Edge Compute in Autonomous UAS

Anderson Innovations — Technical Whitepaper

Abstract

Autonomous mapping aircraft are frequently architected around the assumption of a reliable downlink. This paper argues the opposite baseline: that orthomosaic generation, radiometric calibration and spatial optimization should resolve on-board or on the ground within minutes of landing, with the network treated as an optional convenience rather than a dependency.

1. Problem Statement

Cloud-first photogrammetry pipelines impose transfer costs that scale with sensor resolution and flight duration. In specialty agriculture, where diagnostics are time-sensitive and rural connectivity is inconsistent, this latency erodes the operational value of the data product.

2. Compute Budget

We characterize the workload across three stages: frame ingest and time synchronization against GNSS/IMU telemetry, feature extraction and pairwise matching, and bundle adjustment with mosaic compositing. Each stage is bounded by a deterministic scheduling budget on constrained embedded compute.

3. Spatial Optimization

Stitching is treated as a sparse nonlinear least-squares problem. Careful problem ordering, incremental linearization and fixed-memory solvers keep the adjustment tractable within microcontroller-class and small SoC power envelopes.

4. Degraded-Link Operation

Perception, tasking and route replanning execute on-board. Link loss is a nominal condition, not a fault: the vehicle continues its survey plan and reconciles state opportunistically when connectivity returns.

5. Conclusion

Edge-resolved mapping shortens the loop between flight and decision, reduces recurring infrastructure cost, and produces a system whose behavior can be characterized and audited end to end.