SSC20-II-04

Very Low Frequency Propagation Mapper (VPM) Experience and Results from the Systems Engineering Cycle of a Small Satellite

Travis Willett Gies, Tanya Rodriguez, Garrett Holt
Millennium Engineering and Integration Company
3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117
twillettgies@meicompany.com

Lee Jasper, Brandon Shirley, Collin Enger, Ryan Jaynes
Space Dynamics Laboratory
3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117

Franklin Hinckley
Metis Technology Solutions, Inc.
3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117

Michael Lopez, Gordon Wilson, Stephen Tullino
Air Force Research Laboratory
3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117

ABSTRACT

The Very Low Frequency Propagation Mapper, an Air Force Research Laboratory 6U CubeSat, is currently in operations complementing the Demonstration and Science Experiments (DSX) satellite by taking coincident measurements with DSX of the inner magnetosphere. This presents an exciting way to have a low-cost platform enhance an existing mission's data set with multiple spatial collection points. The Small Satellite Portfolio (SSP) had overall responsibility for the vehicle, ground system, and mission design; by enabling the team to make technical and programmatic decisions on their own, the team has been able to overcome many hurdles in short timeframes. Further, the team was constructed with a diverse set of skills to handle the many complexities of space systems. Finally, the team, and SSP as a whole, recognize that changes to the system and mission are not only expected but desired as mission maturity is gained. This paper outlines a selected set of issues and challenges that occurred, the ways the team dynamically handled the situations, and lessons learned for systems that are constrained in both cost and capability (e.g. small satellites).

INTRODUCTION

The[1] Very Low Frequency (VLF) Propagation Mapper (VPM) satellite, an Air Force Research Laboratory (AFRL) 6U CubeSat, is designed to augment the Demonstration and Science Experiments (DSX) satellite's VLF transmission experiment being flown by AFRL. Small satellites, especially Cube Satellites (CubeSats), present an opportunity to enable simultaneous observations of the same phenomena in fairly low-cost, short timeframes. The ability to augment a much larger mission's data set temporally or spatially is essentially unprecedented within the Air Force and an incredible use-case for these platforms. VPM not only enhances the DSX mission, it experiments with the concept of using small satellite platforms as a key architectural element of a greater mission implementation.

AFRL's Small Satellite Portfolio (SSP) has been designing, integrating, and operating spacecraft that have explored the utility of small satellites in Air Force applications for about eight years. With each iteration of spacecraft, SSP has been able to raise the baseline capabilities demonstrated on the CubeSat platform by implementing lessons learn from previous spacecraft – VPM largely pulled lessons learned from the GEARRS 1, GEARRS 2, and SHARC missions. The GEARRS satellites demonstrated that global telemetry monitoring was possible using a Globalstar transmitter designed for terrestrial use [3, 4]. VPM was able to implement this technology to continuously beacon vehicle state of

[1] AFMC-2020-0258

health telemetry that provides operators with valuable insight into the spacecraft globally (and more often than line of sight passes). The SHARC spacecraft had many similarities to VPM and was pivotal in forming SSP spacecraft design philosophies. VPM was able to implement lessons learned from SHARC to improve vehicle telemetry collection, fault handling and automation of uplinking/downlinking to the vehicle. By leveraging these lessons, SSP was not only able to build a spacecraft to support the VPM science mission, but was also able to explore design, execution, and philosophical approaches for flight and ground systems.

The primary, and unique, schedule driver for the VPM mission was to be operational simultaneously with the DSX spacecraft. Due to the nature of the science mission, VPM was required to launch within six months of the STP-2 Falcon Heavy on a separate launch vehicle. This drove a rapid and reactive development schedule with numerous external variables. Despite the challenges, VPM launched to the International Space Station on 5 December 2019 and began operations upon deployment on 1 February 2020. Operations are expected to last for at least six months in tandem with DSX with goal objectives lasting 12 months (limited currently by frequency license).

Small Satellite Portfolio Mission Design

The VPM mission was initially conceptualized in 2012 and has gone through several iterations of the mission design cycle before finally reaching full authority to proceed in January 2018. Once the DSX launch became solidified, so did the need and timeline for VPM. Even though VPM is relatively "low-cost, short timeframe" it is still a complex system requiring significant resources. The mission was kicked off with an initial acquisition strategy planning the use of commercial CubeSat hardware to support an in-house AFRL developed payload. This hybrid approach allowed SSP to internally have the objective to assess the capabilities available in the CubeSat market while the SSP engineering team could react to changes in schedule and launch vehicle using creative engineering solutions.

From the outset, the team had a defined scope and prioritized improvements to capabilities (building off of the previous missions); this allowed the team to operate and make decisions fairly independently within the well-defined bounding-box. Organizational risk posture was defined early and was decided to be fairly accepting. Trust was placed in the mission team to iteratively prototype solutions to a multitude of problems without inducing much process on the iteration/decision making. The engineers were encouraged to design and test the system with modularity and flexibility in-mind such that procurement risks could be taken on various "black-box" subsystems. By carefully envisioning each decision's impact on the highest-level mission objectives, rapid convergence on a solution was possible without significant programmatic burden [5, 6]. VPM was one of the use-cases that led to the development of the agile-inspired system engineering and mission assurance approach currently being formalized in SSP [7]. The team was enabled to try various execution paths in the system design and also to plan for operations; once in operations rapid response to issues and small, iterative, improvements to capabilities was not only expected, but implemented, given that the system design and the "rules for engagement" in operations allowed for open ended experimentation on orbit.

The remainder of the paper will discuss in some greater detail the above approaches, the issues experienced and how they were handled, and outline the process from design through operations.

VPM DESIGN AND INTEGRATION

SSP missions generally start out with a set of high-level mission objective(s) from which minimal and full mission success criteria are derived. The system design which satisfies these criteria is an open trade space that entails, among other things, how the mission team should put the bus or space asset together. Rather than a rigid set of requirements, the success criteria are an ongoing conversation with the mission stakeholders as the system evolves and grows in maturity. Procurement options include the full spectrum of buy-or-build that seeks to leverage the commercial satellite hardware and software market whenever possible. Due to the rapidly evolving nature of the market and new entrants to capability offerings, many components were assessed and adapted to meet the desired performance. This evolution was expected based on portfolio experience with SHARC and other CubeSat missions. A critical choice in system level design was to emphasize modular design (encapsulated functionalities, not plug-and-play) in both hardware and software to permit the lowest possible impact when a system required modification or replacement. By maintaining a team of engineers with hardware and software design skills, the team could rapidly prototype and integrate custom solutions to overcome gaps or interface mismatches in the commercial hardware. **All critical electrical interfaces were handled with a mission-specific interface card designed and populated in-house**. This design and the associated engineering skill sets allowed new board designs to be prototyped, spun and tested within days or weeks rather than the usual procurement timelines of multiple months. Leveraging the team's ability to buy or build solutions and evolve the mission success criteria to match, the mission altered and improved hardware and

software interfaces and functionality well into the testing phase.

The VPM team went into development planning to have some hardware modification, and saw the flight software as the key interfacing element to enabling this process. This meant one of the first design decision points in the mission development process, before any hardware was delivered, was how to address the flight software implementation. The VPM team opted to take a modular and configurable approach for the flight software architecture that could flex with both hardware and software design/testing outcomes to still achieve success. When presented with challenges in the commercially supplied small satellite design-space, it is highly valuable to have the team and programmatic mentality to accept what cannot be changed, and change what can.

Modular Design

Having loosely established mission success criteria, the team rapidly identified a Minimum Viable Product (MVP) or set of capability that the flight software, hardware, and ground system needed to provide. Small (~5 person) teams were then stood up and co-located as close together as feasible to encourage cross-pollination [7]. Rather than a full requirements derivation, the MVP approach allowed the hardware, software and ground teams to take their minimum functionality and decompose internally to the level they required. This MVP incorporated lessons learned from previous missions: fault handling, parallel development, engineer-led operations, etc. While greater than the absolute bare-minimum set of functionalities, it was determined to be the absolute lowest bar of capability to be worth delivering to orbit.

To marry with the hybrid hardware solution, the teams selected the Radiant flight software. Radiant, at its core, is a modular reusable open-network software development framework targeted at space systems and space systems development. This is SSP's first time using Radiant for a mission, and the first time a Radiant-based space system is delivered to orbit. Significant consideration was put into the flight software trade, but it was done early and decisively by the team and SSP leadership. This proved to be an enabling decision throughout the development and testing process. It should be noted that Radiant core capability exceeded MVP in many cases, but required developing hardware-specific applications for all the hardware components, with some reuse from SHARC (e.g. GPS App). Radiant's modularity and network interfacing of the flight software provided simple integration with AI&T ground software tools allowing the team to "test as you fly" through a COSMOS[2][9] interface from Day 1. It was also a tremendous asset to have the source code accessible to the team, and due to the fact that SDL utilizes Radiant on multiple missions, the code base also had developers from other missions outside SSP contributing.

The flight software rapid development strategy included several other attributes that yield dividends across multiple SSP missions:

1. Leverage agile software approaches.
2. Utilize a development framework for integrating hardware and payloads via serial, Ethernet, Spacewire, USB, and I2C.
3. Streamline process of supporting new/existing spacecraft hardware buses.
4. Grow "library" of examples and hardware support.
5. Leverage automatic code generation for command and telemetry messaging and COSMOS integration.

Flight Software Design

Radiant is based on SSM [9-13], which is a modular reusable open-network software development framework. SSM was the middleware layer for the flight software used on SHARC and had been an enabling technology for that mission. Radiant encompasses a set of core flight software services that ride on top of SSM as the Core Flight Software (FSW) layer shown in Figure 1, the inter-process communication fabric in this case is network based.

[2] https://www.ball.com/aerospace/programs/cosmos

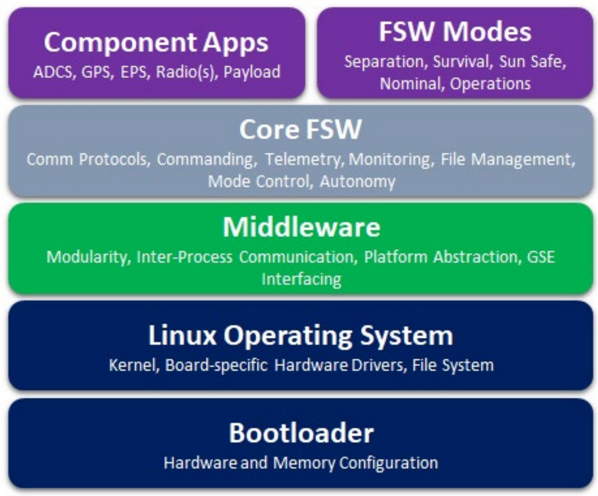

Figure 1: Flight software layers.

This middleware layer allows all the core and mission-specific Apps and services to talk to each other without being tightly coupled. Modularity enables reuse, and good abstraction and encapsulation enable modularity [1, 2]. The *Middleware* and *Core FSW* layers can be reused almost wholly across missions with bug fixes and new features being added. The *Component Apps* can be reused when using the same hardware allowing for a library of hardware to be built up. This allows developers the ability to leverage as much existing capability as possible going into each mission, to rapidly prototype, and to see what doesn't work.

With a core in place and a modular approach for interfacing with each hardware system the flight software became the enabling technology for meeting hardware milestones and providing support for integration testing as needed. This helps the entire system be resilient to change and promote reuse of the support Apps if hardware is reused on other missions.

Modularity and the GPS Subsystem

The modular architecture, with well-defined interfaces for connecting components, allows for parallel development efforts and decouples changes or failures in hardware from dependent or interrelated components from affecting one another. This resiliency to change is illustrated by GPS subsystem modifications made during VPM's development cycle. The first generation of commercial GPS interface board requires an I2C interface for local control and a UART direct to the GPS receiver. This adapter board suffered a latent electrical failure during test. The VPM team opted to simply remove the carrier board and develop a bare-minimum interface board that did not require a management software interface. While software development time was lost with the removal of the I2C carrier board and the board's associated software module, the rest of the transition took place within two weeks. The FSW team absorbed a great deal of the failed board's functionality through additional configuration, e.g. startup and telemetry processing. However, no additional changes to the overall flight software were needed because of low coupling, good modularity, and good encapsulation. Further, the broad skillset hardware team easily handled the design and fabrication of an interface board.

OPERATIONS

Much like AI&T, the SSP team expected to encounter anomalies and prepared to resolve or mitigate them (for instance, the operations team was not trained to specific procedures but focused on key activities so that desired outcomes were known and the team could adapt and problem solve to achieve success). Anomalies on orbit require more creative solutions than the AI&T floor but the same understanding and careful troubleshooting as hands-on debugging. For VPM, this opportunity presented itself immediately after deployment. This section will discuss some of the anomalies seen on VPM during early operations and how the team responded to overcome the anomalies.

Deployment and designing for a tumble

VPM's initial deployment CONOPS was to automatically deploy the solar panels, detumble, and go into a 3-axis stabilized sun pointing mode without operator intervention. This level of automation is not uncommon, but it is a higher risk acceptance stance than many missions take. The decision to allow the vehicle to achieve this high level of functionality was based on the premise that should issues arise, VPM has implemented SSP's design principles for surviving an uncontrolled attitude state:

1. Spacecraft power systems shall be capable of generating power to operate in a safe mode while attitude is uncontrolled.
2. Telemetry shall be received by the operator in a tumble without fine knowledge of spacecraft position.
3. Radio uplinks shall be able to close in a tumble [6].

Thus a higher functionality, higher risk operation, resides on the backbone of a resilient design feature.

In order to implement the first design principle, VPM implemented a set of spacecraft modes below the nominal operations modes: sun safe, survival, and phoenix modes. With each demotion of mode, both power consumption and spacecraft capability were reduced. The second design principle was accomplished

using a Globalstar transmit-only modem to beacon out spacecraft state of health telemetry. The third design principle was implemented using patch antennas with a simple antenna switching algorithm to achieve a near-omnidirectional antenna beam pattern.

VPM utilized all three implementations of designing for a tumble immediately following kickoff. On the first day of operations, using the Globalstar beacon, the team was able to determine that when VPM was deployed, it successfully executed the hazardous operations timer and deployed the solar panels. However, the spacecraft was not de-tumbling and was power negative. VPM quickly transitioned to survival mode, but the spacecraft was still not power positive in the tumble with even the bare minimum of functional avionics and software running. For the next 10 days, the spacecraft existed in a power negative state transitioning between the off-nominal spacecraft modes.

Spacecraft power systems should survive a tumble

The EPS subsystem on VPM was sized to be significantly larger than required by the nominal operational power environment (driven by payload operations and attitude state) to account for a tumbling scenario. VPM was constrained to a 'flower petal' solar panel configuration (Figure 2) because of the many antennas and externally mounted sensors, so there were many tumble scenarios that VPM could be power negative (i.e. the entire "back" hemisphere of the spacecraft's field of view).

Figure 2: The VPM spacecraft with solar panels and antennas deployed.

This led to implementing automatic mode transitions based on battery voltage. The default mode following the hazardous operations period was sun safe mode; this mode had all of the VPM bus subsystems powered. In the case of a low battery, the system would fault into survival mode; this mode would power off the ADCS and GPSR subsystems to conserve power, but retain the ability to command the spacecraft and transmit beacons to Globalstar. In the case that the spacecraft was still not power positive, the system would fault into its phoenix mode; this mode is a charge only mode where the EPS subsystem cuts power to the reset of the system until high enough battery voltage threshold has been achieved (e.g. a state of charge hysteresis allowing for some time for the vehicle to attempt to recover, achieve ground contact, etc.).

There were three main factors that allowed the spacecraft to become power positive enough to allow tumble recovery:

1. Changing the default boot mode to survival allowing for longer up-time
2. Seasonally increasing illumination period
3. A very-slowly precessing tumble

When the spacecraft was booted back up after the phoenix mode charge period had completed, it would nominally set the spacecraft to sun safe mode and remain there for at least one hour – this was to allow the ADCS time to attempt to continue its detumble mode. This meant that the spacecraft spent an hour of its operable time futilely trying to de-tumble the spacecraft which depleted the batteries much quicker than anticipated. Thus, the first command that was sent was to update the default boot mode to survival mode; this change gave the VPM team much higher likelihood of having a line of site pass where the spacecraft was in an operable mode.

The other large factor in survival mode turning into a power positive mode was an increasing illumination period; when VPM deployed it had the shortest time in sun for the orbit. For the first 10 days of operations, VPM rebooted 53 times, but the time spent in phoenix mode was slowly reducing. On the tenth day, the operators were able to upload and execute a command sequence to test a fix in the ADCS subsystem. Less than 5% of total system momentum was dissipated by the magnetic torque rods in this test of approximately 1 hour. From then on, the satellite was power positive in its tumbling survival mode. This even slightly reduced tumble rate likely enabled the EPS maximum power point tracker

algorithm to track power more effectively and the vehicle was recovered.

Telemetry should be able to be received in a tumble without fine knowledge of spacecraft position

The Globalstar beaconing feature on VPM (Figure 3) was designed based on experience from the GEARRS and SHARC satellites; both of those satellites showed a relatively high efficiency of beacon receipt in a tumble. The VPM design did not account for the change in Globalstar modem though; GEARRS and SHARC both used the STX2 modem, but VPM used the newer STX3 modem. The new modem only has an effective radiated power of 96mW which is approximately 1.6 dB less than the STX2 modem[3].

GEARRS 2 demonstrated a beacon efficiency of around 75% efficiency in a tumble of approximately 6 degrees / sec from an elliptical orbit of 350 km x 700 km at 55 degrees inclination [3]. VPM is demonstrating a beacon efficiency of around 30% in a controlled spin of 0.4 - 0.8 degrees /sec in a 460 km, circular orbit at 52 degrees inclination.

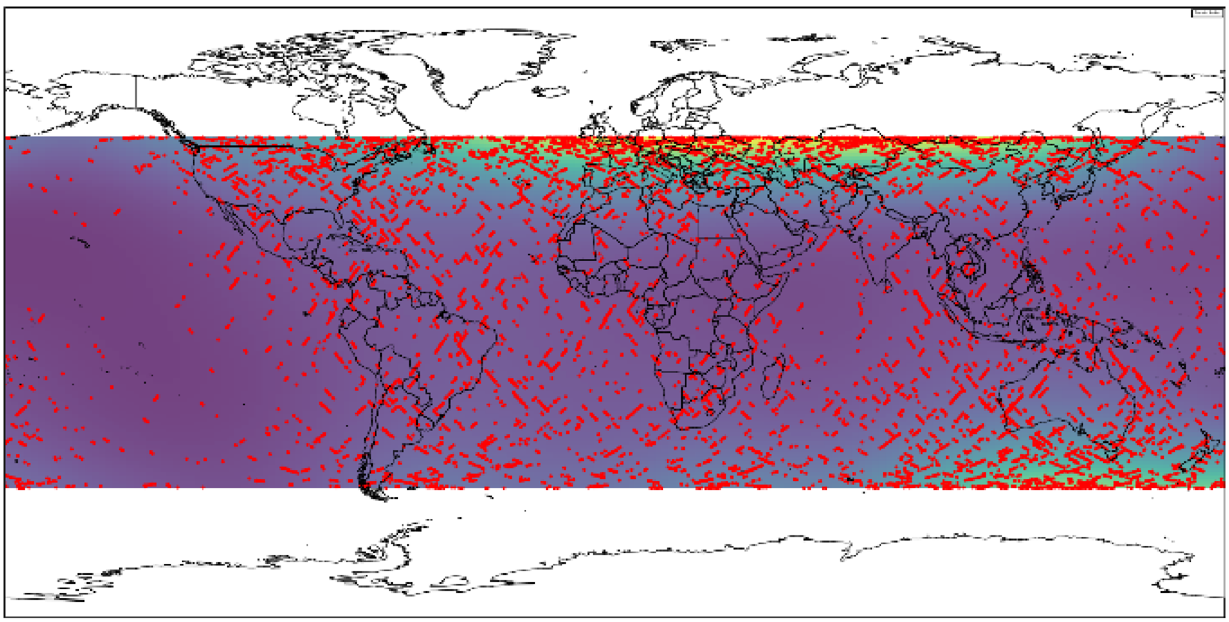

Figure 3: Received "complete" beacons for VPM, demonstrating global reception.

The beacon efficiency was significantly worse in the high rate tumble – this made predicting when the spacecraft was in an operable mode very difficult for operators. In order to better understand the phasing of spacecraft mode over time, the VPM team obtained a Globalstar report that outlined when partial beacons were received (Figure 4).

[3] https://fccid.io/L2V-STX2-1; https://fccid.io/L2V-STX3

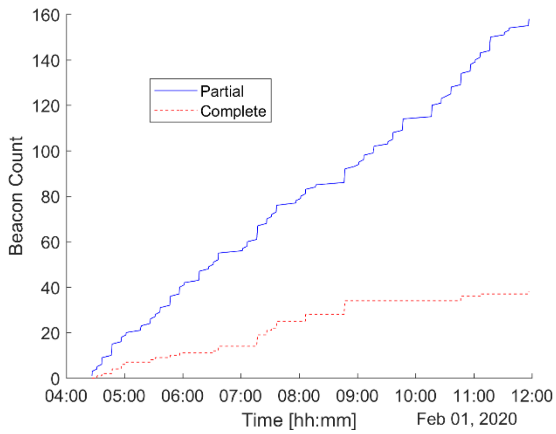

Figure 4: Beacon reception for partial beacons (< 36B) and full beacons. Reception rate for this period was 24%.

This report confirmed that VPM was alive but tumbling too fast to get all four required packets of a beacon reliably, but could still get individual messages through the Globalstar constellation. In the tumble following deployment, VPM was having less than 5% success on complete beacons, but approximately 20% success on receiving a partial beacon before the phoenix mode cycle started. Both the partial and complete beacons proved to be valuable despite the significantly reduced rate at which they were received. Partial beacons could not be interpreted, but the mere presence of a partial beacon helped to determine the power state of the spacecraft, especially when paired with complete beacons. In addition, complete beacons verified that the solar panels were fully deployed, subsystem power draws were nominal, temperatures were nominal, and surprisingly, spacecraft position, velocity, and time were occasionally received.

Radio links should be able to close in a tumble

There was significant, tested, margin on the telemetry, tracking, and command (TT&C) link budget, and the spacecraft could close link in the fast tumble. There were many compounding issues working against the team while tumbling on Day 1. CubeSats often deploy in large clusters and VPM was no different with 18 new objects in its deployment group. This required significant conversation with the antenna operator throughout LEOPS to identify the object and track the uncertain two-line elements (TLEs) generated in the first days and weeks. Coupled with the uncertainty of spacecraft power and the potential for seeing one of the known radio bugs, anxiety was high that communications would work, despite confidence in the link budget. After several days, beacons where received with valid GPS position, velocity, and time, and the VPM team worked closely with the ground station engineers and antenna operators to fence the spacecraft and apply a shift in time to the TLE being used.

This more accurate ephemeris and continuous improvements to the ground software removed variables, allowed confidence to build, and gave the team the chance to send one crucial command, then another, and another until the system stabilized. Note that originally the plan was for the Globalstar beacon to transmit the spacecraft's GPS position once the system had detumbled. This nominally would have allowed the team to quickly and accurately point ground antennas without requiring early, ambiguous TLE use. With the intermittent GPS signal and the shift to booting to survival mode, where GPS was unpowered, this was not generally possible.

Tumble root cause and on orbit repair

When attempting to identify the root cause of the tumble, the VPM team had very little access to telemetry; furthermore, once the cause was identified, the team had to work around having spotty commanding to the spacecraft. As the TT&C link became more reliable (e.g. the orbital position was understood and power generation marginally improved), the team was able to debug the attitude anomaly: an incorrectly reporting magnetometer. During final vehicle assembly, the externally mounted sun sensor/magnetometer was short circuited due to the package being installed upside down. When this error was found during AI&T, the team did its due diligence to ensure the sun sensor was not damaged (it was still functional), but was unaware that the package also included a magnetometer. The team documented the failure and proceeded with assembly. Given the damaged magnetometer was required by the autonomous detumble algorithm, action needed to be taken on orbit to recover the spacecraft from the tumble. The damaged sensor had caused the system to induce a tumble of approximately 15 deg/sec in two axes, and an estimated 40 deg/sec in the third. Through discussion with the ADCS provider, a path forward was identified. The ADCS had the option to change the magnetometer used by its algorithms, so if the secondary magnetometer was still functional, then VPM could detumble. The solution was relatively simple given the flexibility of the flight software and ADCS system, but made difficult by the spacecraft state of health. All Radiant hardware applications included a raw pass-through command to send any new commands that might become necessary on-orbit. A command sequence was uploaded to the spacecraft that would execute a series of raw ADCS commands. The team opted to create two command sequences: the first would be used to verify the secondary magnetometer's functionality, and the second would permanently set the secondary magnetometer as the default. These command sequences were quickly

implemented and tested against simulated hardware as VPM did not have a flatsat against which to verify the commands. Once the command sequences were ready for upload, the VPM team was able to upload the command sequences, verify the secondary magnetometer functionality (Figure 5), and begin the detumble process in just three days.

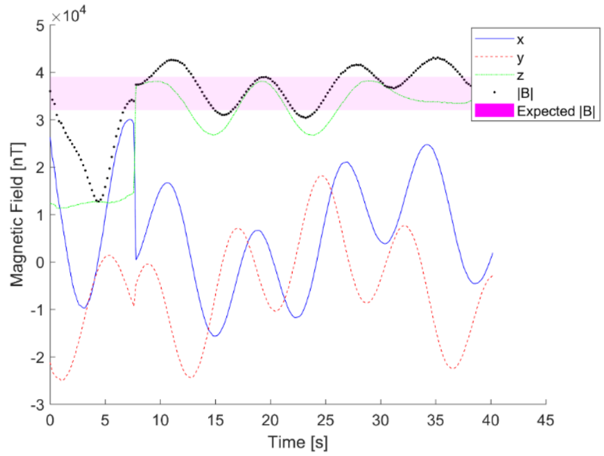

Figure 5: Broken magnetometer (first 7 seconds) and undamaged magnetometer. The magnetometer measuring the field "B" was notibly below expected values for VPM's location in orbit, indicating damage. The undamaged magnetometer shows magnitude near expectations. (Magnetic field vector components and magnitudes shown in the plot.)

Given the power negative state of the spacecraft, it was expected that multiple power cycles of the system would be required to fully detumble and begin sun pointing, but the spacecraft was able to accomplish this with no phoenix mode reboots (Figure 6).

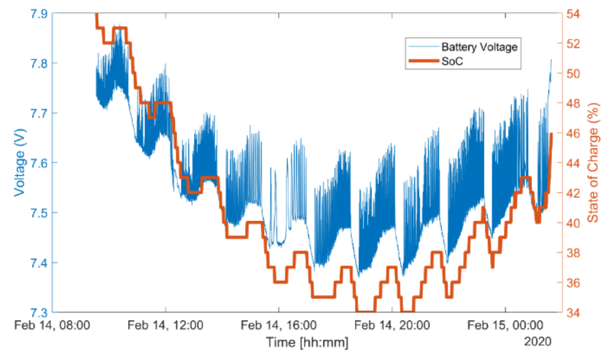

Figure 6: VPM battery voltage/state of charge as detumble occurred. Phoenix mode would occur at 6.5V on the batteries.

That was surprising as it took 16 hours for the satellite to fully detumble (Figure 7), which was significantly longer than the expected duration of a spacecraft survival mode power state. This was likely due to a combination of an increasing illumination period and naturally varying tumble state. The VPM team's design choices, planning for anomalies, allowed the spacecraft to be recovered from this potentially mission ending event. In the end, the system was one integration error away from achieving an aggressive, automated on-orbit commissioning sequence.

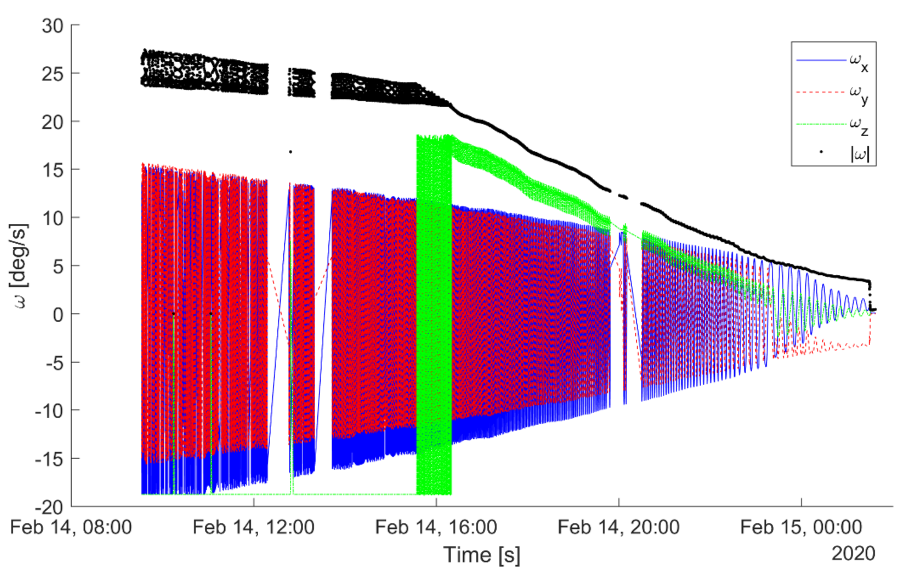

Figure 7: Measured body rates as detumble occurred over 16 hours. Note z-axis is initially saturated. Full three axis control recovered under 2 deg/s tumble rates, right at the end of the timeframe.

OUTCOMES AND FUTURE OUTLOOK

Outcomes

VPM has been very useful to SSP for providing some demonstration and validation of designs as well as practices. Having a team enabled to make corrections and iterations quickly is pivotal. Further, having hands-on hardware as early as possible enables interface debugging and time with higher and higher levels of integration provides confidence in system behaviors. Finally, doing these things and accepting that operations, especially commissioning, is partially an extension of test activities where system behaviors are being discovered and potentially debugged, has led to continual improvements in how the VPM mission is accomplished.

One key example of this mentality continuing through operations is the fact that the team, due to COVID-19 impacts, shifted operations from an on-base operations center to the capability to securely operate the spacecraft from their homes. The majority of the data chain had been deployed into cloud infrastructure and with several quick alterations, and assessment/confirmation of complying with security needs, the VPM team has made it such that the spacecraft can be flown from anywhere there is an internet connection. **This is the first time an Air Force owned and operated spacecraft has used commercial ground service providers, cloud-based operations, and remotely operated a spacecraft**. Given the demonstration from this mission, and the versatility it has provided in adverse conditions (a pandemic), future missions are likely to be operated this way, as security postures allow.

However, there are always numerous elements of the program and SSP's approach that could have been better. Some of the most important take-aways from SSP's approach are as follows:

- **Survive a Tumble:** Designing spacecraft to operate through a tumble, as was the design guidance for SSP missions [6], is not only too driving to the design but it is often not done with vendor supplied spacecraft buses. Designing to survive a tumble, with a well-tested phoenix mode (that can use the charge circuitry to charge batteries even if the vehicle is off), has proven incredibly important for VPM. This is likely to become a standard feature.
- **Test the "in-between" times:** Day in the life testing was, correctly, focused on specific elements (e.g. mission success criteria and key behaviors) but omitted testing the vehicle's operation when it was simply standing by. The behavior between events was not well known until the vehicle was left to its own devices in orbit. Why should it be expected that a vehicle that has never been left on its own to work properly in orbit on its own? In the future, SSP will likely consider things like total up-time vs. max single duration on vs. time spent operating between human input. Focusing on understanding behaviors, not just debugging functionality, is important.
- **There is never enough COMM testing:** SSP aims to test all of its systems with an end-to-end (ground and space segment) over the air, long range communications test. This verifies the link budget and demonstrates that the radio configurations are correct. Similar to the previous bullet, the flight radio behavior was never tested for longer duration, non-transmitting, operations. VPM also had a condensed version of this test (not over the air, and compatibility testing for ground system). Because this system is effectively the most important to do anything else with the spacecraft, it should receive commensurate attention and time.
- **Modular design with expectation that interfaces are immature:** Expecting that interface maturation will occur during the system integration process (hopefully on a flatsat instead of the flight unit(s)). Enabling the team to quickly add to, or modify, the interfaces to bring the system together accelerates the overall integration process to meet design intent.
- **Flight software/ground software configurability is fundamental to agility:** Making sure that the mission team has access to, and can configure as needed, the software systems allow flexibility to uncertainty and the ability to achieve objectives as capability comes on-line. When integrating "black-box" subsystems into an integrated system, modular software allows changes to occur, minimizing schedule impacts compared to fixing the black-box. Further, a configurable system allows flexibility to put off lower priority objectives to later, producing an agile experimentation platform.

Beyond the formulated take-aways, there are multiple on-going issues and performance studies that the team continues to assess. The planned non-volatile storage, an industrial grade SD card, failed several weeks into operations and telemetry was remapped to RAM. The radio also has demonstrated challenging, intermittent behavior that has been partially corrected through fault protection improvements within the flight software; root cause or more permanent corrections other than "reset" are still being assessed. Finally, SSP is considering how to better end-to-end verify the critical elements involved in the detumble process which would improve confidence in achieving higher level functionality more quickly in the mission. Sensor/actuator functionality and performance testing may be a reasonable step or even full Helmholtz cage verification; complexity and accuracy of ground based testing for these activities are in debate.

Future Outlook

A continual thread of development in the Small Satellite Portfolio is the concept of path-agnostic communications. At the mission level, SSP is not relying on a single path, network or radio to deliver mission data to the user. By exploring commercial and government ground networks and satellite to satellite communications, overall geographic and temporal vulnerability can be reduced. VPM contributed to this Portfolio vision by incorporating the evaluation of commercial satellite ground networks as a mission objective. This process presented unique challenges and opportunities as a US Air Force owned and operated mission but has been extremely successful to date. VPM's demonstration has opened the possibility for commercial ground to play an integral role in the larger TT&C architecture of the Small Satellite Portfolio. The commercial ground networks are tied to a cloud-based operations system designed and built in tandem with the VPM mission. A significant systems engineering endeavor in its own right, this cloud-based operations suite allowed near instantaneous reaction to the COVID-19 crisis, allowing secure, remote access to the operations center from off-site locations. A fully remote, virtualized operations center, while not an initial requirement to meet the DSX augmentation mission, has been an example of reactive, agile development with lasting value.

VPM is currently being operated remotely with significantly reduced need for operators, but improvements are continuously being identified and integrated into the flight and ground systems to eliminate the need for an operator. The flight system has several key features that ground automation is able to utilize to replace operator tasks: automatic telemetry downlinking and command queue management on a per command basis (instead of a per pass basis). Automatic telemetry downlinks are managed by the flight system based on the spacecraft mode and are easily configurable. Automatic downlinks prioritize telemetry generated when the link is active and fill the remaining bandwidth with stored telemetry; this allows automation to perform verification steps quickly while downlinking stored state of health telemetry. The command management of the flight software allows for individual commands or command sequences to be added or removed from the schedule; all the telemetry required to identify individual commands is able to be requested. The ground system is currently using these features to automatically synchronize ground schedules to the spacecraft by adding or removing individual commands from the schedule. Each command in the ground schedule is given a time-to-live where the synchronization automation will attempt to add it to the satellite command queue and ensure the command stays in the queue until it has been executed. These features have helped reduce VPM operations staffing from around 17 FTE in the initial month of operations, to less than one FTE since that time, even while the vehicle was still being commissioned. Further automation features are actively being developed and can be quickly implemented and iterated on due to the abstract, modular frameworks of both the flight and ground systems.

Acknowledgements

The authors wish to thank the entire Small Satellite Portfolio for their relentless dedication to the team and the mission throughout this development. This effort was performed by an integrated team of military, civilian, UARC and contractor personnel. This work and the mission behind it could not have been accomplished without the guidance of Dr. David Voss and Kyle Kemble, both of who are now with Space Security and Defense Program (SSDP). Mr. Stan Straight and Mr. Harold Vern Baker provided multiple reviews and insights throughout AI&T and Operations to the team that supported various elements discussed in the paper.

REFERENCES

1. Clyde, S. and J. E. Lascano, "Unifying definitions for modularization, abstraction, and encapsulation as a step toward foundational multi-paradigm software engineering principles," in Proceedings of the Twelfth International Conference on Software Engineering Advances, IARIA. Athens, Greece: International Academy, Research, and Industry Association (IARIA), October 2017, pp. 105–113. [Online]. Available: https://www.thinkmind.org/index.php?view=article&articleid=icsea 2017 5 20 10076.

2. Shirley, B.L., “Developing and Securing Software for Small Space Systems,” (2019). All Graduate Theses and Dissertations. 7544. https://digitalcommons.usu.edu/etd/7544

3. Voss, H. and J. Dailey and J. Crowley and B. Bennett and A. White, “TSAT Globalstar ELaNa-5 Extremely Low-Earth Orbit (ELEO) Satellite,” Proceedings of the SmallSat Conference, Logan, UT, Aug 2-7, 2014. Paper No. SSC14-WK-6.

4. Dailey, J. and H. Voss and A. White and S. Brandle, “Globalstar Communication Link for CubeSats: TSAT, GEARRS1, and GEARRS2,” Proceedings of the SmallSat Conference, Logan, UT, Aug 10-13, 2015.

5. Jasper, L. and L. Hunt and D. Voss and C. Jacka, “Defining a New Mission Assurance Philosophy for Small Satellites,” Proceedings of the SmallSat Conference, Logan, UT, Aug 4-9, 2018. Paper No. SSC18-WKII-05.

6. Clements, J. and T. Murphy and L. E. Z. Jasper and C. Jacka, “Tailored systems engineering processes for low cost high risk missions,” Proceedings of the SmallSat Conference, Logan, UT, Aug 3-8, 2019. Paper No. SSC19-WKII-08.

7. Jasper, L. and B. Braun and L. Hunt, “New Constraint-Driven Mission Construct for Small Satellites and Constrained Missions,” Proceedings of the IEEE Aerospace Conference, Big Sky, MT, Mar 7 – 14, 2020. Paper No. 2.0409.

8. Melton, R. and J. Thomas, “How, What, Where, and Why to Open Source Small Satellite Software,” Proceedings of the SmallSat Conference, Logan, UT, Aug 3-8, 2019. Paper No. SSC19-WKVI-06.

9. Christensen, J. H., “Space plug-and-play architecture networking: A self-configuring heterogeneous network architecture," (2012). All Graduate Theses and Dissertations. 1422. http://digitalcommons.usu.edu/etd/1422

10. Lyke, J. and D. Fronterhouse and S. Cannon and D. Lanza and T. Bryers, “Space plug-and-play avionics," Proceedings of the AIAA 3rd Responsive Space Conference, AIAA. Los Angeles, California: AIAA, April 2005.

11. Kief, C. J. and B. Zufelt and S. R. Cannon and J. Lyke and J. K. Mee, “The Advent of the PnP Cube Satellite," Proceedings of the Aerospace Conference, 2012 IEEE, March 2012, pp. 1-5.

12. SPA Committee, Space Plug-and-Play Architecture Standard: Networking (AIAA S-133-2-2013). American Institute of Aeronautics and Astronautics, Inc., January 2013.

13. Christensen, J. H. and D. B. Anderson and M. E. Greenman and B. D. Hansen, “Scalable network approach for the space plug-and-play architecture," Proceedings of the Aerospace Conference, 2012 IEEE, March 2012, pp. 1-10.