

Why is Space Cybersecurity Unique?

Rajiv Thummala

Sibley School of MAE, Cornell University
rkt34@cornell.edu

Eric Rice

Jet Propulsion Laboratory, California Institute of Technology
eric.b.rice@jpl.nasa.gov

Gregory Falco

Sibley School of MAE, Cornell University
gfalco@cornell.edu

Abstract—As space systems become critical infrastructure, they have attracted increasing attention from the cybersecurity community. This paper argues that securing spacecraft requires a mission-centric cybersecurity paradigm that treats mission continuity and availability as first-order design and security primitives, rather than adapting practices from terrestrial systems. We identify seven constraints that shape the space security problem: mission-specific designs that prevent standardization, physics that couples software to irreversible orbital dynamics, permanent loss of hardware access after launch, communication gaps that mandate autonomous decisions, environmental degradation of electronics, tight subsystem dependencies that enable cascading failures, and governance pressures that constrain feasible security architectures. None of these dimensions is unique in isolation, but their simultaneous presence and coupling produces a distinct security problem and design space.

I. INTRODUCTION

Tens of thousands of broadband low Earth orbit (LEO) satellites are expected to be launched by the end of this decade [1]. Spacecraft have shifted from isolated scientific instruments to critical components of global communications, navigation, defense, and commerce. At the same time, offensive cyber capabilities against space infrastructure are maturing [15] and policy documents now explicitly call out space system cybersecurity as a national priority [2].

Cyber operations offer strategic advantages over kinetic actions in the space domain that shape both adversary incentives and defensive requirements [19]. Unlike kinetic attacks, cyber attacks can enable disabling spacecraft with fewer immediate consequences on the orbital environment, avoiding long-term degradation of shared orbital regimes and the political and operational costs associated with debris creation. Cyber effects can often be achieved without permanent destruction of the target, preserving deniability and enabling plausible non-repudiation, as attribution in space cyber incidents is complicated by shared infrastructure, distributed ground segments, and the difficulty of forensic validation under delayed and incomplete telemetry. The technical and economic barriers

to cyber operations are also substantially lower than those required for kinetic attacks. Cyber operations can be conducted remotely, leverage commercially available hardware and software expertise, and scale with comparatively modest resources. These characteristics make cyber operations a comparatively accessible and attractive means of influencing space missions, reinforcing the need for cybersecurity architectures that prioritize resilience, detection, and mission continuity over assumptions of physical exclusion or deterrence.

This raises a critical question: are conventional cybersecurity practices sufficient for hardening space missions, or do they require a mission-centric cybersecurity paradigm?

Space missions optimize risk toward preserving vehicle health and availability to execute mission objectives. Mission continuity and availability are therefore first-order design and security primitives, not secondary operational concerns. Much of existing cybersecurity research approaches satellites as conventional embedded systems, adapting techniques and standards from enterprise IT and industrial control systems (ICS) with minimal customization. This paper argues that such adaptation is insufficient. Space missions combine constraints that are individually familiar from other domains but, in aggregate and interaction, reshape the security problem. We organize these pressures into seven mission-level principles: mission-specific heterogeneity, physics coupling between software and orbital dynamics, immutability arising from lack of physical access, mandatory autonomy under intermittent and delayed communication, a harsh space environment that continuously perturbs hardware and state, tight interdependence among subsystems, and governance and economic forces that constrain feasible security architectures.

We dissect each of the seven principles in the following sections to illuminate their security implications and how their combination produces a problem space that differs from terrestrial cyber-physical security in important, non-separable ways.

II. LIMITS OF TRANSFERRING ENTERPRISE AND ICS SECURITY MODELS TO SPACE MISSIONS

A. Mission-Driven Risk Optimization and Deterministic Computing

Enterprise IT systems are general purpose and emphasize throughput, flexibility, and rapid patching. Space missions

prioritize deterministic behavior, bounded latency, and high reliability to protect vehicle health and mission availability. Flight software runs on real-time operating systems with carefully managed timing, limited instrumentation, and strict safety envelopes [20]. Logging, monitoring, and dynamic reconfiguration are constrained by bandwidth and power budgets, so security mechanisms must align with mission continuity rather than assume always-on visibility or rapid human intervention. These mission-driven choices produce a compute stack that is structurally different from enterprise IT and conventional embedded systems.

B. Heterogeneity Across Mission Architectures

Cyber-physical systems on Earth, while mission specific, benefit from significant reuse. For example, industrial control environments in power or water utilities frequently deploy large fleets of the same programmable logic controller (PLC) and remote terminal unit (RTU) classes, follow recurring zoning patterns, and rely on vendor hardening guides and certification schemes. Once a secure configuration, monitoring profile, or patching practice is validated for one facility, it is often transferable, with limited modification, to many others in the same sector.

In contrast, spacecraft architectures exhibit structural heterogeneity across the entire stack. Each spacecraft is a bespoke integration defined by a vector of design drivers: mission objective, orbital regime, radiation environment, communications geometry, mission-unique geometric constraints for safety and performance, autonomy requirements, payload criticality, and size, weight, power, and cost (SWaP-C) constraints. Platforms at opposite ends of this design space differ by orders of magnitude in mass, power, lifetime, and risk tolerance. A short lived 3U CubeSat in LEO may be built around commercial off-the-shelf (COTS) components with minimal redundancy and a single ground station. A long lived navigation satellite in medium Earth orbit (MEO) or a cislunar relay may use radiation hardened parts, redundant subsystems, and a global ground network. They are built by different organizations, using different component sets, under different qualification standards, and they assume different operational concepts.

A CubeSat may employ symmetric key cryptography [3] with pre-shared Advanced Encryption Standard (AES)-128 keys stored in electrically erasable programmable read-only memory (EEPROM), omit key rotation due to limited computation and brief mission lifetime, and accept commands from a single ground station. A high value navigation satellite may implement asymmetric cryptography with hardware security modules, hierarchical key management, redundant backups, and periodic updates from multiple authorized stations across a fifteen year mission. A cislunar satellite may use time-bound authentication tokens validated autonomously onboard, with an onboard certificate authority that manages trust relationships while accommodating delay-tolerant networking latencies. These are architecturally incompatible designs, not configurable variants. Command authentication illustrates how this heterogeneity precludes a common baseline. The design

space for space link security spans symmetric and asymmetric cryptography, pre-shared and dynamically managed keys, autonomous trust validation, and delay-tolerant authorization schemes, as surveyed in the CCSDS Space Link Security Green Book [4].

These mission-specific architectures imply that security baselines must be mission-tailored, and that policy guidance should emphasize profiles and assurance objectives tied to mission continuity rather than uniform checklists. Designing for availability and continuity becomes part of the security architecture, not a post hoc operational concern.

III. PHYSICS AND ORBITAL MECHANICS

Conventional cyber-physical security treats the physical process as the target of compromise but not as a determinant of the security context itself [6]. In industrial control systems, vehicles, or medical systems, physics governs the effects of malicious commands, yet network reachability, authentication structure, and access control are largely independent of the physical system state. Orbital mechanics and attitude dynamics couple directly into the security model, so the spacecraft physical state becomes an input to the attack surface and to the feasible defensive mechanisms [5][8].

A. Orbital Dynamics and Contact Geometry

Orbital motion sets basic communication geometry. For a satellite in a circular orbit with semi major axis a around a central body with gravitational parameter μ , the mean motion and period are:

$$n = \sqrt{\frac{\mu}{a^3}}, \quad T = \frac{2\pi}{n} = 2\pi\sqrt{\frac{a^3}{\mu}}.$$

The Earth central angle ψ_{\max} over which a ground station can see a satellite at altitude h satisfies:

$$\cos \psi_{\max} = \frac{R_E}{R_E + h},$$

where R_E is the Earth radius. The corresponding contact duration with one station is approximately:

$$t_{\text{vis}} \approx \frac{2\psi_{\max}}{n}.$$

Reachability and contact duration are therefore functions of orbital radius design, not of administrator policy. As the Earth rotates and the satellite progresses along its orbit, the set of ground stations with line of sight, sufficient elevation, and adequate link budget changes deterministically over time. The space system attack surface is time-varying and state dependent: during one pass only an authorized ground station may be in view, during a later pass that same footprint may include both a trusted site and a hostile uplink. Figure 1 illustrates how geometry bounds when communications can occur, which in turn determines both when interactive security functions can execute and when attacks may be attempted.

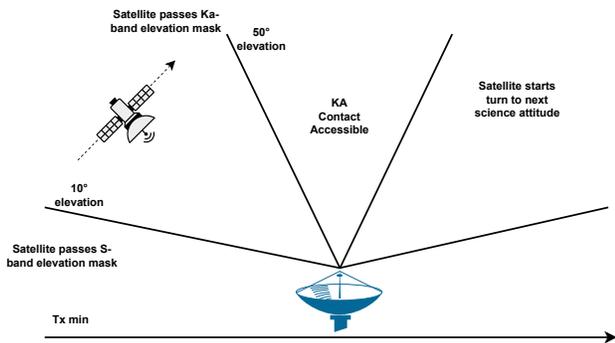


Fig. 1. A satellite is only able to communicate with ground resources while it is in view of a ground asset, subject to constraints imposed by both the telecommunications design and the orbital configuration. The figure illustrates an example vehicle equipped with both S-band and Ka-band links, each with a distinct elevation mask, θ_{mask} . The vehicle is accessible to ground assets on a given frequency band only for a duration that depends on the maximum visibility angle, ψ_{max} , for that band and the orbital velocity of the vehicle.

B. State Estimation, Sensor Geometry, and Security Logic

Attitude and sensor geometry introduce a second layer of coupling. High gain antennas and narrow field of view sensors impose strict pointing requirements, so small errors in attitude can break communication or corrupt measurements. State estimation combines orbital knowledge, via GPS or GNSS measurements or ground-uploaded ephemerides, attitude measurements, inertial measurements, and other environmental sensing to produce an estimate $\hat{x}(t)$ of the true state $x(t)$, where

$$x(t) = \begin{bmatrix} r(t) \\ v(t) \\ q(t) \\ \omega(t) \end{bmatrix},$$

with position r , velocity v , attitude q , and angular rate ω . Autonomy logic, actuator limits, safe mode triggers, and policy checks such as permitting station keeping only within a certain orbital slot are evaluated against $\hat{x}(t)$, not against the unknown true state. An adversary that biases any component of the measurement set can steer $\hat{x}(t)$ away from $x(t)$, and in doing so manipulate the security logic itself. Sensor spoofing affects safety and quality of control in both terrestrial and space systems. In space missions, however, corrupted knowledge of vehicle state can propagate into mission-ending consequences if not detected within limited contact and recovery windows.

C. Irreversibility, Propagation, and Detectability

Orbital dynamics and communications availability drive which security mechanisms can be applied to the communications link, and when they can be applied. For a LEO satellite with $T \approx 90$ minutes, typical single site visibility t_{vis} is on the order of several minutes. All interactive security operations, including session establishment, mutual authentication, command validation, key updates, and initiation of software uploads, must complete within these intervals while competing with mission data downlink and health telemetry. Antenna

acquisition and signal lock consume a significant fraction of t_{vis} , leaving limited time for any protocol that requires multiple round trips. In deep space, light time delay dominates. The one way propagation time

$$\tau = \frac{R}{c}$$

for range R and speed of light c is measured in minutes to hours, which eliminates interactive challenge response and forces pre-authorization of extended autonomous behavior. The content and duration of that pre-authorization envelope are dictated by orbital mechanics rather than by conventional cybersecurity design preference.

The same dynamics that constrain communication also shape how attacks manifest and how detectable they are. The linearized relative motion between two nearby trajectories can be described by Clohessy–Wiltshire equations for circular reference orbits. For example, along track separation $y(t)$ in the Hill frame evolves approximately as

$$y(t) \approx \frac{2\Delta v}{n} (1 - \cos nt)$$

for a small along track velocity impulse Δv at $t = 0$. A small unauthorized Δv can therefore generate kilometer scale along track separation over weeks, while remaining comparable to cumulative effects of drag, J_2 perturbations, and station keeping activity. Small attitude or orbital perturbations that remain within environmental variability may be operationally irrelevant if mission margins can absorb them. From a mission-centric perspective, security concern arises not from the existence of an attacker, but from whether induced deviations consume margins, shorten mission lifetime, or force protective responses that disrupt mission objectives.

Orbital mechanics also imposes irreversibility and propagation. Velocity changes can only be counteracted by expending finite propellant, so many consequences of compromise are effectively irreversible once fuel, thermal margins, or power reserves are exhausted. Misalignment can lead to battery depletion that cannot be recovered. Unauthorized thrusting can move a satellite into an orbit that is geometrically or energetically unreachable with remaining fuel. Collisions generate debris fields, and the population of fragments $N(t)$ can grow as a function of time and collision rate, with models such as Kessler’s suggesting the possibility of self reinforcing cascades that increase debris density faster than it can be mitigated [7]. In contrast, industrial systems permit shutdown, isolation, repair, and replacement.

Physics therefore makes security design inseparable from mission dynamics. Defensive strategies must account for time-varying reachability, state-dependent authorization, and irreversibility, and policy should treat orbital context as a first-class input to access control and recovery planning.

IV. ACCESS

The distinguishing feature of spacecraft relative to terrestrial systems is not simply distance, it is immutability. After launch, the hardware stack, trust anchors, and low level security

architecture become physically unmodifiable for the remainder of the mission [9]. For most spacecraft, there is no realistic path to swap boards, attach probes, or rework the security architecture once the vehicle leaves the launch pad.

A. Hardware Immutability and Frozen Trust Anchors

Hardware security mechanisms are largely fixed at integration and launch. The presence, placement, and trust boundaries of cryptographic accelerators, secure boot logic, hardware security modules, and radiation-hardened processors are determined prior to deployment and cannot be altered in orbit. While some parameters, such as keys or supported algorithms, may be updated through software if explicitly designed for, all such updates remain constrained by the original hardware security envelope. After encapsulation and launch, there is no practical means to replace a hardware security module, introduce a new secure element, or retrofit a fundamentally different root of trust. This constraint may change in the future as on-orbit servicing, assembly, and manufacturing (OSAM) missions [10] enable scaled hardware modifications in orbit.

B. Software Update Constraints and Nuance

Software modification is possible but tightly constrained by both link budget and flight software architecture. Consider a CubeSat with uplink rate $R = 9.6$ kb/s and daily contact $t = 600$ s. The theoretical upper bound on daily upload volume is

$$D_{\max} = \frac{R \cdot t}{8} \approx \frac{9.6 \times 10^3 \cdot 600}{8} \approx 7.2 \times 10^5 \text{ bytes,}$$

on the order of one megabyte per day before protocol overhead and competition with mission traffic. Contemporary small satellites typically fly a monolithic flight software image or a tightly coupled set of statically linked tasks under a real-time operating system [21]. The bootloader expects a single verified image in nonvolatile memory, with limited or no support for dynamic linking, module hot swapping, or on-orbit repartitioning of storage. Updating a nontrivial subsystem therefore requires transferring either a full image or a large, carefully constructed differential update relative to the existing onboard image, staging it in flash, and performing an all or nothing swap at reboot. Given that full images are often tens of megabytes, such an update spans many passes and remains vulnerable to interruption, link errors, and power transients.

These architectural choices interact with the asymmetric risk profile. A failed patch on a terrestrial embedded device typically bricks hardware that can be replaced or reimaged. A failed in-orbit update can strand the spacecraft in an unbootable state with no physical recovery path. To mitigate this, operators introduce conservative mechanisms such as dual image banks, staged activation, rollback timers, and strict ground validation, but these add complexity, consume scarce memory, and are not universally implemented, especially on cost-constrained platforms. In all cases, software updates remain bounded by the original hardware trust anchors: fixed cryptographic accelerators, secure boot chains, and key storage

layouts cannot be altered post-launch, so even successful on-orbit reprogramming cannot upgrade the underlying root of trust.

C. Incident Response Without Physical Access

Diagnosis and recovery of incidents are similarly constrained. Telemetry for forensics is sparse and delayed, and safe mode behavior is intentionally conservative, prioritizing vehicle survival over detailed observation [18]. Recovery often consists of command resets, mode transitions, or selective subsystem power cycling rather than deep inspection. These constraints limit post-incident evidence collection and make it difficult to validate whether recovery actions restored a trusted state or merely preserved availability.

Consider the February 2022 cyberattack on Viasat’s KA-SAT network [15], which illustrates how incident response for space systems is shaped by the absence of immediate physical access and the scale of remote impact. While the attack did not affect the space segment itself, it exploited trusted ground-segment access to issue legitimate management commands that remotely wiped the flash memory of tens of thousands of user terminals, rendering them inoperable. From an incident response perspective, forensic validation required the physical recovery and shipment of compromised terminals from multiple countries before low-level analysis could be performed using JTAG interfaces and modified bootloaders [15]. Even then, responders faced uncertainty about the provenance and integrity of the recovered evidence due to the delay between compromise and inspection [15].

Immutability makes long horizon threat prediction a primary design constraint. A spacecraft launched in 2025 with hardware support for RSA 4096 may appear conservative at deployment. Over a fifteen year mission, advances in cryptography, implementation attacks, or practical quantum capabilities can erode that margin. Without a path to replace hardware roots of trust or to introduce fundamentally different primitives, such as post-quantum schemes, the vehicle remains bound to its original choices. Terrestrial systems can rotate hardware, deploy new accelerators, and retire insecure devices.

Access constraints therefore require security design to assume fixed trust anchors, limited forensic visibility, and risky update paths. Defensive policy should emphasize pre-launch assurance, robust rollback strategy, and explicit mission tradeoffs between availability and the depth of post-incident recovery.

V. AUTONOMY

A. Why Autonomy Is Mandatory

While access immutability fixes the hardware and trust anchors for the mission lifetime, autonomy concerns how security decisions are made during operations when humans cannot be in the loop. Spacecraft operate under communication regimes that make extensive onboard autonomy unavoidable. A typical LEO satellite completes an orbit in approximately 90 minutes, with per station visibility often limited to 10–15 minutes [12]. Geostationary orbit (GEO) platforms maintain

continuous line of sight, but incur one way propagation delays of roughly 240 ms [13], which constrains tight closed loop control of fast dynamics. As a result, spacecraft must function without operator input for most of each orbit in LEO, for extended intervals in GEO, and for weeks or longer on interplanetary trajectories. During these windows spacecraft must autonomously execute the mission, maintain attitude and orbit, operate payloads, manage power and thermal states, and respond to anomalies, assuming long ground-in-the-loop response times. Security mechanisms are subject to the same constraint.

B. Autonomous Security Decision-Making Under Uncertainty

Terrestrial environments, even when highly automated, generally retain the possibility of timely human escalation or physical intervention. Space systems do not. Authentication outcomes, anomaly detections, and threat responses often occur when no ground station is in view. Consider a spacecraft approaching an imaging target that requires a maneuver within a five minute window. A command arrives but fails cryptographic validation. There is insufficient time for a full human in the loop investigation across multiple passes. The onboard logic must decide whether to reject the command, accept it under constrained parameters, or transition to a more conservative mode that preserves safety at the cost of losing the opportunity [18]. That decision must be made autonomously using pre-programmed rules that encode tradeoffs between security risk, mission value, and operational context.

Anomaly detection exhibits similar pressures. Telemetry such as a sudden increase in processor load, unexpected internal traffic, or a new process identifier may indicate malware, a software defect, or a benign configuration change [14]. On the ground, security operations centers can aggregate logs, correlate events across systems, and consult subject matter experts over minutes to hours. A spacecraft has limited instrumentation, constrained telemetry, and seconds to minutes in which to act during time critical phases. An onboard intrusion prevention system must choose between continuing operations while risk remains unresolved, or applying mitigations that may interrupt payload activities or affect control loops. False negatives allow a potential compromise to persist through an entire contact gap. False positives forfeit observation opportunities and may consume scarce resources such as propellant or power margins.

Ground directed response is quantized by orbital dynamics. If an attack or anomaly is first detected near the start of a communications blackout, telemetry describing the event will not reach the ground until the next visibility window, after a delay on the order of one orbital period T_{orbit} . Analysis and preparation of corrective commands add processing time t_{analysis} , and execution typically waits for a subsequent pass. A simple lower bound on response time is

$$t_{\text{response}} \geq T_{\text{orbit}} + t_{\text{analysis}} + T_{\text{orbit}},$$

which is measured in hours for LEO and longer for higher orbits. During this interval, the spacecraft must continue

operating under whatever pre-programmed security posture and countermeasures were in place at the moment of detection.

C. Autonomy as Both Attack Surface and Defensive Opportunity

Autonomy expands the attack surface by increasing the authority and complexity of onboard decision logic, and by concentrating security-relevant choices in software that is difficult to validate exhaustively. At the same time, autonomy is a defensive opportunity because it can encode mission-specific policy, enforce conservative safety envelopes, and apply local anomaly responses when the ground is unavailable. These dual roles mean autonomy must be treated as a security-critical subsystem with explicit assumptions about decision thresholds, failover behavior, and authority limits.

Autonomy in space security is therefore not an optimization choice, it is a structural requirement imposed by latency, contact geometry, and the inability to intervene physically on demand. Security design must focus on predictable autonomous behavior and governance should set expectations for when onboard policy may trade mission performance for continuity and safety.

VI. SPACE ENVIRONMENT

The space environment imposes continuous, unavoidable physical stress on spacecraft electronics. Ionizing radiation, thermal cycling, electromagnetic interference (EMI), and constrained power availability act as a persistent source of faults rather than rare exceptions. As a result, mechanisms traditionally treated as reliability features, error detection and correction, watchdog timers, power cycling, and thermal management, become security critical. They determine whether cryptographic material remains intact, authentication logic behaves correctly, and security services remain available.

A. Radiation Effects on Security-Relevant State

Ionizing radiation is the dominant source of random faults. Galactic cosmic rays, solar particle events, and trapped belt particles induce single event upsets in storage and logic, invert bits in configuration registers, and can trigger latchups. A single bit flip in a stored AES key, a key index, or a command authentication flag can silently corrupt cryptographic state and cause valid traffic to be rejected or malicious traffic to be accepted [22]. Over years, total ionizing dose increases device error rates [23] and can push cryptographic hardware toward marginal operating regimes. Error detecting and correcting codes, memory scrubbing, hardened state encodings, and periodic key verification, often framed as reliability techniques, are therefore essential to prevent environmental faults from manifesting as undetected key corruption or altered control flow in security relevant state machines. Table I summarizes representative failure modes and their security impact.

B. Thermal, Power, EMI, and Survivability Constraints

Thermal and power conditions further constrain security. LEO platforms experience repeated transitions between sunlight and eclipse, with large temperature excursions each orbit.

TABLE I
EFFECTS OF RADIATION-INDUCED SINGLE EVENT UPSETS ON CRYPTOGRAPHIC OPERATIONS

SEU location	Bit-level effect	Cryptographic state change	Resulting security behavior
AES key register	Single key bit k_i flips so $b' = b \oplus 1$	Key becomes $K' = K \oplus \Delta_i$, where Δ_i has a single bit set	Decrypt and MAC checks fail, valid traffic is rejected, and the spacecraft may trigger fault protection due to apparent anomalies
Key index or key selector register	Index field changes from j to j'	Device selects incorrect key $K_{j'}$ for authentication or encryption	Authenticated commands fail, or, if a default/test key is selected, malicious traffic may be accepted as valid
Authentication mode or policy flag	Control bit toggles (e.g., auth_required : 1 \rightarrow 0)	Authentication state machine transitions into an unintended branch	Checks are bypassed or previous failures interpreted as success, enabling execution of unauthorized commands
Replay/nonce counter	Counter bit flips $ctr' = ctr \oplus \Delta$	Monotonicity or uniqueness property violated	Replayed commands appear fresh, or desynchronization drops valid traffic, degrading mission operations

Thermal gradients and cycling accelerate degradation of interconnects and can produce intermittent faults in devices such as hardware security modules and secure boot components. EMI from onboard radios, power converters, and payloads can induce transient faults that mimic or mask cyber anomalies [24]. At the same time, battery capacity and solar array output vary with temperature and aging, forcing dynamic load shedding. When available power drops below a threshold, the system must prioritize essential attitude and power control over ancillary functions. If cryptographic acceleration, continuous monitoring, or logging are categorized as lower priority loads, they will be reduced or disabled precisely when the platform is under the greatest environmental stress. Power management policies and thermal control loops thus implicitly define which security functions are guaranteed and which degrade under duress.

C. Fault Recovery and Security State Preservation

These environmental effects couple spacecraft fault detection, isolation, and recovery (FDIR) tightly to cybersecurity. FDIR mechanisms such as watchdog resets, autonomous power cycling of subsystems, and configuration rollbacks must preserve security invariants across faults and recovery actions. After a reset, boot logic must reestablish a known good cryptographic state rather than enter a permissive or unauthenticated mode. Power cycling a processor to clear a latchup must not erase integrity information while leaving corrupted configuration in place. Error correcting code (ECC) and scrubbing policies must be tuned so that key stores and authentication code paths receive at least the same level of protection as other safety critical data.

Environmental constraints therefore bind reliability engineering to security design. Defensive strategy should explicitly budget for radiation effects on security state, and policy should require that FDIR and power management preserve authentication and integrity guarantees during stress.

VII. INTERDEPENDENCE

Spacecraft architectures exhibit extreme interdependence among subsystems, creating cascading failure pathways

wherein compromise of a single component can precipitate total mission loss. Rather than modular isolation, mission success depends on integrated performance across tightly coupled, multi-disciplinary, multi-physics behaviors. This architectural convolution transforms localized security breaches into systemic failures, as no subsystem operates in isolation.

A. Integrated Performance and Cascading Failure

Consider the fundamental dependencies in a representative satellite architecture. The attitude determination and control system maintains spacecraft orientation, pointing solar arrays toward the Sun and communication antennas toward Earth. This system depends on the onboard computer for command processing, the power subsystem for actuator operation, and the communication subsystem for ground commands. Conversely, the power subsystem requires proper solar array pointing from attitude control to generate electricity, thermal control to prevent battery degradation, and the onboard computer to manage charge cycles. The communication subsystem demands attitude control for antenna pointing, power for transmission, and the onboard computer for protocol execution. The onboard computer itself requires power, thermal regulation to prevent processor failure, and communication links for command reception and telemetry transmission.

This circular dependency structure means compromise of any subsystem degrades or destroys others [16]. An attacker gaining access to the attitude control system can mispoint solar arrays, depleting batteries and disabling power-dependent subsystems including communication, computation, and thermal management. Battery depletion forces the spacecraft into safe mode or causes permanent shutdown. Even if power is later restored, thermal extremes experienced during the blackout may have irreversibly damaged electronics. Alternatively, compromising the onboard computer enables manipulation of power management logic, deliberately draining batteries during eclipse periods when solar generation is unavailable, yielding the same cascading failure without directly attacking attitude control.

B. Coupled Thermal and Propulsion Dependencies

Thermal control dependencies create additional vulnerability pathways. Most spacecraft lack active cooling, relying instead on passive radiators, heaters, and thermal coatings to maintain component temperatures within operational ranges. The thermal subsystem depends on power for heater operation and on attitude control to orient radiators correctly. Compromising either enables thermal attacks. Disabling heaters during eclipse allows batteries to freeze, permanently reducing capacity. Misorienting radiators causes processors to overheat and throttle or fail. Once thermal limits are exceeded, physical damage occurs within minutes to hours, rendering subsequent security responses futile. The spacecraft may remain in orbit but ceases to function.

Propulsion systems demonstrate similar interdependence. Station-keeping maneuvers require the onboard computer to calculate burn parameters, the attitude control system to orient the spacecraft correctly, the power system to energize valves, and the communication system to receive maneuver commands from ground control. An adversary need not directly access propulsion hardware. Corrupting the orbit determination algorithm in the onboard computer suffices to generate incorrect burn calculations. Mispointing the spacecraft during a burn wastes fuel and potentially places the satellite on a hazardous trajectory. Delaying valve commands causes late ignition, again wasting fuel. As propellant is finite and cannot be replenished, even minor inefficiencies accumulate over the mission lifetime, shortening operational duration or forcing premature decommissioning.

C. Communication Boundary and Asymmetric Redundancy

Communication subsystem compromise serves as a particularly effective attack vector due to its external interface [16]. Ground stations transmit commands and receive telemetry, making the communication link the primary cyber boundary. Successful command injection or protocol exploitation grants adversaries access to the onboard computer's command interpreter [17]. From this foothold, attackers can issue commands to any subsystem: reorient the spacecraft, modify power allocations, reprogram autonomous sequences, or upload malicious flight software. The communication subsystem functions as the gateway through which all other subsystems become accessible, yet it must remain externally reachable to fulfill its mission function [17]. Hardening communication security without degrading functionality presents a fundamental design tension.

Redundancy, a standard terrestrial mitigation, is constrained in spacecraft by mass and cost budgets. Critical components such as reaction wheels, star trackers, and transponders may have redundant units, but full subsystem duplication is rare. A CubeSat typically flies with minimal redundancy if at all. Larger satellites may have dual-redundant flight computers but single-string power converters, propulsion valves, or attitude sensors. Partial redundancy creates asymmetric failure modes: losing a redundant computer enables graceful degradation, but

losing the sole power converter causes immediate mission termination.

D. Autonomy and FDIR as Amplifiers

Autonomy exacerbates interdependence vulnerabilities. Automated FDIR routines reside in the onboard computer and depend on telemetry from sensors distributed across all subsystems. An attacker manipulating sensor data can induce false FDIR triggers, causing the spacecraft to unnecessarily swap to backup hardware, enter safe mode without cause, or ignore genuine faults [17]. Spoofing a temperature sensor to report overheating disables heaters and freezes batteries. Spoofing a battery voltage sensor to report full charge prevents recharging during sunlight periods, eventually depleting reserves. FDIR becomes an attack amplifier rather than a defense.

E. Domain-Specific Attack Surfaces

Space-specific attack surfaces are shaped by access patterns that differ from terrestrial IT. Command paths are time-bounded, link budgets are narrow, and ground and relay infrastructure are integral to mission operations. As a result, adversaries can target ground networks, uplink protocols, crosslinks, and payload commanding in ways that are tightly coupled to orbital schedules and mission phases. The SPARTA taxonomy [17] provides a grounded catalog of these tactics and techniques, highlighting how the space domain blends cyber, physical, and operational entry points.

Interdependence complicates isolation assumptions under mission conditions, but does not invalidate existing reliability and fault management practices. Space vehicles are inaccessible, tightly integrated, and extensively tested prior to launch, which limits attack vectors as much as it creates them. In practice, the primary digital interface to the vehicle remains the command path implemented through flight software. Many failure modes induced by malicious behavior are already mitigated by conservative reliability design and FDIR mechanisms intended to handle unknown faults. Advancing secure design for space systems therefore requires identifying where cybersecurity integrates with, reinforces, or stresses existing reliability practices, rather than reimagining reliability engineering as a separate security discipline.

VIII. GOVERNANCE

Spacecraft development operates under governance frameworks that create fundamental tensions between security requirements and operational realities. Three governance challenges compound the technical constraints already established: interoperability with legacy controls designed for previous threat landscapes, cost pressures driving adoption of commercial off-the-shelf components with inadequate security assurance, and security requirements that hinder innovation by imposing restrictions incompatible with rapid development cycles. These governance constraints create a disconnect between security research and operational practice, and they directly shape mission continuity by defining which security tradeoffs are acceptable. This discussion is synthesized from industry

and government findings documented in the NSC and Office of the National Cyber Director 2025 report on space system cybersecurity [2].

Interoperability with legacy controls presents immediate operational barriers. Spacecraft missions span decades, creating fleets of satellites designed under different security paradigms operating simultaneously. A constellation operator managing satellites launched in 2010, 2020, and 2025 confronts incompatible command authentication protocols, disparate key management infrastructures, and heterogeneous cryptographic algorithms across the fleet. Modern satellites implementing post-quantum cryptography must maintain backward compatibility with legacy satellites using RSA-2048, forcing operators to maintain dual security infrastructures. Space and cyber experts broadly agree that legacy space systems, many operating decades beyond expected design life, were generally built without cybersecurity in mind [2]. These legacy systems use legacy coding languages including C and C++, which drives continued use of these languages across the industry despite known vulnerabilities. Industry broadly recommends government take a forward-focused rather than retroactive approach to space system cybersecurity requirements [2].

Cost pressures fundamentally constrain security design choices. Security hardware such as dedicated cryptographic processors, hardware security modules, and tamper-resistant enclosures add mass, consume power, and increase procurement costs. A radiation-hardened processor with formal security certification may cost 100 times more than a commercial equivalent, forcing mission planners to choose between mission capability and security robustness. Commercial off-the-shelf components offer significant cost savings but lack the security validation, supply chain provenance, and long-term support guarantees of space-qualified hardware. This economic pressure makes COTS components increasingly prevalent in cost-constrained missions. Industry representatives noted that smaller firms perceive a resource trade-off between space mission objectives and cybersecurity, with demonstrating viability to investors being the primary motivation in product design or operations rather than prioritizing cybersecurity [2]. However, these same satellites operate in shared orbital regimes and communicate through ground station networks potentially interconnected with critical systems, creating systemic risk propagation.

Security requirements imposed by regulatory bodies and government customers frequently hinder innovation by demanding security controls incompatible with modern development practices. Space system development teams see cybersecurity as detracting from core space missions and an impediment to technical innovations required for fielding competitive and operationally viable products [2]. Organizations silo aerospace and computer engineering workforces, preventing collaboration on cross-cutting challenges like space system cybersecurity. The research-practice disconnect manifests most acutely in cryptographic agility requirements. Security researchers advocate for cryptographic agility enabling satellites to transition between algorithms as threats evolve,

but implementing agility requires additional code complexity, increased attack surface, and non-trivial validation burden. Industry representatives expressed that large elements of the space systems ecosystem operate without cybersecurity sensors tailored for the on-orbit environment, and that developing lightweight sensors suitable for resource-constrained environments would require significant research and funding [2]. The additional complexity contradicts aerospace engineering principles favoring simplicity and heritage software with extensive operational history. Mission teams prioritize flight heritage over theoretically superior but operationally unvalidated security mechanisms.

These governance challenges create systemic vulnerabilities that technical security solutions cannot address. A perfectly secure satellite design remains vulnerable if budget constraints force adoption of insecure COTS components, if legacy interoperability requirements mandate weak cryptographic protocols, or if security requirements prevent deployment of innovative operational practices. Space cybersecurity governance must evolve to accommodate mission diversity, enable rapid security response, and align economic incentives with security outcomes. Governance choices therefore function as security design inputs, and policy should explicitly weigh mission continuity and availability against compliance burden and cost.

IX. CONCLUSION

This paper argued that space mission cybersecurity is not a straightforward extension of terrestrial cyber-physical security, but is shaped by seven constraints that, while individually familiar from other domains, combine to create a distinct security problem. Terrestrial security frameworks that rely on standardization, physical access, human oversight, iterative refinement, and strong isolation boundaries cannot be applied wholesale to this setting. Security for space missions must instead be derived from first principles for each mission, validated pre-launch against decade-scale threat projections, embedded as autonomous decision logic that is aware of dynamics and contact geometry, and designed for graceful degradation under simultaneous environmental stress and systemic cascade risk. Recognizing these structural constraints is a prerequisite for developing methods, tools, and governance models that treat mission continuity and availability as foundational design objectives rather than operational afterthoughts.

X. ACKNOWLEDGMENTS

Contributions to this work were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

REFERENCES

- [1] Z. M. Kassas, "Ad Astra, Navigation with Megaconstellation LEO Satellites," *IEEE Aerospace and Electronic Systems Society Distinguished Lecture*, online, Nov. 13, 2024. [Online]. Available at <https://ieeae-aess.org/event/lecture/ad-astra-navigation-megaconstellation-leo-satellites>
- [2] National Space Council and Office of the National Cyber Director, "Space System Cybersecurity: Space Industry Perspectives," *Executive Office of the President*, Washington, D.C., Jan. 2025.

- [3] C. Greenwood, B. Griswold, and K. Simmons, "Using Symmetric Encryption to Increase the Security of CubeSats," in *45th COSPAR Scientific Assembly*, vol. 45, July 2024, p. 2529. [Online]. Available: <https://ui.adsabs.harvard.edu/abs/2024cosp...45.2529G>
- [4] Consultative Committee for Space Data Systems (CCSDS), "Space Link Security," CCSDS 350.5-G-2 Green Book, July 2022. [Online]. Available: <https://ccsds.org/Pubs/350x5g2.pdf>
- [5] M. Calabrese and G. Falco, "Physics-Informed Satellite Cybersecurity," in *Computer*, vol. 57, no. 5, pp. 106-109, May 2024, doi: 10.1109/MC.2024.3374009. keywords: Satellites;Computer hacking;Intrusion detection;Computer security,
- [6] C. Neuman, "Challenges in Security for Cyber-Physical Systems," in *Proceedings of the DHS Workshop on Future Directions in Cyber-Physical Systems Security*, 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:15983288>
- [7] D. J. Kessler, N. L. Johnson, J. C. Liou, and M. Matney, "The Kessler Syndrome: Implications to Future Space Operations," *Advances in the Astronautical Sciences*, vol. 137, no. 8, p. 2010, 2010. Univelt, Inc.
- [8] K. Sun, Q. Zheng, and R. Jin, "Secure Control for Spacecraft Relative Dual Quaternion Systems with Denial-of-Service Attacks," *IEEE Transactions on Aerospace and Electronic Systems*, early access, pp. 1-11, 2025, doi: 10.1109/TAES.2025.3591386.
- [9] G. Falco, L. Korth, P. Custer, R. N. Schofield and C. Pockock, "How to Scrub a Launch: Spaceport Cybersecurity," 2023 IEEE 9th International Conference on Space Mission Challenges for Information Technology (SMC-IT), Pasadena, CA, USA, 2023, pp. 56-67, doi: 10.1109/SMC-IT56444.2023.00015. keywords: Industries;Systematics;Regulators;Space missions;Complexity theory;Ground support;Information technology;spaceport cybersecurity;spaceport security;spaceport;launch facilities;launch security;space cybersecurity;launch cybersecurity,
- [10] N. G. Gordon and G. Falco, "Reference architectures for autonomous on-orbit servicing, assembly and manufacturing (OSAM) mission resilience," *2022 IEEE International Conference on Assured Autonomy (ICAA)*, Fajardo, PR, USA, 2022, pp. 124-128, doi: 10.1109/ICAA52185.2022.00024.
- [11] Liu qian, Song Ningning and Zeng Wenlu, "Research of centralized multicast key management for LEO satellite networks," International Conference on Cyberspace Technology (CCT 2013), Beijing, China, 2013, pp. 577-582, doi: 10.1049/cp.2013.2154. keywords: LEO satellite networks;centralized;multicast key management;one-way function tree,
- [12] W. Harris, R. Blake, D. Woods, R. Thompson, and J. Stjernevi, "Mission Planning for Constellations," in *SpaceOps 2002 Conference*, Houston, TX, Oct. 2002. American Institute of Aeronautics and Astronautics. [Online]. Available: <https://arc.aiaa.org/doi/10.2514/6.2002-T3-04>, doi: 10.2514/6.2002-T3-04
- [13] T. Katoh, T. Saito, T. Yamashita, K. Yamakawa, M. Matsumoto, and H. Watanabe, "A Conceptual Study of the GEO-LEO High Data Rate Satellite Communications System," in *17th AIAA International Communications Satellite Systems Conference and Exhibit*, Yokohama, Japan, Feb. 1998. American Institute of Aeronautics and Astronautics. [Online]. Available: <http://arc.aiaa.org/doi/10.2514/6.1998-1260>, doi: 10.2514/6.1998-1260
- [14] Y. E. Çiloğlu and Ş. Bahtiyar, "A New Anomaly-Based Intrusion Detection System for MIL-STD-1553," 2023 10th International Conference on Recent Advances in Air and Space Technologies (RAST), Istanbul, Turkiye, 2023, pp. 1-6, doi: 10.1109/RAST57548.2023.10197927. keywords: Military standards;Space vehicles;Space technology;Intrusion detection;Machine learning;Military aircraft;US Department of Defense;MIL-STD-1553;intrusion detection;anomaly;avionics;machine learning,
- [15] N. Saunders, R. Thummala and G. Falco, "Space Cybersecurity Incident Response Framework: A Viasat Case Study," 2025 IEEE Aerospace Conference, Big Sky, MT, USA, 2025, pp. 1-15, doi: 10.1109/AERO63441.2025.11068784. keywords: Satellites;Protocols;Shape;Forensics;Software;Hardware;Geomagnetic storms;Cyberattack;Standards;Resilience,
- [16] G. Falco, R. Thummala and A. Kubadia, "WannaFly: An Approach to Satellite Ransomware," 2023 IEEE 9th International Conference on Space Mission Challenges for Information Technology (SMC-IT), Pasadena, CA, USA, 2023, pp. 84-93, doi: 10.1109/SMC-IT56444.2023.00018. keywords: Space vehicles;Codes;Satellites;Space missions;Space technology;Power system protection;Software;space cybersecurity;ransomware;satellite ransomware;satellite attack;satellite hi-jack;NASA;flight software;core Flight System;flight software security;space ransomware,
- [17] The Aerospace Corporation, *SPARTA: Space Attack Research & Tactic Analysis Matrix*, Available at: <https://sparta.aerospace.org/>, Accessed: 14 November 2025.
- [18] R. K. Thummala, G. J. Falco, B. E. Schake, D. L. Pollock, D. J. Melander, C. P. Banh, B. Pendleton, and R. A. Procell, "Mission aware cyber safe mode for spacecraft," Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States), Tech. Rep., 2025.
- [19] N. Tsamis, B. Bailey, and G. Falco, "Translating space cybersecurity policy into actionable guidance for space vehicles," in *ASCEND 2021*, 2021, p. 4051.
- [20] J. Curbo and G. Falco, "Testable Cyber Requirements for Space Flight Software," in *2025 IEEE Aerospace Conference*, 2025, pp. 1-20, doi: 10.1109/AERO63441.2025.11068629.
- [21] E. Birrane, K. Bechtold, C. Krupiarz, A. Harris, A. Mick, and S. Williams, "Linux and the Spacecraft Flight Software Environment," in *21st Annual AIAA/USU Conference on Small Satellites*, Logan, UT, 2007, Paper SSC07-XII-3.
- [22] V. Bandeira, J. Sampford, R. Garibotti, M. G. Trindade, R. P. Bastos, R. Reis, and L. Ost, "Impact of radiation-induced soft error on embedded cryptography algorithms," *Microelectronics Reliability*, vol. 126, Art. no. 114349, Sep. 2021, doi: 10.1016/j.microrel.2021.114349.
- [23] M. Liu, X. Xu, C. Zeng, and C. Xiong, "A study on the influence of dose rate on total ionizing dose effect of anti-fuse field programmable gate array—The irradiation damage is attenuated at low dose rate," *Frontiers in Physics*, vol. 10, Art. no. 1035846, Oct. 2022, doi: 10.3389/fphy.2022.1035846.
- [24] A. Loveless, L. T. X. Phan, R. Dreslinski, and B. Kasicki, "PCSP00F: Compromising the Safety of Time-Triggered Ethernet," in *2023 IEEE Symposium on Security and Privacy (SP)*, San Francisco, CA, USA, 2023, pp. 3193-3208, doi: 10.1109/SP46215.2023.10179318.