

Performance Analysis and Electromagnetic Compatibility of a Novel Wideband Radio Frequency Remote Sensing Payload

Jacob von Chorus, Erica Peterson, Nicolaas Handoyo, Nathan Cole, Robert E. Zee
 Space Flight Laboratory, University of Toronto Institute for Aerospace Studies
 4925 Dufferin Street, Toronto, Ontario, Canada, M3H 5T6; +1 (613) 889 - 0781
 jvonchorus@utias-sfl.net

ABSTRACT

The increase in scale, complexity, and sensitivity of small satellite radio frequency payloads presents challenges in spacecraft level environmental performance testing. The Space Flight Laboratory is developing a novel wideband radio frequency payload for use on multiple satellites as part of a distributed remote sensing system. Qualification of this payload at the spacecraft level is complicated by the range of frequencies requiring analysis, the variety of received signal types, and having to qualify the payload on multiple satellites with differing configurations. This paper presents the system level radio frequency performance testing framework developed to efficiently qualify this new payload consistently in different bus configurations. The goals of this framework were to reliably determine payload receiver performance with frequencies ranging from VHF to X-band, evaluate the impacts of electromagnetic interference, and automate the electromagnetic compatibility and performance test processes such that they could be efficiently run on multiple satellites. Ultimately, this framework has yielded the ability to characterize the performance of a complex wideband radio frequency payload, and efficiently scale that characterization to a fleet of spacecraft.

INTRODUCTION

Electromagnetic compatibility (EMC) is a concern for any electronic device, and small satellites are no exception. EMC refers to the ability for multiple electronic devices in a system to operate simultaneously without excessive performance loss due to electromagnetic interference [1]. Electromagnetic interference is an interaction between two devices, where a culprit device conductively or radiatively transfers energy to a victim device. Radio receivers are particularly susceptible to electromagnetic interference, as they are designed to be sensitive to electromagnetic signals.

Small satellites often contain at least one radio receiver for telecommand uplink. Due to the critical function of these receivers, their system-level performance must be validated through EMC testing—the process of proving that a system is free from excessive electromagnetic interference.

Remote sensing or communication payloads in particular may introduce numerous receivers in addition to telecommand devices. Historically, conventional radio architectures limited the number of channels a small satellite could support, so the scale of the resulting EMC testing effort was manageable. New software defined radio (SDR) based payloads increase the number of channels that can be supported, and can result in a system requiring a large system-level EMC testing effort.

Conventional Radios

Conventional radio receivers use dedicated filters, mixers, and demodulators for each radio frequency (RF) channel of interest [2]. In the context of small satellites, this limits the number of channels a mission can receive because supporting numerous channels would conflict with the goal of keeping the small satellite simple and its part count low [3] [4]. System-level EMC testing, as a result, is then limited to the telemetry and command links, and one or two payload receivers. AISSat-1, -2, and -3 along with NorSat-1, -2, and -3, for example, carried Automatic Identification System (AIS) receivers for detecting the location messages broadcasted by maritime vessels [5] [6] [7] [8]. NorSat-2 and -3, the most recent of the listed missions, carried a VHF data exchange radio and a ship navigation radar detector respectively. These missions carried two to three conventional receivers, and an attempt to add more would increase system complexity and volume.

Figure 1 shows an example conventional radio receiver architecture. Each component is a dedicated integrated circuit on a printed circuit board. The antenna is connected to a band pass filter (BPF), followed by a low noise amplifier (LNA), and potentially several gain stages [9]. A mixer downconverts the signal to an intermediate frequency (IF) where it is filtered and amplified before being demodulated. This architecture does not lend itself to adaptability of frequency, bandwidth, and modulation.

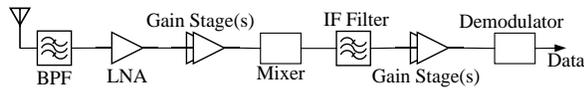


Figure 1: Conventional radio receiver architecture.

Software Defined Radios

There is interest in extending the capability of a single satellite’s payload to support a larger set of channels. For communication systems, heavy use of S-band and X-band frequency spectrum is putting pressure on mission designers to make better use of available bandwidth [10]. Systems that can adapt frequency, bandwidth, and modulation to make best use of link conditions can offer increased performance compared to a fixed system. However, an adaptive system is challenging with conventional radio architectures. For remote sensing payloads, the ability to receive on multiple channels increases the comprehensiveness of an observation. NorSat-3, mentioned previously, is an example with its navigation radar detector as a secondary means of identifying maritime traffic. Detecting additional signals emitted from such targets—communication channels, search radars, weather radars could further augment such a mission’s ability to find its targets. Signals of opportunity are another compelling reason to support multiple receive channels on a payload. Instead of actively transmitting a signal which is reflected off a target, a mission using signals of opportunity detects a reflected signal from a different non-cooperative satellite [11]. Such illuminators could be global navigation satellite system (GNSS) or communication satellites, and they could be used for applications such as soil moisture measurement, ice cover measurement, vessel detection, and wind speed measurement. The advantages are making use of frequencies that would otherwise not be licensed for scientific purposes, and avoiding the need to carry a large, power-intensive transmitter, especially important for small satellites. A mission intended for signals of opportunity benefits from being able to receive on multiple channels because this enables more illuminators to be used, which leads to more observation opportunities.

Software defined radio (SDR) technology, already popular in terrestrial cellular applications, are an alternative to conventional radios gaining traction for space applications because of their flexibility. Figure 2 shows an example SDR architecture [12]. Dedicated hardware, consisting of antennas, bandpass filters (BPF), and low noise amplifiers (LNA), are still typically required in the frontend, especially when supported frequencies differ significantly. However, the mixer which downconverts the signal from a tunable frequency, the lowpass filter which prevents aliasing, and the analog to digital converters (ADC) which

digitize the signals are common to all supported channels. Furthermore, SDRs typically reside on a single integrated circuit making the device especially compact, ideal for small satellites. Demodulation and other processing is performed in the digital domain using software running on field programmable gate arrays (FPGA) or processors. These devices are highly flexible, being able to reconfigure frequency and modulation through software while in-flight. SDRs typically break a signal into in-phase and quadrature (90° phase offset) branches. The resulting complex samples retain both amplitude and phase information for use by the processing software.

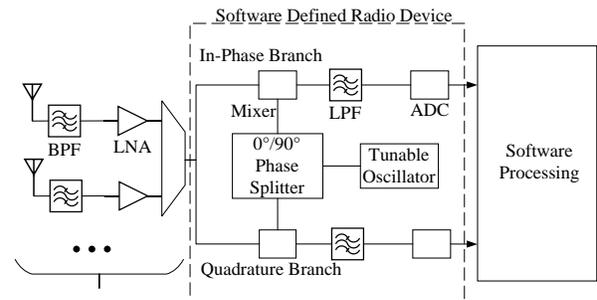


Figure 2: SDR receiver architecture.

Software Defined Radio Payload and EMC Challenges

The Space Flight Laboratory (SFL) is developing a technology demonstration mission to demonstrate a remote sensing payload that uses SDR technology to observe an array of signals transmitted from targets, and signals of opportunity reflected off targets. The in-house designed payload uses a similar architecture to that shown in Figure 2, where a set of frontends for each frequency band is connected to an RF switch, which is connected to a monolithic SDR. Processing is done by both FPGA and processor. In total, the payload can observe on twelve channels. Adding to this scale, the payload will fly on three satellites with different configurations to further enhance observation capabilities. Compared to the missions mentioned above, this jump from one or two payload frequencies to twelve is enabled by the flexibility of SDRs.

As discussed above, system-level EMC is a necessary consideration for any satellite system with radio receivers. The flexibility of SDRs presents a significant challenge to EMC testing because each channel supported represents another channel that is at risk of electromagnetic interference, and which must be tested to ensure compatibility. The large number of channels supported by SFL’s new remote sensing payload meant the scale of EMC testing required for this mission was much larger than previous programs. Additionally, the payload is flying on three satellites with different

configurations, so testing processes needed to be repeatable. Using existing EMC testing methods would have extended the mission development schedule beyond its target. This motivated the creation of a new EMC testing framework to quickly and consistently complete system-level EMC testing with the integrated payload.

EMC Testing Framework

To complete the EMC testing plan for SFL’s upcoming technology demonstration mission, a new EMC testing framework has been developed. Its purpose is to measure RF system performance on each payload receive channel, and to determine the effects of electromagnetic interference at the system level. To accomplish this at scale, hardware reuse and automation of testing and results analysis, are used heavily. This framework and its application to this technology demonstration mission are the subject of this paper. While this work was motivated by a single mission, the EMC testing framework is widely applicable to any mission using SDR payloads to support a large number of receive channels.

BACKGROUND

Electromagnetic interference within a satellite system has traditionally been identified in four categories, where an incompatibility can occur if there is a transfer mechanism from a culprit device to a victim device [1]:

1. **Conducted emissions:** energy generated by a culprit device transferred through electrical conductors
2. **Conducted susceptibility:** energy received by a victim device through electrical conductors
3. **Radiated emissions:** energy generated by a culprit device transferred over the air
4. **Radiated susceptibility:** energy received by a victim device over the air

Traditionally, these are tested individually. For each unit, emissions are measured across a frequency sweep using current probes or a measurement antenna. Susceptibility is measured by introducing conducted or radiated power at some specified level, and checking unit functionality across an identical frequency sweep. These tests, performed on units and subsystem assemblies, ensure that everything in the spacecraft is quiet to some requirement, and is immune to conducted and radiated fields to some requirement. With conservatively set requirements, the spacecraft is known to be safe from electromagnetic interference issues. MIL-STD-461E standardizes this approach to EMC testing [13].

The standardized EMC testing process is inefficient for small satellites. It is time consuming and seeks to show that no possible electromagnetic interference could

occur at any frequency between any two units. However, it ignores whether an interference transfer mechanism exists within the system, and whether receiver performance is actually degraded enough to not meet requirements. On previous missions, SFL has instead tested EMC by measuring the performance of receivers at the system level with all spacecraft hardware powered and operating. This proves that in an operational state, a receiver will meet its performance requirements with all sources of interference present. For a telecommand receiver, as an example common to most missions, this means determining at what received power level the maximum allowable bit error rate occurs, and comparing that to the expected worst-case power on-orbit.

Performance oriented EMC testing is executed with fully integrated satellites so that all potential electromagnetic interference transfer mechanisms are present. This works well for small satellites because if an EMC issue is found, the system is typically simple enough for the source of the problem to be identified and solved. For the single-channel payloads discussed so far, this same testing methodology could be applied by using payload output products to gauge performance at different power levels. This means test setup and post-processing of results is channel dependent. As an example of this type of analysis, Figure 3 shows the performance measured for a dedicated automatic dependent surveillance-broadcast (ADS-B) receiver—a system by which aircraft transmit location and identification messages [14]. This was captured with all spacecraft hardware operating so it represents the worst case for potential electromagnetic interference. It shows decoding performance and allows margin to be directly read off the plot. While this method is effective, there are challenges in scaling to many channels because the test signal and performance metrics are payload specific. For EMC testing of SFL’s new payload, a more generalized version of this performance-based approach was desired.

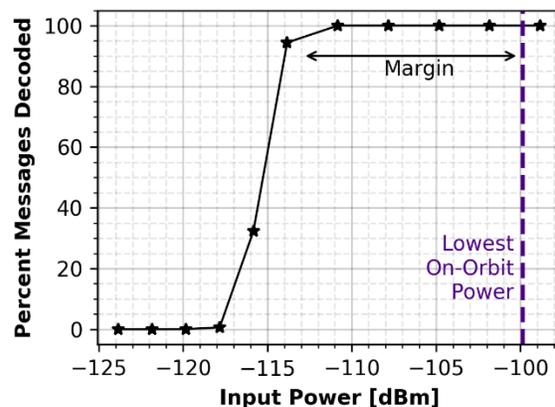


Figure 3: Performance of dedicated ADS-B receiver measured using decoding success percentage.

EMC TESTING FRAMEWORK OVERVIEW

The system-level EMC testing framework is designed around the performance-based approach used on previous missions, where test signals are fed to the payload to evaluate performance. Its goals, meant to scale to many channels, are:

1. Measure performance on radio receive channels so that it can be compared with lowest expected on-orbit RF power.
2. Use common performance metrics for each signal type of interest.
3. Quantify performance degradation due to electromagnetic interference.
4. Automate test process, post-processing, and results reporting.

Signal to noise ratio (SNR) is used as the performance metric for every signal type analyzed. It is readily measured, is valid for any signal, and can be specified as a requirement to ensure a receiver can perform its function. For example, a communication link could use the required carrier to noise ratio before the detector for this purpose. Measuring noise alone to estimate SNR, while simpler to setup, is insufficient as it is a single power measurement that could be influenced by variations in system gain that result from factors such as temperature. SNR, being a ratio of two measurements, is immune to such variations.

Electromagnetic interference is quantified by comparing performance between two test setups and between two spacecraft operational states. The two test setups are:

1. **Cabled testing** – The test signal is transmitted to the payload over a coaxial cable so that only conducted interference is present.
2. **Over the air testing** – The integrated spacecraft is placed in an anechoic chamber where flight antennas are used, so both conducted and radiated interference are present.

Spacecraft operational states are categorized as:

1. **Quiet** - Only always-on and necessary payload hardware are powered. Minimizes interference.
2. **All-up** - All possible spacecraft hardware is powered and operating. Maximizes interference.

Cabled testing in the quiet state provides best case receiver performance since radiated interference is blocked, and conducted interference is minimized. Over the air all-up testing is used to determine whether or not the payload meets performance requirements as this is both the worst case, and closest to on-orbit conditions

where most spacecraft hardware would be operating. Cabled all-up performance is compared with cabled quiet to look for the effects of conducted electromagnetic interference. Over the air quiet provides a baseline for radiated emissions.

Figure 4 shows a hypothetical measurement of conductive and radiative electromagnetic interference. The x-axis shows input power, and the y-axis shows SNR. The highest solid black line is cabled quiet and is the best performing as anticipated. Over the air quiet is slightly worse due to including a baseline amount of radiative interference. Total overall interference is the difference between the over the air all-up and cabled quiet lines. Total conductive interference is the difference between the cabled all-up and quiet lines. Lastly, total radiative interference is the difference between the over the air and cabled all-up lines.

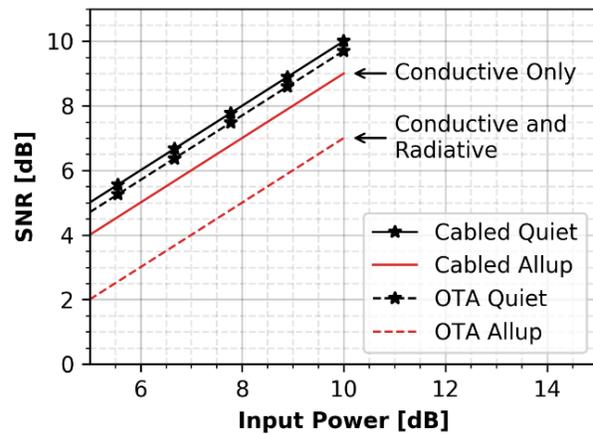


Figure 4: Electromagnetic interference results for illustration.

Test Setups

The two test setups are shown in a basic form in Figure 5. An external SDR is used to generate test signals because it can switch frequency, waveform, and bandwidth without changing test hardware. The external SDR either transmits through an antenna in an anechoic chamber, or directly feeds the spacecraft payload's input using a cable. An anechoic chamber is a faraday cage—a volume completely surrounded by conductive walls—which shields the interior from external electromagnetic interference. Absorptive foam on the interior walls dissipates power radiated by the spacecraft so it does not reflect back, mimicking a space-like environment.

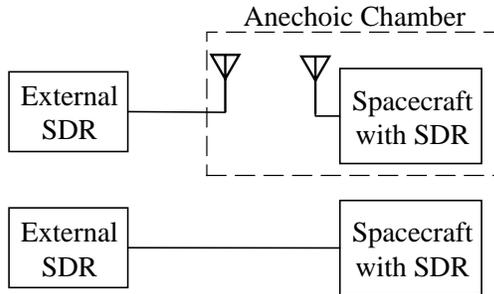


Figure 5: Over the air (top) and cabled (bottom) simplified test setups. Components between external SDR and spacecraft not shown.

Data Capture and Post-Processing

Data is captured by the payload receiver which uses an SDR to save raw complex samples in a capture file. The SDR does not perform any signal-specific processing. Referring to Figure 2, the digital samples from the ADCs are saved. Following a test, a capture is run through post-processing software that determines SNR. In calculating SNR, it also determines signal power and noise power. Lastly, it plots the resulting performance curves similar to Figure 4 so performance margins and electromagnetic interference can be determined.

The next two sections discuss data capture and post-processing in detail.

DATA CAPTURE

The test setup hardware is designed to be flexible so it can remain unchanged between channels. Shown in Figure 6, the setup consists of both RF hardware components and control computers. The RF components (solid lines) consist of an SDR to generate the required frequency and waveform, an electronically controlled variable attenuator to sweep a range of signal input powers, and a spectrum analyzer to measure power from the SDR to determine how much input power the spacecraft receives. There are three computers within the setup. The SDR control computer configures the variable attenuator over a serial connection, configures RF parameters on the SDR, and streams digital waveform samples to the SDR for transmission. The second computer, the payload computer within the spacecraft, controls the on-board SDR and saves captured samples to files in on-board storage. Finally, the main computer orchestrates the entire process. It communicates to the SDR control computer which variable attenuator settings to configure and what signal to transmit, it tells the spacecraft when to capture, it downloads captures from spacecraft storage, and then logs the data to storage in the form of a spreadsheet.

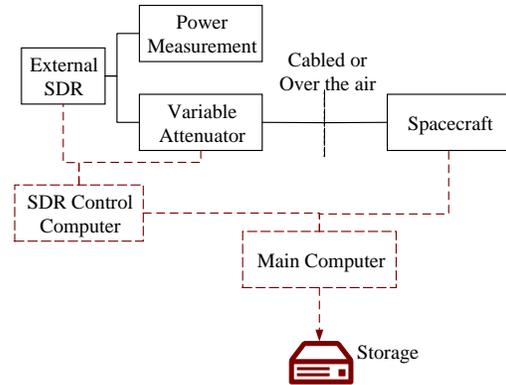


Figure 6: Test setup consisting of RF chain and control computers.

The control scheme enables automation since a power sweep—a set of test points taken at different input powers to establish a trendline in the results—can be configured by a human operator on the main computer, and then the computer handles control of all the other equipment. Prior to this framework, often the variable attenuator, the external SDR, and spacecraft's SDR were manually controlled, resulting in a significant workload for the operator.

Communication between the main computer and SDR control computer is performed over a local-area-network. Communication to the spacecraft can be performed over ethernet or the telemetry and command radios.

Coordinating three computers all running different software is a non-trivial task. A method of flow control is required that is portable between different platforms, whether they be running Python, C++, or Bash scripts. To that end, a standard Linux utility called netcat is used. Netcat establishes basic networking connections where one computer acts as a server listening, and the other as a client making those connections. This is done through scriptable command line calls to this utility. The computer requiring flow control will execute this utility to create a server which will halt program flow as it listens for a connection. A second computer controlling flow will invoke netcat to connect to this server as a client, and then immediately close the connection. The closure results in the first computer continuing on with its script. While very simple in execution, this method has proven effective and easy to implement on both the ground-side computers and the spacecraft's payload computer. Finally, the socket connection offers an opportunity to pass data. For example, the main computer can pass along required frequency to the SDR control computer, and the spacecraft can pass along the name of a newly produced capture file to the main computer for logging.

To assist in automation of the test process, pre-defined test profiles are used. Defined for each channel, these profiles specify frequency, bandwidth, transmit and receive sampling rates, the type of waveform being used, and its associated properties. These parameters tell the testing system how to configure the external SDR and spacecraft payload. Additionally, they are saved as metadata with each captured sweep, so the post-processing software can use them to automatically determine how to process the related captures.

Tests are conducted in the form of a power sweep to produce results like Figure 4. The sweep is run automatically, where the operator specifies which profile to test, and which attenuation settings to use, and then the computer runs a capture for each. Attenuation settings are based on what input power levels the operator wishes to test. To determine this, the insertion loss of the RF chain between the external SDR, and the spacecraft is measured with a vector network analyzer (VNA). Then a spectrum analyzer power measurement is taken with the signal transmitting. The RF chain insertion loss and variable attenuator setting are subtracted from this power measurement to find the payload input power. Typically input power values are based on mission requirements, and so variable attenuation settings to produce desired input powers must be solved for. Rearranging for variable attenuation, the settings used in a test can be found from desired input power values as shown in (1).

$$L_{vari-atten} = P_{sdr} - L_{il} - P_{in} \quad (1)$$

Where $L_{vari-atten}$ is variable attenuation in dB, P_{sdr} is power measured by the spectrum analyzer in dBm, L_{il} is insertion loss of the RF chain in dB, and P_{in} is input power to the payload in dBm.

POST-PROCESSING

Post-processing of capture data was performed to determine receiver performance. This is accomplished by first determining the signal power and noise power present in the captured samples. There are two concepts of signal power that must be differentiated. The first is raw power directly measured from captured samples. This value will always contain the addition of noise power, as noise is always present. Since this measurement includes more than just signal, it is referred to as capture power. The second concept is signal power where the addition of noise has been corrected for. Such a correction is an estimate since noise cannot be undone, only its average influence removed. This is referred to as signal power from this point onwards.

Capture and noise power are calculated in similar ways by performing a power spectral density (PSD) estimate

on raw time domain samples, and integrating the result over the required bandwidth. Capture power is measured from samples taken during test signal transmission, and noise power from samples during which the spacecraft was not receiving anything. A number of methods can be used for the PSD estimate [15]. A modified periodogram is used in this work. This method performs a Fast Fourier Transform (FFT) on the time-domain samples, and then scales the result to units of W/Hz. As opposed to a basic periodogram, the modified periodogram uses a windowing function which gradually tapers the ends of the capture to zero. An FFT assumes an infinitely long capture duration. Anything less than that will, without appropriate windowing, result in spectral leakage from signal peaks to adjacent frequencies. A tapered windowing function suppresses this effect. The other PSD estimators discussed in [15] apply additional processing to make peaks in the spectra more visually distinct. However, since the output of this PSD is integrated in software, this is unnecessary. Equations (2) and (3) show the calculation of capture and noise power respectively, where B is channel bandwidth in Hz, S_{xx} is the modified periodogram of signal containing samples in W/Hz, S_{nn} is the modified periodogram of samples with no signal in W/Hz, and f is frequency in Hz. Since the complex samples are captured after being mixed down to baseband, they are centered on 0 Hz.

$$P_{cap} = \int_{-B/2}^{B/2} S_{xx}(f) df \quad (2)$$

$$P_{noise} = \int_{-B/2}^{B/2} S_{nn}(f) df \quad (3)$$

Signal power is calculated from capture and noise power by assuming noise is uncorrelated with signal. Correlation is typically an issue in terrestrial applications where reflected copies of the signal can arrive at the receiver at different times. That is not the case on-orbit. This means noise power adds incoherently to the received signal, so only scalar addition is required [16]. Equation (4) shows how signal power is determined from capture power. From this, SNR is: P_{signal}/P_{noise} . Power and SNR are presented in dBm and dB respectively, calculated using $10 \log_{10}(\cdot)$.

$$P_{signal} = P_{cap} - P_{noise} \quad (4)$$

Isolation of Signal and Noise

The above methods require that signal-containing samples be isolated from those containing only noise. Rectangular waveforms have proven effective for this purpose. The pulses in the waveform carry the signal, while nothing is transmitted during the troughs between pulses. The pulses can be of constant amplitude for a radar source, but can also contain more complicated waveforms such as spreading codes or message bursts.

Capture power is determined from samples in the pulses, and noise from the trough samples. Lastly, since the waveform is periodic it can be transmitted continuously, where the capture is executed at any time. This is opposed to precisely coordinating the external SDR and payload capture to begin capture at a specific instant, with respect to the test waveform. The former is simpler as it does not depend on precise timing, and it avoids lengthy external SDR startup and shutdown times between data points.

Since captured signal phase is random, this must first be determined. The frequency and duty cycle of the rectangular waveform are known, so a synthetic rectangular waveform with these parameters can be cross-correlated with the capture. The time delay of the cross-correlation peak is the capture phase, it represents

where the synthetic waveform best overlaps the captured pulses. Figure 7 shows a section of a captured radar pulse waveform. The dashed red line is the synthetic waveform. It can be seen that the cross-correlation peak position on the x-axis corresponds with the amount this waveform must be shifted to find the captured pulses. Figure 8 shows a similar result, but with a negative SNR. Despite the signal being 20.5 dB below the noise floor, the cross-correlation method still produces a well-defined correlation peak, so the underlying signal pulses are recoverable with their phase being known precisely. The signal in Figure 8 consisted of a spreading code. Following phase determination, the pulse and trough signals are readily isolated.

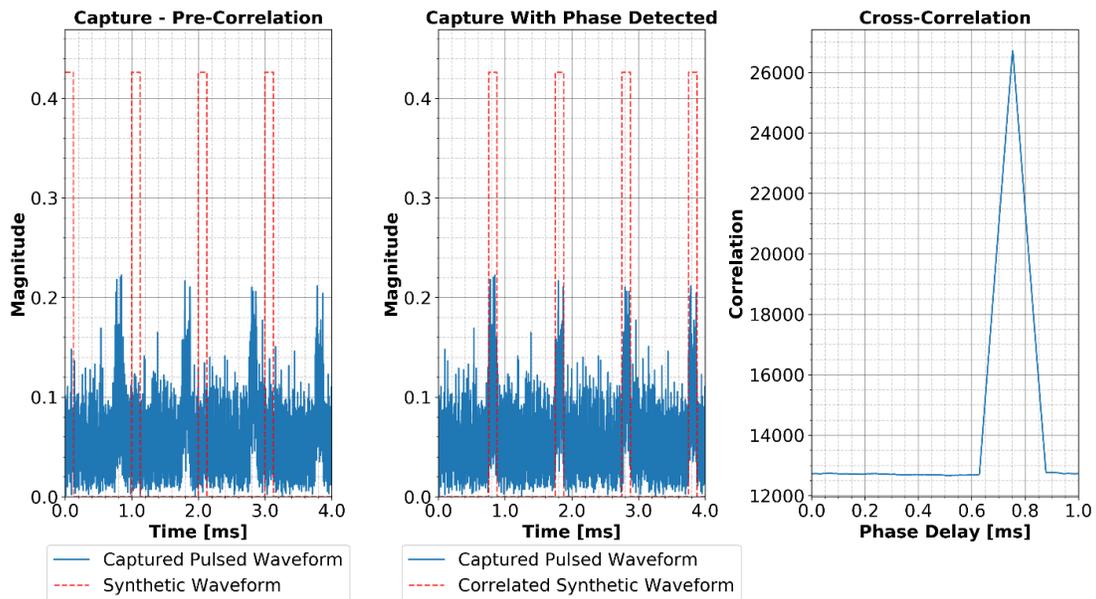


Figure 7: Phase determination of a rectangular waveform through cross-correlation.

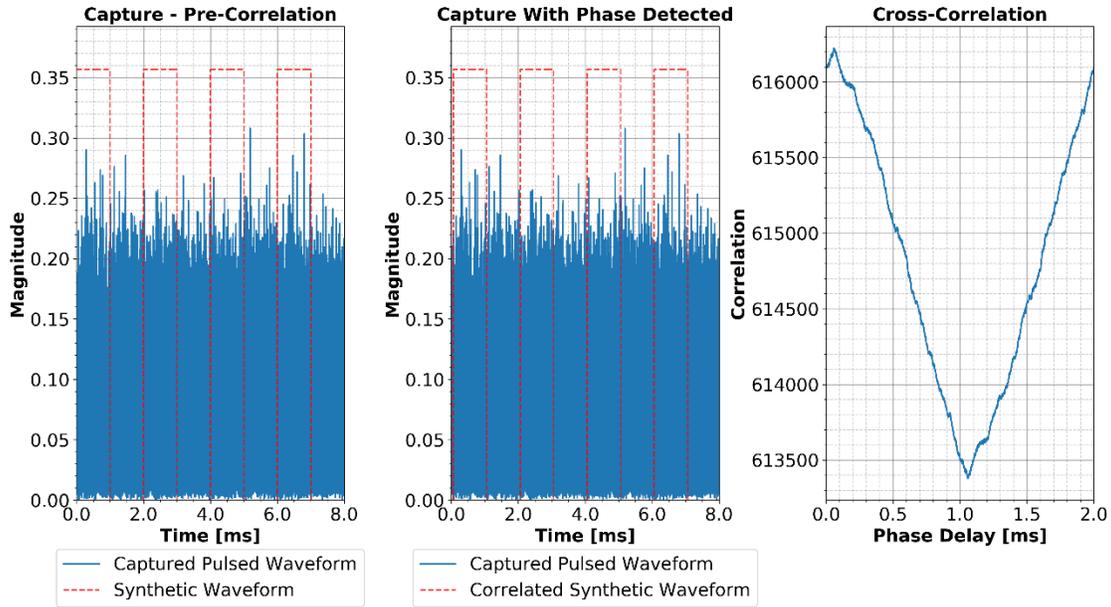


Figure 8: Phase determination with negative SNR (-20.5 dB).

Noise Power Results and Local Oscillator Leakage

Noise power is measured directly as discussed above using isolated noise samples from the troughs between peaks. For this method to work, these samples need to be free of all influence from the transmitting SDR. To avoid transients from the pulses, samples in the troughs near the pulse edges were ignored. However, this proved insufficient during initial testing.

Figure 9 shows the noise spectrum of a pulsed signal. The spectrum has a 1 MHz bandwidth, but only the centre 200 kHz is shown. A large peak is visible in the centre, and its magnitude during testing correlated with the transmitted signal power, implying it is not a spacecraft electromagnetic interference issue. Furthermore, it was always found to be present in the centre regardless of the channel being tested. The peak was the result of the local oscillator in the external SDR leaking into the transmit path, and being amplified by the transmitter, even when it was not supposed to be transmitting any signal. The effect on measured noise power is shown by the solid black line in Figure 10 which increases with input power. Noise is in-band power in the absence of signal, so the correlation with input power further proves that it was caused by a ground support equipment issue. Noise power should have no correlation with input power.

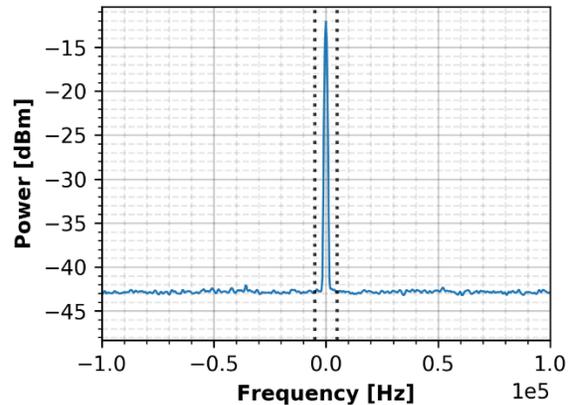


Figure 9: Noise spectrum with local oscillator leakage showing. Dashed vertical lines are area to ignore.

The local oscillator leakage problem is solved in post-processing by ignoring some bandwidth in the centre of the capture spectrum when calculating noise power. In Figure 9, the spectrum between the vertical dashed lines is excluded. Noise power is calculated by integrating the rest of the bandwidth, and then applying a correction factor to account for the missing bandwidth. This correction assumes the average noise power spectral density beyond the excluded zone holds in the centre. Equation (5) shows the updated noise power calculation with the local oscillator cutoff and correction factor included.

$$P_{noise} = \frac{B}{B-B_{lo}} \left(\int_{-B/2}^{-B_{lo}/2} S_{nn}(f) df + \int_{B_{lo}/2}^{B/2} S_{nn}(f) df \right) \quad (5)$$

B is channel bandwidth in Hz, B_{lo} is the bandwidth around the local oscillator to be ignored in Hz, S_{nn} is the noise power spectral density in W/Hz, and f is frequency in Hz. Equation (5) is used in place of (3). Figure 10 shows the noise measurement result with different bandwidth exclusions. A number of values were tested and they all show the influence of the transmitter being completely removed from the noise measurement, as represented by the flatness of the measurement curves. Since the corrected noise power measurement shows no correlation with signal power, it can be concluded that using the troughs between pulses to measure noise is valid with this correction, and the external SDR is not introducing a signal or transients that cannot be easily accounted for. The drawback of removing a larger bandwidth is that legitimate noise from the spacecraft could be ignored. The 100 kHz curve shows this as it is offset from the 2, 10, and 20 kHz curves, implying it excluded captured electromagnetic interference within this frequency range. So, it is desirable to select the lowest local oscillator cutoff bandwidth sufficient to reject the oscillator leakage. On the noise power plot, this produces a horizontal line when plotted against input power. Typically, this was the 2 kHz cutoff.

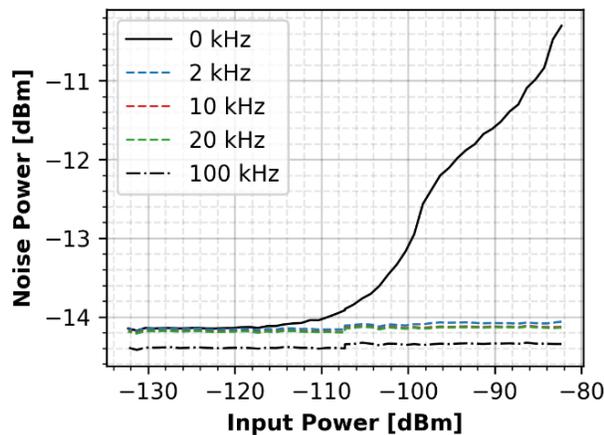


Figure 10: Cutting off bandwidth around local oscillator. Each line corresponds with a different amount of bandwidth about the centre being ignored.

Signal Power and Signal to Noise Ratio

Capture power did not require any special treatment, it is simply calculated with (2). Then it and the corrected noise power are used to calculate signal power using (4). The noise, capture, and signal power results of an S-band test are shown in Figure 11. The noise power, corrected for the local oscillator, is shown as the flat black line. Capture power, the dashed blue line, converges with noise power as input power is reduced. Since capture power includes noise, it will never measure below the noise floor. The sloped purple line is signal power. It is measurable below the noise floor, showing that subtracting noise power from capture power is an effective way to recover signal power. The accuracy for this measurement is shown by signal and capture power converging at high power where noise is less significant, and the signal power linearly decreasing with input power at a slope of 1 dB/1 dB. The dotted line is an ideal 1 dB/1 dB line superimposed over the signal power measurement, showing how precisely it tracks theory. Only 20 dB below the noise floor does it begin to diverge.

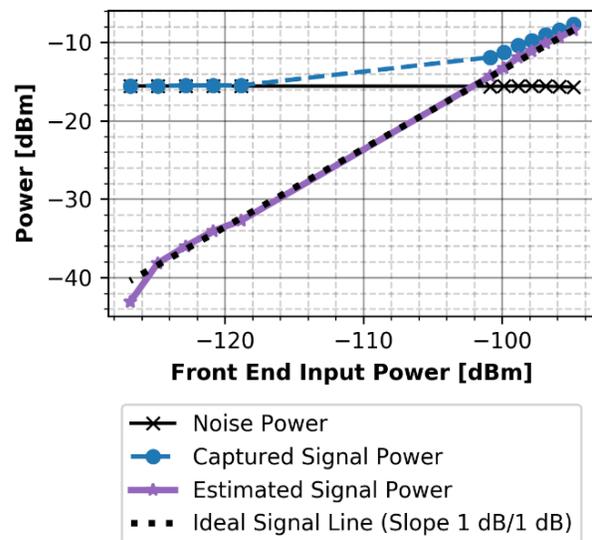


Figure 11: Noise, capture, and signal power measured from an S-band power sweep.

With reliable signal and noise power measurements, SNR is readily calculated, shown in Figure 12. This result can be compared to Figure 11 as the difference between the signal and noise power lines. It should be noted that since signal ideally correlates one-to-one with input power, and noise is constant, the SNR line should have a 1 dB/1 dB slope with respect to input power.

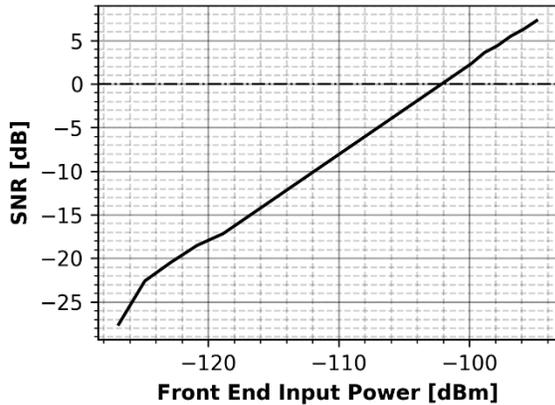


Figure 12: SNR calculated for an S-band power sweep.

Spectrum Plots

In addition to SNR, the post-processing software generates spectrum plots for each capture. These serve as a secondary qualitative means to assess electromagnetic interference, especially useful when investigating performance issues uncovered by the quantitative test. Two spectrum plots are saved per capture: one of the entire capture including signal, and one of only the noise samples. Plots are produced using Welch's method which splits samples into a number of smaller segments, takes the periodogram of each, and then averages the individual results [15]. This reduces random noise in the frequency domain result to allow peaks to be more readily seen, good for qualitative analysis. Figure 13 is an example of a capture spectrum showing the result of a binary phase shift keying (BPSK) modulated signal. This plot shows the signal's main lobe and two sidelobes. It also shows an unexpected peak at approximately -4 MHz that could be investigated. The noise floor rolls off with frequency due to a bandpass filter in place before the receiver whose passband was less than the capture bandwidth.

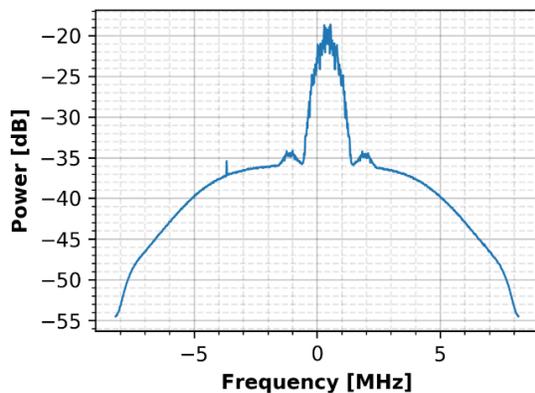


Figure 13: Captured spectrum of a binary phase shift keying (BPSK) modulated signal.

INPUT POWER CALIBRATION

The x-axis of the SNR plots, payload RF front end input power, is determined in one of two ways. The method used by SFL, and alluded to in Figure 6, is to measure the power transmitted by the external SDR, and subtract all of the RF chain losses between this measurement point and the payload RF front end's input. In the cabled setup, this loss consists of coaxial cables, the variable attenuator's fixed insertion loss, and the variable attenuator's setting. Measuring this loss is easily and accurately performable with a network analyzer, prior to connecting the test setup to the payload. So, the input power for cabled testing can be reliably determined.

Over the air testing is more complicated. Loss between the spectrum analyzer and the payload include cables and the variable attenuator like above, but also two antennas, free space, a polarization loss factor, and a pointing loss factor. Accounting for the latter factors is done using antenna gain measurements for the spacecraft and source, a measurement of free space distance to calculate path loss, and an estimate of polarization and pointing loss. This calculation is less precise than the cabled equivalent.

To verify the calculation of input power, a second means of calculating input power is to use the capture itself. The samples from which signal power is determined are measurements of voltage by the ADC, so this measurement really represents the input power to the ADC. Taking the calculated signal power from the capture, and subtracting gain from the RF frontend input to the ADC will result in input power. This gain was measured for each channel during unit level testing of the new payload, and so values were readily available.

Typically, the two measurements of input power agreed closely for the cabled setups, as expected. They did not always agree for the over the air tests. When they varied by more than 1-2 dB, the assumptions of antenna gain, pointing error, and polarization loss were looked at with more scrutiny and adjusted if necessary. In these circumstances, it is necessary to ensure the disparity is due to bad calibration, and not a deficiency in the test signal or the payload receive chain. Improving input power calibration in the over the air test setups is an open item of future work.

OTHER POST-PROCESSING TECHNIQUES

Post-processing of periodic pulses has been presented but other methods of separating signal and noise are available. One used in testing the new payload is meant for digital messages where breaking the signal into pulses would make the message undecodable. The solution is to coordinate the external SDR and the payload using GPS time. The payload would be

commanded to start a capture at some instant in time, and the external SDR would be commanded to start transmitting a few seconds after that instant. Then in post-processing, the noise samples are those first few seconds before transmission, and everything after is signal. This method is advantageous in post-processing as it is simple to separate noise and signal samples, and also allows the captured message to be decoded, but it slows and complicates the test process. The main computer determines the capture start and transmission start times and communicates these to the payload computer and external SDR. The external SDR requires a lengthy startup time before it can begin transmission so these startup times must be far enough in the future to allow for this. Periodic pulses on the other hand allow the signal to be transmitted continuously allowing captures to take place quickly, and for transmit power to be constantly monitored. The timed transmission start method does not allow power to be monitored since the transmission bursts are too brief.

Due to the slower runtime, and inability to monitor transmit power, the periodic pulse method is recommended unless there is a particular need to decode the capture's content outside of this EMC framework.

RESULTS

The SNR outputs from post-processing can be plotted to compare quiet and all-up performance across cabled and over-the-air testing. Figure 14 shows an example SNR plot captured by SFL. Two results are shown, a cabled all-up test and an over-the-air all-up test. Additionally, the minimum SNR requirement and lowest expected on-orbit input power are plotted as horizontal and vertical lines. Critically, this result shows that the worst case, over-the-air all-up, meets the performance requirement with a 5 dB margin in this example. This alone can be sufficient to say the channel passes electromagnetic compatibility at the system level. Nonetheless, the cabled result is colinear with the over-the-air line, meaning no performance is lost due to radiative electromagnetic interference. Although not shown here, previous testing had shown that conductive interference was not present.

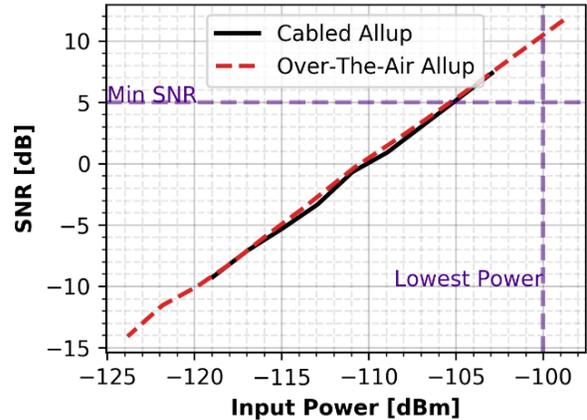


Figure 14: SNR results of a well-performing receiver.

In contrast, Figure 15 shows the performance results from a channel that did see significant radiative interference, noted by the difference between the quiet and all-up results. Only the over-the-air results are shown, the cabled results were colinear with the over-the-air quiet result. A performance drop was discovered when a communication radio was transmitting. When the transmitter was powered off, while otherwise maintaining the all-up state, performance reverted to that of the quiet test as shown by the blue line in the upper-right. This means the performance drop was entirely due to radiative interference from the transmitter. Illustrated on the plot are the SNR requirement and the lowest expected receive power. Without the performance drop, there is a healthy performance margin, while with the transmitter, there is no margin.

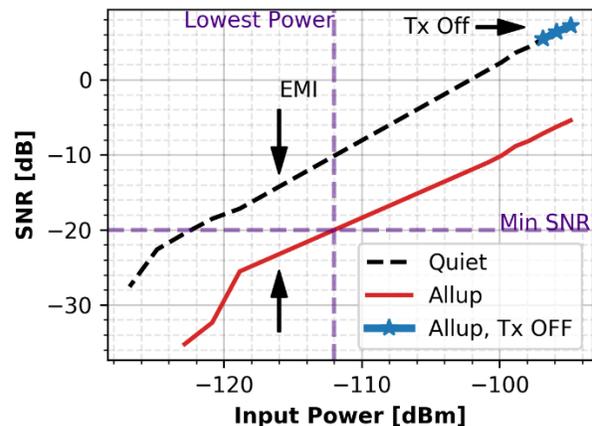


Figure 15: SNR performance of a channel showing electromagnetic interference due to transmitter.

The noise spectrum plots provide a secondary means to confirm the transmitter as the source of interference, shown in Figure 16. The noise spectrum is clean when the transmitter is off, and contains significant ripple

when it is on. The roll-off of each spectrum is due to a bandpass filter in-line with the receiver and is expected.

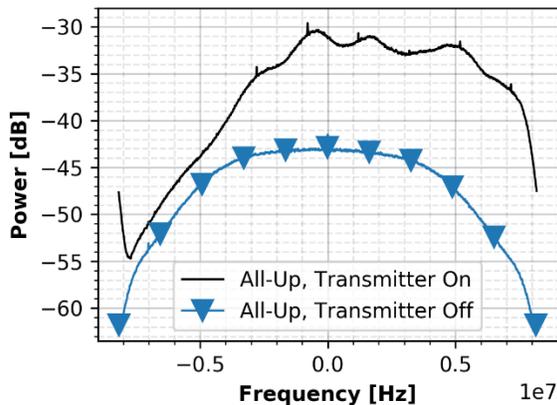


Figure 16: Noise spectrum with and without transmitter.

The solution to this electromagnetic incompatibility is to keep the transmitter off during observations on this channel. Since the concept of operations never has this channel being used within sight of a ground station, this solution does not require any compromises to be made. If the transmitter were to be required to observe on this channel, then extra filtering in the transmit path would likely have been required. This example illustrates how the EMC testing system was applied to a channel on the payload SFL is developing, and how the results can be used.

INTERACTIVE ANALYSIS AND RESULTS GENERATION

Automation of data capture, post-processing, and results display enabled the creation of an interactive analysis feature. Through a single command specifying the test profile and a single variable attenuation setting, an operator is able to start a process that runs a capture, downloads the result, post-processes it, and displays the SNR and spectrum plots. This process skips the data logging that is used to create power sweeps like in Figure 15. The purpose of interactive testing is to help find the source of an interference issue. It completely removes the burden of operating the test setup from the operator, allowing them to focus on setting up the spacecraft's operational state, and then looking at the result from a subsequent capture. An operator can run a number of tests with different sets of spacecraft hardware powered until the source of interference is identified. When this system was used at SFL, interactive analysis proved useful for this type of debug. It was also useful as a quick means to verify a test setup, before running the power sweeps.

EMC TESTING FRAMEWORK WITHIN SFL

Having been applied to the spacecraft forming SFL's upcoming technology demonstration a few times, this EMC testing framework has proven its utility. For example, three sets of cabled tests have been executed. The first took place before the creation of this framework; the test process was largely manual. Data capture took place over several months and did not include any post-processing, so there was no immediate way to verify captures.

The second cabled test took place with a preliminary version of the new EMC testing framework. Automation of data capture resulted in the data capture process being completed in two weeks. Post-processing took an extra week because automatic data handling to pass captures to the post-processing software did not yet exist.

The final iteration of cabled testing with the completed EMC testing system took one week and all post-processing and results display was completed within this time span. Passing of capture data to the post-processing software was fully automated at this stage, so results could be viewed within minutes of completing a test. The months of effort saved in application of the EMC framework have more than paid for the development effort.

CONCLUSION

The flexibility of SDR payloads has led SFL to develop a technology demonstration mission to showcase the use of an SDR as a means of observing multiple RF emissions and signals of opportunity. To support the scale of the system-level EMC testing effort required by this mission, an EMC testing framework has been developed and applied to multiple satellites containing the SDR payload. Results from these tests have informed a number of decisions to mitigate sources of electromagnetic interference. In its final iteration, the framework has reduced the time required to run a set of tests on every payload RF channel from months, to a week, demonstrating its ability to handle EMC testing at scale. Overall, the framework has been able to prove the technology demonstration mission's SDR payload is ready for flight.

More broadly, this framework is applicable to any SDR payload. Since an SDR receiver processes digital samples in software, they can generally be reconfigured to save out raw complex samples as required by this framework. The external data capture test setup is not spacecraft or payload specific, being an SDR, so it is also generally applicable. Lastly, the post-processing software uses SNR as a performance metric which is applicable to any signal, being a ratio of power.

The work presented here will enable future missions that seek to observe or communicate on many channels, and missions that wish to deploy such a payload on multiple satellites.

ACKNOWLEDGEMENTS

The authors would like to thank the staff and students at the Space Flight Laboratory for their mentorship, support, and contributions throughout this project.

REFERENCES

1. T. Clark, M. McCollum, D. Trout and K. Javor, "Marshall Space Flight Center Electromagnetic Compatibility Design and Interference Control (MEDIC) Handbook," NASA Marshall Space Flight Center, Huntsville, AB, 1995.
2. N. Franconi, S. Sabogal and A. George, "A Novel RF Architecture for Simultaneous Communication, Navigation, and Remote Sensing with Software-Defined Radio," in *Small Satellite Conference 2020 Proceedings*, Logan, UT, 2020.
3. P. Stibrany and K. Carroll, "The Microsat Way in Canada," in *ASTRO 2000 - 11th CASI Conference on Astronautics*, Ottawa, 2000.
4. M. Rast, G. Schwehm and E. Attema, "Payload-Mass Trends for Earth-Observation and Space-Exploration Satellites," European Space Agency, Noordwijk, 1999.
5. UTIAS Space Flight Laboratory, "AISSAT-1, -2, and -3," 2023. [Online]. Available: <https://www.utias-sfl.net/aissat-1-2-and-3/>. [Accessed 08 04 2023].
6. UTIAS Space Flight Laboratory, "NorSat-1," 2023. [Online]. Available: <https://www.utias-sfl.net/norsat-1/>. [Accessed 08 04 2023].
7. UTIAS Space Flight Laboratory, "NorSat-2," 2023. [Online]. Available: <https://www.utias-sfl.net/norsat-2/>. [Accessed 08 04 2023].
8. UTIAS Space Flight Laboratory, "NorSat-3," 2023. [Online]. Available: <https://www.utias-sfl.net/norsat-3/>. [Accessed 08 04 2023].
9. Ames Research Center, "9.0 Communications," in *State-of-the-Art of Small Spacecraft Technology*, Moffett Field, National Aeronautics and Space Administration, 2023.
10. M. Maheshwarappa and C. Bridges, "Software Defined Radios for Small Satellites," in *NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, Leicester, 2014.
11. J. L. Garrison, J. R. Piepmeier and R. Shah, "Signals of Opportunity: Enabling New Science Outside of Protected Bands," in *2018 International Conference on Electromagnetics in Advanced Applications (ICEAA)*, Cartagena, Colombia, 2018.
12. T. Collins, R. Getz, D. Pu and A. Wyglinski, *Software-Defined Radio for Engineers*, Norwood, MA: Artech House, 2018.
13. Department of Defense, "Requirements for the Control of Electromagnetic Interference Characteristics of Subsystems and Equipment," 1999.
14. W. Richards, K. O'Brien and D. Miller, "New Air Traffic Surveillance Technology," Boeing, 2010.
15. Mathworks, "Nonparametric Methods," 2023. [Online]. Available: <https://www.mathworks.com/help/signal/ug/nonparametric-methods.html>. [Accessed 03 2023].
16. J. Overvest, F. Jansen, F. Laghezza, F. Uysal and A. Yarovoy, "Uncorrelated Interference in 79 GHz FMCW and PMCW Automotive Radar," in *2019 20th International Radar Symposium (IRS)*, Ulm, Germany, 2019.