Temporal Risk on Satellites

Shiqi Liu
George Mason University
sliu38@gmu.edu

Kun Sun
George Mason University
ksun3@gmu.edu

***Abstract*—Satellite vulnerabilities change over time as orbits shift, power margins tighten, and the space environment deteriorates. However, most cybersecurity risk frameworks still treat threats as static. In practice, the same exploit can be far more damaging during a critical maneuver than during routine operations. We propose a temporal risk assessment framework that makes time an explicit axis in satellite security analysis. It extends existing adversary behavior taxonomies with a five-dimensional temporal capability model and estimates exploitation difficulty across distinct temporal windows of a mission. Rather than producing a single risk score, the framework outputs a series of time-indexed likelihood–impact matrices. It discretizes missions into operationally meaningful time windows and environmental bands to show when systems are most exposed. This view helps operators avoid scheduling sensitive operations in high-risk periods and align defensive resources with a threat landscape that shifts over time.**

I. Introduction

Satellite systems [1] now underpin global communications, navigation, and Earth observation, making them critical infrastructure. As these systems grow more sophisticated and more interconnected, their attack surface expands. Their predictable orbits, limited physical protection, and complex supply chains create persistent opportunities for adversaries [2]. The growing use of Commercial Off-The-Shelf (COTS) hardware and inexpensive ground station services further lowers barriers to entry. These trends make it easier and cheaper to interact with and attack satellite systems [3]. Recent incidents, such as the 2022 KA-SAT cyberattack [4] that disrupted communications across Europe, show that satellite security threats are not hypothetical.

Prior work has emphasized spatial and architectural aspects, but the *timing* of attacks can matter as much as *what* attackers do and *how* they do it. Satellite operations are governed by time-dependent phenomena. Orbital mechanics determines when ground stations can communicate with spacecraft. Mission-critical operations such as orbital maneuvers, payload activations, and data downlinks run within narrow time windows. Adversaries who understand these temporal patterns can time operations to increase impact. Recent studies [5] show that space-specific effects, such as radiation-induced bit flips, can interfere with defenses such as Control Flow Integrity (CFI). During South Atlantic Anomaly (SAA) passages or geomagnetic storms, elevated single-event upset (SEU) rates increase the residual risk of mitigation failure. Complementing these radiation-induced effects, a recent study on energy-drain attacks in satellite Internet constellations also supports this view. Low Earth orbit (LEO) satellites alternate between sunlight and eclipse and rely on batteries during eclipse. The depth of battery discharge has a major effect on lifetime. The STARMELT attack [6] sends crafted traffic through a victim satellite to keep its communication subsystem in a high-power state instead of allowing it to sleep. It exploits predictable orbits and access patterns to align traffic with eclipse periods and to maximize battery drain.

Existing space cybersecurity risk frameworks, such as SPARTA [7], are widely used for threat modeling and for mapping techniques to NIST SP 800-53 controls [8]. However, current approaches treat risk as essentially *time-invariant*. Each technique receives a single risk score that implicitly assumes its feasibility and impact remain constant throughout the mission. The absence of an explicit temporal dimension in risk assessment creates three gaps: *(i) Incomplete likelihood estimation*. Existing models do not capture adversaries' temporal capabilities or how the difficulty of executing an attack varies across time windows. *(ii) Impact underestimation*. Static risk scores cannot reflect how system vulnerability changes with operational context. *(iii) Misaligned defenses*. Without temporal risk profiles, defenders cannot prioritize protection during high-risk windows or schedule sensitive operations outside predictable vulnerability periods.

We present a *temporal*[1] *risk framework* for satellite cybersecurity that treats timing as a primary concern in threat assessment. Building on SPARTA's Notional Risk Scores (NRS), we extend the standard 5×5 matrix over likelihood and impact into a family of matrices indexed by time. Each matrix evaluates risk for specific mission time windows and environmental bands (e.g., eclipse vs. non-eclipse, geomagnetic storm severity levels). The framework fits into existing risk management workflows. Analysts can add temporal considerations on top of established ATT&CK-based assessments. By making explicit *when* attacks are feasible, high-impact, or difficult to attribute, the framework gives operators a time-aware view of risk that current practices lack.

[1] In this paper, we use the term *temporal* to refer to the dependency on physical time, orbital mechanics, and mission schedules. This is distinct from the *Temporal Metrics* in CVSS, which refer to the lifecycle status of a vulnerability (e.g., exploit maturity or patch availability).

Workshop on Security of Space and Satellite Systems (SpaceSec) 2026
23 February 2026, San Diego, CA, USA
ISBN 978-1-970672-02-2
https://dx.doi.org/10.14722/spacesec.2026.23050
www.ndss-symposium.org

II. Background & Related Work

A. Attacks Against Satellite Systems

Satellite attack surfaces span the ground, link, user, and space segments [3], [9]. Adversaries often gain a foothold in the ground segment through ground networks, Ground Station as a Service (GSaaS) platforms, or operational tooling [10], [4]. The compromise can then cascade to space assets. Effective defense therefore requires a holistic analysis of end-to-end mission control and data flows instead of treating the spacecraft in isolation.

Within the link and user segments, over-the-air jamming and spoofing [11], [12], [13] remain prominent, and the modem and terminal layers are frequent weak points. Teardown and measurement studies that focus on low Earth orbit terminals and user equipment [14], [15], [16] have revealed design and implementation defects. They also show concrete paths to exploitation.

On the ground segment, falling ground station costs and the emergence of GSaaS [17], [18] have eroded physical and economic barriers. It is now easier for attackers to interact with assets in orbit and to trigger software vulnerabilities. Empirical analyses of firmware from satellites in orbit [19], [10] show that some systems lack protection of telecommand channels and lack access control. They also exhibit multiple software defects. These weaknesses show that relying on obscurity alone is fragile.

Isolation and recovery capabilities in orbit are therefore central to resilience. Small satellites increasingly run full operating system stacks, for example Linux-based payload computers. Recent work [20], [21] has begun to evaluate application sandboxes as a practical isolation layer for CubeSat-class missions and to document feasibility and engineering challenges for integration in orbit. In parallel, proposals for independent health monitoring and out-of-band telemetry channels [22], [23] highlight the prevalence of subsystem faults. They also show the importance of post-incident diagnosis and recovery for limiting the blast radius of cyberattacks.

B. MITRE ATT&CK in Space

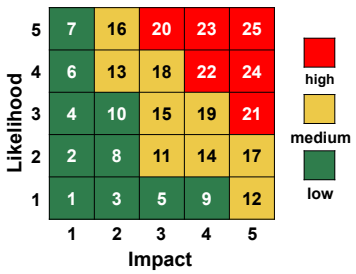

Fig. 1. Risk Matrix Representation [24].

MITRE ATT&CK [25] is a matrix-based knowledge base that provides a common lexicon of adversary tactics, techniques, and procedures (TTPs). It spans the Enterprise, Mobile, and Industrial Control Systems matrices and is widely used for adversary emulation and detection coverage assessment. In recent years, researchers have begun to adapt ATT&CK-based TTP matrices to the space domain. The Aerospace Corporation SPARTA framework [7] defines a space cyber tactics and techniques matrix. Its Notional Risk Scores (NRS) place each technique in a 5×5 matrix over likelihood and impact, as shown in Figure 1. Likelihood is estimated from adversary motivation, exploitation difficulty, and capability tier. Impact covers mission disruption, loss of control or availability, safety effects, financial loss, and national security consequences. Cell values from 1 to 25 are grouped into low, medium, and high and drive the selection of mapped countermeasures and NIST SP 800-53 controls [8].

III. Methodology

Building on the SPARTA Notional Risk Scores summarized above, we now describe how to extend the static 5×5 risk matrix into a temporal form. For each technique, we first examine how likelihood and impact vary across the relevant time windows. We assign 1–5 levels for likelihood and impact to each window.

For likelihood assessment, we apply time-dependent adjustments to *Exploitation Difficulty* and *Adversary Capabilities*. These adjustments reflect temporal accessibility and adversary timing proficiency. *Motivation* retains its original meaning. It is primarily driven by system criticality and the resulting attractiveness of the target. Favorable time windows and exploitation costs may influence how an adversary prioritizes targets. They do not change the intrinsic appeal of the target. It is worth noting that in traditional static frameworks [7], analysts often implicitly account for narrow time windows by lowering the overall likelihood score. While this produces a valid time-averaged assessment, it can mask periods of acute vulnerability.

Impact is also modulated by time. Time windows influence the realized consequences and severity of a successful attack. This effect depends on the mission and on the specific program. Time windows do not change the impact scoring framework itself. The category scheme (I1–I5) and its semantics, such as mission interruption, partial degradation, or loss of mission, remain unchanged. Time determines which level in that framework a specific attack scenario should map to. For example, the same jamming attack might be assessed as I2 (recoverable data delay) during routine operations but as I5 (battery depletion that terminates the mission) during an eclipse-critical period. Both ratings use the same impact framework. The difference arises from how vulnerable the system is in the respective time windows.

A. Temporal Adversary Capabilities

Existing space cybersecurity risk frameworks [7], [26] use a seven-tier attacker model to quantify adversary capability. The tiers range from script kiddies to the most capable state actors. Each tier is described along seven generic dimensions, such as vulnerability discovery and exploitation. To remain

TABLE I
MINIMAL TEMPORAL CAPABILITY SETS M_T FOR THE SEVEN ATTACKER TIERS.

Tier	Representative adversary type	M_T	Intuitive description
1	Script kiddies	$\{A\}$	Uses off-the-shelf tools in rough time windows, with no real prediction, sensing, or precise timing.
2	Hackers for hire	$\{A, F\}$	As above, plus basic forecasting of access windows using public orbit and pass-prediction tools.
3	Small hacker teams	$\{A, F, S\}$	Add continuous sensing (e.g., spectrum monitoring) to refine when the target is accessible and responsive.
4	Insider threats	$\{A, F, S, Y\}$	Combine high-quality internal state data with the ability to align actions to precise operational time slots.
5	Large well-organized teams	$\{A, F, S, Y, H\}$	Possess the full temporal capability set, including deliberate timing for statistical and operational hiding.
6	Highly capable state actors	$\{A, F, S, Y, H\}$	Same temporal capability set as Tier 5, but applied at greater technical depth and operational scale.
7	Most capable state actors	$\{A, F, S, Y, H\}$	As Tier 6, further extended to coordinated, multi-theater, multi-constellation campaigns.

comparable to these frameworks, we keep the seven-tier categorization. Instead, we map *time* back into attacker capability itself. We introduce a temporal mastery profile $\langle A, F, S, Y, H\rangle$ (actuation, forecasting, sensing, synchronization, and hiding) as an additional threshold for tier assignment. For each attacker class we define the minimum required set of temporal capabilities, as shown in Table I. This preserves the generality of the original tiers while embedding temporal proficiency into attacker capability metrics. It also lets us compare timing-driven adversarial operations across tiers.

Actuation (A) captures whether an adversary can cause effects on the target within a given access window. It includes RF power shaping and directed jamming at the waveform layer and malicious uplinks at the protocol and command layers. In extreme cases, it also covers physical actions such as laser dazzling. The exploitable window is bounded by the visibility window. This window is set by practical engineering constraints such as line-of-sight (LoS), antenna pointing, and radio frequency (RF) chain availability.

Forecasting (F) turns observations into time-localized predictions. It applies simplified perturbation models for orbital propagation (SGP4 for near-Earth and SDP4 for deep-space objects) together with link geometry calculations. These models estimate future access windows, eclipse ingress and egress times, and link budget extrema. Forecasting also uses space weather and radiation models to anticipate SAA crossings and geomagnetic storm intervals.

Sensing (S) provides near-real-time awareness of system state. It includes passive monitoring of unencrypted beacons and telemetry, spectrum and power side-channels, and optical or radar observations. It also covers high-quality, low-latency telemetry from compromised ground assets or from attacker-controlled ground stations. These sources can be fused with third-party orbital data and space weather feeds to produce time-aligned situational awareness.

Synchronization (Y) is the ability to align actions with a chosen window in time. It depends on a stable time base and a well-characterized end-to-end delay budget. In practice, this includes disciplining clocks via GPS or NTP, modeling and compensating propagation and processing delays, and phase-locking actions to onboard timing or ground-station handovers. Without Synchronization (Y), even strong forecasting and sensing capabilities rarely produce reproducible effects.

Hiding (H) aims to make induced effects hard to distinguish from noise in statistics. It keeps observable outcomes within expected statistical bounds and mimics natural failure signatures such as SEU-like event rates. It also times actions to coincide with space weather activity, Sun outage periods, or maintenance windows. A high level of Hiding (H) assumes a baseline model of nominal and anomalous behavior and competent statistical modeling.

B. Temporal Exploitation Difficulty

Following the SPARTA framework [7], we retain *Exploitation Difficulty* as a component of technique likelihood. We decompose it into static and temporal components. The static component corresponds to the original notion in SPARTA. It captures code complexity, required tooling, and non-temporal preconditions. The temporal component, which we call *Temporal Exploitation Difficulty (TED)*, is orthogonal to this static part. TED measures how hard it is to exploit a technique at the *right time*, assuming the attacker already knows how to execute the technique in principle. We distinguish two sources of temporal difficulty:

Intrinsic properties of the time window. These properties come from the system's orbital visibility and mission operations timeline. They are independent of any specific adversary. Key factors include the duration of each usable window, how often such windows occur, how predictable their timing is, and whether the opportunity is one-shot or repeatable. For example, routine downlink passes offer frequent and predictable windows with considerable slack. In contrast, windows driven by geomagnetic storm activity or contingency operations are less predictable in practice. Intrinsic properties of time windows can often be analyzed with system-level tools. Orbit propagation and link geometry models [27] capture access and eclipse windows. Radiation and SAA models [28] capture regions with enhanced particle flux. Statistical models capture the frequency and duration of rare combined states.

Scenario-specific timing constraints. These constraints arise from the concrete attack objective. They do not change the underlying visibility or mission windows. Instead they restrict how those primitive windows can be used in practice. For example, an adversary may need to align an uplink with a narrow command window or overlap with a critical mission phase such as a maneuver or payload operation. Such constraints shrink and correlate the set of usable windows for that scenario. Intuitively, scenarios with high TED often require a richer temporal capability profile $\langle A, F, S, Y, H\rangle$ on the adversary side. We keep this separation explicit: TED is evaluated

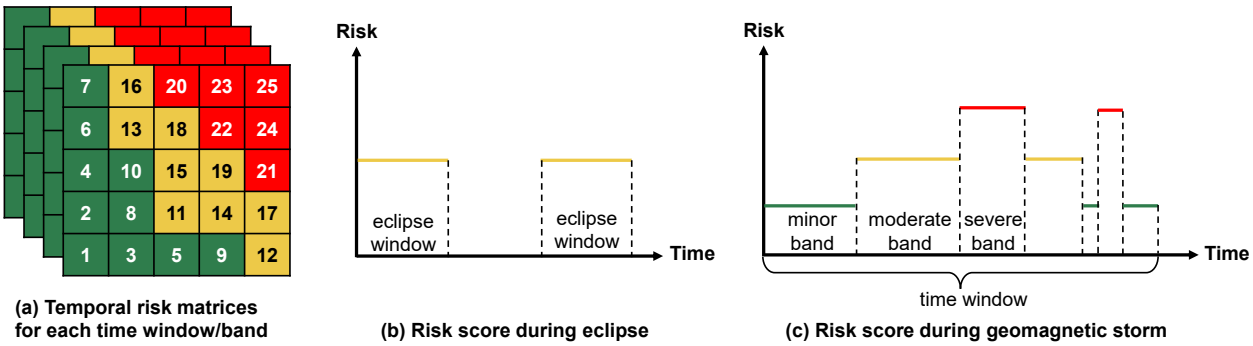

Fig. 2. Temporal Risk Matrices and Time-Resolved Risk Across Windows and Bands.

from the system and the scenario, while $\langle A, F, S, Y, H \rangle$ is used to determine which attacker tiers can satisfy those timing requirements.

Combining motivation, temporal capabilities, and time-varying exploitation difficulty yields the likelihood score. As in NRS [7], this likelihood is a semi-quantitative, expert-elicited rating rather than a closed-form expression. The difference is that elicitation now explicitly accounts for temporal considerations.

C. Temporal Risk Matrix

In real-world systems, both likelihood and impact are time-dependent. Time should not be viewed merely as an external parameter that scales a fixed risk value computed under the assumption of time-invariant likelihood and impact. To stay compatible with the semi-quantitative 5×5 risk matrix, we avoid a full three-dimensional risk surface. Instead, at the modeling level we treat likelihood and impact as functions of time, $L(t)$ and $I(t)$. We then project this time-resolved analysis back onto a two-dimensional grid by assigning discrete scores from 1 to 5 across the mission timeline. The resulting temporal risk matrix is best viewed as a family of 5×5 matrices indexed by time, as shown in Figure 2(a), rather than as a single time-agnostic picture of risk.

To keep the number of risk-matrix snapshots tractable, we discretize the mission into time windows, such as passes, mission phases, or other critical intervals. Beyond recurring orbital events, distinct lifecycle phases act as unique risk windows: the Launch and Early Orbit Phase (LEOP) heightens TT&C vulnerability during orbit raising and stabilization, while Decommissioning reduces vigilance and relaxes security, changing the system's defenses. In the simplest cases, for example eclipse versus non-eclipse as shown in Figure 2(b), the relevant time windows are naturally discrete and internally homogeneous. A single (L, I) level can then be assigned to each window. For environmental episodes such as geomagnetic storms or South Atlantic Anomaly transits, as shown in Figure 2(c), risk can vary within the episode as a function of storm intensity or radiation flux. In our framework, we further discretize the underlying intensity, for example storm level or particle flux, into a small number of bands and treat each band as a separate time window with its own (L, I) level. Each time window or time-severity band still maps to a well-defined cell in the 5×5 risk matrix, while the underlying analysis remains time-resolved and environment-aware.

Likelihood varies across time windows and bands through the combined effect of motivation, static exploitation difficulty, TED, and the minimal attacker tier. Once the mission has been discretized, we assign an ordinal likelihood level from 1 to 5 to each window or band. For a fixed technique and attack objective, we treat likelihood as piecewise constant within homogeneous window classes. Windows of the same type, such as routine passes with similar geometry or eclipse intervals with similar temporal characteristics, share the same likelihood level. Likelihood changes when the system moves between window classes, for example from routine operations to eclipse-critical or high storm conditions, or when the attack objective itself changes. Impact varies across windows and bands; it reflects the consequences if the attack succeeds in a given window.

IV. Case Study

We apply our temporal risk framework to STARMELT [6] as a concrete case study. Under the SPARTA framework [7], STARMELT maps to *EX-0013.01: Valid Commands*, where an attacker issues legitimate (protocol-valid) actions in a way that progressively depletes onboard resources such as stored battery energy.

For high-criticality satellite systems, SPARTA yields a static NRS of 25 (High) by taking worst-case Likelihood and Impact. In LEO operations, however, the same technique can have very different consequences depending on the power regime: during sunlight the platform is typically power-positive, whereas during eclipse it operates on battery discharge. Treating the entire orbit as uniformly "High" encourages always-on throttling even when margins are available, while a single "average" score can understate the short eclipse interval where resource depletion is most consequential. To make this dependence explicit without leaving the 5×5 matrix, we score Sunlight and Eclipse as separate windows.

A. Likelihood Analysis

We assess the likelihood of the STARMELT attack based on Temporal Adversary Capabilities and TED. Internet Con-

stellations (SICs) are critical infrastructure, so we assume the adversary's motivation remains high.

Temporal Adversary Capabilities. To execute STARMELT, an attacker must possess a specific Minimal Temporal Capability Set (M_T). We assume an adversary with at least Tier-4 temporal capabilities (per Table I), as executing STARMELT requires the following capabilities:

- STARMELT relies on predicting when the victim satellite enters and exits eclipse and on understanding constellation topology/routing dynamics. This requires capability F (e.g., ephemeris propagation and lighting prediction) and capability S (sensing system connectivity and exploiting exposed topology/routing logic).
- STARMELT employs "Victim Stalking" and "Victim Crossing" strategies, coordinating geo-distributed bots to inject traffic aligned with fast orbital motion and window transitions. This requires capability Y to align ground asset activation with satellite phase timing.
- STARMELT also requires capability A. The attacker must be able to generate protocol-conformant traffic and inject it over the victim's relevant links at a high enough rate. The goal is to sustain congestion and keep power draw elevated.

The required capability set is thus $M_T = \{A, F, S, Y\}$. For this case study, we treat M_T as constant across sunlight and eclipse windows, since generating protocol-conformant traffic and coordinating timing do not directly depend on illumination. Any illumination-driven operational changes (e.g., routing policy changes) are instead reflected through TED.

Temporal Exploitation Difficulty. We distinguish scenario-specific constraints from intrinsic window properties.

a) Intrinsic properties: The eclipse window is shorter than the sunlight window, which reduces temporal slack. While the attack vector (protocol-conformant traffic injection) can be initiated on demand and does not require a dedicated warm-up phase, effective execution still depends on reaching a sufficient injection rate and maintaining it long enough within the window. A shorter window therefore provides fewer retries and less margin for timing error, moderately increasing TED compared to sunlight.

b) Scenario-Specific Constraints: The primary constraints for STARMELT depend on the attack vector. Ground-to-satellite link (GSL) saturation is geometry-limited, requiring bots to remain within the victim footprint to maintain LoS, while inter-satellite link (ISL) saturation is topology-limited, requiring traffic to traverse feasible routing paths as connectivity evolves. These constraints do not directly depend on illumination, but illumination-driven operational behavior can further shrink the effective usable portion of each window; we account for such effects as scenario-specific constraints when evaluating TED.

Because TED is moderately higher during eclipse due to reduced slack and a potentially smaller effective window, we rate the likelihood as Very High (5) in sunlight and High (4) in eclipse.

B. Impact Analysis

Unlike Likelihood, the operational Impact is heavily modulated by the satellite's power state. We align our scoring with standard NASA risk consequence criteria [29].

Sunlight Window. During sunlight, the solar arrays generally cover the bus load and allow the battery to recharge. In STARMELT's illuminated-case evaluation [6], sustained traffic flooding did not increase battery depth-of-discharge (DoD), consistent with a net-positive energy balance. Under illumination, the consequences are primarily non-energy, driven by higher onboard computation. This fits the *Minor Degradation* category: performance degrades but stays within system margins. We therefore score Impact as 2 (Low).

Eclipse Window. The satellite relies entirely on battery discharge. STARMELT shows that keeping the satellite out of low-power sleep during eclipse increases discharge and accelerates battery wear, reducing effective battery life (over repeated eclipse intervals, up to ∼76% over 1.5 years in the evaluated setting) [6]. With tight power margins, this may trigger undervoltage protection and safe modes; in the worst case, it can disrupt attitude control and communications. We therefore treat eclipse consequences as *Mission-Threatening* and score Impact = 5 (Very High).

C. Operational Value

TABLE II
TEMPORAL RISK COMPARISON (HIGH CRITICALITY SYSTEM)

Assessment Mode	Likelihood	Impact	Risk Score
Static Baseline (SPARTA)	5	5	**25 (High)**
Temporal (Sunlight)	5	2	**16 (Medium)**
Temporal (Eclipse)	4	5	**24 (High)**

Table II highlights the temporal split: eclipse carries the highest risk, whereas sunlight is materially lower. In practice, this supports stricter controls during eclipse and less restrictive non-critical throttling during sunlight to preserve throughput, without relaxing baseline safeguards.

V. DISCUSSION

Validation and calibration. Our framework remains semi-quantitative. Time-resolved likelihood and impact levels are still elicited from experts. An open problem is how to calibrate and validate such temporal risk models. Public space cyber incident data are scarce and incomplete. They rarely reveal how outcomes depend on timing. The community still lacks standardized time-aware red-team exercises, simulations, and benchmarks. These artifacts would help test whether temporal risk profiles produced by our approach align with periods when systems would be most vulnerable.

Temporal attack-defense games. We currently treat temporal risk as a property of a given system and adversary, without modeling strategic interaction between attackers and defenders who both choose when to act. An open problem is to embed temporal risk assessment into explicit attack-defense games over time, where attackers select windows to exploit and

defenders decide when to harden, monitor, or reschedule operations.

Scalable compositional modeling. Our temporal abstractions are currently defined for a single spacecraft and its environment. Future missions increasingly involve constellations, shared ground infrastructure, and multiple services. Enumerating all combinations of time windows, environmental bands, assets, and services does not scale. However, this complexity also offers an opportunity for defense through *temporal diversity*: since satellites in different orbital planes enter high-risk windows (e.g., SAA transit) at staggered times, future frameworks could enable *risk-aware routing* that dynamically directs traffic away from vulnerable assets. To achieve this, an open problem is how to aggregate risk at the orbital plane level to build constellation-wide assessments without losing necessary temporal structure.

VI. Conclusion

In this paper, we argue that security risk for space systems cannot be treated as a static property of individual techniques. It must be evaluated in light of when attacks are feasible and how operations and environment change over time. We complement existing frameworks with a time-sensitive view that augments attacker tiers with a temporal mastery profile and separates static from timing-related difficulty. We then represent risk as a family of likelihood–impact matrices parameterized by operational windows and environmental bands. These elements let operators tie risk assessment more directly to orbital dynamics, mission timelines, and space weather.

Acknowledgment

This work was supported by ONR grant N00014-23-1-2122.

References

[1] SpaceX, "Starlink: Reliable high-speed internet from space," https://www.starlink.com/, 2025, accessed: 2025-11-16.

[2] J. Vanlyssel, E. Sobrados, R. Anwar, G. Roman, and A. Anwar, "Spychain: Multi-vector supply chain attacks on small satellite systems," *arXiv preprint arXiv:2510.06535*, 2025.

[3] J. L. C. Remy, E. Ear, C. Chang, A. Feffer, and S. Xu, "SoK: Space infrastructures vulnerabilities, attacks and defenses," in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2025, pp. 1028–1046.

[4] Viasat Inc., "KA-SAT network cyber attack overview," Viasat Corporate Blog, Mar. 2022, https://www.viasat.com/perspectives/corporate/2022/ka-sat-network-cyber-attack-overview/ (accessed 2025-11-16).

[5] J. Willbold, T. Cloosters, S. Wörner, F. Buchmann, M. Schloegel, L. Davi, and T. Holz, "Space RADSIM: Binary-agnostic fault injection to evaluate cosmic radiation impact on exploit mitigation techniques in space," in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2025, pp. 1047–1063.

[6] Y. Zhang, Q. Wu, Z. Lai, Y. Deng, H. Li, Y. Li, and J. Liu, "Energy drain attack in satellite internet constellations," in *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)*. IEEE, 2023, pp. 1–10.

[7] E. Ear, B. Bailey, and S. Xu, "Towards principled risk scores for space cyber risk management," *arXiv preprint arXiv:2402.02635*, 2024.

[8] Y. Kurii and I. Opirskyy, "Analysis and comparison of the NIST SP 800-53 and ISO/IEC 27001:2013," in *Proceedings of the 2nd International Workshop on Cyber Protection of Information and Telecommunication Systems (CPITS 2022)*, ser. CEUR Workshop Proceedings, vol. 3288. CEUR-WS.org, 2022, pp. 21–32.

[9] R. Peled, E. Aizikovich, E. Habler, Y. Elovici, and A. Shabtai, "Evaluating the security of satellite systems," *arXiv preprint arXiv:2312.01330*, 2023.

[10] J. Willbold, M. Schloegel, M. Vögele, M. Gerhardt, T. Holz, and A. Abbasi, "Space odyssey: An experimental software security analysis of satellites," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2023, pp. 1–19.

[11] C. Tibaldo, H. Sathaye, G. Camurati, and S. Capkun, "GNSS-WASP: GNSS wide area SPoofing," in *USENIX Security 2025*, 2025, pp. 7039–7058.

[12] 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 (WiSec '23)*. ACM, 2023, pp. 341–352.

[13] G. Oligeri, S. Sciancalepore, and R. Di Pietro, "GNSS spoofing detection via opportunistic IRIDIUM signals," in *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '20)*. ACM, 2020, pp. 42–52.

[14] L. Wouters, "Glitched on earth by humans: A black-box security evaluation of the SpaceX starlink user terminal," in *Black Hat USA 2022*, 2022, conference talk and whitepaper.

[15] J. Willbold, M. Schloegel, R. Bisping, M. Strohmeier, T. Holz, and V. Lenders, "VSAsTer: Uncovering inherent security issues in current VSAT system practices," in *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '24)*. ACM, 2024, pp. 288–299.

[16] J. Smailes, E. Salkield, S. Köhler, S. Birnbach, and I. Martinovic, "Dishing out DoS: How to disable and secure the starlink user terminal," *arXiv preprint arXiv:2303.00582*, 2023.

[17] Amazon Web Services, "AWS Ground Station," https://aws.amazon.com/ground-station/, 2018, accessed: 2025-11-16.

[18] Microsoft Azure, "Azure Orbital Ground Station," https://azure.microsoft.com/en-us/products/orbital, 2022, accessed: 2025-11-16.

[19] T. Scharnowski, F. Buchmann, S. Wörner, and T. Holz, "A case study on fuzzing satellite firmware," in *Proceedings of the 2023 Workshop on the Security of Space and Satellite Systems (SpaceSec'23)*. San Diego, CA, USA: Internet Society, 2023. [Online]. Available: https://www.ndss-symposium.org/wp-content/uploads/2023/06/spacesec2023-230707-paper.pdf

[20] G. Marra, U. Planta, P. Wüstenberg, and A. Abbasi, "On the feasibility of cubesats application sandboxing for space missions," *arXiv preprint arXiv:2404.04127*, 2024.

[21] N. Yadav, F. Vollmer, A. Sadeghi, G. Smaragdakis, and A. Voulimeneas, "Orbital shield: Rethinking satellite security in the commercial off-the-shelf era," in *2024 Security for Space Systems (3S)*. IEEE, 2024, pp. 59–69.

[22] A. Young, C. Kitts, M. Neumann, I. Mas, and M. Rasay, "Initial flight results for an automated satellite beacon health monitoring network," in *Proceedings of the 24th Annual AIAA/USU Conference on Small Satellites*. Logan, UT, USA: AIAA/USU, 2010, sSC10-XII-1.

[23] E. J. Wyatt, M. Foster, A. Schlutsmeyer, R. Sherwood, and M. K. Sue, "An overview of the beacon monitor operations technology," in *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, 1997.

[24] The Aerospace Corporation, "Space Attack Research & Tactic Analysis (SPARTA)," https://sparta.aerospace.org/, 2025, accessed: 2025-11-16.

[25] The MITRE Corporation, "MITRE ATT&CK®," https://attack.mitre.org/, 2025, accessed: 2025-11-16.

[26] B. Bailey, "Cybersecurity protections for spacecraft: A threat based approach," The Aerospace Corporation, Tech. Rep. TOR-2021-01333-Rev A, Apr. 2021, https://aerospace.org/paper/cybersecurity-protections-spacecraft-threat-based-approach.

[27] B. Rhodes, "Skyfield: High precision research-grade positions for planets and earth satellites generator," https://ascl.net/1907.024, 2019, astrophysics Source Code Library, record ascl:1907.024.

[28] D. M. Sawyer and J. I. Vette, "Ap-8 trapped proton environment for solar maximum and solar minimum," NASA Goddard Space Flight Center, Tech. Rep. NSSDC/WDC-A-R&S 76-06, 1976.

[29] National Aeronautics and Space Administration, "Nasa risk management handbook," National Aeronautics and Space Administration, Tech. Rep. NASA/SP-2011-3422, Nov. 2011, version 1.0. https://www.nasa.gov/wp-content/uploads/2023/08/nasa-risk-mgmt-handbook.pdf.