely for the second deletion event.

Data integrity

The Cadet radio appears to have an inherent bit-flip problem. This behavior was observed during ground testing in a laboratory environment but got worse once on orbit. The Cadet radio has an error correction protocol that can fix the first-bit flip and detect the second one in each row of data. However, more than 2 bit-flips per row occurs very frequently. We have found empirically that data has to be downloaded several times to fill the gaps and provide better data coverage. This affects the data budget requiring more ground passes to download the same data.

Uploading

Dellinger has the capability to add, delete or replace applications and relative time sequences from the ground. This feature is very useful and helped Dellinger across the entire operations phase, particularly in recovering from potentially mission-ending situations.

Uploads are quite challenging, however, due to limitations in the Cadet radio's ability to upload.

Each radio uplink command size is 101 bytes, but more than half are used for headers (PRN, Radio Header, and CCSDS header). Such headers make the effective upload for each command only 48 bytes. The radio is half duplex, meaning that it cannot send and receive data at the same time. Also, the radio is in receive mode for 0.1s every second, while the rest of the time it is in idle mode. Considering each message takes 0.084 seconds (101bytes in 9600bps), we need to send each message at least ten times to make sure we are hitting the listening window of radio. In addition, each command includes a delay of around 8 seconds to increase the success rate of each command. In general, a command needs to be sent several times to be successful. Sending a command 3 times was typically enough to have a high success rate during uplink. In summary, the radio upload speed is theoretically 9,600 bps, but in real world operations, we were only able to upload 48 bytes (384 bits) every 24 seconds, with an effective average upload rate of 16 bps.

RESILIENCY VS. RELIABILITY.

Dellingr was envisioned to test the reliability of CubeSats, CubeSat subsystems, and CubeSat development efforts for NASA science missions, and from that perspective, it was a success. We now have a wealth of valuable lessons-learned that we are applying to future CubeSat missions to increase mission reliability, and these lessons span FSW to systems engineering to MOC operations.

Although Dellingr was not a reliable mission, it was, and continues to be, resilient. Just on the ACS side alone, Dellingr lost GPS and the IMU, had a noisy FSS and a bug in the FSW that caused uncontrolled spin-ups, and an I2C issue that required turning off the reaction wheels. Yet despite these failures and anomalies, we were still able to obtain to maintain a stable power positive orientation through B-dot and obtain INMS ion data. At the time of this writing, we are still working to commission the neutral side of the INMS instrument, but there is every indication that this will occur.

A typical spacecraft mission has sets of requirements that it has to meet to be termed 'successful'. Dellingr did not carry such requirements, but rather had loosely defined objectives. A primary objective was to obtain good INMS ion and neutral spectra. Given a large expenditure of resources, we could probably recover some level of ram-pointing through some major software changes. But the team instead focused on obtaining INMS data through a slow roll of the spacecraft, such that it periodically rotates the INMS aperture in the ram direction. This would be sufficient to obtain the data

necessary to validate the instrument. If we had held the mission to specific Level 1 requirements, e.g., this would not be considered success. But this flexibility in mission implementation is required for resource constrained CubeSats. In the end, only 3 systems cannot fail: power, communication, and C&DH. And we suggest that it is critically important for CubeSats to focus limited resources on building an "architecture of resiliency" into the design, with a narrowly focused reliability where needed. This flexibility, when smartly incorporated into the mission design, greatly increases the odds of mission success, even if that path to success was not the one initially envisioned.

KEY LESSONS LEARNED

Dellingr on-orbit operations provided a wealth of lessons-learned. We list these below, in no particular order:

- A daily reset is a good failure detection and correction solution if the concept of operations allows. This minimizes the development of algorithms and code to detect and react to certain events in addition to overhead for the on-board computer to monitor such telemetry points.
- Patching software is a common theme over the course of the Dellingr mission operations. The ability to replace almost any code driving operations is a vital capability that can potentially save missions.
- Uploading is challenging due to limited capability. A conscious effort needs to happen during the development of the software to minimize uploads. Examples include table-based parameters and the ability to change parts of a table instead of having to upload a new table when changing one of such parameters, and small application sizes
- Adding as much telemetry as possible and creating filter tables is a good way to improve the ability to troubleshoot. During nominal operations of Dellingr, the amount of bus telemetry points recorded and their cadence can be lowered to minimize downloads. In case of anomalies or during commissioning, the filter table can be set to acquire selected telemetry and at a faster rate as needed.
- Consider using redundant systems when having reliability concerns if budget allows. Sometimes it is less expensive to fly a redundant system than attempting to improve the reliability of a single component. Dellingr flew fine sun sensors from two separate providers.
- With the increased risk posture of CubeSats, one should plan in more contingency schedule time and

engineering staff for checkout and corrective actions on-orbit than a traditional mission.

- Plan to invest more time into data reduction, analysis and trending to help identify subtle hardware problems and take corrective action before it turns into a larger problem.
- Don't design your FDC for only handling failed components but also misbehaving hardware that works intermittently.
- Add hardware diagnostics to your software to help isolate problems.
- Add more flexibility to your software in its ability to reconfigure between different ACS control modes and science instrument usage to gain more resiliency in your overall system to perform its mission.
- Having a full flatsat available is preferable, as is extensive flight software testing with hardware in the loop and a physics simulator. When selecting hardware components, generally the fewer devices that share a communication bus the better.
- Sensor and actuator hardware should match the flight computer's processing power; overshooting specifications can cause as many problems as under shooting.
- Relatively low-cost lessons for flight software are to ensure that low level hardware interfaces are available to ground control software. This will likely be advantageous for development and testing. Dellinger used a separate diagnostic mode to develop low level drivers and perform tests, which accelerated development early on but is now useless for diagnosing issues in flight. Using a flight software framework that allows for software to be updated in flight, such as NASA's open source core Flight System, is essential for enabling fixes after launch.

CONCLUSIONS

Despite multiple hardware anomalies and failures Dellinger continues to operate at the time of this writing. We are commissioning the INMS neutral mode, which involves activation of an ionization filament. The key to Dellinger's continued operation is that the most critical subsystems – communication, power and C&DH – still function. With a limited budget and tight schedule, we made the strategic decision to focus most of the extensive ground testing on ensuring the reliability of those three subsystems, and it has paid off for the mission. The ACS system was the last one to be fully integrated and tested; it was also the most difficult and costly subsystem to perform end-to-end tests on to uncover hardware or software issues. The team strategy for reducing ACS risk and controlling test cost was to focus on sun-pointing mode reliability. If Dellinger could sun-point and stay power positive, then we could use our

ability to upload software patches to correct pointing or other issues in the control system software. However, we had not expected that so many hardware issues would affect Dellinger. Still, this on-orbit flexibility was the lifesaving feature that allowed us to continue operations and work around these discrete ACS component issues as well. The bottom line is that this “architecture of resiliency” is always a good thing to build into your CubeSat mission, and the low costs of subsystem components such as sun sensors makes it easy to add in extra hardware.

Acknowledgments

Dellinger would not have been successful without the support of GSFC Center management and the many scientists and engineers who worked the project. We are grateful for their unwavering support and dedication.

References

1. Kepko, L., C. Clagett, L. Santos, B. Azimi, D. Berry, T. Bonalsky, D. Chai, M. Colvin, A. Cudmore, A. Evans, S. Hesh, S. Jones, J. Marshall, N. Paschalidis, Z. Peterson, J. Rodriguez, M. Rodriguez, S. Sheikh, S. Starin, and E. Zesta, *Proceedings of the AIAA/USU Conference on Small Satellites*, Session 3: Next On The Pad 1, SSC17-III-06, 2017.
2. Stickler, A. and K. Alfriend, K. “An elementary magnetic attitude control system,” *Journal of Spacecraft and Rockets*, 13(5), 282–287, 1976.