UAV Flight System Architecture for Industrial and Commercial Applications

— with Energy Aware Design Principles and Renesas Industrial MCU Reference Implementation

***Levan Bidzishvili**, System Architect, TAA / UX*

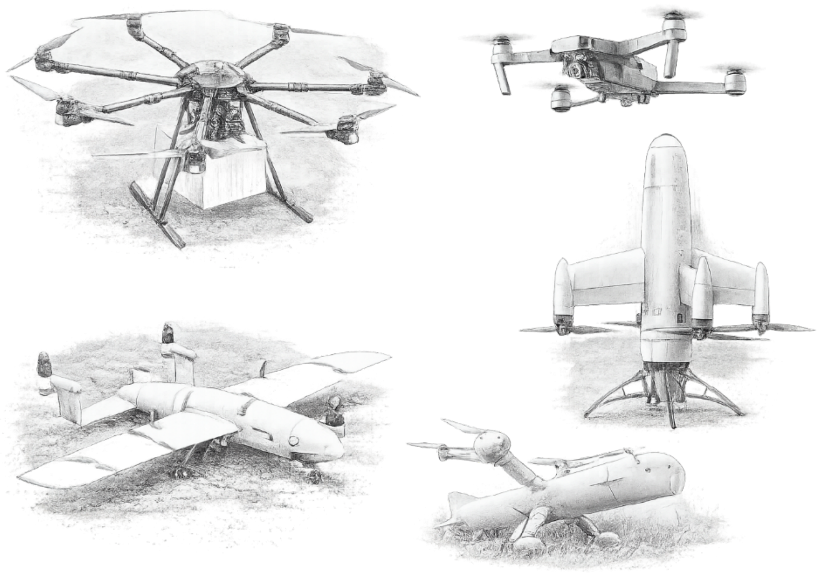

Executive Summary

Unmanned Aerial Vehicles (UAVs) used in industrial, commercial, and defense environments are no longer simple remotely piloted platforms. They are evolving into autonomous, energy-constrained cyber-physical systems that must operate reliably in complex environments, often without continuous access to satellite navigation signals.

This white paper presents a Renesas-enabled UAV reference architecture, derived from the block diagram shown in Figure 1, that addresses these challenges through deliberate system partitioning. The architecture combines intelligent battery management, robust power distribution, deterministic real-time control, distributed actuation, and AI-enabled perception. Rather than assuming ideal operating conditions, the design explicitly supports GPS-degraded and GPS-denied scenarios, where vision-based navigation and sensor fusion can match or even exceed GPS performance.

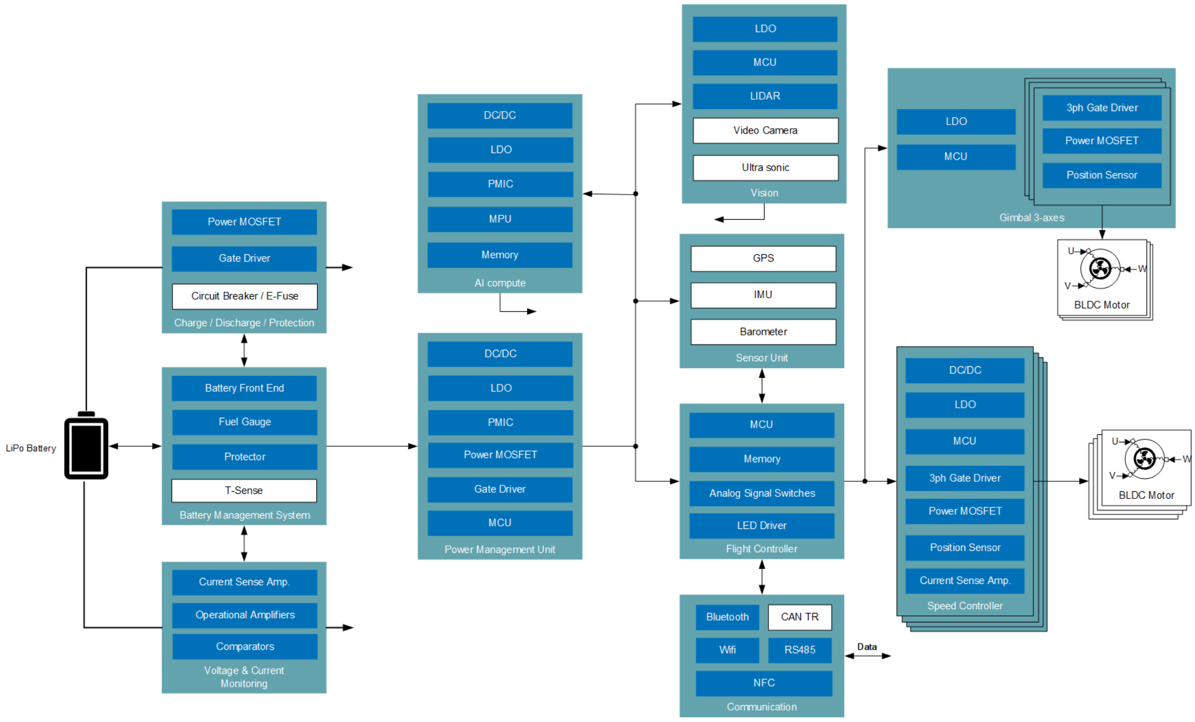

Figure 1: UAV Drone system architecture

Architecture Driven by Physics and Mission Reality

Electric UAV flight is fundamentally constrained by physics. Vertical take-off, hover, and aggressive maneuvering demand high instantaneous power, while endurance and range are limited by total stored energy. These constraints dominate system behavior far more than individual software features.

The architecture shown in Figure 1 reflects this reality. Energy storage, power distribution, flight control, perception, and actuation are separated into well-defined subsystems that cooperate but do not collapse into a monolithic design. This separation allows the system to scale in payload, compute capability, and autonomy while preserving safety and predictability.

Crucially, navigation is treated as a multi-source problem. GPS is present, but it is not the single point of dependency. Vision, inertial sensing, LiDAR, and ultrasonic ranging are all first-class inputs to the flight controller.

Battery Management: From Energy Storage to Energy Awareness

At the foundation of the system lies the Battery Management System (BMS). In professional UAVs, the battery is not merely an energy reservoir; it is a safety-critical subsystem whose behavior directly influences mission success.

The BMS in the reference architecture continuously monitors voltage, current, and temperature, and derives not only the remaining energy but also the battery's ability to deliver power at a given moment. This distinction is essential, because high discharge rates during hover or ascent can significantly reduce usable energy and accelerate aging.

By reporting real-time battery capability to the flight controller, the system enables predictive behavior. Instead of reacting to critical battery thresholds, the UAV can adapt its flight profile early, reducing aggressiveness, limiting hover time, or initiating a controlled return. Energy thus becomes an active system variable, not a passive constraint.

Power Management as the Electrical Backbone

Between the battery and the rest of the UAV sits the Power Management Unit (PMU). In demanding flight phases, current draw changes rapidly as motors respond to control commands and environmental disturbances. Without proper isolation and regulation, these transients can destabilize sensitive electronics.

The PMU in this architecture provides regulated power domains, controlled sequencing, and isolation between high-current motor electronics and low-noise sensing and control circuits. This design ensures that navigation sensors, AI compute, and control loops remain stable even when propulsion demand fluctuates sharply.

Such robustness is a prerequisite for autonomous operation in industrial and defense environments, where electrical noise, temperature variation, and load dynamics are unavoidable.

Deterministic Flight Control as the System Orchestrator

The Flight Controller, implemented with an industrial Renesas MCU, acts as the real-time orchestrator of the UAV. Its primary responsibility is not raw computation power, but deterministic behavior.

The flight controller executes stabilization loops, state estimation, and mission logic with predictable timing. It consumes navigation and perception results from other subsystems, rather than processing raw sensor data itself. This separation allows the controller to remain responsive and certifiable, even as perception and AI workloads evolve.

From an architectural perspective, the flight controller is the decision authority, while AI and vision subsystems act as advanced observers.

Distributed Actuation and Safety Through Redundancy

The propulsion system is implemented as a set of distributed speed controllers, each with its own MCU, gate drivers, power MOSFETs, and current sensing. This distribution reduces load on the flight controller and improves fault tolerance.

If a motor or controller degrades, the system can redistribute thrust and maintain controlled flight. This approach reflects lessons learned from multi-rotor VTOL systems, where redundancy and graceful degradation are essential for safety.

AI Compute and Perception as a Dedicated Domain

The block diagram explicitly separates AI Compute from the flight controller. This reflects a key design principle: perception and inference are computationally intensive and non-deterministic, while flight control must remain deterministic and bounded.

The AI compute domain, powered through its own regulators and memory, processes data from cameras, LiDAR, and ultrasonic sensors. It performs feature extraction, pattern recognition, depth estimation, and environment understanding. The output of this domain is not raw imagery, but structured information such as relative motion, obstacle maps, and navigation cues.

By isolating AI compute, the architecture allows perception capabilities to scale independently of flight control, without compromising real-time safety.

Navigation Beyond GPS: AI Vision in Degraded and Denied Environments

Why GPS Alone Is No Longer Sufficient

GNSS technologies such as GPS work well in open environments, but their limitations become evident in real-world missions. Signal reflections in urban canyons, attenuation indoors, interference near metallic structures, and intentional jamming in defense scenarios all degrade accuracy or cause complete loss of positioning.

For UAVs operating close to infrastructure, inside facilities, or in contested environments, relying solely on GPS is increasingly impractical.

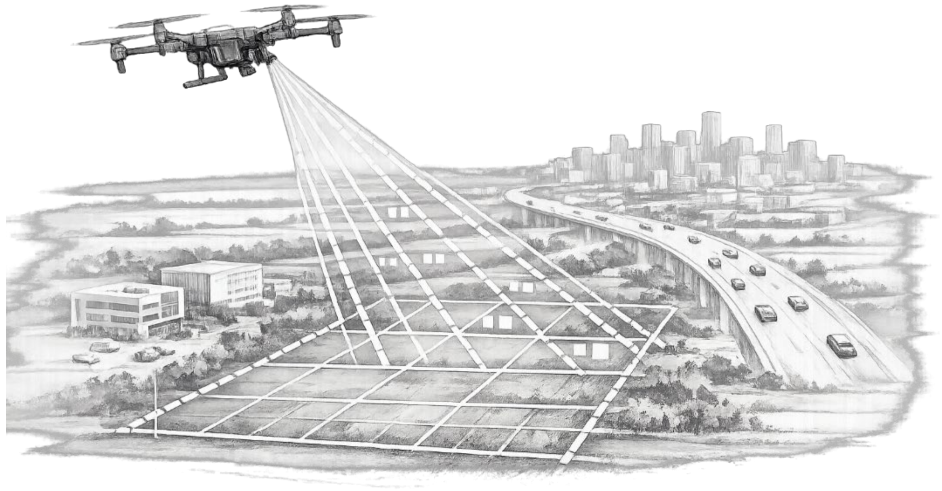

Figure 2

Vision-Centric Navigation as an Equal or Better Alternative

The architecture shown in Figure 1 enables navigation based on local perception rather than global satellite signals. Cameras combined with AI pattern recognition can estimate motion by tracking visual features across successive frames. When fused with inertial data, this approach provides continuous and smooth position estimation even in the absence of GPS.

LiDAR adds precise geometric depth information that is largely independent of lighting conditions, while ultrasonic sensors provide robust short-range altitude and proximity measurements, particularly valuable during take-off, landing, and indoor flight.

In many scenarios, this combination offers higher local accuracy than GPS, especially for relative positioning and obstacle avoidance.

Sensor Fusion as the Key Enabler

No single sensor is sufficient in all conditions. The strength of the architecture lies in sensor fusion. Visual perception excels in textured environments but degrades in poor lighting. LiDAR remains reliable in darkness but adds cost and power consumption. Ultrasonic sensors are simple and robust but limited in range. The IMU provides high-rate motion data but drifts over time.

By fusing these inputs, the AI compute domain produces navigation estimates that remain stable even when individual sensors become unreliable. The flight controller receives these fused results and adapts its control strategy accordingly.

Seamless Transition Between GPS and Vision-Based Navigation

Because GPS is treated as one input among many, the system can transition seamlessly between navigation modes. In open environments, GPS contributes absolute positioning. As the UAV enters a

GPS-degraded zone, vision and inertial sensing naturally take over without abrupt mode switching or operator intervention.

From the perspective of the flight controller, navigation remains continuous. Only the weighting of input sources changes.

Renesas Single-MPU UAV Architecture: Coordinated Real-Time Control and AI Perception on RZ/V2H

A central design goal of the proposed UAV platform is to increase capability and autonomy while reducing system complexity. Traditional advanced UAV designs often rely on multiple processors—separate devices for flight control, vision and AI processing, and safety supervision. While functional, this approach increases board complexity, power consumption, and integration effort.

Renesas addresses this challenge with a single-MPU UAV architecture based on the RZ/V2H (Figure 3), which integrates real-time control, high-level processing, AI acceleration, and safety supervision into one device. Rather than merging all workloads into a single execution context, the RZ/V2H uses a heterogeneous multi-core architecture that allows each function to run in the domain best suited to its timing, performance, and safety requirements.

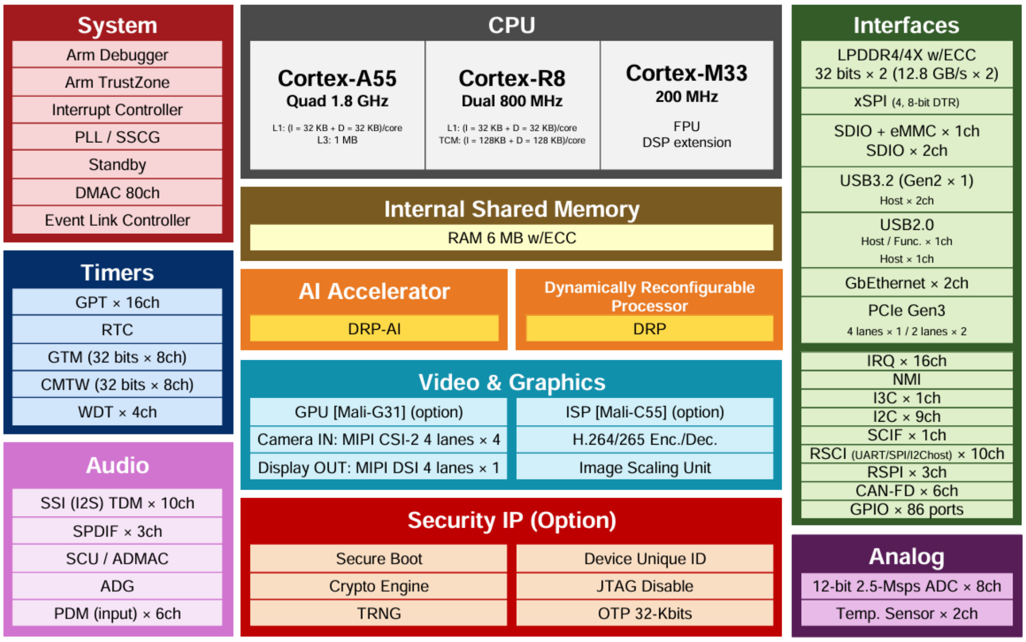

Figure 3: RZ/V2H MPU - architecture

Real-Time Flight Control on Cortex-R8

The Cortex-R8 dual-core subsystem is dedicated to the Flight Management Unit (FMU). It executes all time-critical and safety-relevant functions, including attitude stabilization, rate control, motor command generation, and immediate failsafe handling. With tightly coupled memory and predictable interrupt behavior, the Cortex-R8 ensures deterministic flight behavior independent of system load elsewhere in the MPU.

The Cortex-R8 typically runs bare-metal or with a lightweight real-time operating system. A strict architectural rule is enforced: only the Cortex-R8 is allowed to generate actuator commands. All higher-level navigation and autonomy decisions are mediated through this real-time authority, preserving safety and control integrity.

Vision, Navigation, and Autonomy on Cortex-A55

High-level perception and autonomy are handled by the quad-core Cortex-A55, usually running Linux. This domain manages camera acquisition via the integrated ISP and MIPI CSI interfaces, LiDAR and ultrasonic preprocessing, visual-inertial odometry, SLAM, obstacle avoidance, and mission planning.

By isolating these workloads on the Cortex-A55, the system can leverage rich software ecosystems and advanced algorithms without impacting flight stability. Navigation outputs are expressed as trajectories, constraints, or mode requests and passed to the Cortex-R8 for real-time arbitration and execution.

Efficient AI Inference with DRP-AI

AI-based perception is accelerated by the integrated DRP-AI, optimized for low-power, deterministic inference. Vision models such as object detection, feature extraction, or semantic segmentation are orchestrated by the Cortex-A55 and offloaded to DRP-AI for execution. The accelerator returns structured perception results rather than raw data, reducing processing load and power consumption while enabling advanced AI capabilities.

Safety Supervision on Cortex-M33

The Cortex-M33 provides an independent safety and supervision layer. It monitors system health, enforces secure boot and trust boundaries, manages watchdogs, and supports emergency control paths. By separating supervision from both real-time control and Linux execution, critical safety functions remain available even in the presence of software faults elsewhere in the system.

LiDAR Sensing and Data Acquisition for Reliable AI-Assisted Navigation

Reliable perception in GPS-degraded or GPS-denied environments requires sensing modalities that provide direct geometric information about the surrounding space. While camera-based AI vision delivers rich semantic understanding, its performance depends on illumination, texture, and contrast. To address these limitations, the proposed Renesas UAV architecture integrates LiDAR as a complementary depth-sensing technology, tightly coupled with AI compute and navigation.

LiDAR (Light Detection and Ranging) measures distance by emitting short laser pulses and detecting the reflected signal from surrounding objects. The fundamental measurement principle is time-of-flight (ToF), expressed as:

$$d = \frac{c * \Delta t}{2}$$

where ***d*** is the distance to the object, ***c*** is the speed of light, and ***Δt*** is the measured round-trip time between emission and reception. This direct physical measurement yields absolute depth information, largely independent of ambient lighting and visual texture.

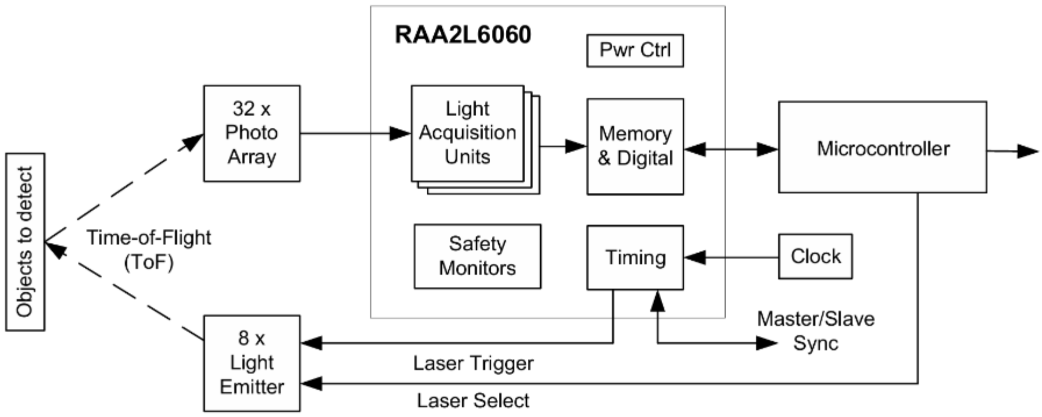

Figure 4: Solid State LiDAR system example

In UAV applications, LiDAR systems perform repeated distance measurements across different directions to build a spatial representation of the environment. Depending on implementation, scanning may be achieved through mechanical beam steering, solid-state electronic scanning, or multi-beam emission. The resulting output is typically a point cloud or depth profile, describing the local geometry of the environment with high accuracy.

Within the Renesas RZ/V2H-based UAV platform, LiDAR data complements camera and inertial sensing by providing reliable short- and mid-range depth measurements. When fused with AI-based vision outputs and IMU data, LiDAR improves obstacle detection, local mapping stability, and navigation robustness, particularly during low-altitude flight, precision landing, and operation near structures or in visually challenging environments.

LiDAR Data Acquisition with RAA2L6060

Accurate LiDAR perception depends not only on optical principles but also on precise signal acquisition and timing resolution. The reflected laser signal is typically weak and must be captured with high fidelity to preserve distance accuracy and measurement consistency.

In the Renesas UAV architecture (Figure 3, Figure 4), the RAA2L6060 LiDAR data acquisition IC plays a central role in this signal chain. RAA2L6060 is designed to interface with LiDAR receiver front-ends, enabling accurate capture of return signals and precise time measurement required for time-of-flight calculation. By converting optical return events into well-defined digital timing and amplitude information, the device forms a reliable bridge between the LiDAR sensor hardware and the digital processing domain.

The data produced by RAA2L6060 is forwarded to the perception and AI compute layers, where it is combined with camera-based features and inertial measurements. This separation of concerns allows LiDAR signal acquisition to be handled by a dedicated, purpose-built device, while higher-level interpretation and fusion are performed on the RZ/V2H processing cores.

LiDAR and Camera-Based Scanning: Complementary Roles

LiDAR and camera-based vision systems observe the environment through fundamentally different mechanisms. Rather than competing approaches, they provide complementary information that, when fused, significantly improves navigation reliability.

Aspect	LiDAR Scanning	Camera / AI Vision
Measurement principle	Time-of-flight distance	Reflected light intensity and texture
Primary output	Absolute depth / point cloud	Image features and semantics
Dependence on lighting	Low	High
Sensitivity to texture	Low	High
Strengths	Precise ranging, obstacle distance, low-light operation	Object recognition, classification, scene understanding
Limitations	Limited semantic information	Depth ambiguity without stereo or motion
Role in UAV navigation	Geometry, collision avoidance, landing support	Context awareness, path planning

In the proposed architecture, camera-based AI vision provides semantic interpretation of the scene, while LiDAR supplies reliable geometric depth. Inertial sensing ensures continuity of motion estimation, and fusion algorithms running on the RZ/V2H combine these inputs into a coherent navigation model.

Conclusion

The UAV architecture presented in Figure 1 and Figure 2 is not a collection of independent blocks, but a coherent, physics-driven and system-partitioned design. By combining energy awareness, robust power management, deterministic real-time control, distributed actuation, and AI-enabled perception within a single heterogeneous MPU, the architecture enables reliable operation in environments where traditional GPS-centric designs are insufficient.

AI vision, LiDAR supported by dedicated data acquisition such as RAA2L6060, ultrasonic sensing, and inertial fusion are not treated as experimental add-ons, but as integral and cooperating navigation sources. When combined and fused appropriately, these modalities can match or exceed GPS performance in real-world, GPS-degraded or GPS-denied conditions.

With the RZ/V2H integrating real-time Cortex-R8 control, high-level Cortex-A55 processing, DRP-AI acceleration, and system supervision in a single device—together with its broad portfolio of industrial-grade MCUs, analog, power, and sensing solutions—Renesas is uniquely positioned to enable the next generation of resilient, intelligent, and autonomous UAV platforms.

References

- Unmanned Aerial Vehicle (UAV) — Renesas application overview: https://www.renesas.com/en/applications/industrial/robotics/unmanned-aerial-vehicle-uav
- RZ/V2H Quad-core Vision AI MPU with DRP-AI3 Accelerator: https://www.renesas.com/en/products/rz-v2h
- RZ/V2H Robotic Development Kit: https://www.renesas.com/en/design-resources/boards-kits/ws125-v2hrdkrefz
- Accelerating UAV Development with PX4 on RZ/V2H Robotics Development Kit: https://www.renesas.com/en/video/accelerating-uav-development-px4-rzv2h-robotics-development-kit
- RAA2L6060 LiDAR Data Acquisition IC: https://www.renesas.com/en/products/raa2l6060
- RAA2L6060KIT LiDAR Starter Kit: https://www.renesas.com/en/design-resources/boards-kits/raa2l6060kit
- Robotics Solutions — AI Vision and Precision Motor Control: https://www.renesas.com/en/applications/industrial/robotics
- Motor and Inverter Control Solutions: https://www.renesas.com/en/applications/industrial/motor-drives-robotics/motor-control-solutions

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu, Koto-ku, Tokyo 135-0061, Japan
https://www.renesas.com

Trademarks

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit: https://www.renesas.com/contact-us

Doc Number: R30WP0017EU0100