SSC20-X-03

A Prototype Virginia Ground Station Network

Zach Leffke, Jonathon Black
Virginia Tech, Hume Center for National Security and Technology
1311 Research Center Drive, Room 2037, Blacksburg, VA 24061; 540-231-4174
zleffke@vt.edu

Kevin Shinpaugh, Ian Harnett, Bryce Clegg, Nick Angle
Virginia Tech, Center for Space Science and Engineering Research (Space@VT)
1341 Research Center Drive, Suite 1000, Blacksburg, VA 24061; (540) 231-1246
kashin@exchange.vt.edu

Chris Goyne, Connor Segal, William ‘Trace’ LaCour, Mike McPherson
University of Virginia, Aerospace Research Laboratory
122 Engineers Way, Charlottesville, VA 22903; 434-982-5355
cpg3e@virginia.edu

Dimitrie Popescu, Samuel Jensen, Jason Harris
Old Dominion University, ECE Department
231 Kaufman Hall, Norfolk, VA 23529; 757-683-5414
dpopescu@odu.edu

Mary Sandy
Virginia Space Grant Consortium
600 Butler Farm Road, Suite 2252, Hampton, VA 23666; 757-766-5210
msandy@odu.edu

Mike Miller
Sterk Solutions Corporation
702 Pine Street Philipsburg, MT, 59858; (406) 859-3105
mlmiller@sterk.space

ABSTRACT

This paper provides a detailed technical description of a prototype ground station network, the Virginia Ground Station Network (VGSN), developed for the Virginia Cubesat Constellation (VCC) mission. Virginia Tech (VT), University of Virginia (UVA), and Old Dominion University (ODU) have each constructed ground stations to communicate with their respective VCC spacecraft. Initially, each university was responsible for commanding its own spacecraft via its own ground station. As the mission progressed, it was decided to network the ground stations and operations centers together to provide backup communications capability for the overall mission. The NASA Wallops Flight Facility (WFF) UHF smallsat ground station was also included in this network. Implementing the VGSN led to the establishment of successful communications with UVA’s Libertas spacecraft via the VT Ground Station (VTGS), demonstrating the utility of collaboration and of the VGSN. This paper provides a technical overview of the VGSN, details concerning signal processing requirements for the mission, a discussion concerning the radio regulatory process as applied to the VCC mission, and plans for future upgrades of the network to continue to support Virginia (and partner institution) small satellite missions.

INTRODUCTION

This paper describes the design and development of the prototype Virginia Ground Station Network (VGSN). The VGSN was initially developed to increase the overall operational capability of the Virginia Cubesat Constellation (VCC) mission. While the VGSN was developed in response to specific operational needs of the VCC mission it is not limited by VCC mission requirements. The network is referred to as a prototype because it is the first of its kind for the Commonwealth of Virginia and to indicate that overall design and implementation is ongoing. The authors have a strong desire to continue to evolve the operational capability, robustness, and security of the VGSN during the ongoing VCC mission lifetime and for future missions. Participation in the network helps reinforce collaborative smallsat efforts among Virginia universities, something that started with the VCC mission itself under the leadership of the Virginia Space Grant Consortium (VSGC). A very high degree of collaboration was achieved by the VCC team from mission concept all the way through mission operations and this collaboration extended to the VGSN development.

While the current primary members of the network are VT, UVA, and ODU, other Virginia institutions, including academic, state, federal, and commercial entities, are encouraged to join the network. This could also include entities located outside of the state but partnered with VA institutions that are members of the network. The ultimate goal of the VGSN is to provide a more robust overall capability for the Commonwealth of Virginia for smallsat operations. The primary focus of this effort is on academic research and education, however economic impact and development is a key goal as well. This will be achieved by pooling ground communications resources via the VGSN, as well as the expertise of the students, research faculty, and academic faculty involved.

This paper is organized as follows. The Introduction section provides introductory and background material. Then a description of the general VCC mission is provided, as the VCC mission was the main driver behind the VGSN development and serves as a backdrop for highlighting specific design details and constraints. A description of the ground stations involved in the network is then included. All of this is necessary background information in order to understand the functionality of the overall network. The next section then describes the overall VGSN design. This is then followed by notable implementation details for the VGSN.

A description of the network connectivity of the respective mission operations centers and tracking stations over the public Internet and the associated security requirements is discussed. A detailed description of the over the air packet protocols and the interfaces between the operations centers and tracking stations are described. The signal processing framework used for additional link layer and physical layer processing and transmitter hardware control are described. Connection to the WFF system and the associated signal processing requirements are described. The requirements for Federal Communications Commission (FCC) and National Telecommunications and Information Administration (NTIA) licensing of the ground stations and spacecraft transmitters is discussed. Finally, the scheduling requirements and real time operations via the VGSN are described.

Network Background

The VCC mission has served as the impetus to develop the VGSN. The overall VGSN design, development, and functional capabilities are described through VCC mission related examples. This paper will often reference VCC specific mission details in order to highlight a specific capability or explain a particular design decision of the VGSN. While VCC serves as the backdrop for describing VGSN capabilities, the network is not specifically limited by the VCC mission requirements and can be relatively easily adapted and expanded to meet new mission requirements.

A large number of ground station networks already exist. Examples from federal organizations include the NASA Near Earth Network (NEN), the NASA/JPL Deep Space Network (DSN), the DoD’s Air Force Satellite Control Network (AFSCN), and NOAA’s National Environmental Satellite, Data, and Information Service (NESDIS), to name only a few. Commercial networks also exist such as the Kongsberg Satellite Services (KSAT) and more recently Amazon is offering ground station services linked directly to their cloud computing infrastructure. In the more public and open domain there exist ground station networks such as the Amateur Radio community’s telemetry collection network for the FOX series spacecraft run by AMSAT-NA. AMSAT-UK has a similar, but distinctly separate system for their FUNCUBE spacecraft. The international Satellite Network of Open Ground Stations (SatNOGS) project is very popular among Amateur Radio and DIY satellite and radio enthusiasts.

With reference to the above existing networks, the VGSN is most similar to the more open networks, employing nearly all open source software and with similar design goals. The network is intended for command and control of smallsats and does not have the same stringent requirements as the more capable and complex networks of the US Federal Government or other commercial providers. One notable difference from the open networks is the ability to actually route spacecraft command uplinks through the network in near real time. This is an important distinction from the receive-only AMSAT/SatNOGS type networks and has very specific implications for regulatory and security requirements. Mechanisms for expanding the VGSN to include connections to external networks (i.e. KSAT, NEN, etc.) are part of the roadmap of the VGSN in order to increase the overall footprint and coverage of the network, with initial work already accomplished with the WFF UHF smallsat ground station (herein referred to simply as the WFF-GS).

VGSN Purpose

The VGSN connects multiple university Mission Operations Centers (MOCs) to multiple tracking stations across the Commonwealth, including NASA's Wallops Flight Facility UHF Ground Station. The overall goal is to pool the collective ground station capability and expertise of the Commonwealth and distribute that capability to multiple institutions.

The primary operational use of the VGSN is to provide a backup communications capability for a given institution for real time mission operations. Should a ground station be down due to maintenance or for other reasons (i.e. staffing), having alternative communications paths is desirable to continue to execute the mission. The VGSN also enables more streamlined and coordinated communications, such as one ground station transmitting while another receives on the same frequency, which helps support activities such as communications troubleshooting.

The VGSN provides centralized resources and signal processing capabilities as needed. For example, when interfacing with the WFF-GS, additional bit level physical layer processing was required for the VCC mission. Rather than perform this processing in an ad hoc manner at separate institutions, all commands and data can be routed through a single central processor that achieves the desired functionality. This saves time and money as it reduces duplication of effort, particularly important for budget constrained mission. This also provides a single portal for connection to external ground stations allowing for simplified scheduling, flexibility, and centralized control and monitoring for security purposes. This external interface capability could be expanded to include other tracking station facilities that are not core members of the VGSN and not necessarily located in the Commonwealth. The addition of this type of capability will largely be driven by mission objectives and funding availability for specific missions.

The VGSN also provides a limited data warehousing capability for any given mission. The data from multiple missions can be stored in a centralized location for data warehousing purposes for the lower level data and telemetry. All parties involved in the mission, whether mission operations or scientific analysis, can be provided access to the network and thus access the data warehouse. This is a function or service that needs careful consideration on a mission by mission basis and is potentially susceptible to scope creep. The core purpose of the VGSN is real time communications with spacecraft and mission operations support. Higher level data processing and scientific analysis for a particular mission should not by default be expected to be part of the VGSN, rather the VGSN should be used as the access point for obtaining the lower level data for higher level processing. This functionality should reside with a Science Operation Center for a particular instrument or mission that at the moment is not included in the VGSN design. This type of functionality could be added if deemed appropriate at a future date and if the appropriate resources are made available to the VGSN team.

Finally, a future goal of the VGSN is to provide mission operations support for institutions in Virginia (and their partners) that may not have the time, funding, or expertise for full mission operations but would like to pursue a small satellite mission. Examples might include high school students building spacecraft or payloads that do not have a ground station capability. The network also provides opportunities for continued collaboration among Virginia institutions of higher learning for more advanced research and experimentation in smallsat communications related work. While VCC served as the driver for the first operational network in Virginia for smallsat communications, it is the hope of the authors to continue to mature the capability for future mission support.

VGSN Terminology

For the purposes of this paper, a Mission Operations Center (MOC) is any computer system that is executing the ground command and control (C2) software for a particular spacecraft. The personnel operating the MOC are expected to be the subject matter experts for a particular spacecraft and are the ones responsible for the overall health and operations for that spacecraft. Also for the purposes of this paper, a tracking station is defined as the systems operating the actual RF subsystems, such as the antenna pointing control, RF front ends, software radio modems, etc. The tracking station operators are responsible for maintaining and operating the RF apertures, and are not necessarily responsible for a particular spacecraft. Collectively, a MOC and a tracking station comprise an institution's overall ground station. The MOC and tracking station functions may reside on a single computer system with everything physically located together, or may reside in separate locations with functionality spread across multiple compute resources. This terminology will be used to delineate where specific functions are occurring on the network.

Finally, this paper will make frequent use of the International Standards Organization (ISO) Open Systems Interconnect (OSI) Model.[1] The ISO OSI Model is the back drop to begin to breakdown the functional assignments across the VGSN and define where specific functions are expected to occur. For the VGSN, a common simplification is used that collapses layers 7, 6, and 5 into a single application layer.

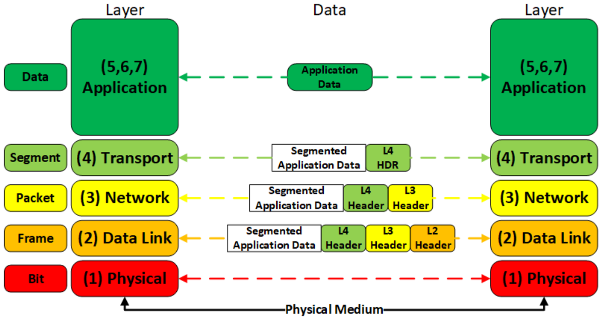

Figure 1. ISO OSI Model depicting logical layer connections, layer datatypes, and encapsulation.

The basic premise of the OSI Model is that logical connections are established between equivalent layers in two communications devices. To move data from one device to another, data is passed down the OSI stack to the physical layer, sent over the communications channel, and then passed back up the layers by the receiving communications device. A process called encapsulation adds information about how to route the data between layers and is used to keep track of the data. Figure 1 shows a graphical representation of the OSI Stack that will be referenced throughout this paper. The concept of sublayers also exists that further breaks down an individual layer. This is not depicted in the figure but will be discussed with reference to the data link layer and the AX.25 protocol in more detail in subsequent sections.

VCC MISSION

Mission Overview

The Virginia Cubesat Constellation (VCC) mission is a collaborative effort between multiple Virginia institutions under the leadership of the Virginia Space Grant Consortium (VSGC) which is serving as the Principal Investigator (PI) institution. The VCC university teams include Virginia Tech (VT), University of Virginia (UVA), Old Dominion University (ODU), and Hampton University. The primary objective of the VCC mission is to provide a hands-on learning experience for undergraduate students in the Commonwealth interested in careers in the space industry. The primary science objective of the mission is to obtain orbital decay data of each cubesat and to either develop new, or validate existing, atmospheric drag models. The constellation consists of three 1U cubesats, each with a slightly different mass and drag profile. Each spacecraft, one per institution, was built by undergraduate students, at VT, UVA and ODU. Hampton University is in charge of mission data analysis and modelling.

The VCC mission is part of NASA's CubeSat Launch Initiative (CSLI) which provides opportunities for small satellite payloads built by universities, high schools, and non-profit organizations to fly on upcoming launches. It is funded by the NASA Undergraduate Student Instrument Program (USIP), the Virginia Space Grant Consortium, and internal research and development funds at the various universities. The Undergraduate Student Instrument Program is managed by NASA's Wallops Flight Facility on Virginia's Eastern Shore.

On July 3rd, 2019 the three cubesats of the VCC mission were deployed to orbit from a Nanoracks deployer onboard the International Space Station (ISS), shown seconds after release in Figure 2, a proud moment for all involved. Over 160 undergraduate students across Virginia have participated in the VCC mission, with support from graduate students and faculty. The three spacecraft in the constellation are named for the three Roman Goddesses on the reverse of the Virginia State Seal; UVA chose Libertas, the goddess of individual liberties; Virginia Tech chose Ceres, the goddess of agriculture; and Old Dominion University chose Aeternitas, the goddess representing eternity. The motto on the reverse of the Seal is *Perseverando* translated in English to *Persevering*, a particularly appropriate motto for smallsat work.

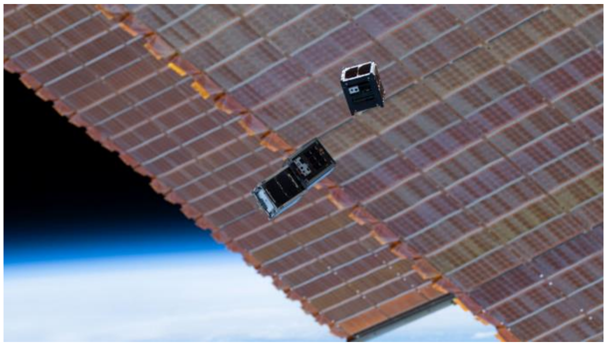

Figure 2. The Virginia Cubesat Constellation cubesats seconds after deployment from the International Space Station.

VCC Communications Systems Overview

The fundamental design of all three spacecraft are nearly identical in terms of communications subsystems. All spacecraft carry the well-known Lithium-II radio from Astronautical Development, LLC (AstroDev). The spacecraft operate in the 401 MHz band under FCC issued experimental (part 5) licenses. Sterk Solutions Corporation provided FCC license application and related services, and served as the single point of contact to the FCC for all three member institutions of the constellation in order to obtain the required FCC licenses. Table 1 gives the spacecraft names, NORAD IDs, owner institutions, FCC call signs, and center frequencies of each spacecraft. The spacecraft frequency assignments are adjacent to each other in the spectrum with 40 kHz separation. All three spacecraft utilize deployable crossed dipole UHF antennas, with UVA and VT using a commercial off the shelf (COTS) solution from Endurosat and ODU using a custom built antenna. All three spacecraft share identical physical and link layer specifications with the exception of the specific center frequency. Each operates with GMSK modulation at a bit rate of 9600 bits per second (bps) using AX.25 Unnumbered Information (UI) frames.[2] At the packet layer and above, each university was allowed to exercise design freedoms to implement their own packet specification (contents of the AX.25 frame information field). Each link is simplex in design, therefore the same center frequency is used by both the spacecraft and the associated ground station, with uplink/downlink transmission direction alternated as necessary.

Table 1. VCC Spacecraft Frequency Information

Spacecraft Name	NORAD ID	Parent University	FCC Call Sign	Frequency [MHz]
VCC-A (Libertas)	44428	UVA	WJ2XMR	401.04
VCC-C (Aeternitas)	44430	ODU	WJ2XOH	401.08
VCC-B (Ceres)	44431	VT	WJ2XMS	401.12

VCC Mission Operations Overview

While the overall science objectives were the same, each university was allowed a significant amount of design freedom in the development of its specific spacecraft. A significant amount of collaboration on the design, component selection, and test processes between VT, UVA, and ODU was achieved in order to ensure that all constellation wide mission objectives would be achieved. Weekly calls between student team leads, and in the height of development and integration activity, often multiple calls per week, were normal. In many cases, similar, often identical, components were used, such as the radio. In other cases, vastly different implementations were used. Examples of different implementations include the flight computers, the flight software architecture, higher layer communications protocols, ground command and control software, and the command authentication mechanisms.

In the original VCC mission concept, each university would be responsible for commanding its own spacecraft via its own ground station. As with the individual spacecraft design freedom, the original objectives also allowed for a significant amount of freedom in the mission operations, with each institution responsible for its own spacecraft. Once the specific science data was obtained and submitted to the mission data warehouse, the operational requirements would have been met for a given institution. Higher level data formatting, normalization, and analysis could then be executed in order to achieve the constellation level science objectives.

The design and operational freedom have direct impact on the design of a unified ground station network and implications for where specific functions should or should not reside on the network. Clear definitions of process boundaries is a strict requirement for the network in order to achieve both flexibility and maintain appropriate levels of security. Where these boundaries are drawn for the VGSN and the associated design decisions are discussed in further detail in subsequent sections. This also has significant implications from a regulatory perspective, as the original licenses followed the single ground station to single spacecraft pattern, but the individual tracking stations needed authority to transmit on any of the three assigned constellation frequencies to support network operations. This will also be discussed in more detail in subsequent sections.

GROUND STATION OVERVIEW

Decision to Re-band the VCC Mission

As originally proposed, the VCC mission was planned for operation in the Amateur Satellite Service spectrum at 435 to 438 MHz, either under Part 97 (Amateur Radio) or Part 5 (Experimental) licenses. As the mission design progressed, it was decided that coordinating Amateur band frequencies for Part 5 (Experimental) operation would be challenging, and the mission scope at this point did not include Amateur activity. Therefore, a decision was made to change the frequencies to operate in the range of 401 MHz, outside the amateur band under Part 5 (Experimental) licenses. Coordination in this range, presented challenges as well as this is spectrum occupied by multiple federal and commercial entities as well as experimental systems, which is discussed more in subsequent sections of this paper.

Unfortunately, by the time the change in frequency occurred, ODU and UVA had already procured ground station antenna and RF equipment for the 435-438 MHz allocation and were forced to modify their systems for operation at the new center frequencies. This resulted in generally functional systems, but with less than optimal performance as the antennas and RF systems were now operating outside the original design parameters. VT had other mission requirements outside of the VCC mission for the already existing Amateur Radio ground station subsystem, and therefore modification for 401 MHz was not an option. Through support from the Aerospace and Ocean Engineering (AOE) department at VT, the VTGS was able to procure new equipment specifically designed for the 401 MHz band and was able to get the systems installed in time for the mission. This situation was a major contributor to the decision to develop the VGSN as a number of ground station related issues presented themselves with the UVA and ODU systems early after deployment of the spacecraft, leading to a constellation level decision to route commands and data through the VTGS and/or WFF systems that had specific support for the 401 MHz band.

VT Ground Station Overview

The Virginia Tech Ground Station (VTGS) began in 2011. Since then, the system has evolved and now includes multiple RF apertures covering multiple frequencies (VHF to Ku-Band). On and near the Blacksburg campus, the VTGS consists of three separate tracking stations and two separate MOCs along with a core network, compute resources, and other IT related resources.

For the VCC mission, a pair of 401 MHz antennas from M2 Antennas Inc., mounted on an Alfa Spid Big-RAS/HR rotator are utilized. The feed points are polarization sense selectable, allowing for switchable RHCP or LHCP polarization. Mast mounted low noise amplifiers, with transmit bypass features, from Advanced Receiver Research (ARR) are utilized. The high power transmit amplifiers are based on surplus equipment and components from Improvised Explosive Device (IED) jammers that were donated to the VT Hume Center and decommissioned by researchers at that organization. An Arduino microcontroller based custom amplifier control board was developed to perform all amplifier control, monitor, alarm (CMA) functions and to interface with the rest of the tracking station systems over low latency Ethernet.[3,4,5] A GNU Radio Out Of Tree Module (OOTM) called *gr-rffe_ctl* was developed to provide direct TX/RX sequencing control from within a GNU Radio flowgraph.[6] An Ettus Research N210 USRP with UBX daughtercard is the backend SDR. The USRP is connected to a GPS disciplined oscillator (GPSDO) for a stable frequency and time reference. An additional server monitors all hardware systems, such as the RF front ends, and antenna pointing control, remote relay control (for power and polarization control), a local weather station (for wind and local lighting alerts), etc.

The VTGS tracking stations are connected to a core ground station network over the VT campus network using gigabit connections and IPSEC VPNs. For the VCC mission, the systems are remotely controlled from console stations in the Space@VT Operations Center. A second Operations Center also exists in the Hume Center's research building, but this is for more restricted access programs. Antenna pointing control is accomplished using custom software in conjunction with free satellite tracking software (gpredict). All VTGS systems utilize the Ubuntu Linux operating system. The VTGS system is remotely accessible over a secure VPN connection for authorized student and faculty operators, allowing operations to be conducted either from the main Operations Centers during normal work week hours or from a separate location as is sometimes necessary for late night or early morning passes. The command and control software selected for the VCC mission (and planned for continued use in future missions) is COSMOS from Ball Aerospace.[7] Full system automation is a major goal on the VTGS development roadmap.

The VTGS core network and servers were utilized to establish the prototype VGSN. A server existing outside the core VTGS LAN on a Demilitarized Zone (DMZ) port of the VTGS firewall has been designated as a *client portal* interface for external VTGS partners. It is this client portal that was utilized to instantiate the prototype VGSN. It is also this server that provided the connection endpoint for the Wallops interface. Finally, the client portal also provide additional physical layer signal processing for the WFF interface. The implementation details for this, heavily based on the use of Docker containerization, are discussed in more detail in subsequent sections.

UVA Ground Station Overview

The UVA tracking station is very similar to the VT tracking station. A pair of M2 Antennas UHF crossed yagis, mounted on an Alfa Spid Big-RAS rotator are used. An ARR mast mounted, bypassable LNA is used, connected to the inside plant RF front end over low loss coax. As mentioned, the most notable difference is the required re-banding of the RF front end systems to move the frequency of operation from the UHF Amateur Satellite Service (~436.5 MHz) down to the 401 MHz band. The UVA team also developed a custom RF front end based on COTS components and their own microcontroller based TX/RX sequencer. An Ettus Research N210 with UBX daughtercard and internal

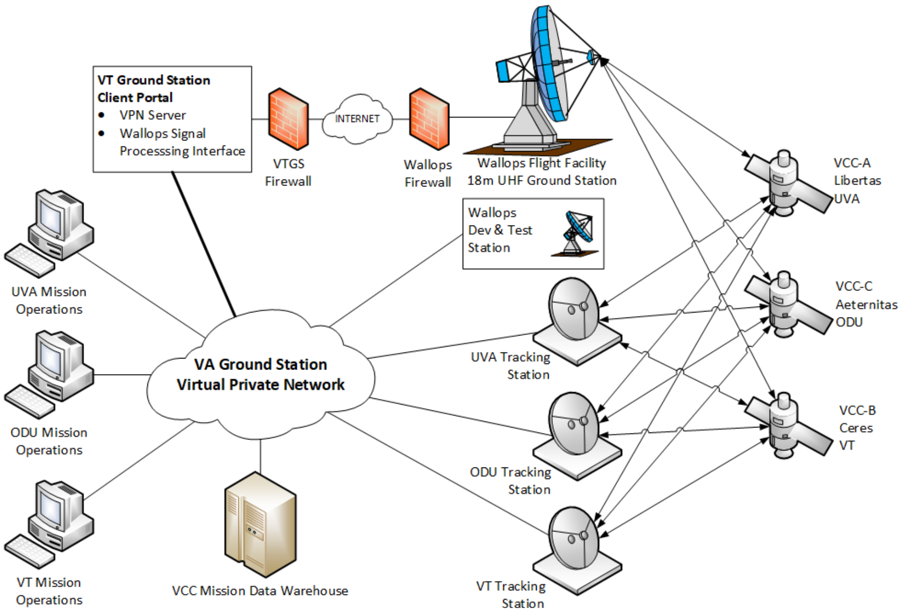

Figure 3. Virginia Ground Station Network overview block diagram.

GPSDO is used for the backend SDR. GNU Radio is utilized for the core signal processing.

The UVA team elected to develop its own custom command and control software for mission operations with their spacecraft. Another notable difference is that the UVA team tends to run their command and control software on the same server that runs the signal processing and other tracking station functions, such as antenna pointing. This is an example of the MOC functionality residing on the same physical system that performs the tracking station functions, which also has implications for network design. Access to the server is via a secured remote desktop application, allowing operations to occur remotely if necessary.

ODU Ground Station Overview

The ODU ground station utilizes a single UHF crossed yagi from M2 Antennas mounted on a Yaesu rotator system. The ARR LNA is a different model than the VT/UVA variants and is integrated into their RF front end system located inside the building. ODU also elected to develop its own custom RF front end and associated sequencing circuitry. Like UVA, ODU was forced to re-band most of their RF systems for operations down in the 401 MHz allocations for the VCC Mission.

For mission operations, ODU also elected to develop custom command and control software and to run that software on the same computer system performing the signal processing functions for nominal operations. The ODU system is not currently remotely accessible, requiring a human operator to be physically present at the control computer to operate the tracking station. The command and control software can be run on a separate computer system however if operating over the VGSN.

Wallops UHF Smallsat Ground Station

The highly capable and well known Wallops UHF smallsat ground station consists of an 18m parabolic antenna and a collection of hardware and software radio backend solutions. Complete description of the WFF system is out of scope for this paper, but important points relevant to the VGSN will be discussed. For the VCC mission, WFF utilized a newly installed commercial software radio solution call *quantumGND* from Kratos Inc. During the operational use of this system (December 2019 to January 2020) for VCC operations, the *quantumGND* architecture did not support AX.25 processing, though AX.25 is on the roadmap for development (and may be complete by the time of this publication).

Therefore, a ‘split signal processing’ architecture was developed where link layer signal processing and some bit level physical layer processing was performed on the VTGS client portal for all VCC users and remaining physical layer processing was performed using the *quantumGND* framework and the rest of the WFF-GS. A connection to the VGSN was established automatically on a pass by pass basis, allowing the split signal processing to be executed in real time, and the specifics of the signal processing and this interface are discussed in more detail in subsequent sections.

NETWORK DESIGN OVERVIEW

The primary purpose of the VGSN is to link any of the Mission Operations Centers to any of the Tracking

Stations for real time smallsat communications and operations. Figure 3 shows the high level overview of the VGSN in its current form. Given the design freedoms at higher layers of the protocol stack for the VCC mission, it is important to clearly define the delineation of which system is responsible for a given function or process. This is accomplished by using well defined interface control documentation. This is important not only for technical reasons, but also for clear definition of the roles and responsibilities of the sub teams and for security purposes.

For example, should a VT student have access to the cryptographic keying material used for authentication on the UVA spacecraft command uplink (and vice versa)? Should UVA students learn how to operate VT COSMOS C2 software, or is that unnecessary additional burden on the UVA team? Perhaps this is OK for an undergraduate educational mission such as VCC which was very collaborative, but future missions should not be required to surrender their C2 software and encryption/authentication material to a remote tracking station on the VGSN. Ideally, UVA students would continue to operate the UVA spacecraft, regardless of which tracking station is utilized on the VGSN and without concern over or loss of direct control over their own mission specific cryptographic material. UVA and VT students were used in the above example, but this general design criteria should hold true for any MOC on the network.

AX.25 and HDLC Link Layer Sublayers

Full discussion of the AX.25 protocol, primarily a link layer protocol, is outside the scope of this paper, however there are a few important points that must be addressed in terms of interfaces over the VGSN. The AX.25 protocol inherits features from the High-Level Data Link Control (HDLC) protocol and sometimes the specific functions can be confusing. The link layer for AX.25 can be broken down into an upper and lower sublayer. For this document, the upper layer will be referred to as the AX.25 sublayer. The lower sublayer (closer to the physical layer) will be referred to as the HDLC sublayer. Collectively these two layers make up the overall AX.25 link layer.

The AX.25 sublayer is responsible for the formation of an AX.25 Frame. More specifically for smallsat missions such as VCC, the Unnumbered Information (UI) Frame type is utilized. This sublayer is responsible for generating the AX.25 packet header, including the source and destination call sign, Control Bits, and PID, as well as forming the information field. The information field of the AX.25 frame contains the higher layer TM/TC data as specified for each spacecraft (different for each for the VCC spacecraft). Figure 4 shows the structure of the AX.25 UI Frame format.

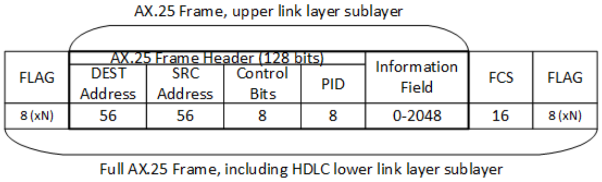

Figure 4. AX.25 Unnumbered Information (UI) Frame, sublayer functions highlighted.

Once the AX.25 frame is formed by the upper sublayer (AX.25 Header and information field) it passes the frame to the lower HDLC sublayer. The HDLC sublayer first computes and appends the Frame Check Sequence, used to detect bit errors at the receiver. Then it performs a process called bit stuffing that inserts 0s in the bit stream to prevent any sequence of 6 contiguous 1s or more from appearing. It then prepends and appends the Flag bytes which are used for frame detection and delineation. Often times multiple flag bytes are appended and prepended, for example to aid in receiver synchronization. The Flag byte value is 0x7E in hexadecimal representations which is 0111110 in binary representation. The six contiguous 1s in the flag byte are why bit stuffing of the AX.25 frame is required in order to prevent false detections should 6 contiguous 1s be present in the AX.25 frame data.

It is between this sublayer boundary, between the AX.25 upper sublayer and the lower HDLC sublayer, that the line is drawn for the VGSN with respect to systems using the AX.25 protocol. All processing from the application layer to the upper AX.25 sublayer are the responsibility of the Mission Operation Centers and their respective command and control software for their mission. The tracking stations are then responsible for the lower link layer (HDLC) processing and actual RF transmission and reception. The TCP/IP protocols are called out as they are used to transport the information over the network.

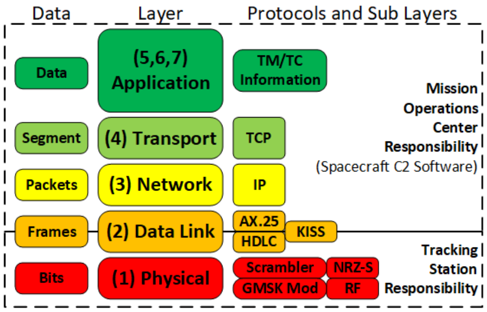

Figure 5. Functional Assignments for VGSN Entities using the AX.25 protocol according to OSI Model.

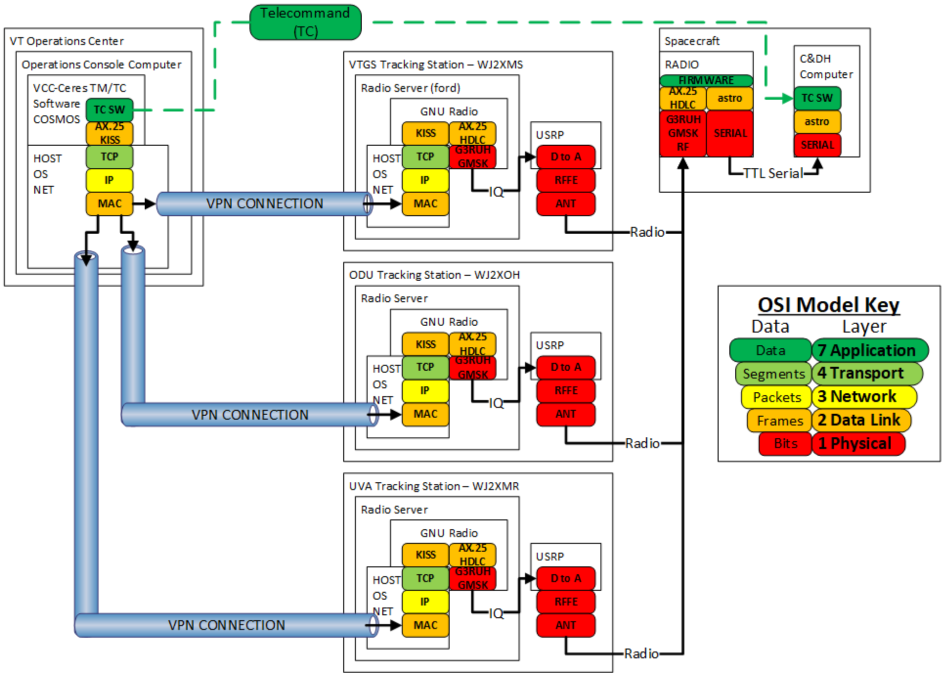

Figure 6. OSI Stack breakdown for the VCC Mission over the Virginia Ground Station Network (VGSN).

KISS Protocol Design Decisions

The 'Keep It Simple Stupid' (KISS) protocol is a companion link layer protocol often utilized with AX.25 frames.[8] Like AX.25, full description of the protocol is out of scope for this paper, but it is a key part of the VGSN support for AX.25 protocols. This protocol was developed by Mike Chepponis, K3MC and Phil Karn, KA9Q in the mid-90s in order to expose more modem functionality to experimenters, usually Radio Amateurs.

The KISS protocol is often used to 'wrap' an AX.25 frame before transport over a given link. This is directly relevant to the previous section discussion about the separation of AX.25 and HDLC sublayers. KISS was developed as a mechanism to move data between the two sublayers, which historically was typically over an RS-232 serial link between a host computer and a Terminal Node Controller (TNC, also known as a modem) and more recently is done over IP network links using either TCP or UDP transport protocols. The modems perform the HDLC and lower layer operations (at the tracking stations), but developers and experimenters should have direct access and control over the AX.25 frame and higher layers (at the MOC in the C2 software).

A number of existing software and hardware modems and higher layer packet decoder/encoder programs make use of the KISS protocol in conjunction with AX.25. For this reason, a design decision for the VGSN was made to support this protocol. This could potentially allow MOCs to use existing software for the packet generation. Similarly, this would also enable the use of existing hardware based TNC at the tracking stations if desired. For the VCC mission, which is all GNU Radio based at the tracking stations, support for both AX.25 and KISS already exists through the *gr-satellites* Out of Tree Module from Daniel Estévez, EA4GPZ, which makes implementation relatively trivial.[9] The C2 software in each MOC on the VGSN must not only generate correctly formatted AX.25 frames, they must also properly wrap that frame using the KISS protocol for transport over the VGSN to the tracking station end point. This is why the KISS link layer protocol is depicted in the middle of the AX.25 and HDLC sublayers in Figure 5 and is the responsibility of both the MOC and tracking station to properly process.

OSI Layer Summary for VCC over the VGSN

The Command and Control (C2) software for a particular spacecraft operates at the application layer of the OSI model, with the general goal of establishing a logical connection with the flight computer on the spacecraft. The 'Data' in this case is the tele-command (TC) data originating at the ground and sent to the spacecraft and telemetry (TM) data generated in the spacecraft and sent to the ground (this includes both telemetry and mission science data even though the acronym references telemetry). The TM/TC data is broken down into smaller and smaller chunks as it moves down the OSI stack for bit level transmission over the radio link and then reassembled as it moves back up the stack on the receiving end of the link.

Figure 6 depicts a nearly complete breakdown of the OSI stack flow for an AX.25 TC packet sent to a spacecraft over the VGSN for the VCC mission. To the left of the figure, the VT MOC and VCC-Ceres C2 software is shown as an example, but this could also be the UVA or ODU MOC running their respective C2 software. The TC data is generated in the MOC C2 software down to the AX.25/KISS layer. The KISS frame is then sent in TCP packets over an IP network connection (VPN connections, discussed more in the next section) to the desired tracking station IP endpoint. The tracking station endpoint is a GNU Radio flowgraph that begins with a TCP socket connection. The KISS frame is received by the flowgraph, and the rest of the link layer and physical layer operations are performed by the tracking station RF systems, culminating with a transmission over the radio channel to the spacecraft. The spacecraft receives the transmission in the Lithium-II radio, performs the necessary physical and link layer operations and sends the TC packet over a TTL serial connection to the spacecraft flight computer, depicted as a Command and Data Handling (C&DH) computer in Figure 6. Note that the AstroDev Lithium-II does not use the KISS protocol and instead has a custom protocol used for transport over the serial connection, denoted as a link layer 'astro' protocol in the figure. While the example given is for a TC packet sent over the uplink path to the spacecraft, the same OSI stack flow for a TM packet generated in the spacecraft and sent over the downlink is used, simply in the reverse direction.

Note that at no time was a tracking station required to have specific knowledge of the TM/TC data at the application layer. This is an important detail. This shows the modularity of the network and the importance of clear definition of process boundaries. This also has important security implications. The VCC spacecraft links, while not actually encrypted, do use encryption techniques for command authentication onboard the spacecraft. The specific techniques will not be discussed in this paper, however it is within the C2 software, when generating a TC packet, that this step is applied. Because of the selected boundaries relative to the OSI model, access to the cryptographic keying material for a particular spacecraft can remain with that spacecraft's MOC and does not need to be distributed outside of the owner institution. Put simply, with reference to the above example, VT does not have to give out its spacecraft authentication keys to UVA or ODU.

Network Connectivity & Security

As the MOCs and tracking stations are spread across the Commonwealth of Virginia (and potentially elsewhere in the future), the public internet is the primary mechanism for connecting all locations. Thus cybersecurity protections are of utmost importance to ensure that only authenticated end users are authorized to send commands to a specific spacecraft. The command sent to the spacecraft and the data and telemetry from the spacecraft should also be secure from potential eavesdroppers on the Internet. Finally, the traffic flowing between MOC and tracking station endpoints should have a degree of trust that the data was not manipulated in transit. Thus, an early design decision for the prototype VGSN was the selection of Virtual Private Network (VPN) technology to provide these security mechanisms. This is also depicted in Figure 6 as secure tunnels connecting MOCs to tracking stations. More specifically, use of the free and open source OpenVPN package was selected. When established and properly configured, the VPN connection form a virtual LAN that ensures proper authentication and provides commercial grade encrypted communications links between endpoints on the network, ensuring protection of the data in transit from eavesdropping and unauthorized data manipulation.

The agreed upon IT security policy within the VGSN team was to generate a pair of institutional VPN credentials, one for the MOC and one for the tracking station. These credentials are generated and maintained in the VTGS core network that runs the VPN server. The credentials were delivered via a secure file transfer mechanism to the technical leads for communications at the member institutions. The authority at the member institution (PI, student team lead, or PI authorized representative) accepts responsibility for protecting the credentials according to their own institution's IT policies and procedures.

This method of generating institutional VPN credentials as opposed to specific user credentials was selected for two reasons, one managerial and one technical. The managerial reason has to do with individual institution IT policies and to minimize unnecessary effort on the VPN server side. So long as the use of VPN client software is allowed, connection to the VGSN avoids conflicting with IT policies at the member institution, such as potentially opening specific ports on firewalls for example. This also avoids undue burden on the VPN server side to generate and maintain credentials for every possible student or faculty member that may need to access the VGSN. Student turnover is a common problem with universities (especially with undergraduate programs), and tracking that for three separate universities could be an administrative problem. Therefore, when a particular operator connects to the network, they do so under their parent institution's VPN credentials.

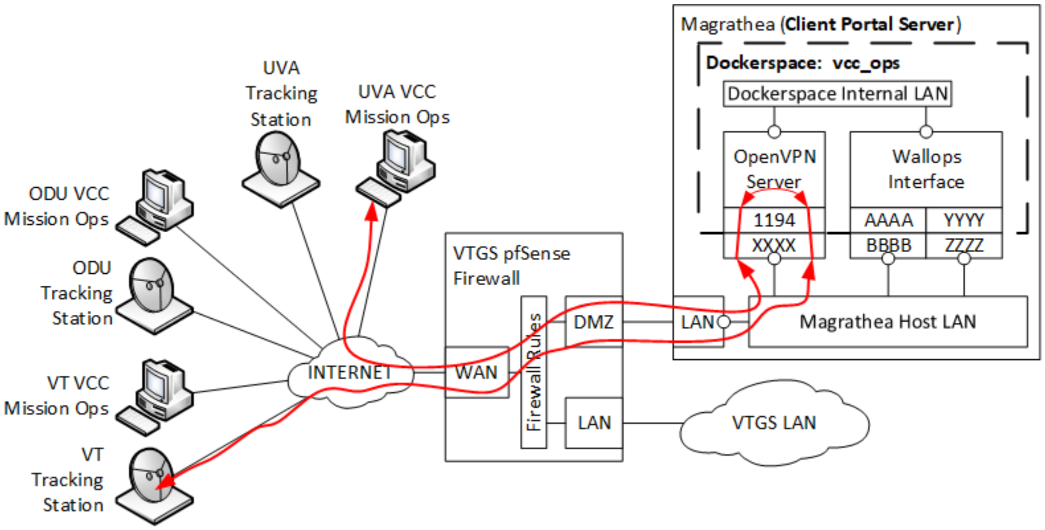

Figure 7. Connecting MOCs to Tracking Stations via the VTGS client portal server using mission specific containerized instances of OpenVPN.

The technical reason for this decision is that the VPN established for VCC Mission Operations uses static IP address assignment within the VPN LAN. Assigning static IPs to 6 fixed sets of credentials (1 tracking station and one MOC at three universities) is easily managed. The static IP addresses on the VPN for the tracking stations are the end points that the C2 software in the MOC connects to for real time command and control operations with their spacecraft. In order for the clients to access each other over the VPN, the server must also be configured to allow traffic between the clients, which does pose a potential security risk should any of the client side credentials for the VPN ever become compromised.

In its current form, any and all traffic between VPN clients is allowed to flow over the VPN, not just the TCP connections between the C2 software and the GNU Radio modems. Should any one of the systems be compromised by an attacker, they would have direct access to any client via the VPN tunnels, bypassing any firewall rules that may be in place at a given institution's higher level IT security team. This threat is somewhat mitigated in that the individual missions are isolated from each other through the use of Docker containers, which will be discussed in subsequent sections. Host level firewall rules for the client systems (such as *ufw* on Ubuntu) can also help mitigate this problem. As with any cyber security threat, layered approaches that balance security with usability is an important consideration.

The use of static IP assignments on the VPN is not optimal, but was a fast solution given timeline constraints for the VCC mission and standing up the prototype VGSN. In the case of the VPN implementation for the VCC mission, all organizations involved are highly trusted due to the collaborative nature of the mission. Should VPNs continue to be utilized, more work will be done that avoids the need for static IP assignments, such as the use of an internal Domain Name Service (DNS) system on the VPN LAN. Additional cybersecurity protections will also be explored, such as server side firewall protections to control the flow of specific traffic between clients. Future work will include exploring alternative techniques that may or may not involve the use of VPNs and may use some other form of secure IP based communications for moving data over the network.

NETWORK IMPLEMENTATION

Firewalls, Client Portals, & 'DockerSpace'

As mentioned, the VTGS core network was utilized to instantiate the first prototype VGSN. More specifically, the VTGS uses a concept of a *client portal* server that is connected to the Demilitarized Zone (DMZ) port of the pfSense based firewall for the VTGS system. Deployment of this server on the DMZ port isolates this public Internet facing server from the internal VTGS LAN and by extension the rest of the VTGS systems connected to the core VTGS network over persistent IPSEC VPN links. This is depicted in Figure 7. To be clear, the core VTGS private network is an independent network from the Virginia Ground Station Network and the VTGS systems are supporting missions other than VCC that must retain a significant amount of program separation. Some of the equipment used does overlap however, as the VTGS ground station had available IT resources to support the prototype VGSN for the VCC mission.

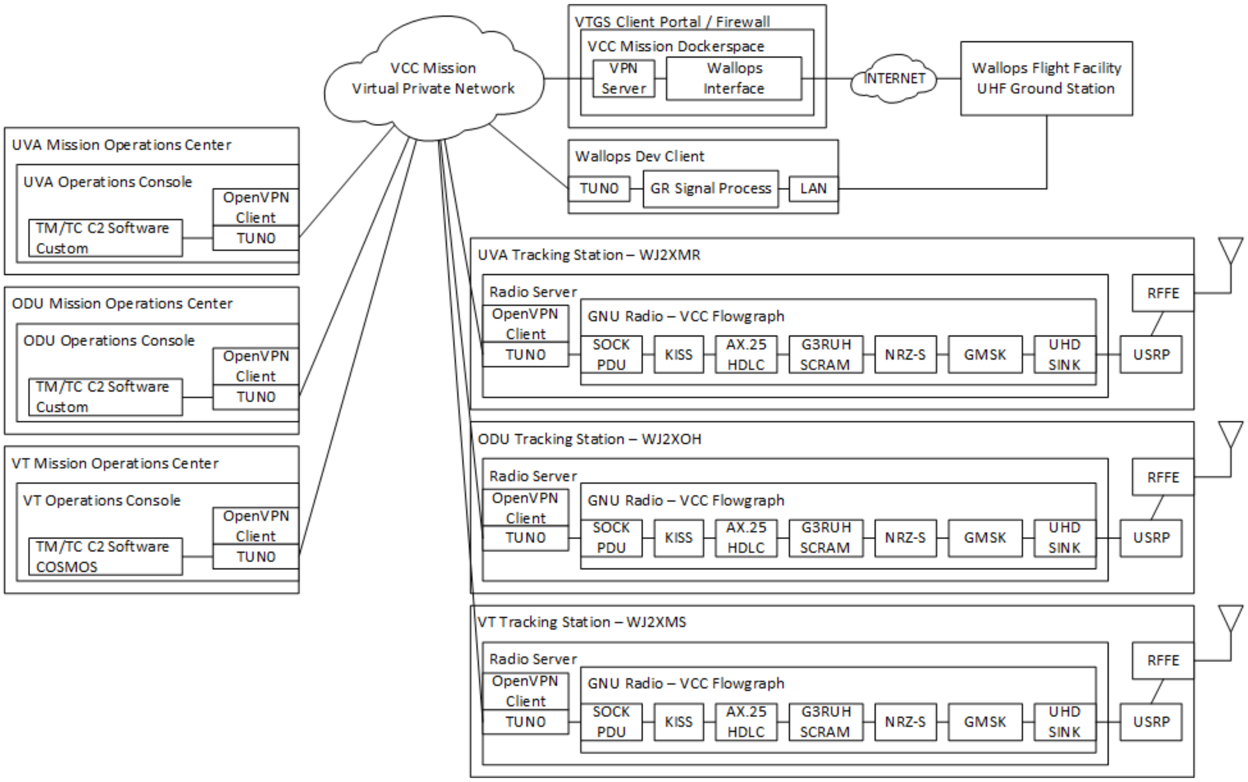

Figure 8. Detailed diagram of the logical connections from MOCs to the tracking stations for the VCC Mission on the VGSN. Additional GNU Radio detail shown.

On the client portal server itself, the VTGS employs the use of Docker for containerization and isolation of mission specific functions. A collection of containers for a specific mission is referred to as a *dockerspace* by VTGS personnel and is depicted by the dashed lines on the client portal server to the right of Figure 7. The specific containers in a dockerspace implementation are connected to each other over isolated Docker networks. Each dockerspace instantiation is mission specific, allowing the client portal server to support multiple missions simultaneously while still maintaining isolation between the specific mission dockerspaces. Common container images can be reused by multiple missions by assigning an instance of them to a new dockerspace, reducing duplication of development effort. An example of this is OpenVPN, which is used on multiple missions, but with different instances running in separate dockerspaces with completely separate certificate authorities and client credentials, preventing the possibility of one mission client connecting to a different mission's dockerspace. A key step in standing up a new mission for the VGSN is the configuration of the dockerspace for that mission on the VTGS client portal server.

The VTGS Firewall also requires some discussion with respect to the VGSN. As mentioned, the client portal server resides on an isolated DMZ port of the firewall, separating its operations from the rest of the VTGS network. Specific non-standard ports are then opened on the public facing WAN port of the firewall. Firewall rules then forward these ports directly to the dockerspace for a specific mission. The only ports that are actually opened for this purpose are related to OpenVPN. A containerized instance of an OpenVPN server is the primary connection mechanism for a given mission. Internal to the Docker container the standard OpenVPN port 1194 is mapped to a non-standard port on the client portal host LAN. This port number is then mapped through the VTGS firewall and is exposed to the public internet. Only clients with specific credentials are then allowed to connect through the firewall to the OpenVPN instance. An example of the UVA MOC connecting to the VT tracking station, a common use case for the VCC mission, through the firewall and VCC dockerspace OpenVPN Docker container is depicted in Figure 7.

University Tracking Station Interface

The primary members of the VGSN include VT, UVA, and ODU, with VTGS resources hosting the core VGSN network resources and running the VCC Mission specific dockerspace services. As discussed in the network design section of this paper, for the VCC Mission, the MOCs execute the satellite specific C2

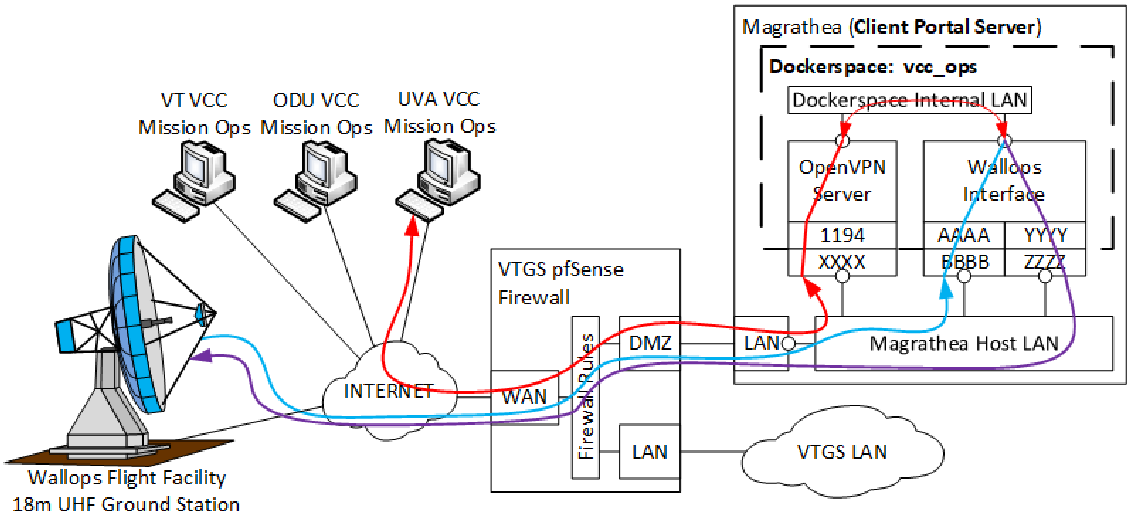

Figure 9. Wallops Flight Facility UHF Smallsat Ground Station Connection the VGSN for the VCC Mission.

software connected to the desired tracking station over secure VPN links. More specifically, AX.25 frames, encapsulated in KISS frames, are sent via the VPN connections over TCP/IP to GNU Radio flowgraphs running on tracking station systems. The GNU Radio flowgraphs then perform additional link and physical layer processing and interact with the spacecraft over the physical RF channel.

Figure 8 shows the logical connections available from MOCs to the Tracking Station endpoints. GNU Radio flowgraphs handle the signal processing. Full description of the flowgraphs utilized for VCC is outside the scope of this paper, but the major steps in the processing are depicted in the figure. Each TUN0 OpenVPN interface is a statically assigned IP address on the VCC VPN. The MOC simply connects to the correct IP address on the VPN, on the correct TCP port established by the Socket PDU server running inside the GNU Radio flowgraph.

Wallops UHF Ground Station Interface

As noted, the Wallops Flight Facility UHF Smallsat Ground Station (simply referred to as WFF-GS in this paper) was used between December 2019 and January 2020 in support of VCC communications troubleshooting efforts. The regulatory details of this are discussed elsewhere in this paper and this section will cover the technical interfaces. Figure 3, Figure 7, and Figure 8 all depict connectivity to the WFF-GS. However, the specific interface mechanisms are not the same as the nominal university tracking station interfaces over the VGSN. In order to interact with the WFF-GS, the VGSN had to provide additional IT and signal processing services in order to both meet the mission requirements as well as the requirements from NASA.

The WFF-GS was not part of the core VCC Virtual Private Network and did not run a VPN client on their system for the VCC mission. This VPN method would have been overly cumbersome for the WFF-GS system operators who are handling multiple missions. Instead, firewall rules at the VTGS and WFF allowed direct connection of TCP clients within the WFF-GS network to reach out to TCP server endpoints.

These connections were established automatically as part of WFF scheduling, with the connection coming up a minute or so before a given pass and disconnecting immediately after the pass was complete. In order to reduce scheduling complexity and to provide increased flexibility for the VCC team, a single endpoint was presented to the WFF-GS within the VCC dockerspace on the VTGS client portal, depicted in more detail in Figure 9. Two TCP ports were utilized for the connection to the WFF-GS, one for uplink depicted in purple, and one for downlink depicted in blue in Figure 9. The dockerspace internal ports were mapped to the client portal host network, which were then passed through the VTGS firewall, over the DMZ to WAN interfaces, using port forwarding rules. This did present a bit of a security issue, however, as mentioned previously, these ports are only opened on the firewall just before a pass of a specific satellite and are closed just after the pass is complete and when not in use in between passes.

The same VCC VPN connection was used to connect the university MOCs to the VCC dockerspace, depicted in red in Figure 9 as an example of the UVA connection. However, instead of connecting back over the VCC VPN to a university tracking station endpoint IP (see Figure 7), a different IP address over the internal dockerspace LAN for the VCC mission is presented to the MOCs,

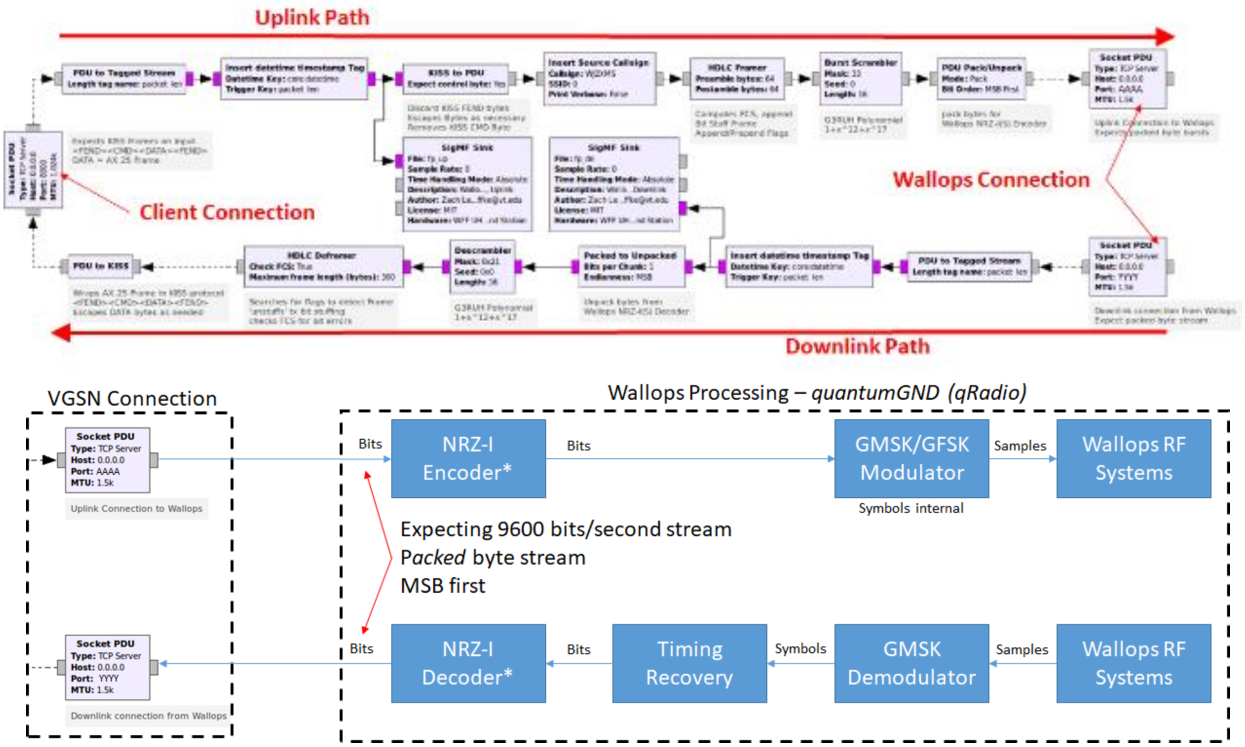

Figure 10. Wallops Interface GNU Radio signal processing running inside a Docker container on the VTGS/VGSN client portal server inside the VCC mission dockerspace (top). WFF-GS physical layer signal processing using the *quantumGND* software radio framework (bottom).

called the 'Wallops Interface' in Figure 7 and Figure 9. From the operators' perspective in the MOCs, they now simply have a 4th option for a tracking station IP endpoint and no modification to their existing C2 software is required to interface with the WFF-GS. The Wallops 'Dev' system shown in Figure 3 and Figure 8 refer to a system used for onsite flatsat testing and is not part of nominal operations.

As discussed previously, the software radio framework used by the WFF-GS for the VCC mission, Kratos' *quantumGND*, did not natively support AX.25 at the time that the station was used by the VCC team. This required additional signal processing to be performed within the Wallops Interface Docker container in the VCC dockerspace. The line between MOC and tracking station responsibility, with respect to the OSI model, shifted lower and actually resides in the physical layer. This *could* have been accomplished by modifying each University's C2 software in the respective MOCs, but this would have led to two instances of the C2 software needing to be maintained at the MOCs, would have broken the agreed upon separation of responsibilities with respect to the OSI model, and would have required triple the effort to implement, which was not feasible with VCC timelines for access to the WFF-GS. The required signal processing was already being performed in the tracking station GNU Radio flowgraphs, so to avoid confusion and to reduce complexity, this signal processing was inserted in a central location within the Wallops Interface Docker container. This container image is essentially an instance of Ubuntu 18.04 with GNU Radio installed to perform the signal processing. This technique dramatically simplified the IT requirements and implementation details for all of the universities and NASA WFF personnel, simply by standing up a new container within the VCC specific dockerspace.

The WFF-GS *quantumGND* system was not capable of handling the AX.25/HDLC protocols as well as the required G3RUH scrambling process at the physical layer. The system was capable of handling the NRZ-S encoding (NRZ-S is a form of NRZ-I encoding, where the S stands for specific conventions commonly used on satellite links) as well as the GMSK modulation and demodulation, along with the required timing recovery on the receive path. Therefore, the 'split signal processing' agreed to between the VGSN and WFF-GS for the VCC mission involved performing all signal processing up to and including the scrambler within the VGSN. At this point data is passed over the TCP connections between VCC/VGSN endpoint and the WFF-GS. This complete process is shown in Figure 10.

The top half of the figure depicts the GNU Radio flowgraph in the VGSN Wallops Interface Docker container performing the required signal processing up to the scrambler and the bottom half of the figure shows the rest of the signal processing completed in the WFF-GS *quantumGND* system.

Network Scheduling & Coordination

A number of scheduling and coordination steps had to be conducted by the VCC team prior to execution of operations with either the VCC tracking stations (VT, UVA, and ODU) or with the WFF-GS. Some of these steps were directly tied to FCC requirements per the experimental licenses of the university tracking stations. The regulatory aspects of this are discussed in a subsequent sections, but this should be understood as part of the operational flow as it impacted the coordination process with NASA's Johnson Space Center (JSC).

The VT team developed an overall VCC Mission network scheduler process using a combination of AGI's Systems Tool Kit (STK) and Python scripts. Typically, this scheduling process would be executed for approximately a one month window. STK was utilized to produce Access and Azimuth/Elevation/Range (AER) reports. This simulation accounted for all four tracking stations on the network and all three spacecraft in the constellation. The STK simulation produces every Access reports for every combination of tracking station and spacecraft as well as AER reports with one second time steps for every pass over every tracking station during the simulation time. The access and AER reports are exported from STK in Comma Separated Value (CSV) format.

A Python script was developed that ingests the STK access and AER report CSV files in as inputs and combines them to produce two outputs. The first output was a CSV file that simply time aligned all passes according to the Acquisition of Signal (AOS) timestamp. The typical result was that for a given spacecraft, there would be four sequential entries representing contact times with each of the tracking stations. This combined and time aligned CSV was retained for scheduling internal VCC/VGSN operations. A second output was also produced based on this combined schedule.

The second output CSV was a *network access* report. This essentially collapsed the four entries of a given spacecraft (one for each tracking station) into a single entry where the ground station was represented as a single VCC-NET entity. The AOS for this entry was the earliest AOS timestamp of the four tracking stations and the Loss of Signal (LOS) entry was the last LOS of the four tracking stations. This produced a network schedule where anytime in between AOS and LOS, one of the four ground stations in the VGSN network may be transmitting to the specific VCC spacecraft during its pass over Virginia.

It was this network pass report that was then delivered to the Spectrum Managers at NASA's Johnson Space Center (JSC). The author, serving as the VCC communications lead, as well as the UVA student team lead served as the primary points of contact with JSC (and WFF) to streamline communications and coordination, under the approval of the Constellation PI from the VSGC. The FCC experimental licenses for the VCC mission required that *every* contact with the VCC spacecraft must be approved by JSC. The reason for this is very straight forward. The frequency band used by the VCC mission could produce an interference condition to the International Space Station, potentially during critical operations such as docking or Extra-Vehicular Activities (EVAs) by the ISS crew. Therefore, in the interest of safety to the astronauts and cosmonauts onboard the ISS and vehicles that may be docking with the ISS, all communications in that band had to be verified by JSC that no potential interference condition would be present during critical operations. Interference conditions could be present when either the VCC tracking stations or the VCC spacecraft had line of sight to the ISS. To satisfy this requirement, the VCC team would deliver the one month network pass report to the JSC Spectrum Managers. The JSC team would redline all communications that were planned by the VCC team that satisfied both criteria; line of sight from a VCC station (ground or spacecraft) to ISS and critical operations underway at ISS. JSC was very accommodating in this process, quick to respond to each coordination request, and would send regular updates during a coordinated window if changes were made to the critical operations schedule of ISS (such as a resupply mission launch delay).

Once the JSC coordination requirement was satisfied, and known 'blackout' windows were clearly defined in the schedule, the JSC approved network schedule and the combined schedule (the first CSV output of the network scheduler) were delivered to all VCC operations teams at each institution. Those institutions were then free to operate their own spacecraft from their own ground stations so long as they adhered to the JSC approved communications windows, the times of which could be found on the combined schedule for their specific tracking station and spacecraft. This was the originally planned type of operations for the VCC mission and no further coordination within the VCC team was required.

Alternatively, if a particular institution wanted to execute passes that required use of a tracking station other than their own over the VGSN, then further internal VCC coordination was required. This was a relatively trivial process due to the highly collaborative nature of the mission and since weekly VCC meetings were held into the operations phase of the mission. The coordination details for operating either with the VT tracking station or the WFF-GS would be determined and the next steps taken. If operating with the VT tracking station, coordination was complete and the operations were scheduled at the respective institutions.

If operation with WFF-GS was desired, a further coordination step was required. With the JSC approved schedule in hand, the VCC team next had to coordinate operations with the WFF-GS scheduling team. Not all access windows in the monthly VCC network pass report were available for VCC operations with the WFF-GS system. This is easily understandable as the WFF-GS operates with multiple spacecraft, under multiple levels of mission priority, and has other mission requirements they must meet outside of the VCC mission. The VCC team would work with the WFF scheduling team and select the candidate passes that the WFF team had deemed available on their schedule for the VCC spacecraft. Once again, the NASA WFF team was very accommodating, assigning high priority to the VCC team early on as the whole point was communications troubleshooting.

The WFF-GS would then send out the confirmed weekly schedule for all WFF-GS passes on the Friday prior to the confirmed week of passes, with pass execution beginning the following Monday. The WFF *quantumGND* operators would then schedule the system to connect to the VCC/VGSN endpoints for those passes and the WFF-GS antenna operators would be on site for the specific pass windows. A WFF-GS antenna Operations dial-in conference line was provided so that the VCC team could communicate directly with the antenna operators during a pass window.

Network Operations

The VT Operations Center would become the VGSN Network Operations Center (NOC) for passes that involved resources on the network, with the author serving as the primary VGSN operator. Zoom sessions for the specific passes were scheduled so that all VCC team members, across multiple institutions, could communicate directly with each other and share screens during the pass. Typically, 10 to 15 minutes before a pass, the VCC operators would all be online, including the respective MOC for a pass and the VGSN operator. Initial system checks would be performed to ensure all systems were ready for the upcoming pass. The primary tracking stations used over the network were either the VT system, with the dedicated 401 MHz equipment, or the WFF-GS.

When using the VT systems, the VGSN operator would also serve as the VTGS operator. Over the Zoom session, the GNU Radio flowgraph GUI would be shared so that the MOC operators (typically students) running the specific C2 software for their satellite could monitor the RF spectrum in real time. During these types of network operations, the VTGS operator would handle all tracking station operations. This included controlling the antenna pointing, interactions with the RF front end control system, using the remote relay systems to turn on the various components, executing the correct GNU Radio flowgraphs, and managing any recorded data after a pass (such as raw IQ recordings, in SigMF format, that are always recorded and stored on the VTGS NAS after a pass).

This type of operation was very common and ultimately lead to the establishment of reliable communications with UVA’s Libertas spacecraft. This was accomplished with the UVA MOC operating over the VGSN through the VTGS tracking station. This was executed during a period of time when the UVA ground station was down for maintenance and repairs that developed early in the mission life after deployment of the spacecraft. These repairs were necessary in part due to the forced re-banding of the UVA GS from the 435-438 MHz band to 401 MHz band as discussed in the Ground Station Overview section. Limited time and funding for the UVA GS repairs, along with the wait period for the license modifications (discussed later) for network operations were the main reason for this delay between deployment and first contact with Libertas.

Unfortunately, the Aeternitas (ODU) and Ceres (VT) spacecraft were generally non-responsive to any communications attempts after deployment by any of the VCC tracking stations, either directly from the owning institution or via the VGSN, leading to use of the WFF-GS for more communications troubleshooting. The specific radio issues are out of scope for this paper, but the decision to utilize the WFF-GS because of these issues had direct impact on the development of the VGSN.

When operating with the WFF-GS, again the author would serve as the VGSN NOC operator. This primarily required connecting into the Wallops Interface Docker container and executing the appropriate flowgraph within the container and generally monitoring that all network links over the VCC VPN and to the WFF-GS were working properly. For a specific pass the VCC team at the respective MOC and the VGSN operator would all get online approximately 15 minutes before a pass.

Initial checks for this included ensuring the appropriate endpoints were entered into the MOC C2 software. This was tested by using a loop back technique to connect the uplink path to the downlink path in place of the TCP connection to the WFF-GS.

Once all systems were confirmed to function properly from the MOC through the VPN and the Wallops Interface Docker container, the VGSN NOC Operator would then dial into the WFF antenna operations line and bridge the connection with the Zoom session. This placed all operators on a single voice loop from the MOC, through the VGSN NOC, all the way to the WFF antenna operators. Approximately 1 minute prior to the start of the pass the WFF-GS systems would connect to the VGSN Wallops Interface, and test packets would be sent from the MOC to the WFF-GS systems. The WFF-GS antenna operators would confirm center frequencies and power levels and then the pass would begin. Throughout the pass, the operators would give callouts at their respective facilities about what they were doing or what was being observed. At the conclusion of the pass, the WFF-GS team would provide a pass report and the VCC/VGSN team would disconnect from the bridge line. The VCC team would then discuss next steps and conclude the pass activity.

Often times, while using the WFF-GS, the VGSN operator would also execute operations in parallel with the VT tracking station as part of the troubleshooting process. This would include operations such as sending Lithium-II radio ‘ping’ commands from the VTGS and listening with the WFF system for responses. This technique was also tried in the opposite direction, with MOC commands uplinked through the WFF-GS and monitoring the downlink with the VTGS. This ‘pitch/catch’ technique was used due to the Lithium-II’s rapid response to a ping command (approximately one millisecond), often times too fast for the TX/RX switching circuitry in the respective tracking stations (tens to hundreds of milliseconds).

The increased Effective Isotropically Radiated Power (EIRP) of the WFF-GS on TX as well as the increased Figure of Merit (G/T) on receive (approximately 20 dB better than the VTGS) were exercised in multiple ways over the course of about a two month window in an attempt to establish communications with the VCC spacecraft. Unfortunately, this was not achieved, with Aeternitas and Ceres remaining silent throughout all attempts. It has not yet been confirmed that this is specifically due to a communications issue (either on the ground or onboard the spacecraft) and troubleshooting is ongoing at the time of this publication. This troubleshooting has also lead to the discovery of some significant issues related to the Lithium-II, requirements for ground based signal processing when using a Lithium-II on orbit, and lessons learned related to integration of the radio with the flight computer and flight software, perhaps appropriate for follow on publications once the troubleshooting work is complete, but currently considered out of scope for the purposes of this paper. However, despite these issues, the Virginia Ground Station Network itself operated as expected and this is an important accomplishment, both for the VCC team, and for the Commonwealth.

REGULATORY

Background

As discussed in the VCC mission overview, the original plan for the constellation was operation in the Amateur Satellite Service between 435-438 MHz, with each satellite operating on a separate frequency. However, plans changed to Experimental (CFR 47, Part 5) operation, in the 401 MHz range. Sterk Solutions Corporation assisted the VCC Mission on FCC licensing and frequency coordination. The individual university teams provided all required technical data to Sterk, working through a single student leader and coordinator. This included data to support the license application, coordination with potentially affected parties, and orbital debris assessment. A separate license for each of the three universities was sought, as the original vision was that each university would operate its own satellite, with its own ground station.

A comprehensive electromagnetic coordination analysis was required, to assure an expectation of no interference with other operations in the 401 MHz range. These operations included NOAA, NASA, and multiple commercial groups. Members from each university made significant contributions to this study, with VT and ODU focusing on the NOAA DCS issue and the UVA team lead handling coordination with the multitude of commercial groups operating in the band. All of this was accomplished with methodology advice from NASA, and coordination by Sterk. With responsive support from FCC, license grants were obtained well before integration and launch of the spacecraft.

Modification for Network Operations

As described previously, the original plan for operations involved each institution operating its own spacecraft and then submitting the collected data to the data warehouse for the science team to access and perform higher level data processing. Therefore, three distinct licenses were obtained, one for each institution. However, after deployment, a number of communications problems presented themselves, both with the ground segment and with the spacecraft. This led to the decision to network the ground stations

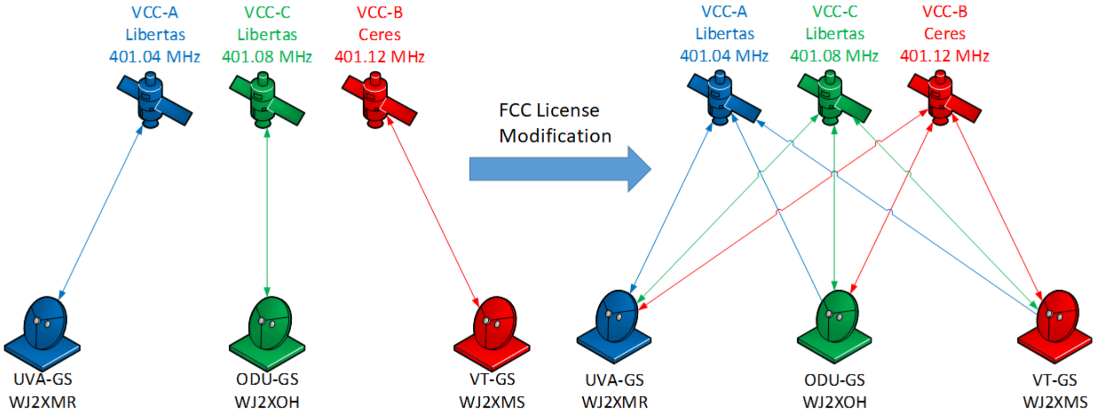

Figure 11. FCC experimental license modification authorizing any VCC ground station to transmit on any frequency assigned to the VCC mission. The satellites' frequencies remain unchanged.

together and to pursue access to the WFF-GS antenna system for more advanced troubleshooting. To evolve from a model of discrete spacecraft – ground station pairs, to a model of "any ground station communicates with any space station", required expansion of the scope of the existing FCC licenses. Special Temporary Authorities, one per satellite, were sought and obtained from the FCC Office of Experimental Technology (OET), to allow any ground station to transmit on any one of the three frequencies of the VCC mission, and to authorize each of the spacecraft to transmit to any tracking stations in the network. Figure 11 depicts the requested modification. Authority to transmit from each satellite to the Wallops Flight Facility station was included, see discussion below.

The decision to apply for experimental STAs rather than experimental licenses was due in part to the urgency to troubleshoot the communications problems with the spacecraft, and to the orbital lifetime of the ODU spacecraft being limited by deployment of its atmospheric drag brake. Because of the success of VGSN operations and the desire to continue to evolve the network, we will seek longer term experimental licenses.

Wallops Flight Facility Licensing

Authority to transmit with the WFF-GS antenna system on the three VCC frequencies was also required. The VCC Team including Sterk, provided the spectrum management team at NASA Wallops, and they were very responsive in obtaining the necessary authority to allow them to support our mission.

CONCLUSIONS & FUTURE WORK

The prototype Virginia Ground Station Network development and operations have been highly successful. Creation of the network led to the accomplishment of significant milestones for the VCC mission, such as the reliable contact of the Libertas spacecraft. Additionally, the development of the network has highlighted an often overlooked, but critical component of a successful smallsat mission, the ground segment and the need for reliable communications. Finally, the network development has demonstrated the overall value of collaboration in academic and research pursuits, an important lesson for young engineers, even if the institutions involved may have rival football teams.

VGSN Improvements

A primary goal of the VGSN is to transition from a prototype demonstration network into a more robust, reliable, and permanent system for smallsat operations and to enable new smallsat related communications research. Given the limited time available to stand up the prototype network, a number of quick decisions were made. A number of key areas are being considered for future implementations. A non-comprehensive discussion of these areas are given in this section to highlight some of the larger design goals.

Security is of utmost concern for the VGSN. This includes both IT security as well as general program security. Participants in the VGSN should be confident that their spacecraft and data are protected and that their overall mission risk posture is better through use of the network, not worse. As mentioned throughout this document, a number of technologies are under review in this area. Another security consideration pertains to restricted access programs in general. VCC was a generally open and educational project. Future missions may have International Trade and Arms Regulation (ITAR), Defense Federal Acquisition Regulation Supplement (DFARS), or have proprietary (commercial) restrictions. Network based solutions to these types of restrictions, if executed appropriately, could represent further value added to new missions.

Another area under review concerns the potential use of cloud solutions for VGSN operations. Existing solutions in this area for spacecraft operations are under review by the team. While some functionality may still need to reside in the specific tracking stations (such as the lower layer operations), some network enabled features may better serve future missions in a cloud deployment. Deployment of the core VGSN network based dockerspaces, as well as consideration for including C2 software in a cloud deployment are being considered for example. This in itself has specific security implications with respect to clear secure boundaries on the network. Cloud solutions in general must also pay close attention with respect to ITAR and DFARS approved solutions.

Increased compatibility with other satellite communications protocols and standards is a major goal of the VGSN. This includes both over the air protocols as well as ground segment protocols, such as the Consultative Committee for Space Data Systems (CCSDS) Space Link Extension (SLE) protocol. Adding this capability will facilitate providing access for VGSN member institutions (and partners) to other ground segment tracking stations should mission objectives require such features or coverage not offered by core participants tracking stations (KSAT coverage of the polar regions for example).

Advanced Smallsat Ground Communications Research

One specific area under consideration is closed loop command routing for real time operations. One of the major feature upgrades for the network would be centralized command routing via a message passing framework as an alternative to direct MOC to tracking station TCP (or other) connections. RabbitMQ is currently the most likely candidate, as it provides a centralized Broker service. All traffic would be routed through the broker, with each message using specific keys to route them to the desired endpoint.

With multiple tracking stations tasked to track a specific spacecraft, as spacecraft responds to uplink commands, the downlink is received by all tracking stations in the footprint. The GNU Radio flowgraphs report Signal to Noise Ratio (SNR) of each received packet, also via the RabbitMQ broker. This SNR metric can be used to update routing tables in a centralized command processor for the mission. As the SNR degrades at one station and increases at another station, the commands from the MOC can be routed automatically to the station with the highest reported SNR as well as the authority to transmit at that station. Multiple metrics could further be used to refine this, such as tracking the dropped packets at each station. This line of research could result in optimized throughputs and overall data downlink for a given mission.

Another area of research has to do with raw IQ sample collection by the network of ground stations. Again consider a single satellite passing over the network. Due to the baselines between the tracking stations, each packet received will arrive at a given tracking station at a slightly different times and slightly different frequencies. These IQ sample streams, either in recorded file form (such as SigMF) or via streaming protocols such as VITA49 or IEEE P1914.3, could then be brought together at a central location for additional signal processing. The earlier stated assumption that certain lower layer functions must reside in the tracking stations may *not* be true strictly speaking, even if they are perhaps more practical with today's technology. Advances in software radio technologies have enabled standardized real time network transport of raw IQ data, such as the VITA49 standard as well as the IEEE P1914.3 standard. Combined with cloud computing resources as discussed earlier, advanced applications for communications in general, but for smallsats in particular, can then be explored. Some of these concepts have been discussed in prior work, and the area remains ripe for future exploration.[10]

Multiple algorithms from various domains of communications and signal processing could be brought to bear and explored for smallsat communications. This includes geolocation techniques, initial orbit estimation and orbit determination, MIMO techniques, digital antenna beamforming and sparse aperture combining, Very Long Baseline Interferometry (VLBI) techniques, and numerous others. As an example, geolocation algorithms such as Time Difference of Arrival (TDOA) and Frequency Difference of Arrival (FDOA) can be applied to geolocate the spacecraft at multiple instances as it passes over the network, and under certain assumptions related to reference frequency and timing qualtiy, initial orbit estimation is then feasible. Additionally, various signal combining techniques, such as maximal ratio signal combining, could be applied to recover packets that were perhaps too weak to be decoded by any one station. These are just a few examples in an area ripe for additional research.

ACKNOWLEDGEMENTS

Virginia Space Grant Consortium (VSGC)

The authors would like to thank the Virginia Space Grant Consortium for its leadership from concept through operations for the Virginia Cubesat Constellation mission. Additionally, the lead author would like to thank the VSGC, specifically its fearless leader and Director, Mary Sandy, for providing funding for dedicated support related to the communications for the VCC mission, ultimately enabling the creation of the

first ground station network dedicated to smallsat operations for the Commonwealth of Virginia.

NASA Wallops Flight Facility

The authors would also like to thank the Wallops Flight Facility team, in multiple departments, for their enthusiastic support of the VCC mission. The authors would also like to extend a special thanks to Steve Bundick for the hours spent setting up, troubleshooting, and validating the signal processing features of the WFF *quantumGND* architecture in conjunction with the VGSN systems, including the time spent with the VCC Team onsite for flatsat testing. The authors would also like to extend a special thanks to the Antenna Operations team, an exceptional collection of engineers who spent a lot of time on the phone helping to brainstorm solutions for the VCC communications issues and providing excellent mentorship to the VCC team. Their experience in smallsat communications is greatly appreciated and invaluable to the smallsat community.

NASA Johnson Space Center

The authors would like to thank the spectrum managers at Johnson Space Center for their timely and straight forward coordination process. The authors look forward to continued work with JSC for the lifetime of the VCC mission.

Virginia Smallsat Data Consortium (VSDC)

The lead author would like to thank the newly formed Virginia Smallsat Data Consortium (VSDC). The VSDC organization is a collaboration between the VT Hume Center and ODU. VSDC is funded by the State of Virginia Research Investment Fund (VRIF) with a general goal of stimulating smallsat work in the commonwealth for academic, research, and economic purposes. Funding provided by the VSDC enabled continued VGSN operations for the VCC mission as well as provided conceptual design and planning time for expansion of the VGSN to transition from a prototype system into a more robust and reliable smallsat asset for the Commonwealth of Virginia.

REFERENCES

1. H. Zimmerman, "OSI Reference Model - The ISO Model of Architecture for Open Systems Interconnection"," *IEEE Transactions on Communications*, vol. 28, no. 4, pp. 425-432, April 1980.
2. "AX.25 Link Access Protocol for Amateur Packet Radio - Version 2.2," July 1998. [Online]. Available: https://www.tapr.org/pdf/AX25.2.2.pdf. [Accessed 1 June 2020].
3. Z. Leffke, "Github, Arduino Code for VHF/UHF PA Control Board," VTGS, Hume Center, June 2019. [Online]. Available: https://github.com/vt-gs/pa_vhf_uhf. [Accessed June 2020].
4. Z. Leffke, "Github, EagleCAD PA VHF/UHF Control Board," VTGS, Hume Center, June 2019. [Online]. Available: https://github.com/vt-gs/eagle_pa_vhf_uhf. [Accessed June 2020].
5. Z. Leffke, "Github, VTGS RF Front End Control Daemon," VTGS, Hume Center, June 2019. [Online]. Available: https://github.com/vt-gs/rffe_daemon. [Accessed June 2020].
6. Z. Leffke, "Github, RF Front End Control GNU Radio OOTM," VTGS, Hume Center, June 2019. [Online]. Available: https://github.com/zleffke/gr-rffe_ctl. [Accessed June 2020].
7. "COSMOS - The User Interface for Command and Control of Embedded Systems," Ball Aerospace Inc., [Online]. Available: https://cosmosrb.com/. [Accessed June 2020].
8. P. Karn, M. Chepponis, "The KISS TNC: A simple Host-to-TNC communications protocol," 1997. [Online]. Available: http://www.ax25.net/kiss.aspx. [Accessed 5 June 2020].
9. D. Estévez, "Github, gr-satellites," June 2020. [Online]. Available: https://github.com/daniestevez/gr-satellites. [Accessed 5 June 2020].
10. S. Hitefield, Z. Leffke, M. Fowler and R. W. McGwier, "System overview of the Virginia Tech Ground Station," 2016 IEEE Aerospace Conference, Big Sky, MT, 2016, pp. 1-13, doi: 10.1109/AERO.2016.7500716.