Engineering a 160-Node Distributed Payload in 3U

Authors

Khurram Hussain[1], Martin Platt[1], Dr Atma Prakash[2]

Corresponding author: Martin Platt (martin@ambasat.com)

[1] AmbaSat Ltd, Northallerton, North Yorkshire, United Kingdom

[2] Teesside University, Middlesbrough, United Kingdom

Abstract

Rapid in-orbit technology validation is increasingly constrained by integration effort and spacecraft resource margins. We present a systems-engineering approach for hosting up to 160 low-power experiment nodes within a standard 3U CubeSat as a contained distributed payload. The architecture treats each node as an independently testable endpoint while the spacecraft provides shared power, timing, command and a controlled data path to the downlink.

To bound peak power and reduce coupled interactions between experiments, we implement deterministic, time-slotted activation using a fixed time-division schedule controlled by the OBC. Only one node is granted electrical activity per time slot (power-gated and bus-enabled), flattening peak demand and bounding peak power and peak thermal load, while integrated thermal dissipation scales with mission activity. We describe the supervisory logic, interface constraints for compliant nodes, and a hardware fault-containment layer that detects bus contention and quarantines non-compliant nodes via per-node switching and watchdog-driven lockout, preserving bus integrity for remaining experiments.

We report representative mass, power and data allocations and verification methods covering tray/node interchangeability, end-to-end command/telemetry integrity, and resilience to injected failure modes (tray control interface fault). The resulting pattern supports high-density, debris-compliant in-orbit validation where modularity, determinism and fault containment are primary drivers.

Introduction

In-orbit validation remains a bottleneck for new sensors and subsystems. Even with commercial off-the-shelf (COTS) spacecraft buses, integrating a novel payload typically requires bespoke mechanical/electrical interfaces, verification effort and operational planning that increase cost, schedule and integration risk. For many emerging technologies, the time from concept to flight can exceed the practical development window in which in-orbit data remains useful.

This paper describes a systems-engineering approach to increasing validation throughput by hosting numerous low-power experiment nodes within a single standard 3U CubeSat, while maintaining bounded and predictable power, thermal and data behaviour. The architecture consolidates spacecraft services at the platform level (power, timing, command and a controlled data path to downlink) and constrains each experiment node to a fixed mechanical envelope and a tightly bounded electrical and operational profile. Deterministic scheduling and hardware-enforced power-gating are used to bound peak load, reduce coupling between experiments and simplify integration at scale.

Prior approaches to high experiment density have included deploying large numbers of ultra-small free-flying devices from a carrier spacecraft. While effective for maximising experiment count, debris-mitigation requirements and the operational burden of tracking and conjunction management increasingly constrain mission profiles that intentionally release large populations of objects. A contained, internally hosted payload model retains the benefits of high experiment density without creating additional orbital objects.

Background

The technological foundations of AmbaSat originated in low-power, distributed environmental sensor networks and end-to-end remote sensing solutions. The same design constraints that drive terrestrial sensor nodes, such as tight power budgets, intermittent connectivity, autonomous operation and robust fault handling, map naturally to small satellite payload experimentation.

Early AmbaSat concepts focused on validating technologies by deploying very small “ChipSat”-class devices as independent spacecraft. As the programme matured, it became clear that evolving debris-mitigation and operational constraints strongly favour architectures that avoid releasing additional objects. This prompted a shift from free-flight deployment to a contained experiment model, in which multiple independent experiment nodes are hosted internally within a single 3U CubeSat.

The resulting approach reduces system-level complexity at the individual experiment node while consolidating critical spacecraft services at the platform level. The 3U spacecraft provides shared power, timing, command and a controlled data path to downlink, while each node is constrained to a fixed mechanical envelope and a bounded operational profile. This enables a “many-experiment” mission within a monolithic external spacecraft, with predictable resource usage and reduced integration overhead compared with repeatedly integrating bespoke payloads into separate spacecraft.

System Overview and Internal Architecture

A fully populated payload stack comprising 20 flight trays and 160 experiment nodes presents an inherent resource-management challenge. If nodes were permitted to operate without coordination, peak electrical demand and coupled thermal interactions could exceed the margins of a standard 3U electrical power system (EPS) and reduce platform availability. The architecture presented here addresses this by treating each node as an independently testable endpoint, while enforcing deterministic operation, shared platform services and hardware fault containment to safely multiplex in-orbit validation.

The 3U spacecraft is built around commercial off-the-shelf (COTS) bus elements for standard spacecraft functions (structure, solar power generation, EPS and UHF communications). This reduces non-recurring engineering effort on the baseline platform and allows development focus to be placed on the payload-hosting functions:

- a customised on-board computer (OBC); and
- a modular internal payload system capable of hosting many nodes with predictable resource usage.

A central design principle is consolidation of spacecraft services at the platform level. The EPS and OBC provide shared, standardised resources for power distribution, timing, command sequencing, data aggregation and downlink buffering. Internally, the payload volume is partitioned into stackable flight trays, each implemented as a discrete sub-assembly capable of hosting eight experiment nodes. Each tray includes an always-on local microcontroller which remains in a listening state until commanded by the OBC to begin its cycle. The tray controller enforces per-node power-gating and slot timing using solid-state load switching and watchdog-driven termination.

The OBC acts as the master supervisor for the payload stack. It issues tray-level commands over a low-bandwidth control interface, orchestrates deterministic activation windows and records execution outcomes. Payload data is collected via a separate internal wireless data plane: nodes transmit a single LoRa telemetry packet directly to a receiver integrated with the OBC during their allocated slot. The OBC timestamps and buffers received packets in non-volatile storage, enabling delay-tolerant operation until a downlink window is available.

Figure 1 summarises the platform architecture and the separation between the wired control plane and internal wireless data plane.

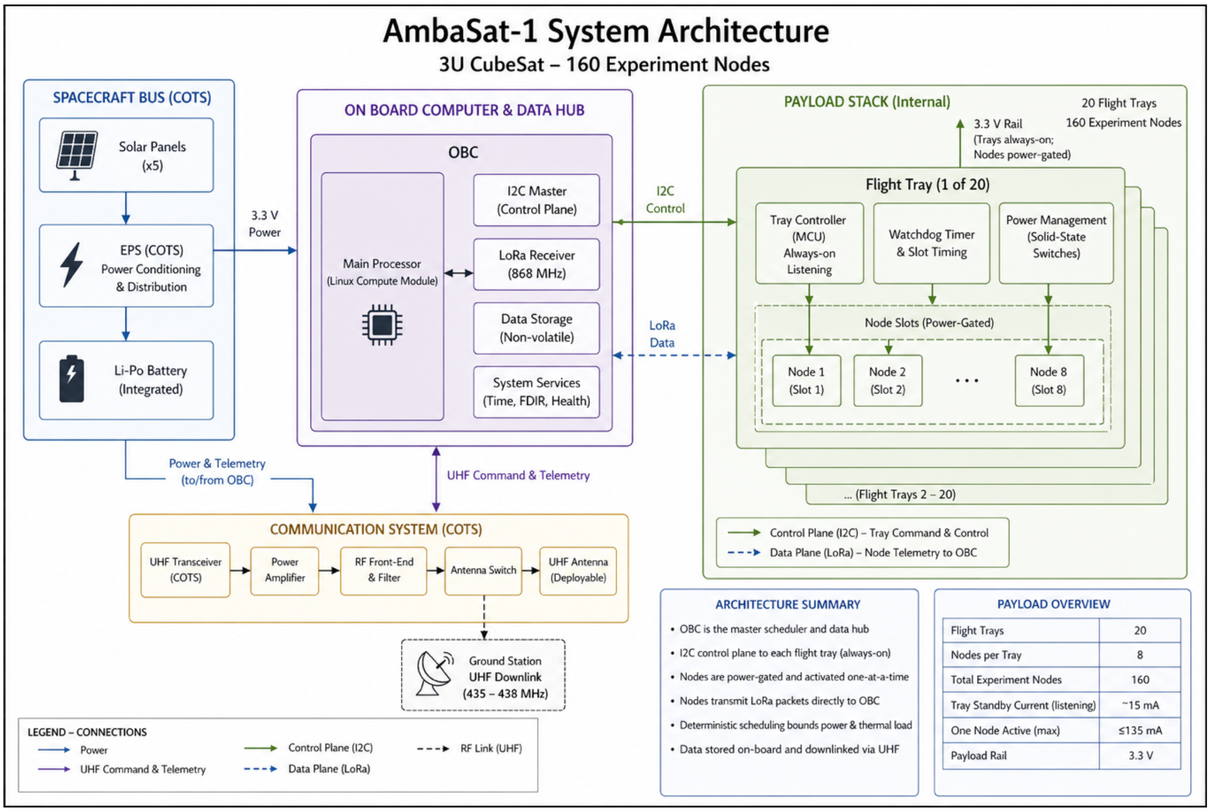

ARCHITECTURE SUMMARY

- OBC is the master scheduler and data hub
- I2C control plane to each flight tray (always-on)
- Nodes are power-gated and activated one-at-a-time
- Nodes transmit LoRa packets directly to OBC
- Deterministic scheduling bounds power & thermal load
- Data stored on-board and downlinked via UHF

PAYLOAD OVERVIEW

Flight Trays	20
Nodes per Tray	8
Total Experiment Nodes	160
Tray Standby Current (listening)	~15 mA
One Node Active (max)	≤135 mA
Payload Rail	3.3 V

Figure 1. AmbaSat 3U contained distributed payload architecture, showing separation between the wired OBC-to-tray control plane and the internal LoRa data plane from experiment nodes to the OBC.

Architecture and Interfaces

The payload architecture is defined by a deliberate separation between the power/control plane and the data plane. This reduces harness complexity and wiring mass, limits the consequences of electrical contention, and simplifies mechanical interchangeability across a high-density internal payload stack.

The control plane is implemented as a wired controller-device link between the on-board computer (OBC) and each flight tray controller. A standard I^{2}C interface is used for tray-level commands and timing synchronisation (with SPI available where additional tray-level hardware support is required). The ^{I2}C link is used only between the OBC and the flight tray controllers; experiment nodes are not connected to I^{2}C and are never addressed as bus devices. Nodes remain electrically isolated and are only ever interacted with via controlled power application. The I^{2}C control plane is designed with bus integrity measures appropriate to the 20-tray stack (e.g., segmentation/buffering and pull-up sizing) and is verified during system-level integration.

The data plane is implemented as an internal wireless network. Each experiment node integrates an 868 MHz LoRa transceiver and transmits a single telemetry packet during its allocated operational window. Nodes transmit directly to a receiver integrated with the OBC, avoiding high-pin-count backplanes and preventing a failed node from contending on a shared wired data bus. The OBC timestamps and buffers received packets in non-volatile storage for later downlink.

The internal LoRa data plane operates at low data rate with high processing gain, providing robustness to multipath and attenuation in the internal spacecraft environment. The payload stack is housed within an open frame 3U structure rather than a sealed enclosure, which reduces shielding effects while still presenting a reflective internal RF environment. The deterministic schedule further reduces collision risk by limiting transmissions to one node per slot. Reception performance is verified during bench testing across the full tray stack to identify any tray-position sensitivity and confirm packet aggregation success under representative integration conditions.

To guarantee physical and electrical interchangeability, compliant experiment nodes must adhere to fixed interface parameters. Mechanically, each node is constrained to a 35 mm × 35 mm footprint and interfaces to the flight tray through a standardised header/pin configuration. Electrically, each node is bound by a per-node power budget and must remain completely passive until explicitly powered. Prior to power application, the node must present a high-impedance state at the tray interface and must not initiate autonomous activity.

Power delivery and timing enforcement are implemented at the flight tray level using solid-state load switching under microcontroller control. Rather than distributing continuous rail voltage to payloads, the tray controller connects and disconnects the power path to each node according to instructions from the OBC. A compliant node must complete its boot sequence, execute its experimental routine, and transmit its single LoRa packet entirely within the allocated time slot.

This architecture provides inherent hardware fault containment. At the end of a node's time slot, the tray controller physically disconnects power, limiting the duration of anomalous behaviour and preventing persistent over-current conditions from propagating into the shared power system. The combination of tray-level power gating and the wireless data plane enables large experiment counts with bounded resource usage and reduced integration coupling between adjacent nodes.

Deterministic Scheduling and Power Management

In this paper, ‘polling’ refers to deterministic, time-slotted activation of trays/nodes (i.e., issuing an enable command), rather than querying nodes for data over the control bus.

To operate a 160-node payload stack within the power and thermal margins of a standard 3U electrical power system (EPS), the AmbaSat architecture uses a deterministic polling regime that bounds peak load and enforces RF silence outside of allocated transmit windows. The on-board computer (OBC) acts as the master scheduler, coordinating the 20 flight trays to ensure predictable power consumption and timing behaviour across the payload volume.

At the start of a collection cycle, the OBC addresses a single flight tray and commands it to begin its local scheduled execution sequence. A fixed execution window is allocated to the selected tray (nominally 320 seconds) to complete its internal cycle before control advances to the next tray. During this window, the tray controller powers its eight hosted experiment nodes sequentially using solid-state load switching. Each node is granted a fixed operational slot of 26 seconds active followed by a 6 second guard interval. The nominal 320s tray window comprises eight 32s node slots (26 s active + 6 s guard; 256 s total) plus ~64s of tray-level overhead to accommodate command/handshake margin with the OBC, tray housekeeping, and contingency time for retries or fault handling without eroding slot timing for subsequent trays.

This sequencing implements a strict “one-at-a-time” rule: only one experiment node is energised and permitted to operate during any slot, which bounds instantaneous payload power demand to the tray overhead plus the maximum permitted draw of a single node. Therefore, peak power is largely decoupled from total experiment count, enabling a fully populated payload stack without a linear increase in worst-case electrical demand. The same scheduling logic also reduces coupled thermal interactions by spreading dissipation in time and across the physical stack rather than allowing concurrent hotspots. Peak thermal load is bounded by the one-at-a-time rule; however, integrated thermal dissipation scales with duty cycle and total mission activity.

Payload telemetry is returned over the internal wireless data plane during the active slot. Nodes transmit a single LoRa (868 MHz) packet to the OBC receiver within their allocated window, supporting deterministic RF activity and simplifying collision avoidance (only one node is expected to transmit during a given slot). The OBC timestamps and buffers successfully received packets in non-volatile storage for downlink when communications opportunities are available.

Fault Detection, Isolation, and Recovery (FDIR)

With an architecture scaled to support 160 independent experiment nodes, isolated faults at the payload level are expected over mission life (e.g., radiation-induced single-event latch-up (SEL), software hangs, or component failures). The AmbaSat platform is therefore designed around a fail-operational principle at the payload layer: anomalous behaviour is intended to be detected and isolated locally so that the remainder of the payload stack can continue scheduled operation.

Fault detection is implemented across both the electrical and logical domains. At the individual node level, detection is primarily electrical, as nodes are intentionally isolated from the control bus and are only connected to the system via controlled power application. The solid-state load switching on each flight tray monitors the current draw of the actively gated node. If a node exhibits an over-current condition (e.g., short circuit or latch-up behaviour), the tray hardware detects the event and removes power within the node's allocated slot.

At the tray level, the on-board computer (OBC) monitors the health of the I^2C control plane used to communicate with flight tray controllers. Each flight tray controller is an always-on ATmega device with a unique address and remains in a listening state until commanded by the OBC to begin its polling sequence. A tray controller anomaly that results in repeated command timeouts, failure to acknowledge, or abnormal bus behaviour is treated as a subsystem-level communication fault and logged by the OBC. In this condition, the affected tray is treated as unavailable for the current cycle, and the OBC advances scheduling to the next tray to preserve overall mission throughput.

Isolation of node-level faults relies on strict, watchdog-driven timing enforcement and hard power-gating. Each node is allocated a fixed time window to boot, execute its routine, and transmit a single LoRa telemetry packet. The flight tray controller starts a local watchdog timer when a node is energised. If the node fails to complete within its window (e.g., software hang), the tray controller terminates the slot by physically disconnecting power, limiting the duration of anomalous behaviour.

When a node fails to deliver its expected telemetry (whether due to over-current isolation or watchdog expiry), the OBC records a failure against that endpoint. If a node accrues three consecutive execution failures, it is marked quarantined and is subsequently skipped during polling cycles, denying it further power application and operational time.

Recovery is implemented in tiers. The first tier is slot-level power cycling: because nodes are de-energised at the end of each slot, the next scheduled attempt provides a cold reboot that can clear transient software hangs or single-event upsets (SEUs). If failures persist, the node is quarantined as described above. For tray-level anomalies, recovery is implemented through bus-level recovery actions at the OBC (e.g., re-initialisation of the I^2C interface and re-trying the command sequence) and graceful degradation at the scheduler level. If a tray remains non-responsive after a defined number of retries, the OBC skips that tray for the remainder of the polling cycle and continues with the next tray.

By combining a wireless internal data plane with per-node power-gating and watchdog enforcement, the architecture limits the scope of failures and reduces the likelihood that a single payload anomaly disrupts operation of the wider payload stack. A severe fault is intended to remain localised to the affected node or, in the worst case, to a single flight tray, allowing continued operation for the remaining experiments under the deterministic schedule.

A sustained RF anomaly (e.g., repeated unexpected transmissions) is handled operationally by skipping the affected node in the deterministic schedule and, if required, disabling its power slot via quarantine logic; the one-at-a-time power-gating limits the duration of any node-originated transmission.

Budgets and Assumptions

A representative resource budget has been developed to confirm that the 160-node payload remains compatible with standard 3U bus margins. The key design feature is that instantaneous payload load is bounded by deterministic scheduling: only one node is energised during any operational slot, using a nominal 26 s active + 6 s guard per node, with a nominal 320 s execution window per flight tray (eight nodes) before advancing to the next tray.

Table 1. – Bounded payload power and data budget (representative)

Parameter	Representative value	Notes
Slot timing	26 s active / 6 s guard	One node per slot
Tray cycle window	320 s	Eight nodes per tray
Flight tray standby ("listening") current	0.75 mA	Tray controller powered, awaiting command
Max current (one tray + one node energised)	130 mA	Worst case during a node slot
Equivalent power (at 3.3 V)	~0.45 W	For 130 mA at 3.3 V
Telemetry per node execution	≤60 bytes	One 868 MHz LoRa packet per slot
Data per full sweep (160 nodes)	≤9.6 kB	160 × 60 bytes, buffered for downlink

Preliminary Bench Measurements

To validate the power and data budgets, early bench testing was conducted on multiple populated flight trays. Measurements confirm that the baseline tray controller draws a steady-state standby current of 0.75 mA whilst awaiting commands. During a simulated execution window, energising a single node yielded a measured peak current of 130 mA at the regulated payload rail (equivalent to ~0.45 W at 3.3 V). Capturing these early hardware-in-the-loop results demonstrates that the deterministic "one-at-a-time" scheduling effectively bounds the maximum instantaneous power draw, providing confidence that the full 160-node configuration will remain within standard 3U bus margins.

Assumptions:

- Only one node is energised per slot (deterministic “one-at-a-time” rule).
- Nodes transmit one LoRa packet during their allocated slot; no node-to-node networking is required.
- The OBC timestamps and buffers packets in non-volatile storage until UHF downlink windows are available.
- Currents are quoted at the regulated payload rail; the corresponding power values assume a 3.3 V rail (scales linearly with rail voltage).

Verification and Validation Programme

The scale of a 160-node distributed payload requires a Verification and Validation (V&V) programme that emphasises repeatability, compliance checking, and fault containment at the payload layer. While conventional SmallSat V&V often concentrates on system-level environmental testing, this architecture benefits from a bottom-up approach that validates:

- deterministic scheduling margins,
- per-node electrical compliance under hard power-gating; and
- robustness of the internal wireless data return under representative integration conditions.

Integration Flow and Compliance Checks

Integration is executed as a hierarchical three-stage progression to isolate variables and avoid compounding errors:

- Node-level compliance: Prior to physical integration, individual experiment nodes undergo electrical and logical compliance checks on a diagnostic test jig that simulates the flight tray interface. This verifies adherence to the 35 mm × 35 mm footprint, confirms a high-impedance/passive state prior to power application, and measures both boot-to-transmit timing and peak current draw to ensure compatibility with the allocated slot window and per-node power budget.
- Tray-level integration: Compliant nodes are mated to the flight tray header pins. At this subsystem level, the tray controller's ability to sequence per-node power, actuate solid-state load switches, and enforce watchdog-driven slot termination is verified independently of the master on-board computer (OBC). This stage demonstrates that node execution is bounded by hardware timing and power-gating at the tray, and that a misbehaving node cannot prevent the tray from advancing to the next slot.
- System-level integration: Fully populated trays are stacked into the 3U form factor and integrated with the spacecraft bus and OBC. A continuous deterministic polling schedule is executed across the payload stack to confirm end-to-end behaviour under representative operating modes, including:
    i. bounded peak load drawn from the shared EPS consistent with the "one-at-a-time" power-gating envelope,
    ii. stable slot and tray timing over repeated cycles; and
    iii. reliable telemetry aggregation by the OBC. Bench-stack runs are used to execute full-sweep cycles and to quantify packet aggregation success rate for compliant, non-quarantined nodes (target >99% over a full operational sweep).

Tray and Node Interchangeability Tests

A core driver of this architecture is modularity and repeatable integration. To validate the plug-and-play intent, blind interchangeability tests are performed during the V&V programme. Randomly selected compliant nodes are swapped between active slots on a single flight tray, and complete eight-node flight trays are physically reordered within the 3U stack.

The system is considered to have passed this phase when the OBC initiates the deterministic schedule, establishes communication with tray controllers over the I^2C control plane, and aggregates the expected sequence of LoRa telemetry packets without manual software reconfiguration. Timing behaviour is assessed by confirming that slot boundaries and tray execution windows remain within defined tolerances following swaps and reordering. This demonstrates that operation is not dependent on a fixed spatial arrangement of nodes or trays, and that the scheduler is agnostic to payload placement within the stack.

Injected Fault Modes

To validate the payload-layer fail-operational principle and the FDIR behaviour described in Section 6, representative hardware and software faults are deliberately injected into the integrated system:

- Hung node (watchdog expiry): A test node is programmed to enter an infinite loop and fail to transmit its packet. The objective is to verify that the tray controller terminates the slot by removing power at the slot boundary, and that the OBC records the failed execution and subsequently skips (quarantines) the node once the configured failure threshold is reached.
- Stuck/failed tray control interface: A fault is introduced that causes a target flight tray to become non-responsive to OBC commands (e.g., repeated command timeouts or abnormal behaviour on the I^2C control plane). The objective is to verify that the OBC detects the condition, performs defined bus-level recovery actions (e.g., re-initialisation and retry), and if the tray remains unavailable, skips that tray for the remainder of the polling cycle while maintaining deterministic scheduling for the remaining trays.
- Over-current event (latch-up simulation): A test node is forced to draw current above the allocated limit to simulate a short circuit or latch-up condition. The objective is to verify that tray-level power switching isolates the node by removing power, and that the event does not cause a sustained voltage sag on the shared EPS or disrupt subsequent node slots.

Acceptance Criteria

Flight readiness is contingent upon demonstrating bounded resource usage, deterministic continuity, and fault containment under representative operating conditions. Key acceptance criteria include:

- Bounded instantaneous load: During deterministic operation (including injected faults), measured current draw at the regulated payload rail remains within the bounded "one-at-a-time" envelope (tray overhead plus the maximum permitted draw of a single node).
- Deterministic continuity: Injected node faults do not prevent progression to subsequent slots; injected tray-level non-responsiveness does not prevent progression to subsequent trays (graceful degradation).
- Fault containment: Over-current events are isolated by power-gating without sustained impact on the shared EPS and without disrupting operation of other nodes/trays.
- Data aggregation performance (bench stack): For compliant, non-quarantined nodes, the OBC achieves a >99% packet aggregation success rate over a full operational sweep during bench-stack testing.

Conclusion

This paper has presented a contained 3U distributed payload architecture for high-density in-orbit validation using 20 flight trays and 160 power-gated experiment nodes. By separating the wired control plane (OBC to tray controllers) from a wireless internal data plane (LoRa telemetry direct to the OBC) and enforcing deterministic one-at-a-time slot execution (26 s active + 6 s guard), the design bounds peak electrical demand and reduces coupling between experiments. Hardware-enforced power switching and watchdog-driven slot termination provide fault containment at node level, while scheduler-level graceful degradation enables continued operation when individual nodes or trays become unavailable. Ongoing work focuses on completing the full verification programme across repeated full-sweep bench-stack cycles, expanding environmental testing where required, and finalising operational constraints and downlink budgeting for flight.

References

[1.] The CubeSat Program, California Polytechnic State University (Cal Poly), CubeSat Design Specification, Rev. 14.1, 9 Feb. 2022.

[2.] International Organization for Standardization, ISO 24113:2019 - Space systems - Space debris mitigation requirements, 2019.

[3.] European Cooperation for Space Standardization (ECSS), ECSS-E-ST-70-41C - Space engineering - Telemetry and telecommand packet utilization, 15 Apr. 2016.

[4.] NXP Semiconductors, UM10204 - I^2C-bus specification and user manual, Rev. 7 (or latest), Apr. 2021.

[5.] United Kingdom, Space Industry Act 2018, 2018 (c. 5).