SSC19-IX-03

The Mobile CubeSat Command and Control (MC3) Ground Station Network:
An Overview and Look Ahead

Giovanni Minelli, Lara Magallanes, Noah Weitz, David Rigmaiden, James Horning, James Newman
Naval Postgraduate School
777 Dyer Road Bu-106A, Monterey, CA 93943; (831) 656-2089
mc3@nps.edu

MAJ Mark Scott, USA, Sean Brady, Chiffon Watkins
United States Department of Defense

Jacob Christensen, Chad Buttars, Ryan Beus, Riley Oakden
Space Dynamics Laboratory

ABSTRACT

The Mobile CubeSat Command and Control (MC3) ground station network is a Department of Defense (DoD)-led effort to build common-use infrastructure supporting communications and mission operations of small satellites for a wide range of US government organizations, contractors, universities, and foreign partners. The network consists of low-cost ground station terminals fielded at participating institutions, providing operators bent-pipe access to their satellites from any location with an internet connection. MC3 currently consists of eight active stations, and three international collaborators. One of the most important aspects of the ground station network has been the diverse community of small satellite users that have come together to share capabilities of mutual interest.

This paper describes the MC3 network and presents an overview of cost-effective future capabilities that will benefit researchers flying experiments on small satellites. Key capabilities include the Satellite Agile Transmit and Receive Network (SATRN) software, flexible software-defined radio architectures, fast-track radio licensing, expanded frequency support, and integration into secure cloud-based infrastructure. The paper also highlights some of the research undertaken at the Naval Postgraduate School (NPS) which utilizes the MC3 network and the satellites it operates as a testbed for advanced concepts. Research topics include optimization of constellation operations, predictive modeling of pass quality, and representative communications experiments flown on high altitude balloons and high power rockets.

INTRODUCTION

Initially fielded in 2012, the primary motivating factor for building the MC3 network was to bring together a handful of institutions involved in small satellite development for the US government. These groups could share hardware of mutual interest and streamline mission operations in a cost-constrained research and development (R&D) environment. Each institution could potentially contribute meaningfully to the federated network, or opt to run its hardware in a stand-alone mode to support local tests or operations. The ground stations were networked together with commercial off-the-shelf (COTS) virtual private network (VPN) devices for secure, cost-effective communications between sites. The distributed VPN architecture allowed external operators to connect to the network from anywhere with internet and fly their satellite using the MC3 stations as a bent-pipe;

dramatically lowering the ground segment cost of these R&D missions.

Early capabilities predominantly utilized hardware for transmitting and receiving to CubeSats in low Earth orbit (LEO) using modest data rates in UHF frequencies while leveraging protocols common to the small satellite community such as AX.25. As of 2019, the network expanded to eight stations, shown in Table 1, and includes active use of both UHF and S-band channels after the addition of several 3-meter parabolic dishes. The stations have also grown to support a wide range of waveforms and protocols by leveraging software-defined radios (SDRs) and open-source/commercially available solutions such as GNU Radio, MATLAB, and LabView. Proprietary or custom software SDR solutions are also possible and have been leveraged by several users. By leaving the radio software implementations unconstrained, the MC3

network can accommodate a wide range of users and radio vendors, provided they operate within the frequencies and bandwidths supported by the antenna apertures and radio licenses of the network.

Table 1:   MC3 Station Locations

Site (Designator)	Location	Capability
Hawaii Spaceflight Laboratory (HSFL)	Honolulu, HI	UHF
Naval Postgraduate School (NPS)	Monterey, CA	UHF / S-band
Space Dynamics Laboratory (SDL)	Logan, UT	UHF/ S-band
University of New Mexico/ Cosmiac (UNM)	Albuquerque, NM	UHF / S-band
Air Force Institute of Technology (AFIT)	Dayton, OH	UHF / S-band
US Coast Guard Academy (USCGA)	New London, CT	S-band
Malabar Transmitter Annex (MLB)	Palm Bay, FL	UHF / S-band
University of Alaska, Fairbanks (UAF)	Fairbanks, AK	S-band

MC3 NETWORK OVERVIEW

Following the small satellite paradigm, the MC3 architecture leverages COTS devices where possible to achieve a low-cost, yet reliable ground segment for small satellite missions. Tying the stations together is the Satellite Agile Transmit and Receive Network (SATRN) software, which is tailored specifically for small satellite mission operations.

SATRN Software

Developed by the Space Dynamics Laboratory (SDL) and owned by the US government, SATRN provides secure bent-pipe communications between operators and their satellites. Figure 1 shows the various components of SATRN including the Client, Server, and GroundSite.

The SATRN software is agnostic to a mission's command and control operations software and is responsible for routing packets over TCP/IP from the remote mission operations center (MOC) using the SATRN Client to the ground station which runs the SATRN GroundSite. A user is still required to develop mission-specific software for operating the satellite, as is often the case regardless. Since TCP/IP is commonplace, interfacing mission-specific operations software to SATRN is typically straightforward.

Well before satellite launch the MOC establishes a VPN tunnel to the MC3 network and configures their satellite in SATRN including parameters such as desired frequencies and network protocols. Once the spacecraft is in orbit, the MOC schedules contacts through the Client interface after loading the spacecraft's two-line elements (TLEs). The desired contacts are arbitrated by the central SATRN Server application. The Server utilizes a simple deconfliction algorithm which accepts or rejects a particular contact based on availability of the ground station, previously scheduled contacts by users of various priority levels, or administrative permission for a mission to use a certain ground site. Research is underway for more complex automated schedule arbitration, discussed later in the paper.

Once the desired pass is underway, a network tunnel opens between the MOC and remote ground station through SATRN. Uplink packets originating in the MOC are sent from the satellite's command and control software to the SATRN Client, from the Client to the GroundSite, from the GroundSite to the radio, and radiated from the antenna. The flow is reversed on the downlink.

During a contact the SATRN GroundSite is responsible for steering the antenna, loading the SDR solution specific for the satellite being serviced, compensating for Doppler, and actuating any peripherals such as relays or power switches which support the contact.
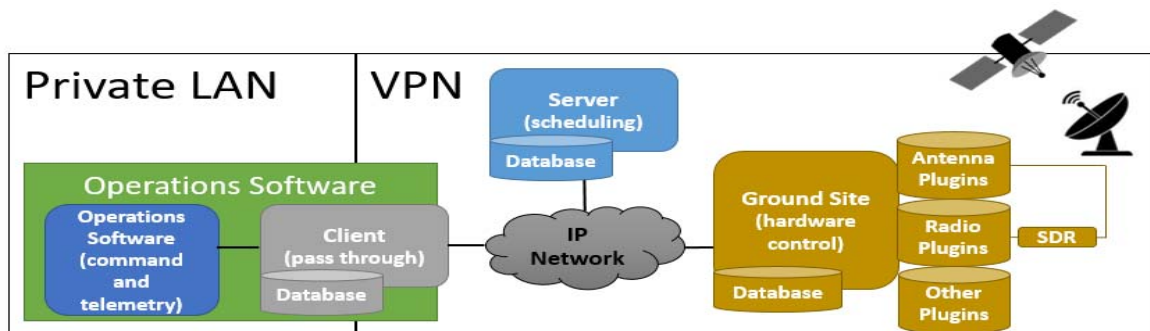

Figure 1: SATRN Architecture

Once the contact is finished, the network tunnel between MOC and GroundSite is closed, and the station is made available for the next contact in the master schedule as managed by the Server.

As part of an effort to increase reliability, resiliency, and scalability, the MC3 network has recently transitioned many of its networking functions to cloud-based services. By hosting both Client and Server on the Cloud, the network benefits from additional attributes such as DoD-approved network accreditation, increased network security, and high uptime. Additionally, MC3 administrators can create arbitrary amounts of virtual private clouds (VPCs) to service various missions, making the solution scalable to the large constellations of small satellites that are projected to launch in the coming years.

Hardware and Licensing

Antenna hardware in use at the ground stations generally consists of a 3-meter diameter parabolic dish housed in a radome, and Yagi antennas mounted nearby, as shown in Figure 2. SDRs are heavily leveraged to maximize the flexibility of waveforms supported by the ground stations.

Figure 2: MLB MC3 Node – UHF Yagis and S-band Dish/ Radome

The hardware is designed to service the frequency ranges licensed to MC3, shown in Table 2. These ranges are also common to many small satellite radios developed across the industry, and leverage dedicated channels reserved for satellite communications such as the Universal S-band (USB). The UHF uplink channel falls within a US government allocation, however the UHF downlink in the 902-928 MHz range is part of the Instrument, Scientific, and Medical (ISM) band. This band, initially a quick alternative for otherwise lengthy spectrum allocation processes, is heavily used by terrestrial devices, often negatively affecting the downlink. ISM small satellite downlinks are therefore strongly discouraged.

Though downlinks in the USB channels are predominantly uncongested, uplinks are shared with the Electronics News Gatherer (ENG) community, which has priority use over the 2025-2110 MHz allocation. Getting the license to operate a communications uplink in this band is increasingly difficult due to the coordination required, particularly in urban environments where spectrum is already maximally congested.

Table 2: MC3 Frequency Ranges

Band	Frequencies	Designator
UHF uplink	449.75-450.25 MHz	12K5F1D 43K0F1D
UHF downlink	902-928 MHz	115KG1D
S-band uplink	2025-2110 MHz	2M00G2D 2M45G1D
S-band downlink	2200-2290 MHz	1M60G1D 2M00G2D 2M45G1D
X-band uplink	7190-7250 MHz	(future)
X-band downlink	8025-8400 MHz	(future)

The MC3 team has been actively engaged with US government frequency regulators and ENG representatives to coordinate the use of S-band channels within this shared range. Using the NPS MC3 station, an experiment was conducted to determine at which power levels a ground station would need to transmit to interfere with a local news station's receivers in the shared USB spectrum [1]. The results demonstrated that though it is indeed possible to have a ground station interfere with local ENGs, an appropriate and logical solution would be to mask transmission elevation angles to mitigate the interference. Given that link margins often preclude low-elevation communications, imposing a transmission keep-out zone below 10 degrees in elevation solved virtually all interference issues. The information from this study was used to certify other MC3 locations at partner institutions which were also located in urban environments.

The MC3 network is adding X-band capabilities starting in 2019. At the time of writing, specifics of the X-band frequency ranges for licensing shown in Table 2 are still being coordinated. The network anticipates

spacecraft with increased data requirements, to include high bandwidth X-band downlinks, to start coming online in 2020.

The goal for all frequency allocations concerning the MC3 network is to fast-track the ground and spacecraft licensing process. By coordinating pre-existing channels, the MC3 team expects to reduce the spacecraft licensing process from up to two years, to several months or less. By working with other spectrum users around the US government, including regulators from the DoD, DoC, and NASA, considerations for small satellite missions' spectrum requirements are now coordinated in tandem with their more high budget counterparts. This offers a comprehensive spectrum use plan for spacecraft regardless of their size and budget. Leveraging this community-based approach, researchers from around the US government can streamline their frequency coordination process by referencing the MC3 license as a starting point for either ground stations or spacecraft.

Operations

As of this writing, the MC3 network supported daily operations for seven CubeSats and expects to add several more by the end of 2019. Clusters of CubeSats are expected to begin using the network starting in 2020 and dozens of anticipated future missions are in various stages of development and expressing interest in MC3.

These satellites belong to a wide range of civil and defense government organizations. The missions supported by the MC3 network to date include: Picosats Realizing Orbital Propagation Calibrations Using Beacon Emitters (PROPCUBE) consisting of three satellites [2], Polar Scout [3] which is a flight of two, Space-based High Frequency Testbed (SHFT), and RSat [4]. The STP-2 mission is planning to launch two satellites to be operated by MC3; NPSAT-1 [5] and FalconSat-7 [6]. Other missions are expected to soon follow.

Supporting the various requirements of each mission offers the MC3 team a unique vantage point for observing lessons learned across programs. Early acquisition after launch can sometimes pose challenges on missions with large numbers of secondary payload deployments due to uncertainties in satellite location. The method described in [3] has worked well for the team in recent years. MC3 users can take advantage of having multiple ground stations by tracking the initial launch vector with some of the antennas, and pointing/radiating inertially at the point of closest approach of the orbital plane with others. The inertial pointing technique typically guarantees several seconds of visibility of the satellite within the antenna beam since the track slowly compensates for the rotation of the Earth, and Doppler shift is minimized while signal strength is maximized due to the closest approach point. Using these techniques usually results in a first contact within 24 hours of launch if a spacecraft is actually responsive.

The MC3 team recommends more work be done to provide secondary situational awareness like satellite-satellite modems broadcasting GPS and basic telemetry even just a few times per day, and incorporating passive identification tags such as radar or laser retroreflectors. If a spacecraft is unresponsive, not only do these provide identification for satellite cataloging, but can help better assess spacecraft attitude. For example, determining whether or not the unresponsive satellite is tumbling can provide insight into whether or not the attitude control system, power system, flight software, or communications channels are functioning.

Lastly, the team recommends that over-air testing between satellite and MC3 ground station be performed well before shipping the satellite for integration. With small teams and limited budgets, schedules are often compressed later in the development cycle leading to reduced end-to-end system testing. A "day-in-the-life" test should not just last one day, but rather should last weeks or months. Ideally, the satellite is "flown" on the ground for extended periods of time using the same software and systems which will be operating the vehicle on orbit. If possible, any umbilical devices are removed from the spacecraft months before shipment and final checks are performed entirely through a flight-like communications chain, as these tests find edge cases likely not considered in short-duration scenarios.

RESEARCH TOPICS

MC3's inherent involvement in academia provides a powerful mechanism for student research involving real satellites and ground stations. This section highlights some of the Master's and Doctoral-level research at NPS of interest to the community, most notably NPS's focus on the automation of operating a large, disaggregated population of small satellites with minimal operator intervention. This includes predictive pass quality modeling, automated deconfliction of ground resources, and minimizing station downtime by automating anomaly recovery processes.

Research is also underway for performing cost-effective experiments utilizing X-band downlinks from high altitude balloon and high power rocket platforms. These demonstration flights provide representative links from near-space to validate design choices for both spacecraft transmitters and ground station receivers, as

well as low-cost communications platforms useful for military applications.

Predictive Pass Quality Modeling

Downlink margins involving small satellites are typically the limiting factor for communications due to reduced onboard power, lower-cost (lower gain) antennas, and restrictions in licensing space-ground transmission power and bandwidth. This makes the downlink particularly susceptible to poor geometries and RF interference on the ground. There may be contacts whose poor performance is counterintuitive; the expectation of data transfer was not met due to factors such as satellite orientation (*i.e.,* passively magnetically stabilized CubeSat antennas pointing inefficiently), and terrestrial obstructions or RF interferers. Spending valuable ground resource time on these contacts may take away from other, more advantageous passes.

Leveraging predictive modeling from historical ground station and satellite performance data allows for additional realism in an autonomous optimization model. Research conducted in [7] mapped downlink performance for the MC3 stations and PROPCUBE satellites. Using these mappings, a predictive model was developed for each satellite/ground station pairing depending on the ground track geometry of the pass. Using the example case shown in Figure 3, the historical downlink rates for PROPCUBEs Flora and Merryweather depended on the initial azimuth angle and percentage of contact time (unless the orbit changes, there is a fixed relationship between initial azimuth, maximum elevation, final azimuth, and pass duration).
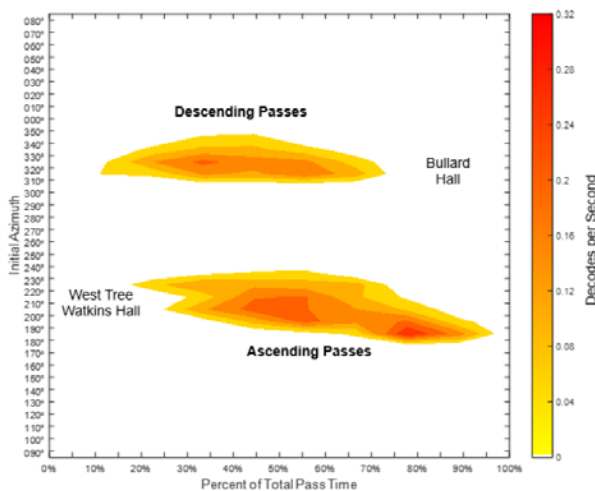
This model would place a strong emphasis on contacting the satellite when, statistically speaking, there is a greater likelihood of successfully downlinking data from it, rather than wasting ground station capacity by tracking a single spacecraft horizon-to-horizon. Collecting these data requires many contacts, and the spacecraft would need to have been in stable operation on the order of weeks or months. Therefore, the model could be advantageous when guiding the day-to-day communications activities of a massive constellation.

Optimization of Autonomous Operations

Motivated by the ambitious goals of small satellite developers to field large quantities of satellites, numbering in the hundreds or thousands, and the associated operational constraints on ground-based networks, techniques in optimal control were applied to maximize the capacity and benefit of line-of-sight communications between satellites and their ground stations [8]. The research aimed to autonomously configure the ground station and slew ground antennas while targeting the objects of greatest computed benefit to the overall mission. This algorithm is most useful for operating a large number of satellites concurrently in view of comparatively few ground stations. The resulting ground-based slew trajectories were occasionally counter-intuitive and difficult to solve by inspection once several satellites were introduced to just a few ground stations, highlighting the algorithm's value when considering hundreds of contacts per day.

The research constructed a benefit value function (BVF) for each satellite potentially in view of a particular ground station. The BVF consisted of both qualitative and quantitative factors such as satellite orbital position, estimated link budget, and mission priority. Figure 4 shows an example of how a BVF would change from acquisition of signal (AOS), time of closest approach (TCA), and loss of signal (LOS) represented as a two-dimensional polar plot of azimuth and elevation angles with respect to the ground station.
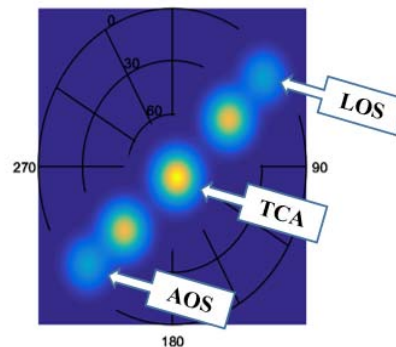

Figure 3: Predictive Pass Model for MC3 station and PROPCUBE satellites

Figure 4: Benefit Value Function for simulated satellite contact

Referencing Figure 4, brighter values depict more benefit and are maximized over the ground station due to the lowest free space path loss in the signal, resulting in the best link margin.

To optimize BVF targeting, a model was constructed to capture real-world kinematic and dynamic parameters for each antenna and any required boundary conditions including the desired planning horizon for the scenario. A numerical solver called DIDO [9] was then used generate optimized antenna slew trajectories based on all of the model's inputs.

When applied to real-world scenarios, the optimized trajectories were loaded into the MC3 system and executed for successfully contacting up to three PROPCUBE satellites simultaneously in view of two antennas. While these satellites all required a single ground-based configuration, future work will extend to distinctly diverse spacecraft which require different configurations. Having each satellite's configuration stored in a database (as is currently done in SATRN), enables on-the-fly transitions between satellites. As ground network tasking increases with more spacecraft, a typical operating method may be to only contact a particular satellite for a few minutes before transitioning to the next one. Due to the proliferation of low-cost ground terminals, another station continues the downlink of the first satellite a few hundred kilometers away for just a few minutes before transitioning to another spacecraft, and so on. The resulting concept of operation resembles a peer-to-peer network more so than a single point-to-point architecture, as is the industry standard today. Such a method requires rethinking satellite tasking, where objectives would be preprogrammed by operators and executed by the system when opportunities present themselves, resulting in a "human on-the-loop" rather than "human in-the-loop" paradigm.

High Altitude Balloons and Rockets

Though CubeSats offer relatively short development cycles in an educational environment, there exist several platforms that can leverage the CubeSat form factor and provide still faster timelines at lower costs. Two such platforms in use at NPS are the high altitude balloon (HAB) and high power rocket (HPR). The HAB platform has been recently used to develop representative space-ground links using X-band transmitters [10, 11]. Thesis students have been developing low-cost transmitters which fit in a 1U volume leveraging the RaspberryPi single board computer and USRP B205mini software defined radio. Additional upconversion, filtering, and amplification is added for X-band transmissions, resulting in a total package that costs approximately $2,500 and is operated by open-source software. Figure 5 shows a 2U HAB payload which hosts a 1U transmitter and 1U bus. The bus provides power to the payload and a secondary communications channel to operators on the ground. 3D printed materials are heavily utilized for structural components given the ability to rapidly iterate designs, light weight, and desirable performance in the near-space environment.


Figure 5: High Altitude Balloon Payload

A 2U payload such as the one shown in Figure 5 can also be manifested on a suborbital HPR to quickly deliver a communications capability above a region of interest [12]. NPS has begun such a test campaign with launches of increasing complexity. The first such launch was attempted in February 2019, carrying a 2U payload for beyond line-of-sight VHF communications relay experiments to 10,000 meters. The rocket disintegrated shortly after launch, losing the payload with it. Since then the 3-meter-long rocket has been successfully flown (shown in Figure 6), and a reflight of the full experiment is scheduled for fall, 2019 [13]. An even more ambitious communications relay experiment will involve two hops between an HPR, a HAB, and MC3 terminal across California. Eventually these relay demonstrations will also involve tasking a satellite on orbit through these other platforms. The work is in support of the Nuclear Command, Control, and Communications (NC3) program which seeks to find alternative methods of maintaining a communications architecture under any conditions. Rapidly standing up overhead relay capabilities for terrestrial and space-based communications would be of substantial benefit.

Figure 6: High Power Rocket Test Flight

CONCLUSION

The MC3 network is capable of providing low-cost bent pipe access to a diverse population of small satellites supporting research applications for the US government and its allies. By taking a community-based approach, the infrastructure development and accompanying research is made available across organizations, minimizing rework and strengthening the quality of the end products. The network has transitioned from proof-of-concept efforts to 24/7 operations supporting multiple stakeholders. A significant challenge ahead will be to make the system resilient for supporting operations of tens, if not hundreds, of satellites in the coming years while maintaining a low-cost, lights-out capability. The following section highlights some upcoming directions for MC3.

A Look Ahead

As more satellites join the network, ground station downtime will have an increasingly greater impact as the system approaches saturation [14]. To prevent downtime in a low-cost system with many single-string design elements, monitoring telemetry from the stations and acting before a problem occurs will be a priority. Given the similarities, the team has elected to operate its ground stations much how it operates its CubeSats.

Without maintaining a 24/7 watch floor, a reliance on automation will be necessary. By having a standard set of telemetry collected at each site, it becomes easier to recognize trends across the entire network.

Machine learning software can provide a low-cost solution to analyze health and status telemetry at each ground station and recognize operational trends. For this reason, machine learning tools such as NASA's Inductive Monitoring System (IMS) are being evaluated. With an automated monitoring solution in place, MC3 operators can focus on repairing issues with ground stations rather than investigating the cause of a failure. Furthermore, the introduction of autonomous monitoring will help drive development toward eventual automated failure resolution.

Defense-related research organizations from the US, UK, Canada, New Zealand, and Australia are also working to stand up the International Small Satellite Command and Control Network (ISC2N). The MC3 network represents the US contribution to the effort. These "Five-Eye" partners employ the same community-based approach adopted for MC3 when creating interfaces and standards that apply internationally between the various research programs.

The team is also looking forward to adding participants from the US Naval Academy, Texas A&M University, and several other DoD organizations later in 2019.

USNA: Jin Kang, PhD; CDR Jeff King, USN; Bob Bruninga

Aerospace Corporation: David Ping; David Harvatin, PhD; Russ Luherson

NIWC Pacific: Marcus Matsumura; Dexter Barit; Noah Acosta; Martin Lindsey, PhD

MEI: Kasandra O'Malia, Katherine Fackrell, Gary Stuart, Heather Powner, Gregg Leisman

JPL: Jim Lux

DTG: John Kay, PhD; Neil Easton; John Ferguson

DSTG: Natalie Stevens; David Lingard, PhD

DSTL: Christine McCullough; Sean Murphy, PhD; Junayd Miah, PhD

References

1. Forbes, A., "Sharing S-Band Communications to Conduct Small Satellite TT&C", Master's Thesis, Naval Postgraduate School, Monterey, CA, Mar. 2018.

2. Bernhardt, P., "The Largest Ionospheric Disturbances Produced by the HAARP HF Facility." [Online]. Available: http://ies2015.bc.edu/wp-content/uploads/2015/05/040-Bernhardt-Slides.pdf

3. O'Malia, K., G. Stuart, K. Fackrell, H. Powner, G. Leisman, S. Wallner, and G. Minelli, "Needle In a Haystack: Finding Two S-band CubeSats in a Swarm of 64 within 24 Hours", AIAA/USU Conference on Small Satellites, Logan, UT, Aug. 2019.

4. Hanlon, N., M. Lange, B. Keegan, J. Roser, D. Wenberg, and J. Kang, "Design Features and Flight Results for the Autonomous Mobile On-orbit Diagnostic System (AMODS)", AIAA Space 2016, Long Beach, CA, Sep. 2016.

5. Sakoda, D., J. Horning, and S. Moseley, "Naval Postgraduate School NPSAT1 Small Satellite", ESA Small Satellite Systems and Services Symposium, Chia Laguna, Sardinia, Italy, Sep. 2006.

6. Anderson, G., M. Dearborn, and G. McHarg, "FalconSat-7: A Photon Sieve Solar Telescope." [Online]. Available: https://www.amostech.com/TechnicalPapers/2011/Optical_Systems/ANDERSEN.pdf

7. Leone, J., "CubeSat Pass Quality Analysis and Predictive Model", Master's Thesis, Naval Postgraduate School, Monterey, CA, Jun. 2018.

8. Minelli, G., "Resource-Constrained Autonomous Operations of Satellite Constellations and Ground Station Networks", Dissertation, Naval Postgraduate School, Monterey, CA, Sep. 2018.

9. Ross, I. M., "A Primer on Pontryagin's Principle in Optimal Control", (2nd ed.), Collegiate Publishers, Carmel, CA, 2015.

10. Lovdahl, B., "Software-Defined Radio Payload Design for CubeSat and X-Band Communications", Master's Thesis, Naval Postgraduate School, Monterey, CA, Dec. 2018.

11. Bischoff, G., "X-Band Software Defined Radio for CubeSat Applications", Master's Thesis, Naval Postgraduate School, Monterey, CA, Sep. 2019.

12. Pross, J., "Filling the Gap: Rocket Delivered Short-Term Expeditionary Beyond Line-of-sight Narrowband Communications Relay", Master's Thesis, Naval Postgraduate School, Monterey, CA, Jun. 2019.

13. Pierce, D., "Development of a Rocket Test Platform Capable of Delivering Standard Dimension Payloads to Near-Space Altitudes", Master's Thesis, Naval Postgraduate School, Monterey, CA, Jun. 2019.

14. Writt, A., "Optimization of CubeSat Ground Stations for Increased Satellite Numbers", Master's Thesis, Naval Postgraduate School, Monterey, CA, Jun. 2018.