# Onboard Science Processing on a Microsatellite with Limited Resources


**Scott A. Miller**
**Southwest Research Institute®**
**6220 Culebra Road**
**San Antonio, TX 78238**
**210-522-4249**
**scott.miller@swri.org**

**Ronnie L. Killough**
**Southwest Research Institute**
**6220 Culebra Road**
**San Antonio, TX 78238**
**210-522-3616**
**ronnie.killough@swri.org**

**Steve W. Cook**
**Southwest Research Institute**
**6220 Culebra Road**
**San Antonio, TX 78238**
**210-522-6322**
**steve.cook@swri.org**

**Christopher S. Ruf**
**The University of Michigan**
**2455 Hayward St.**
**Ann Arbor, MI 48109**
**734-764-6561**
**cruf@umich.edu**



*Abstract*—**The National Aeronautics and Space Administration (NASA) Cyclone Global Navigation Satellite System (CYGNSS) mission aims to understand the coupling between ocean surface properties, moist atmospheric thermodynamics, radiation, and convective dynamics in the inner core of tropical cyclones (TCs). The mission is comprised of eight microsatellites (µSats) in low-earth orbit (LEO) at an inclination of 35 degrees.**

**The mission faces unique challenges in hardware and software design to satisfy mission restrictions: the small size of the µSats implies a low power budget for telemetry downlink bandwidth, science data processing, and Attitude Determination and Control (ADC) processing. Additionally, the LEO path of each µSat implies shorter ground passes that limit the downlink time. To accommodate these constraints, creative and efficient hardware and software designs are required.**

**This paper discusses how downlink and power limitations will be accommodated by the design of the CYGNSS Spacecraft hardware and software.**

**Each µSat contains a Delay Doppler Mapping Instrument (DDMI) which receives direct signals from Global Positioning System (GPS) satellites as well as GPS signals scattered by the ocean surface. The direct signals pinpoint the location of the µSat, while the scattered signals respond to ocean surface roughness from which wind speed is derived. The science data derived from these direct and scattered GPS signals, Delay Doppler Maps (DDMs), will be used to provide information about and improve forecasts of the intensity of TCs.**


## TABLE OF CONTENTS



## 1. INTRODUCTION

The design of the hardware and software for the Cyclone Global Navigation Satellite System (CYGNSS) mission faces unique challenges in satisfying mission restrictions: the small size of each microsatellite (µSat) implies a low power budget for telemetry downlink bandwidth, science data processing, and µSat Attitude Control and Determination (ADC). Additionally, the Low Earth Orbit (LEO) path of each µSat implies shorter ground passes that limit the downlink time. To accommodate these constraints, creative and efficient hardware and software designs are required.

This paper discusses how downlink and power limitations will be accommodated. The paper describes (1) the science data being produced by the mission, (2) the drivers of constraints for downlinking the science data, (3) what can be done about the bandwidth constraint, and (4) what can be done about the Central Processing Unit (CPU) power constraint.

## 2. SCIENCE DATA PRODUCTION

CYGNSS accomplishes its science goal using a Delay Doppler Mapping Instrument (DDMI) on each µSat [1, 2, 3]. In total, the DDMI consists of the Delay Mapping Receiver (DMR) electronics unit, a Right Hand Circularly Polarized (RHCP) antenna on the zenith side of the spacecraft for receiving the direct Global Positioning System (GPS) signal, and two Left Hand Circularly Polarized (LHCP) antennae on the nadir side of the spacecraft for receiving the scattered GPS signal from the surface. Figure 1 illustrates the GPS signal transmitted by a GPS satellite and its reception by the antennae on a CYGNSS µSat.


978-1-4799-1622-1/14/$31.00 ©2014 IEEE

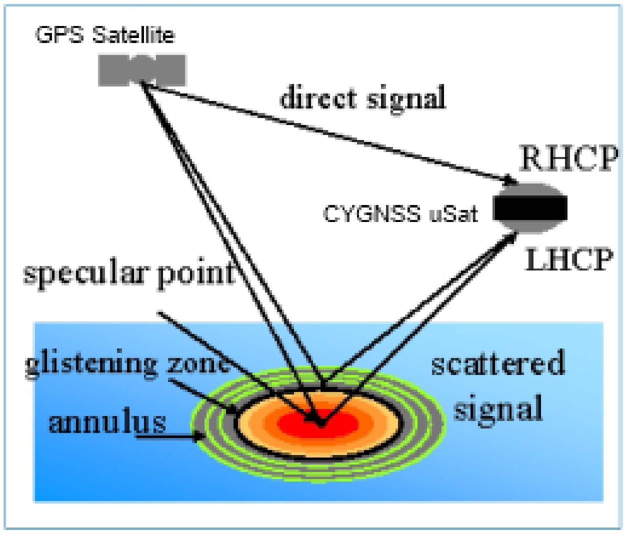


**Figure 1 - GPS Signal Propagation and Scattering**

The direct GPS signal provides a coherent reference for the coded GPS transmit signal and is used to pinpoint the location of the µSat. The scattered signal contains detailed information about the ocean roughness, from which local wind speed can be derived. The location on the ocean's surface where the angle of incidence is equal to the angle of reflection of the transmitted GPS signal is referred to as the *specular point*, and will correspond to a sample of highest power recorded by the DDMI. The annulus region surrounding the specular point that produces the majority of the reflected GPS signal received at the µSat antennae is referred to as the *glistening zone*.

Onboard the µSat, the DDMI takes the received reflected GPS signal and generates a type of scattering image referred to as a Delay Doppler Map (DDM), shown in Figure 2. The DDM is the primary science data output of the DDMI.

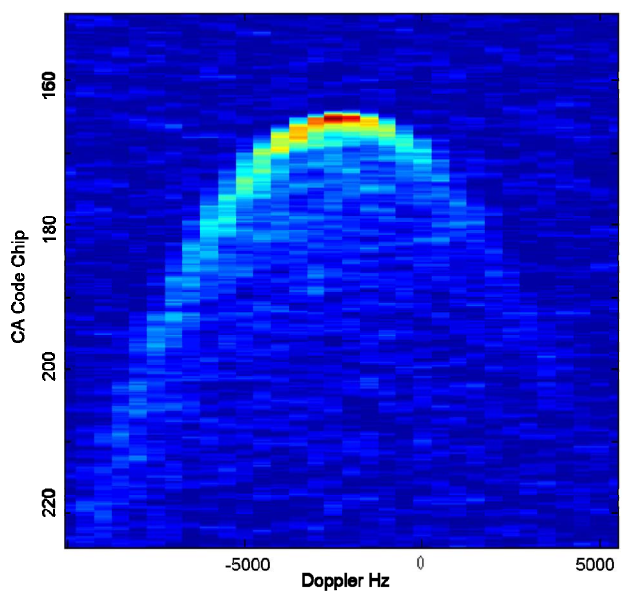


**Figure 2 - Delay Doppler Map**

The coordinates of a DDM are Doppler shift and time delay offset. Each pixel of the DDM is obtained by cross-correlation of the received signal with a locally generated replica time delay and Doppler shift. An open-loop tracking algorithm allows each DDM to be processed by predicting the position of the specular reflection point from the known positions of the receiver and GPS transmitter. Available hardware resources allow generation of four simultaneous DDMs per second.

## 3. Constraints

Required global coverage for the CYGNSS mission is provided by eight µSats dispersed in a 500 km, 35° circular orbit. In this LEO, each µSat in the CYGNSS constellation will only make ground contact with a Universal Space Network (USN) station for an average of 420 seconds (7 minutes) every two days. See Figure 3 for an illustration of the CYGNSS ground coverage.

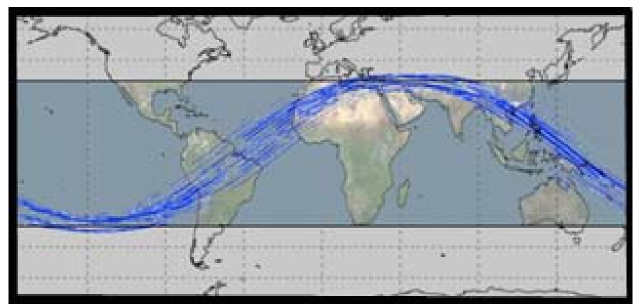

**Figure 3 - Ground Coverage of CYGNSS Constellation over 90 min (one orbit)**

Because each observatory in the CYGNSS constellation is a µSat, limited space is available for solar array placement (Figure 4). The limited space available for solar arrays equates to a limited amount of power for the avionics.

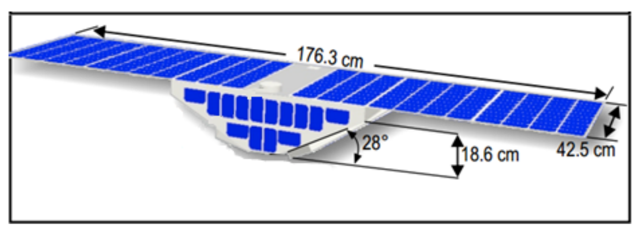


**Figure 4 - CYGNSS µSat Dimensions**

A limited amount of power for the avionics equates to a limited amount of Radio Frequency (RF) energy able to be produced by the transceiver for data transmission. It is currently anticipated that the transceiver for the CYGNSS mission will be capable of downlinking data at a maximum rate of 3.5 Mbps. Total downlink throughput of the CYGNSS constellation is also driven by the LEO of the constellation, meaning short ground contact times.

In addition to a limited amount of power available for transmitting data via RF, a limited amount is available for

978-1-4799-1622-1/14/$31.00 ©2014 IEEE

the CPU running the Flight Software (FSW). It is currently anticipated the CPU will be clocked at 12.5 MHz of the available 100 MHz the processor is capable of running.

## 4. ADDRESSING BANDWIDTH CONSTRAINTS

As with any mission, one goal of the avionics system is to downlink the maximum number of bits to the ground in the least amount of time possible. This goal is especially important on CYGNSS, with the limited downlink rate of 3.5 Mbps and the ground contact time of 420 seconds per pass.

Given the ground contact time and downlink rate, calculations show a maximum of 184 megabytes of data can be downlinked during one ground pass. This 184 megabyte set must include the Spacecraft Housekeeping (H/K) telemetry and all science data collected and stored since the last ground contact.

To address the constraint of the limited ground contact time and downlink rate, several design concepts are being incorporated into the FSW of the µSat:

1. Data Reduction and Compression of the Science Data
2. Compression of the H/K and Engineering Data
3. Minimizing the H/K and Engineering Data

Each of these solutions for addressing downlink constraints are discussed in the following paragraphs.

*Data Reduction and Compression of the Science Data*

As discussed earlier, the DDMI primarily produces one type of Science Data: DDMs. Each DDM in its raw form consists of 2,560 pixels (128 delay rows by 20 Doppler columns) with 32 bits per pixel. The DDMI will transfer four DDMs per second, which translates to a rate of 0.3125 Mbps. The DDMs will be received by the onboard Spacecraft FSW for further processing and eventual downlink. If DDM production continued at a rate of 0.3125 Mbps for two days, 6.592 gigabytes of data would need to be stored and downlinked. This 6.592 gigabytes of data equates to 3668% more than the maximum of 184 megabytes that can be downlinked in one ground pass.

To reduce the size of the DDM dataset, a data reduction and compression algorithm produced by the CYGNSS science team will be incorporated into the Spacecraft FSW. For each raw DDM, the algorithm identifies a subset of pixels that contain mostly reflected signal (i.e. glistening zone) and identifies pixels that contain mostly noise (i.e. noise floor). From these subsets, the algorithm generates two datasets for eventual downlink that are further compressed by reducing the number of bits represented by each pixel.

The glistening zone dataset of a DDM contains the pixels that have captured the majority of the reflected GPS signal. To identify the glistening zone pixels, the algorithm first creates a low pass filtered DDM from the raw DDM and then finds the maximum valued pixel within the filtered DDM. This maximum valued pixel represents the estimated specular point of the glistening zone. This filtered version of the DDM is used to reduce sensitivity to noise in finding the computed specular point. Once the specular point pixel is found, a predetermined number of pixels around the computed specular point, which represent the glistening zone, are selected, as shown in Figure 5. The glistening zone pixels are further compressed by selecting the nine most significant bits that represent scattered signal for each glistening zone pixel before creating the output glistening zone dataset for downlink.

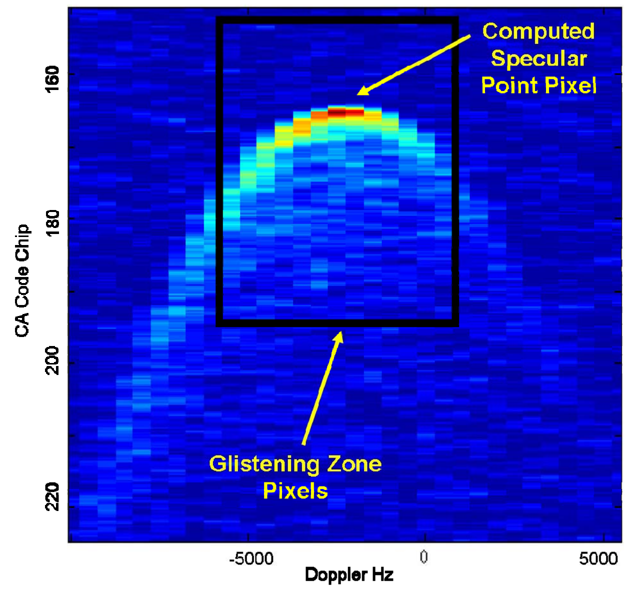


**Figure 5 - DDM with Glistening Zone Pixels Identified**

To estimate the system noise contribution to ocean scattered signal within the glistening zone, the noise floor of each DDM must also be sampled. This is done by selecting a predetermined number of rows (e.g. four) above the computed specular point pixel as shown in Figure 6, summing the pixel values in each row into one value per row, and for each summed row value selecting the 12 most significant bits that represent noise before creating the output noise floor dataset for downlink. Each individual row is downlinked to verify the absence of ocean scattered signal as the noise rows approach the specular point.

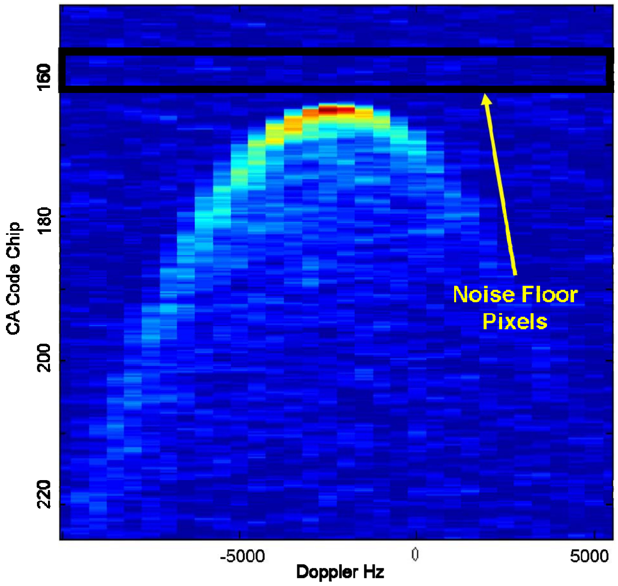


**Figure 6 - DDM with Noise Floor Rows Identified**

After the compressed glistening zone pixel dataset and the compressed noise floor dataset are generated from each raw DDM, they are packetized and stored in onboard flash memory for eventual downlink. Since both the glistening zone and noise floor datasets are bit sample decimated, the compression algorithm is considered lossy. It is currently anticipated that the DDM compression and decimation algorithm will reduce the size of the DDM science data by a ratio of 47:1.

Although a full DDM is not captured for later downlink when using this algorithm, the compressed dataset produced is more than adequate to meet the spatial resolution and sampling requirements of the mission. A compressed DDM restricts the region of the ocean scattering to the portion within 50 km of the specular point, and four DDMs will be gathered and compressed at a rate of 1 Hz.

If necessary for diagnostic purposes, the mission software will support a mode that turns off the compression and stores full DDMs for later downlink.

*Compression of the H/K and Engineering Data*

To reduce the size of the Spacecraft H/K and Science Meta Data packets, a standard software compression library [4] may be used to further reduce the amount of bits to eventually downlink. Given the anticipated stream of data representing nominal H/K and Science Meta Data and the compression library chosen, a compression ratio of 3:1 is expected.

While from a bandwidth perspective it is advantageous to have the data as compressed as possible, doing so adds complexity to ground processing software and also reduces the onboard Spacecraft CPU utilization margin. Initial experimentation has shown that compression may take up to 0.038 milliseconds of CPU time per byte compressed.

*Minimizing the H/K and Engineering Data*

Additionally, to make the most of the limited bandwidth, efforts are being made in the design phase of CYGNSS to minimize the size and rate of real-time Spacecraft H/K packets produced that will inform the Missions Operations Center of the spacecraft health and status. Less bandwidth used by real-time telemetry allows for the majority of the data downlink bandwidth to be used for science data.

To accommodate for the reduced amount of H/K that will nominally be recorded to onboard Flash memory, a wrap-around First in First out (FIFO) buffer may be used to record a history of critical spacecraft telemetry. This FIFO buffer may be downlinked in the event of an anomaly.

## 5. ADDRESSING POWER CONSTRAINTS

As discussed earlier, the size of the μSat and the corresponding space for Solar Arrays equates to less power available for the CPU. Although the CPU can be clocked up to a speed of 100 MHz [5], it is planned to be clocked at 12.5 MHz to conserve power.

To address the constraint of the processor being clocked at 12.5% of its maximum clock speed, several design concepts will be incorporated into the design of the hardware and software of the μSat:

1. Multi-processing capabilities of the dual-core processor will be used for concurrently running two sets of CPU instructions.
2. Science data collection and downlink operations within the FSW will be mutually exclusive.
3. Ability to change the CPU speed will be possible via a Field Programmable Gate Array (FPGA) during Integration and Test (I&T) operations.

Each of the above solutions to addressing the power constraint of the CPU is elaborated on in the following paragraphs.

*Multiprocessing and Dual-Core Processor*

To use the multi-processing capability of the dual-core processor, the FSW will consist of two Operating System (OS) executables running in an Asymmetric Multiprocessing (AMP) configuration [6]. Each OS executable running one core of the CPU will be responsible for a different set of tasks. For example, one OS executable may be responsible for spacecraft health monitoring and command/telemetry processing, and the other OS executable may be responsible for running the ADC algorithm and science data processing. Each OS executable will run on one core of the CPU. Due to the CPU cores sharing resources (Interrupt Controller and On-Chip Peripheral Interfaces), it has been found that the amount of power needed to run in a dual-core configuration is 80% of what it would require to run in a single-core configuration at twice the dual-core clock frequency. See Figure 7.

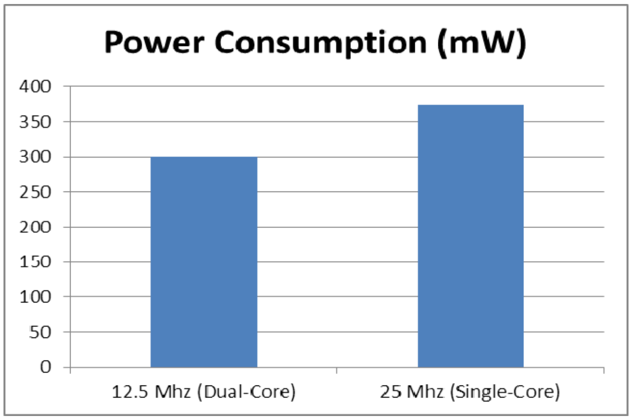


**Figure 7 - Dual-Core vs. Single-Core**

While the use of a dual-core does save power, it comes at the price of complexity. Since the chosen OS does not support Symmetric Multiprocessing (SMP), multiple copies of the OS must be stored in volatile and non-volatile memories for one version of the FSW. See Figure 8 for the anticipated partitioning of the onboard non-volatile Magnetoresistive Random Access Memory (MRAM).

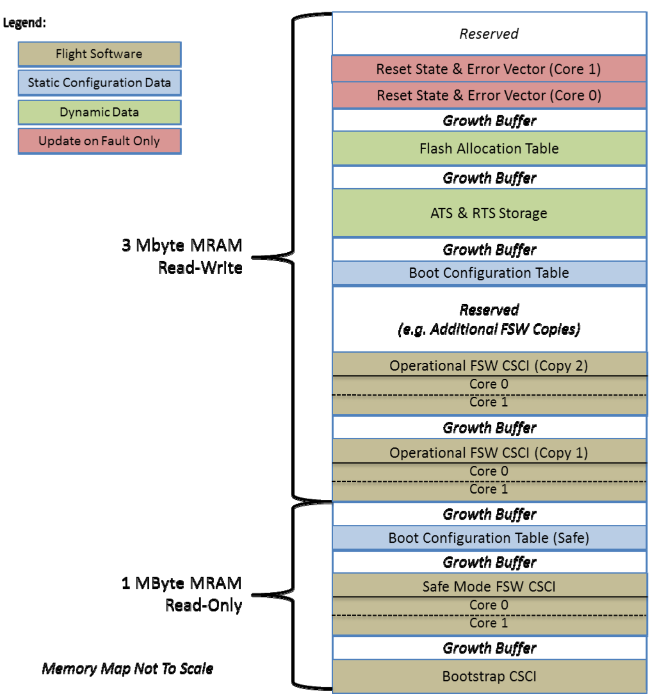


**Figure 8 - Memory Map**

*Mutual Exclusion of Science Collection and Downlink*

While most missions would require that science data collection executes while a downlink is taking place, the relatively short time of the downlink (420 seconds every two days) makes it reasonable to design the system such that science data collection and a downlink operation be mutually exclusive. This will prevent the µSat FSW from needing to accommodate downlink operations while simultaneously devoting CPU cycles to processing and storing science data.

*CPU Clock Speed Modification*

While the current design allows for an appropriate level of margin in the processing capability, additional margin could be attained by changing the clock speed of the processor. However, this increase of processor speed comes at a cost to the power margin. Assuming the processor runs both cores, each 1 Mhz increase in clock speed will increase the power consumption by 24 mW [5].

The onboard FPGA is designed such that clock dividers and multipliers may be used to change the CPU clock speed. This ability will allow the possibility of increasing the speed of the CPU during I&T operations if it is found that the FSW needs to handle more processing than anticipated.

## 6. CONCLUSIONS

As the miniaturization of electronics continues, more technologies become available making µSats and other small-size commercial applications possible. The Size, Weight and Power (SWaP) savings in a µSat or a miniaturized commercial product are large and beneficial, but come at a price of increased engineering complexity and constraints.

On CYGNSS, the smaller size of the µSat compared to a conventional science satellite equates to less power than normally available. Less power equates to less margin available for downlink bandwidth and CPU speed, which creates design challenges to be overcome by those responsible for the design and implementation of the hardware and software of the µSat. This paper offers solutions to common mission constraints for µSat avionics architectures.

978-1-4799-1622-1/14/$31.00 ©2014 IEEE

## REFERENCES


[1] R. Rose, C. Ruf, D. Rose, M. Brummitt, A. Ridley. The CYGNSS Flight Segment; A Major NASA Science Mission Enabled by Micro-Satellite Technology. IEEE Aerospace Conference, Big Sky, Montana, 2013.

[2] C. Ruf, S. Gleason, Z. Jelenak, S. Katzberg, A. Ridley, R. Rose, J. Scherrer, V. Zavorotny. The NASA EV-2 Cyclone Global Navigation Satellite System (CYGNSS) Mission. IEEE Aerospace Conference, Big Sky, Montana, 2013.

[3] M. Unwin, P Jales, P. Blunt, S. Duncan, M. Brummitt, C. Ruf. The SGR-ReSI and its Application for GNSS Reflectometry on the NASA EV-2 CYGNSS Mission. IEEE Aerospace Conference, Big Sky, Montana, 2013.

[4] P. Deutsch and J. Gailly. ZLIB Compressed Data Format Specification version 3.3. Request for Comments (RFC) 1950.

[5] Dual-Core LEON3-FT SPARC V8 Processor GR712RC Data Sheet. Aeroflex Gaisler. February 2013, Issue 2.0.

[6] D. Edell. Multi-Core Processing in Flight Software. Flight Software Workshop 2012. San Antonio, Texas, 2013.


## BIOGRAPHIES

*Scott Miller is a Senior Research Engineer in the High Reliability Software Section at Southwest Research Institute. He has eight years of experience in the areas of embedded systems, data analysis, and computer networking. He has worked as a developer and System Engineer for A-10 military jet avionics and for the Instrument Suite (IS) for the NASA Magnetospheric Multiscale (MMS) Mission. He has a B.S. in Computer Engineering from Texas A&M University. He currently serves as the Project Manager for the CYGNSS Spacecraft Flight Software.*

*Ronnie Killough is Director of the Communications and Embedded Systems Department at Southwest Research Institute (SwRI). He oversees several areas of technology including network-centric systems, tactical communications, cyber security, wireless sensors, smart energy systems, and high-reliability software. He is a 1987 graduate of Angelo State University and a 1990 graduate of Texas A&M University. He has bachelors and masters degrees in computer science, with minor concentrations in electrical engineering, accounting, and finance.*

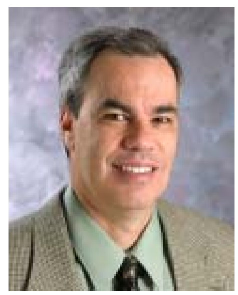

*Steve Cook is a Senior Research Analyst in the High Reliability Software Section at Southwest Research Institute. His background and expertise are in software architecture, parallel and distributed computing, and scientific application development. He has a M.S. in Computer Science from Texas A&M University and a B.S in Geophysical Engineering from the Colorado School of Mines. He is currently serving as a developer for the CYGNSS Spacecraft Flight Software.*

*Chris Ruf received his B.A. degree in physics from Reed College, Portland, OR, and the Ph.D. degree in electrical and computer engineering from the University of Massachusetts, Amherst. He is currently Professor of Atmospheric, Oceanic, and Space Sciences and Director of the Space Physics Research Laboratory at the University of Michigan and Principal Investigator of the NASA EV-2 CYGNSS Mission. He has worked previously at Intel Corporation, Hughes Space and Communication, the NASA Jet Propulsion Laboratory, and Penn State University. He has published in the areas of satellite microwave radiometry and atmospheric, oceanic, land surface, and cryosphere retrieval algorithms. Dr. Ruf is a Fellow of the IEEE, and a member of the American Geophysical Union (AGU), the American Meteorological Society (AMS), and Commission F of the Union Radio Scientifique Internationale. He has served on the editorial boards of AGU Radio Science, the IEEE Transactions on Geoscience and Remote Sensing (TGRS), and the AMS Journal of Atmospheric and Oceanic Technology. He is formerly the Editor-in-Chief of TGRS.*