

Security Risks of Adaptive Coding and Modulation in Space Systems

Edd Salkield
University of Oxford
edd.salkield@cs.ox.ac.uk

Sebastian Köhler
University of Oxford
sebastian.kohler@cs.ox.ac.uk

Simon Birnbach
University of Oxford
simon.birnbach@cs.ox.ac.uk

Ivan Martinovic
University of Oxford
ivan.martinovic@cs.ox.ac.uk

Abstract—Adaptive Coding and Modulation (ACM) space protocols are currently being implemented and standardised to maximise data throughput on satellite links in the presence of signal fading and radio interference. These systems use an uplink feedback channel, allowing the ground segment to influence the selection of modulation and coding parameters to suit the current channel conditions. However, no current academic work considers the security implications of physical-layer attackers targeting this uplink channel.

In this paper, we introduce the *uplink-assisted attacker* which hijacks the currently unauthenticated ACM feedback mechanisms, using only cheaply available equipment, to select ill-suited communication parameters and prevent the channel from responding to radio interference attacks on the downlink. Our results show the high impact of this attack class: an uplink-assisted noise jammer can cause a 50% frame error rate at 11.5 dB less average power than a noise jammer alone, and up to 16.9 dB if higher modulation and coding parameters are supported. Uplink-assisted spoofing and bandwidth restricting attackers are also shown to be more effective than their counterparts which attack the downlink alone.

Unfortunately, these issues cannot be resolved by cryptographic authentication alone, especially where an attacker can pose as one of multiple terminals reporting channel quality. We therefore conclude with a discussion of countermeasures to prevent and detect this form of attack, and draw out lessons learned for secure ACM design.

I. MOTIVATION

Adaptive Coding and Modulation (ACM) is a promising method for improving the throughput of downlink space communications data links, and is seeing increasingly widespread interest for both New Space systems such as Starlink, and for standardisation and deployment by national space agencies [1]–[4]. Very recent CCSDS recommendations and standards in particular have been published to define and standardise the behaviour of ACM [5], [5]–[7]. Alongside increasing the throughput of the channel, ACM is being explored for its anti-jamming security properties; the robustness of the channel can be automatically adjusted when radio interference is detected.

As a result, current research is looking into the deployment of ACM in the next generation of government satellites, targeting increased throughput for high data rate telecommunication payloads [8]. The technology was explored in NASA’s SCaN testbed, in 2016, when an ACM-capable payload was deployed on the ISS; these protocols are now being standardised and




Fig. 1: An illustration of the uplink-assisted attacker, which abuses the Adaptive Coding and Modulation system, as compared to a conventional attacker. The conventional attacker is illustrated in dashed red, and a) transmits RF interference within the vicinity of the victim terminal, to b) affect decoding. The uplink assisted-attacker is in solid blue and instead 1) overshadows the uplink to select an advantageous Modulation and Coding, and then 2) transmits RF interference tailored to the selected parameters, to 3) affect decoding with significantly lower signal power, by factor Δ .

developed for deployment in near-future missions such as OPS-SAT 2 [4], [9], [10].

In parallel, additional work has explored extending DVB-S2 with ACM modes through mechanisms such as the DVB-RCS2 feedback channel, and a number of proprietary implementations such as the CDM-625 modem and Iridium [2], [11], [12]. However, to date the SCaN testbed remains the most fully documented real-world implementation, and is therefore the primary subject of our study.

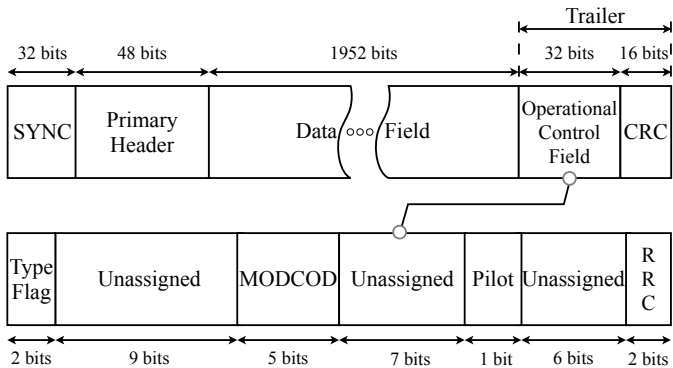


Fig. 2: Feedback Packet Structure. Note that it is unauthenticated.




Fig. 3: The DVB-S2 Frame Structure. Adaptive and Variable Modulation and Coding are supported since the PLHeader contains the MODCOD parameters of the subsequent XFEC Frame. The PLHeader is the most resilient part of the communication, using a QPSK encoding and strong error correcting code.

ACM achieves these performance improvements, in contrast to Static Coding and Modulation (SCM) systems, by matching the parameters of the protocol to the radio channel conditions. This mechanism requires a secondary uplink *control channel* which is used to select the parameters; when channel conditions are poor, a more robust error correcting code and modulation scheme is chosen, and vice versa. If implemented correctly, this scheme allows for continuous near-optimum utilisation of the available channel, approaching the theoretically maximum Shannon capacity even under fading conditions or interference [13]. Furthermore, these systems can be backward-compatible with *Variable Modulation and Coding* protocols such as DVB-S2, which self-report the current modulation and coding parameters as shown in Figure 3.

However, the security risks presented by the uplink control channel remain largely uninvestigated, even though the uplink remains unauthenticated across all documented satellite systems, and can therefore be directly abused by the attacker as illustrated in Figure 1. To the best of our knowledge, the security of ACM is considered only in two works: an ESA internal report, and an analysis of the security of satellite modems [8], [14]. In the latter paper, it was shown that for the given proprietary protocol, the attacker can spoof the downlink messages which specify the MODCOD, and therefore disconnect the terminal. However, neither paper investigated the impact of attacking the ACM uplink in terms

of the availability of the downlink.

There are also avenues for abuse even of authenticated uplinks without violating cryptographic operations; the attacker can exercise control over the chosen data rate by selective jamming of the uplink, or behaving as a legitimate party in a multi-user system. This also weakens the downlink to jamming, alongside other attacks such as overshadowing [15].

Furthermore, the increased prevalence of software-defined radio hardware, which itself makes on-board ACM feasible, has also lowered the cost of implementing this attack. Therefore, unless the security issues relating to ACM systems are resolved sufficiently quickly, the next generation of planned ACM satellites will unintentionally be at increased risk of jamming.

II. CONTRIBUTIONS

In this work, we introduce the *uplink-assisted attacker*, demonstrating for the first time how the control channel can be abused to weaken satellite systems to radio interference attacks targeting the downlink. In Section III we assess the related work, drawing particular attention to security papers on ACM in the related application of LTE. We provide a threat model in Section IV, relating the requirements for hijacking the uplink in combination with a downlink interference attack to a realistic attacker budget.

In Section V we outline the approach and theoretic gains achievable by uplink assisted attackers with the three objectives of denying service, degrading service, and hijacking the downlink. These are evaluated through software simulations described in Section VI; we discuss the results in Section VII. We finally conclude with a discussion of lessons learned and countermeasures in Section VIII.

III. RELATED WORK

One of the major purposes of employing ACM in radio systems is increasing the resilience of the channel to varying noise levels. In a security context, particular focus has been on how the channel can adapt and respond to apparent noise on due to radio interference and jamming [16], [17]. Specifically, research has shown that ACM channels are more resilient than static modulation and coding channels of equivalent average data bandwidth in a satellite context. However, to the best of our knowledge no current work considers how the ACM feedback channel can be hijacked to weaken the downlink to other forms of RF attack.

The most related work is within LTE, where there has been initial discussion on vulnerabilities in computing Channel Quality Indicators (CQI) for ACM. In particular, by jamming specific parts of the communication, an adversary can lead the system to select less resilient modulation and coding parameters [18].

Furthermore, it has been shown that LTE can be prevented from adapting to the correct channel conditions by jamming the LTE *physical uplink control channel*; this channel is used to report CQI [19], [20].

More recently, Zheng et al. have shown that deep learning methods for selecting MODCOD parameters are vulnerable to adversarial ML attacks [21]. This indicates that great care must be taken to mitigate adversarial channel manipulation when designing these systems.

IV. THREAT MODEL

The objective of the attacker is to decrease the power requirements of performing a radio interference attack on the downlink by hijacking the ACM control channel on the return link. Specifically, the adversary aims to weaken the downlink channel to jamming and overshadowing interference by selecting less resilient modulation and coding parameters, and by preventing the ground terminal from successfully requesting more resilient parameters.

The attacker therefore requires two radio transmit chains, both for the up- and downlinks; these transmitters do not necessarily have to be co-located. We assume that the attacker has access to commercially available equipment in order to assemble these transmitters, which includes two software-defined radios, upconverters as required, suitable amplifiers, and high gain antennas. We also assume that the downlink attacker can maintain presence in the vicinity of the victim ground segment.

A. Required equipment

In order to assess the requirements of performing uplink-assisted attacks, we now provide an example transmitter system to attack NASA’s SCaN ACM testbed, which is the focus of our evaluation.

Interestingly, both the uplink feedback channel and downlink communications channels operate within the S-band; however only the feedback channel is given precisely as 2216.5 MHz [4], [22]. As a result, similar hardware can be used for generating the signals for both channels.

A signal of the correct frequency can be directly generated by certain software-defined radios, which makes a subsequent upconverter not required. We take the HackRF as an example, which has a frequency range of 1 MHz to 6 GHz and is available for ~\$500 at the time of writing.

Finally, an amplifier and antenna is required to transmit the signal at a sufficient effective isotropic radiated power (EIRP) to overshadow the victim channel. Against the return link, our evaluation in Section VII demonstrates that this is possible with 0.2 dB relative attacker-to-victim power; therefore an attacker with an equivalent amplifier and transmitter as the legitimate ground segment is capable of executing the attack from anywhere within the satellite’s coverage region.

This hardware is also cheaply available: a 2.4 m tracking dish antenna is available from RF Hamdesign¹, and can be combined with an amplifier, such as the Mini Circuits ZHL-5W-2GX+. The requirements of overshadowing the downlink communication in this frequency band are well known, and are explored in more detail in existing work [23].

¹<https://www.rfhamdesign.com/>

V. ATTACK THEORY

We proceed to consider the three primary objectives that a physical-layer attacker can hold: denying service, degrading the service, and spoofing. In particular, we consider how control over the uplink channel assists an attacker in achieving each of these objectives.

For the Denial of Service and spoofing attacks which correspond to radio interference attacks, we measure the benefit achievable by the attacker in terms of a reduction in required attacker-to-victim power ratio. For degrading service, we calculate the reduction in bit rate that is possible under the different victim modes.

A. Jamming: denial of service

To deny service, we consider a jamming adversary; this attacker emits an interfering RF signal within the bandwidth of the victim signal whilst in the vicinity of a receiver. The jammer achieves denial of service once the attacker’s signal causes a sufficient bit error rate in the victim receiver.

The bit error rate caused in the receiver is ultimately determined by the modulation and coding parameters of the victim signal, which is adjusted in ACM systems to match the channel conditions within a margin of tolerance [24]. For example, in the presence of poor signal, a sparser modulation structure such as 8-PSK or QPSK can be selected, alongside a stronger LDPC code rate such as 1/4.

1) *Uplink-assisted jamming*: The fundamental advantage of an ACM channel against a jammer is that since the modulation and coding (MODCOD) parameters are selected to match the channel conditions, the attacker is required to deny service against the strongest available MODCOD, with the lowest data rate. In the NASA SCaN testbed implementation, this was achieved by defining a bit error rate known as Quasi-Error-Free (QEF), which corresponds to a packet error rate of $1e-5$ [4], and requesting MODCOD parameters to achieve QEF in the current channel conditions. Therefore, as the attacker begins to deny service, so the receiver selects more robust MODCOD parameters; this process continues up to the limit of QPSK 1/4, the most robust MODCOD parameters. Therefore, to achieve complete denial of service, the attacker’s signal must be of sufficient power to overcome QPSK 1/4 modulation and coding; we explore the case of a reduced error rate in Section V-B.

However, by exercising control over the uplink channel, the attacker can perform a physical-layer downgrade attack. The attack proceeds by first overshadowing the uplink control channel on the satellite, and selecting the weakest available modulation and coding parameters: for DVB-S2 this corresponds to 32-APSK 9/10. As a result, the victim communication is downgraded to the least resilient MODCOD, and the attacker only has to transmit a signal of sufficient power to overcome the downgraded communication. Due to the uplink overshadowing, the victim receiver’s requests to switch to a more resilient MODCOD are ignored.

TABLE I: The attacker-to-victim ratio, P_a/P_v , required to either degrade communication performance below quasi-error-free, or cause a given reduction in bit rate, against a representative selection of DVB-S2 modulation and coding parameters. The potential gain that the attacker achieves is given by the difference between P_a/P_v of the chosen MODCOD, and QPSK1/4.

MODCOD	P_a/P_v [dB] to jam		Bit rate	
	Absolute	QPSK1/4	Loss factor	Threshold [dB]
QPSK 1/4	2.4	0.00	9.09×	-4.0
QPSK 3/4	-4.0	-6.38	2.99×	-5.5
QPSK 9/10	-6.4	-8.77	2.49×	-7.9
8-PSK 3/5	-5.5	-7.85	2.50×	-6.4
8-PSK 3/4	-7.9	-10.3	2.00×	-9.0
8-PSK 9/10	-11.0	-13.3	1.66×	-12.7
16-APSK 2/3	-9.0	-11.3	1.68×	-10.2
16-APSK 3/4	-10.2	-12.6	1.50×	-11.0
16-APSK 9/10	-13.1	-15.5	1.25×	-16.0
32-APSK 3/4	-12.7	-15.1	1.20×	-13.1
32-APSK 9/10	-16.0	-18.4	1.00×	–

2) *Calculating the attacker advantage:* We now measure the potential impact of the uplink-assisted DoS attack as compared to conventional noise jamming by considering the difference in attacker-to-victim power ratio between the two approaches. This difference is given by considering the power required to deny service to a QPSK 1/4 channel, which is selected under conventional jamming, and the power required to deny service against a channel with attacker-controlled MODCOD parameters.

We do this by considering the DVB-S2 specification, which gives the error performance for DVB-S2 under different modulation and coding parameters [25]. In particular, this is given as E_S/N_0 – the signal-to-noise ratio, per modulated symbol, to achieve a quasi-error-free (QEF) service level.

We can use this to calculate P_a/P_v , the attacker-to-victim power ratio required to degrade service below QEF, by considering N_0 as the noise power spectral density induced by the attacker rather than thermal noise within the receiver. This gives:

$$P_a/P_v = -E_s/N_0$$

The resulting values of attacker-to-victim power ratio required to deny service are given in Table I.

It can be seen that the potential advantage of this approach is significant: by selecting the weakest available MODCOD parameters of 32-APSK 9/10, an uplink-assisted noise jammer can violate QEF at -18.4 dB as compared to the conventional noise jammer. The attack is still effective if the attacker has only limited control over the selected MODCOD; selecting QPSK3/4 yields instead a theoretic -6.38 dB difference.

B. Jamming: degrading data rate

Instead of aiming to completely deny service, another viable objective is simply to degrade the rate of the received data, resulting in a decrease in the volume of data downlinked per pass. Attacks of this form rely upon the property that more

resilient MODCOD parameters ultimately result in a lower data rate; sparser constellations encode fewer bits per symbol, and more robust error correcting codes require more parity bits.

A conventional jammer can exploit this by transmitting interference at a carefully selected power level; the receiver detects the interference and selects via the control channel a sufficiently robust MODCOD. However, the chosen parameters are of sufficiently low data rate that the desired application can no longer be supported and the attacker’s objective is achieved.

The advantage of the uplink-assisted attacker is that a low data rate MODCOD can be selected without requiring any interference on the downlink.

1) *Calculating the attacker advantage:* We quantify the potential advantage that an uplink-assisted attacker can realise in selecting a lower data rate, with respect to the bit rate achieved at each of the selectable MODCOD parameters. In particular, we calculate the *loss factor* – the factor of how slow a selected bit rate is relative to the bit rate prior to the attack.

The actual bit rate of the signal is determined by the symbol rate alongside the MODCOD parameters, which we tabulate in Table I. These values are calculated by considering the number of available bits relative to the number of total bits in DVB-S2 under the normal frame length of 64800. We also provide the power required for a noise jammer to cause this bit rate through degrading channel quality.

It can be seen that by overshadowing the uplink, the attacker can cause up to a $9.09\times$ reduction in data rate if the victim was previously in MODCOD 32-APSK9/10. This would have required attacker-to-victim power -4.03 dB for a conventional jammer. This corresponds to a significant power saving, and increased stealthiness: the attacker does not need to transmit any interference at the terrestrial receiver.

C. Overshadowing: spoofing

The final attacker objective that we investigate is overshadowing. Against an ACM system such as DVB-S2, the attacker can always select the most resilient modulation and coding to transmit and overshadow the downlink; in DVB-S2 this is QPSK1/4. However, in contrast to a conventional overshadowing attacker, the uplink-assisted attacker can also control the modulation which is then overshadowed.

For the uplink-assisted attacker, there is very limited potential in this approach. As shown in Figure 3, the PLHeader of the DVB-S2 packet is always $\pi/2$ QPSK modulated; the selected MODCOD parameters affect only the subsequent XFEC Frame. Therefore, regardless of the selected MODCOD, the attacker’s QPSK signal must compete with the QPSK signal.

Therefore, as described in previous work on the requirements of overshadowing, the attacker-to-victim power ratio we would have expected to successfully overshadow is -0.1 dB, regardless of the selected MODCOD [23]. Despite this, as we discover in Section VII, there is a marginal gain to be had by the uplink-assisted attacker in this case; even though

the PLHeader MODCOD can't be changed, the XFEC Frame MODCOD can be controlled.

VI. EXPERIMENT DESIGN

To assess the feasibility of uplink-assisted attacks we firstly evaluate the requirements of overshadowing the SCaN testbed uplink protocol, and then quantify the gains achievable by the attacker for each objective described in Section V. These evaluations consisted of physical-layer software simulations. All of the source code is available online at [redacted for review].

We now provide an overview of how the data for these scenarios was generated, and the results measured.

A. Uplink implementation

For our evaluation, we implemented a GNU Radio software encoder and decoder of the NASA SCaN testbed feedback channel. This was selected since it is the most publicly documented ACM return channel structure, was intended as a testbed for future ACM protocols, and is built of well-established satellite signal processing components [4]. We proceed to describe this transmit pipeline, which is illustrated in Figure 4.

B. Transmit/receive pipeline

In the NASA SCaN testbed feedback, the MODCOD parameters are selected by an adaptive algorithm which is run at the ground station. This algorithm selects, in particular, the modulation type, coding, pilot enable, and SRRC filter rolloff factor. These parameters are encoded ad-hoc within the *Operational Control Field* trailer of AOS Space Data Link Protocol Transfer Frames [26], as illustrated in Figure 2. The data field is left unassigned so that additional functions such as telecommand can operate in tandem with ACM. Although the precise location of this information within the frame structure varies per mission, we note that the security properties are identical for any unauthenticated structure.

Each Transfer Frame has a 16 bit Frame Error Control Field, which contains parity bits used for CRC error detection, and is decoded with a Viterbi Decoder. The new MODCOD parameters are only accepted if the CRC decodes with a syndrome of zero, which indicates no errors. An Attached Sync Marker is then appended to identify the start of each frame.

The frames are then modulated onto the carrier wave using the Glenn Goddard TDRSS (GGT) waveform under Mode F, as illustrated in Figure 4. This mode is a Single Access service; it is intended for a single groundstation-satellite link [22].

The frames are firstly scrambled by XORing with a pseudo-random bit sequence. Forward error correction is provided with a rate 1/2 convolutional code. The precise convolutional code used is unspecified in the report; we therefore verified through internal discussion that the standard CCSDS code as used on the Voyager program was adopted, which has rate 1/2 and $k = 7$ [27].

The waveform is then modulated with BPSK and a pulse shaping filter applied. The chosen data rate is 155.346 kbit/s,

on a 2216.5 MHz carrier [4], [22]. Since the pulse shaping filter is also unspecified, we have picked the standard root-raised-cosine (RRC) filter, with excess bandwidth factor $\alpha = 0.2$.

1) *Method*: By implementing the above transmit chain in GNU Radio, we generated clean I/Q sample data of the GGT waveform. To assess the bit error rate, we encode a pseudo-random data stream within the GGT waveform, which can be compared against after being received.

We used this to generate two sets of clean sample data: one representing the attacker and the other the victim signal. These signals are added together in I/Q space, and the combined signal is fed into a GNU Radio decoder. We consider the phase effects of the two desynchronised signals being generated from different reference oscillators out of scope for this work.

We assess the bit error rate by comparing the input and output byte sequences, as we vary the power of the attacker signal relative to the victim signal.

C. Downlink implementation

To compare the gains achievable by an uplink-assisted attacker to a conventional jammer or spoofing attacker, we set up a DVB-S2 software receiver using `gr-dvbs2rx` [28].

The data was generated by encoding a video signal into DVB-S2 packets, using `dvbs2-tx` to generate the I/Q sample data for each QPSK and 8-PSK MODCOD in Table I. Future work should consider including APSK modulations; however, these modes are unsupported by `gr-dvbs2rx` at the time of writing.

Throughout the evaluation we enabled DVB-S2 pilot symbols; these increase the resilience of the communication and thus increase the difficulty of jamming and overshadowing attacks.

We finally calculated the Attacker to Signal power ratio for jamming and overshadowing attacks, in both the conventional and uplink-assisted cases.

For the conventional attacker, we implemented the SCaN testbed adaptive algorithm: the reported signal strength measurement from `gr-dvbs2rx` is used to select the correct mode from a lookup table. The table uses the quasi-error free E_s/N_0 levels from the DVB-S2 standard document [25]. We added a further margin of 4 dB, based upon the given range of 0 dB to 9 dB from the SCaN experiment [4].

For the uplink-assisted jammer, the adaptive algorithm was not enabled, allowing the attack to proceed with the victim being held at a particular MODCOD.

For the jamming scenario, we mixed the signal with additive white Gaussian noise filtered by a root-raised cosine filter with $\alpha = 0.2$. For the overshadowing scenario we instead mixed the victim with a competing DVB-S2 signal at QPSK 1/4; the attacker maximises their advantage by selecting the most resilient MODCOD.

VII. EVALUATION

We now assess the security risks of Adaptive Coding and Modulation systems against satellite communications, through the experiments described in Section VI.




Fig. 4: Feedback Waveform Encoding.

A. Uplink overshadowing

We look firstly at the power requirements of overshadowing and jamming the uplink feedback channel, which enables the attacker to select the desired modulation and coding parameters for the downlink, and prevent the legitimate ground stations from doing so. The results are seen in Figure 5, which shows how the bit error rate induced in the satellite receiver varies as the attacker power increases. The attacker-to-victim power P_a/P_v is the relative power ratio of each signal, as they appear incident upon the receiver.

Since the uplink protocol is designed to be communicated over very long distances, a strong error correcting code (ECC) is used in combination with a QPSK modulation. This results in a very steep digital cliff effect for both jamming and overshadowing attackers, where the bit error rate suddenly changes at -1.3 dB for an AWGN jammer, and 0.2 dB for an overshadowing attacker. Since the chosen ECC is very resilient, there is only a 1.5 dB difference between the power levels required to jam and to overshadow; we address in Section VIII how different MODCODS could be selected to increase the difficulty of overshadowing further.

The fundamental takeaway is that, to perform this attack, the adversary requires transmitter hardware that has an EIRP at least 0.2 dB W greater than the legitimate ground station.

B. Downlink attacker

With the ability to overshadow the uplink feedback channel, an RF attacker against a SCaN-style ACM system can arbitrarily select the modulation and coding to their advantage. A less privileged attacker may instead be able to influence the MODCOD selection by acting as one of a number of terminals reporting signal quality within the satellite beam.

We therefore quantify the gain that an uplink-assisted attacker can realise against a conventional attacker, under different assumptions about the MODCOD selected by the victim downlink.

These experiments are run using the setup described in Section VI.

1) *Jamming: Denial of service:* The results of the jamming analysis are shown in Figure 6, which shows the frame error rate induced at the ground station receiver under the 8-PSK and QPSK victim modulations, comparing the uplink-assisted jammer to a jammer under adaptive modulation and coding. The frame error rate threshold we have chosen for denial of service is 50%.

It can be seen that the attacker realises the highest gain of -10.6 dB against 8-PSK with convolutional code rate $9/10$;

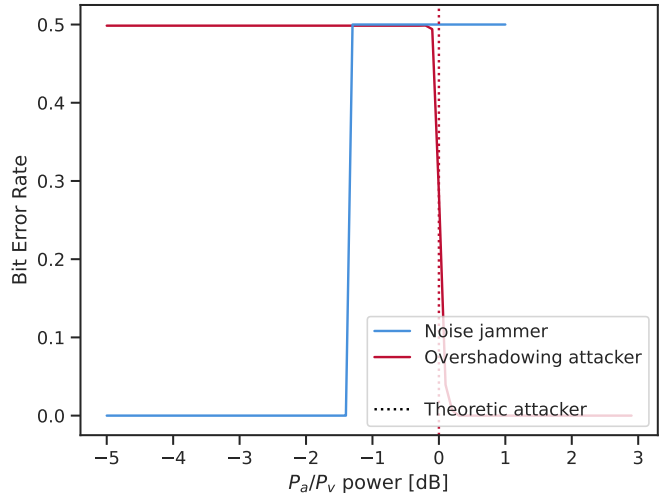


Fig. 5: The bit error rate of the feedback packets under both a noise jammer and overshadowing attacker. It can be seen that the noise jammer denies service to the uplink at -1.3 dB. To spoof the MODCOD request messages involves overshadowing the existing message, and instead requires only 0.2 dB.

this is only slightly lower than the theoretically predicted value of -13.3 dB from Section V. We note that even if the attacker can only partially control the selection of the MODCOD, a reduced but significant gain can still be realised. For example, the measured gain is -4.95 dB as compared to the -6.38 dB theoretic value, even if the MODCOD can only be switched to QPSK3/4.

The vertical dotted line represents the threshold at which a theoretic noise jammer causes sufficient bit errors to be induced to violate the quasi-error free condition, which were tabulated in Table I. It can be seen that this threshold is very close to the point at which the frame error rate climbs. A further study is required to evaluate the gains realisable against the least resilient MODCODs, such as 32-APSK9/10, against which the QEF condition is violated at -16.0 dB; with our experiments showing that QPSK1/4 is jammed to 50% frame error rate at 0.9 dB, the expected gain would be -16.9 dB. A further point of investigation would be the new VL-SNR mode for DVB-S2, which introduces a MODCOD even more resilient than QPSK9/10; this would increase the relative gain of the attacker further [29].

2) *Jamming: Degrading data rate:* We now extend the analysis to considering the alternative attacker objective of

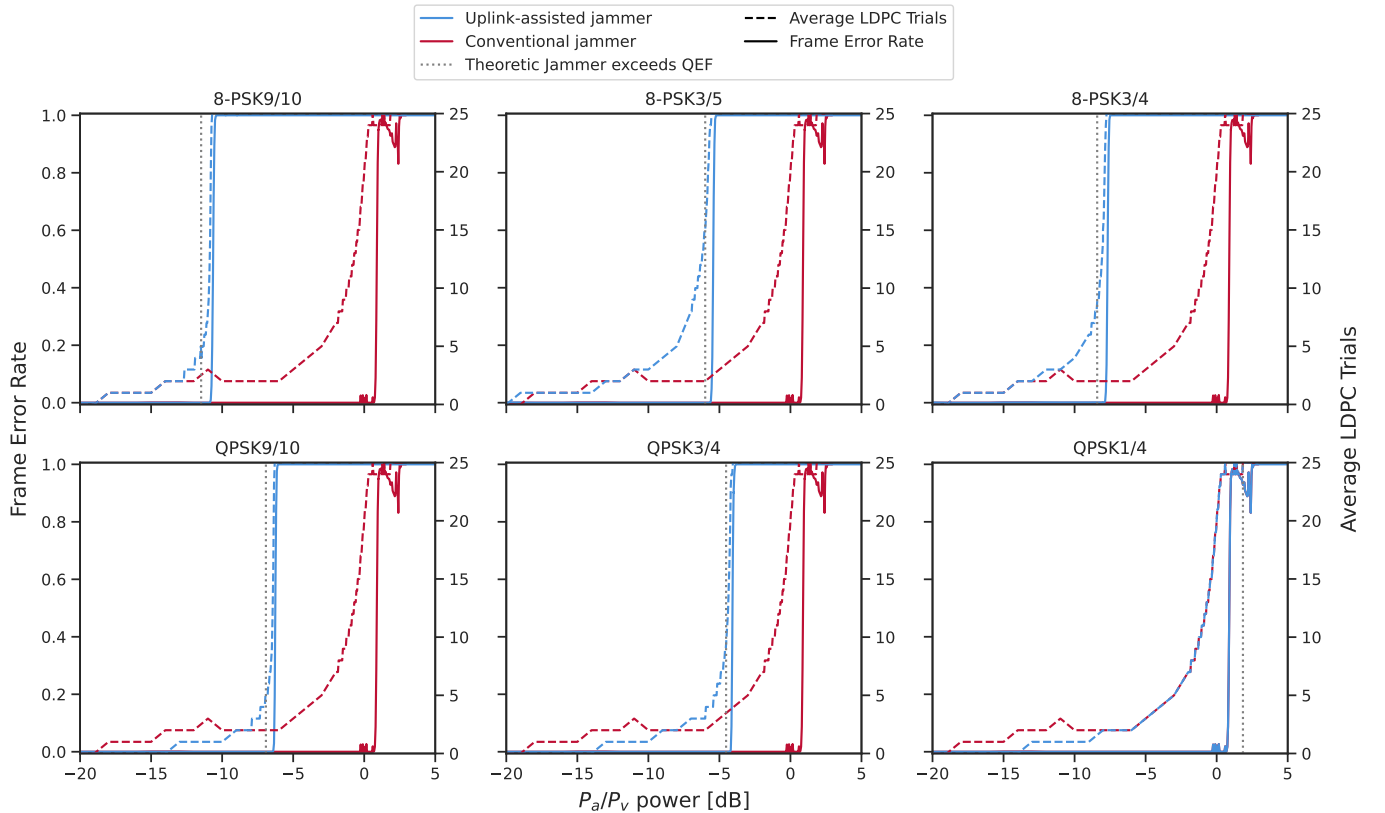


Fig. 6: The number of error correcting trials reported by `gr-dvbs2rx` as the P_a/P_v power ratio increases. As the conventional jammer power increases, the receiver continues to select more resilient error correcting codes, only failing at $P_a/P_v \approx 0$. In contrast, the uplink-assisted jammer which can fix the MODCOD parameters at 8-PSK9/10, yields a -10.6 dB gain over its conventional counterpart.

degrading the data rate of the downlink, rather than denying service outright. One motivation for this objective is making distributed denial of service attacks easier such as ICARUS, by reducing the volume of data required to saturate a link [30].

This objective is particularly relevant in ACM channels, where MODCODS with lower data rates are chosen in response to increases in jammer power.

To understand the advantage of the uplink-assisted attacker in this scenario, we assess the equivalent power of a noise jammer which degrades channel quality to select a MODCOD of a given error rate.

The results can be seen in Figure 7, which graphs how the bit rate received at the ground station varies as conventional jammer power increases. We note that the bit rate decreases stepwise as different MODCODs are chosen.

3) *Overshadowing*: We finally assess the advantage that an uplink-assisted spoofing attacker realises, if the attacker can influence the selected MODCOD parameters.

Although from the theoretic analysis in Section V we do not expect a particular gain since a successful overshadowing attack involves overshadowing both the inner data portion and the more resilient PLHEADER illustrated in Figure 3.

However, our results show that selecting weakened MOD-

CODs for the XFEC Frame has a slight effect in the uplink-assisted attacker's favour. In particular in Figure 8 we illustrate the P_a/P_v power required to bring reduce the frame error rate of the attacker's spoofed packets. From these results it can be seen that if the attacker aims for a frame error rate of 0.1, then 8-PSK3/4 is overshadowed at -0.7 dB, as compared to the -0.1 dB theoretic value, for a -0.6 dB gain.

VIII. DISCUSSION

The results shown in this work for the uplink-assisted attacker have serious implications for the security of adaptive modulation and coding systems. In particular, we have shown that if an attacker can influence the selection of the modulation and coding parameters of the channel, then the resilience of the channel can be significantly degraded against a jamming adversary by up to -16.0 dB. Although these gains can only be realised by an attacker with arbitrary control over the selected MODCOD, we have also shown that an attacker that can only modestly affect the selected MODCOD can still outperform the conventional jammer by 4.03 dB.

Furthermore, we have shown that existing feedback channels, such as the one implemented for the SCan testbed, remain unauthenticated; therefore the full potential of uplink-assisted attacks can be realised against these systems.

We now seek to discuss the implications of this attack type for existing satellite systems, and outline important considerations for the design of secure ACM.

A. Cryptographic authentication

The first and most obvious solution to an overshadowable uplink channel is cryptographic authentication: through signing the uplink messages and including a sequence number and timestamp, overshadowing and replay attacks can be effectively mitigated. Although necessary for a secure implementation, authentication is by no means sufficient: the attacker can still affect the chosen parameters by selectively jamming the uplink control messages; we later go on to discuss secure behaviour in this context.

Nevertheless, for the NASA SCaN protocol, this might be implemented enabling optional SDLS transport for the AOS Space Data Link [26]. This would be an important step, given that no publicly-available proposal for satellite ACM suggests authentication.

Further challenges with authentication arise in the context of multi-user uplink systems, such as Starlink or Iridium. In this case, the modulation and coding is instead chosen by an algorithm which takes into account the reported signal strength from all terminals within its beam; with uplink hardware, the attacker can participate in this by behaving as a legitimate authenticated party. Depending on the implementation, an attacker may be able to jam the uplink messages of all other terminals or get a majority of other terminals under their control, and thereby precisely control the selected MODCOD.

Furthermore, we acknowledge that mission operators may be apprehensive about uplink authentication due to its potential impact on reliability: this scheme necessarily involves rejecting unauthenticated messages.

B. Uplink hardening

Recent work on satellite spoofing has indicated mechanisms for hardening communication links to overshadowing, even if cryptographic authentication is not desired [23]. In particular, a sparser constellation might be considered which, despite being less resilient to jamming, is significantly more resilient to spoofing.

C. Detection mechanisms

Due to the above difficulties in supporting authenticated ACM feedback channels, and since attackers could still influence the chosen MODCOD in the multiple access setting, detecting this form of attack is of increased importance.

One approach would involve monitoring for inconsistencies between the requested and provided MODCOD parameters at the ground station. This relies upon a design feature specific to DVB-S2, in which the currently active MODCOD parameters are encoded within the downlink waveform, allowing ground stations to receive data at MODCODs other than the one that was requested.

The result of this is that, even when the XFEC Frame is undecodable, the MODCOD parameters from the PLFRAME are

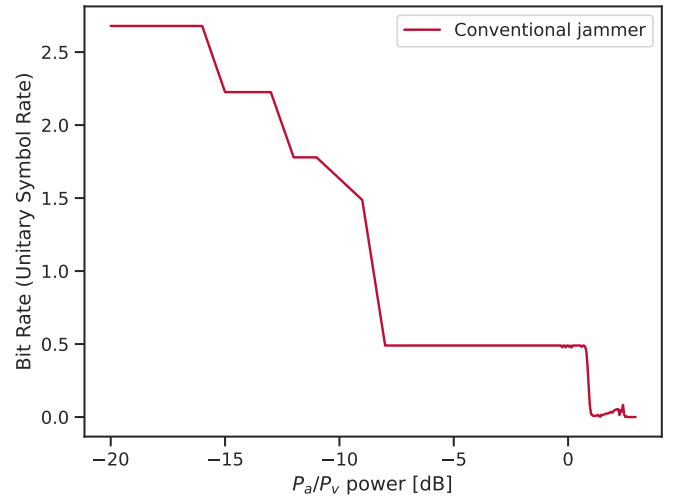


Fig. 7: The bit rate, relative to a symbol rate of 1, of the victim downlink signal as the P_a/P_v power ratio increases. It can be seen that the bit rate degrades stepwise as different MODCOD parameters are selected. Eventually service is completely denied at $P_a/P_v \approx 0.2$, and the bit rate drops to zero.

still recoverable. Therefore, unless the ground station is under sufficiently high jamming to also destroy the PLFRAME, an alert can be raised when inconsistencies between the reported signal level and the chosen MODCOD are detected.

D. Secure implementation

We finally turn our attention to potential pitfalls in the secure implementation of ACM, drawing upon lessons learned from the SCaN implementation.

1) *Measuring signal quality*: The SCaN testbed uses RSSI, a measure of received signal strength, in order to look up and request the correct MODCOD parameters. The idea is that since the receiver noise remains constant, RSSI allows the signal-to-noise ratio to be calculated. However, this approach is not resilient to radio interference, which appears as an increase in RSSI strength. This has been used in a similar way in LTE [18].

2) *Error correction*: To ensure that the MODCOD parameters are only adjusted when the satellite is certain of them being received, the SCaN behaviour is to reject any frames with bit errors. The unfortunate side effect of this is decreased resilience to jamming: the attacker only has to introduce one bit error through the convolutional code to deny service.

A cryptographic authentication scheme would provide a better guarantee of correct message delivery, whilst still being compatible with a concatenated outer error correcting code.

3) *Failing safe*: The SCaN implementation currently does not behave safely in the presence of uplink jamming. In particular, if no uplink packets are correctly received then the communication remains in the last selected MODCOD. This allows an attacker to select an advantageous MODCOD, and then continuously jam the uplink at lower power to prevent it from being changed. This can be resolved by defaulting to




Fig. 8: Comparison of the received quality of the attacker’s spoofed data when overshadowing the downlink, against a variety of victim MODCODS. The attacker transmits the most resilient QPSK1/4 waveform. It can be seen that the least resilient victim MODCOD parameters require ~ 1 dB less power to overshadow than QPSK1/4.

more resilient parameters in periods where no uplink packets are received.

IX. CONCLUSION

We have demonstrated that unauthenticated ACM feedback channels pose a major risk to the resilience of satellite communications against RF attackers. This is because the adversary can exploit the mechanism used to upgrade the robustness of the channel, and instead select communication parameters that make further downlink attacks easier; in our evaluation, it was sufficient to exceed the received signal power of the legitimate uplink by only 0.2 dB.

Through the simulations in this study, we have established that the gains for a jamming attacker are significant: in our simulations, the uplink assisted counterpart can deny service at -10.6 dB relative to its conventional counterpart. We have also shown that overshadowing attackers succeed at -0.6 dB relative to a conventional overshadower, and that attackers seeking only to reduce the data rate can cause up to a $9.09\times$ reduction in bit rate by only overshadowing the uplink.

These results not only draw attention to the dangers of unauthenticated ACM feedback channels, but also apply in the context of authenticated multi-user ACM systems. This class of attack can therefore only be fully mitigated in the design of the protocol itself. We have therefore discussed the steps involved in hardening the uplink to these attacks, implementing suitable detection mechanisms, and securing the implementation.

To prevent the next generation of satellite systems being vulnerable to these attacks, satellite operators must act quickly to research, standardise, and deploy hardened ACM mechanisms.

REFERENCES

- [1] F. Vieira, *Quality-of-service provision for satellite systems implementing adaptive physical layer*. Universidade do Porto (Portugal), 2008.
- [2] R. Miller, “Adaptive Coding and Modulation (ACM) in the CDM-625 Advanced Satellite Modem,” Tech. Rep. MSU-CSE-06-2, Comtech EF Data, June 2009.
- [3] S. Zhang, G. Yu, S. Yu, Y. Zhang, and Y. Zhang, “LSTM-Based Adaptive Modulation and Coding for Satellite-to-Ground Communications,” *Journal of Beijing Institute of Technology*, vol. 31, no. 5, pp. 473–482, 2022.
- [4] J. Downey, D. Mortensen, M. Evans, J. Briones, and N. Tollis, “Adaptive coding and modulation experiment using nasa’s space communication and navigation testbed,” in *34th AIAA International Communications Satellite Systems Conference*, p. 5736, 2016.
- [5] CCSDS 131.2-B-2, “FLEXIBLE ADVANCED CODING AND MODULATION SCHEME FOR HIGH RATE TELEMETRY APPLICATIONS,” *Recommendation for Space Data System Standards*, February 2023.
- [6] CCSDS 431.1-B-1, “VARIABLE CODED MODULATION PROTOCOL,” *Recommendation for Space Data System Standards*, February 2021.
- [7] CCSDS 130.12-G-2, “CCSDS PROTOCOLS OVER DVB-S2—SUMMARY OF DVB-S2 – SUMMARY OF DEFINITION, IMPLEMENTATION, AND PERFORMANCE,” *Report Concerning Space Data System Standards*, June 2023.
- [8] ThalesAlenia Space, ESA, “Definition and End to End System Analysis for the Use of ACM Techniques in the 26 GHz Data Downlink in Future Earth Observation (EO) Missions,” tech. rep., 2016.
- [9] J. A. Downey, D. J. Mortensen, M. A. Evans, and N. S. Tollis, “Variable Coding and Modulation Experiment Using NASA’s Space Communication and Navigation Testbed,” tech. rep., 2016.
- [10] D. Evans, “OPS-SAT-2 Optical Space Lab,” tech. rep., 2021.
- [11] H. Bischl, H. Brandt, T. de Cola, R. De Gaudenzi, E. Eberlein, N. Girault, E. Alberty, S. Lipp, R. Rinaldo, B. Rislow, *et al.*, “Adaptive coding and modulation for satellite broadband networks: From theory to practice,” *International Journal of Satellite Communications and Networking*, vol. 28, no. 2, pp. 59–111, 2010.
- [12] X. Wang, H. Li, and Q. Wu, “Optimizing adaptive coding and modulation for satellite network with ML-based CSI prediction,” in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1–6, IEEE, 2019.
- [13] C. Shannon, “The zero error capacity of a noisy channel,” *IRE Transactions on Information Theory*, vol. 2, no. 3, pp. 8–19, 1956.
- [14] R. Bisping, J. Willbold, M. Strohmeier, and V. Lenders, “Wireless Signal Injection Attacks on VSAT Satellite Modems,” in *33rd USENIX Security Symposium (USENIX Security 24)*, August 2024.
- [15] E. Salkield, S. Birnbach, S. Kohler, R. Baker, M. Strohmeier, and I. Martinovic, “Firefly: Spoofing Earth observation satellite data through radio overshadowing,” 2023.
- [16] K. Pinyoanuntapong, M. Goswami, A. B. Habib, H. M. Kwon, and K. Pham, “Boundaries of signal-to-noise ratio for adaptive code modulations,” in *MILCOM 2016 - 2016 IEEE Military Communications Conference*, pp. 132–137, 2016.
- [17] J. Peng, Z. Zhang, Q. Wu, and B. Zhang, “Anti-jamming communications in uav swarms: A reinforcement learning approach,” *IEEE Access*, vol. 7, pp. 180532–180543, 2019.
- [18] R. M. Rao, *Perspectives of Jamming, Mitigation and Pattern Adaptation of OFDM Pilot Signals for the Evolution of Wireless Networks*. PhD thesis, Virginia Tech, 2016.
- [19] M. Lichtman, J. H. Reed, T. C. Clancy, and M. Norton, “Vulnerability of LTE to hostile interference,” in *2013 IEEE Global Conference on Signal and Information Processing*, pp. 285–288, 2013.
- [20] M. Lichtman, T. Czauski, S. Ha, P. David, and J. H. Reed, “Detection and Mitigation of Uplink Control Channel Jamming in LTE,” in *2014 IEEE Military Communications Conference*, pp. 1187–1194, 2014.
- [21] K. Zheng and X. Ma, “Designing learning-based adversarial attacks to (mimo)-ofdm systems with adaptive modulation,” *IEEE Transactions on Wireless Communications*, vol. 22, no. 9, pp. 6241–6251, 2023.
- [22] D. T. Chelmins, “Glenn Goddard TDRSS Waveform 1.1. 3 On-Orbit Performance Report,” tech. rep., 2014.
- [23] E. Salkield, M. Szakály, J. Smailes, S. Köhler, S. Birnbach, M. Strohmeier, and I. Martinovic, “Satellite Spoofing from A to Z: On the Requirements of Satellite Downlink Overshadowing Attacks,”

- in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pp. 341–352, 2023.
- [24] G. Giambene and S. Kota, “Cross-layer protocol optimization for satellite communications networks: A survey,” *International Journal of Satellite Communications and Networking*, vol. 24, no. 5, pp. 323–341, 2006.
- [25] E. S. EN, “ETSI EN 302 307 V1.1.1: Digital Video Broadcasting (DVB),” *Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications (DVB-S2)*, European Telecommunications Standards Institute, Valbonne, France, vol. 3, 2005.
- [26] A. S. D. L. Protocol, “CCSDS 732.0-B-2 Recommendation for space data system standards, blue book, issue 2,” *Washington DC, USA: National Aeronautics and Space Administration*, 2006.
- [27] S. Butman, L. Deutsch, and R. Miller, “Performance of concatenated codes for deep space missions,” *The Telecommunications and Data Acquisition Progress Report*, pp. 42–63, 1981.
- [28] I. Freire and R. Economos, “gr-dvbs2rx: a Software-Defined DVB-S2 Receiver based on GNU Radio,” Mar. 2023.
- [29] “White Paper on the use of DVB-S2X for DTH applications, DSNG & Professional Services, Broadband Interactive Services and VL-SNR applications,” tech. rep., March 2015.
- [30] G. Giuliani, T. Ciussani, A. Perrig, and A. Singla, “ICARUS: Attacking low earth orbit satellite networks,” in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pp. 317–331, 2021.