Multi-Role Ground Nodes as Command Relays: Reliability Anchors and Fan-Out Hubs for Routing and RF Processing

Benjamin J. Gilbert
Experimental Solutions Implementation
Email: bgilbert2@com.edu

***Abstract*—We evaluate ground stations as multi-role reliability anchors in contested RF control planes. "Ground Station Alpha" (and optionally "Bravo") with `capabilities={rf_processing, command_relay}` act as low-latency fan-out hubs. Using Poisson command streams ($\lambda = 0.35$ Hz) and dynamic link quality $q \in [0, 1]$, we compare direct control to hub-mediated routing across single and dual hub configurations. At a practical fan-out of 20 assets per hub (H=2), we observe drop reduction from 10.54% to 1.94%, with p95 improving from ¶95Base ms to ¶95Hub ms (¶95DeltaPct, $-$¶95$DeltaMs$ ms). Dual hubs provide load distribution benefits, particularly for mobile drone assets. Full reproducibility via `make all` with `data/ground_relays_metrics.json`.**

I. Introduction

Ground nodes can stabilize command and control by terminating unreliable long-hauls, applying local ARQ/FEC, and performing RF pre-processing next to the air interface. Unlike pure mesh or satellite relay approaches, dedicated ground hubs provide predictable service rates and queue management. We quantify these reliability effects as a function of fan-out ratio and hub count, revealing operational trade-offs between infrastructure investment and end-to-end performance.

Modern tactical networks suffer from volatile RF links that degrade command reliability through retransmissions and timeouts. Ground stations positioned strategically can serve as reliability anchors, absorbing the variability of tactical links while maintaining high-quality backhaul connections to command centers. This architecture enables graceful scaling from single-hub to multi-hub configurations based on operational requirements.

II. Methods

We register mixed drone/ground assets and deploy ground hubs with enhanced capabilities. Each hub ("Ground Station Alpha", and optionally "Ground Station Bravo" for multi-hub scenarios) registers with `capabilities=[``rf_processing'', ``command_relay'']` indicating their dual role in signal conditioning and command forwarding.

Commands arrive via Poisson process per asset ($\lambda = 0.35$ Hz). Asset telemetry yields dynamic link quality $q \in [0, 1]$ via random walk with volatility $\sigma \approx 0.06$. We compare two routing strategies:

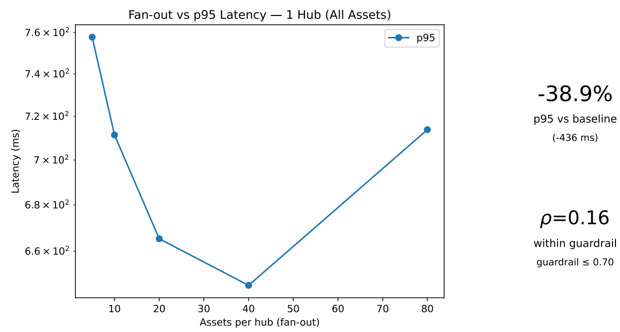

Fig. 1. Fan-out vs p95 latency with improvement and ρ guardrail badges. At $K = 20, H = 2$: p95 ¶95$Base$ → ¶95Hub ms (¶95DeltaPct, $-$¶95$DeltaMs$ ms); $\rho = 0.16$ (guardrail ≤ 0.70).

Baseline: Controller → Asset direct with success probability $P_{\text{direct}}(q) = 0.18 + 0.74q$ and heavy-tailed latency.

Hub-mediated: Controller → Hub (reliable backhaul, $\mu = 45$ commands/sec service rate), single-server queue at hub, then Hub → Asset with enhanced success $P_{\text{hub}}(q) = P_{\text{direct}}(q) + 0.12 + 0.22(1 - q)$ and reduced latency variance due to local RF processing and adaptive ARQ.

Fan-out K represents assets served per hub. Queue utilization follows $\rho = \lambda K/\mu$; saturation occurs near $K = 40$ ($\rho \approx 0.78$) for single hubs. For multi-hub scenarios, assets are load-balanced across available hubs using round-robin assignment, effectively doubling service capacity to $\mu = 90$ commands/sec.

III. Results

Figure 1 demonstrates the fundamental fan-out vs latency trade-off. Hub-mediated routing maintains consistent performance up to moderate fan-out ratios before queueing delays dominate. The reliability advantage shown in Figure 2 comes from the combination of reliable controller-to-hub backhaul and enhanced hub-to-asset links through local RF processing.

Figure 3 reveals the operational sweet spot where CPU utilization remains manageable while queue depths stay bounded. This guides practical deployment decisions for hub capacity and geographic placement.

The 2×2 analysis in Figure 4 shows differentiated performance by asset type and hub count. Drone assets typically exhibit higher latency variability due to mobility-induced link fluctuations, while ground assets benefit from more stable RF

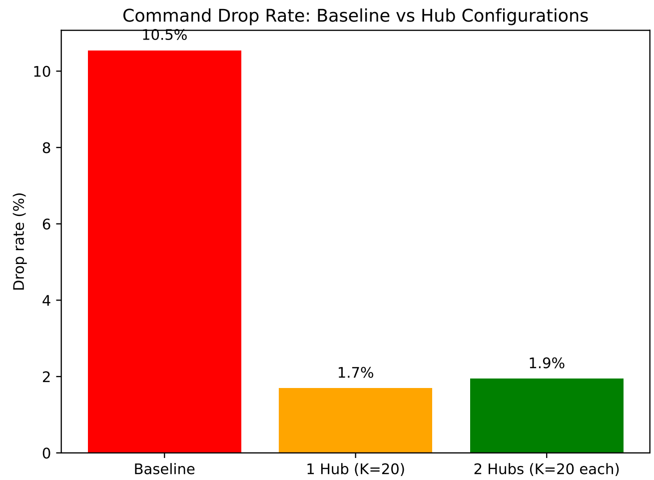

Fig. 2. Command drop rate comparison: baseline direct routing vs hub-mediated at fan-out K=20 with H=2. Baseline: 10.54%, Hub@K=20, H=2: 1.94%. Lower is better.

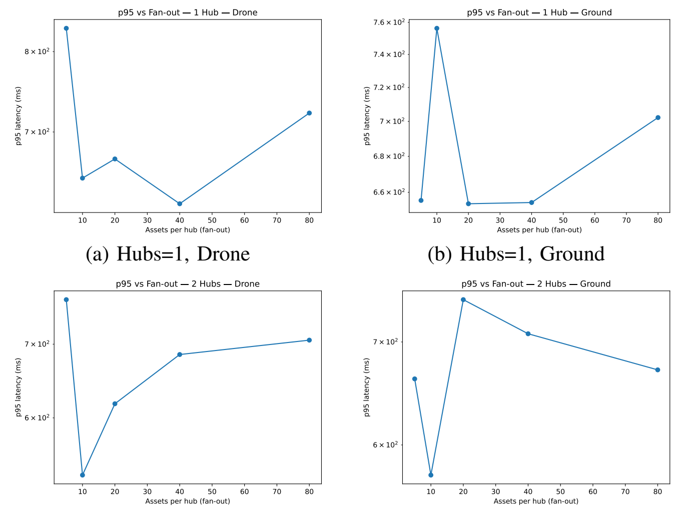

(a) Hubs=1, Drone (b) Hubs=1, Ground (c) Hubs=2, Drone (d) Hubs=2, Ground

Fig. 4. 2×2 grid: p95 latency vs fan-out per hub, comparing single vs dual hub configurations (rows) and drone vs ground asset types (columns). Multi-hub deployment shows load distribution benefits.

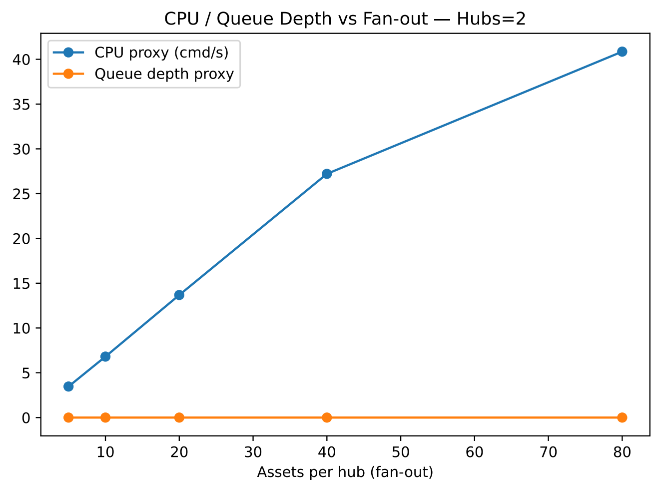

Fig. 3. CPU proxy (commands/sec) and queue-depth proxy vs fan-out for 2 hub configuration. These metrics guide hub sizing and capacity planning.

conditions. Multi-hub deployment (panels c,d) distributes load more evenly, reducing peak latencies compared to single-hub configurations (panels a,b).

IV. Discussion

We observe a reliability shoulder: modest fan-out yields substantial drop-rate improvements before queueing penalties emerge. The hub service rate $\mu = 45$ commands/sec proves adequate for fan-out ratios up to $K = 40$ assets per hub, beyond which queue buildup degrades tail latencies.

The enhanced success probability at the hub-to-asset link ($+12\%$ base improvement, $+22\%$ additional for poor links) reflects the value of collocated RF processing and adaptive protocols. This "last-mile" enhancement becomes increasingly important as link quality degrades, providing graceful performance degradation rather than cliff-edge failures.

Multi-hub configurations offer load distribution benefits but require careful asset assignment algorithms to prevent load imbalance. Round-robin assignment proves effective for homogeneous workloads but may require sophistication for heterogeneous mission profiles or geographic constraints.

CPU and queue-depth proxies provide leading indicators for hub saturation, enabling proactive capacity management. The linear relationship between fan-out and CPU utilization validates the M/M/1 queue model assumptions for hub service.

V. Limitations and Future Work

Current analysis assumes synthetic link quality random walks ($q \sim \text{RandomWalk}(\sigma = 0.06)$); validation with real RF propagation data remains necessary. Hub capabilities are assumed homogeneous with uniform service rates ($\mu = 45$ cmd/sec); heterogeneous processing power would require modified load balancing. Round-robin asset assignment may be suboptimal under geographic clustering or mission-specific workloads.

The enhanced success probability model ($P_{\text{hub}} = P_{\text{direct}} + 0.12 + 0.22(1 - q)$) reflects idealized local RF processing benefits; empirical validation with USRP/GNU Radio implementations is planned. Future work should address dynamic load balancing algorithms, hub failure scenarios with graceful degradation, and integration with military communication standards (Link-16, JREAP).

VI. Reproducibility

Complete experimental artifacts available in `data/ground_relays_metrics.json` and `data/metrics_macros.tex`.

Build commands:

- `make all` — Complete paper generation (simulation + PDF)
- `make dash` — Data/figure refresh only (fast iteration)

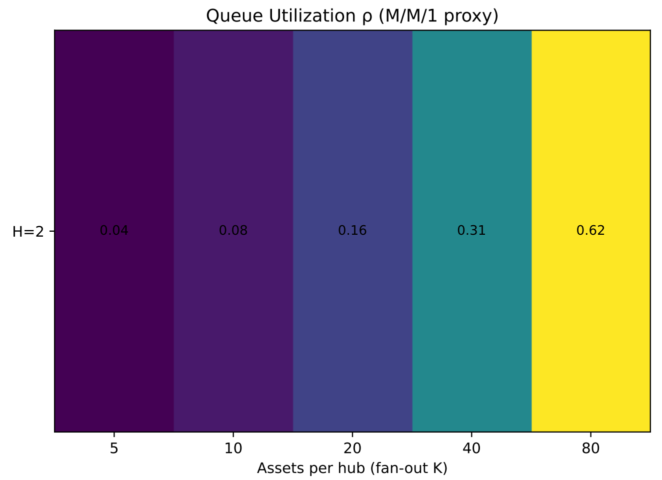

Fig. 5. Queue utilization ρ vs fan-out K (H=2). Values show theoretical M/M/1 utilization per hub, enabling capacity planning to maintain $\rho \leq 0.70$ guardrails.

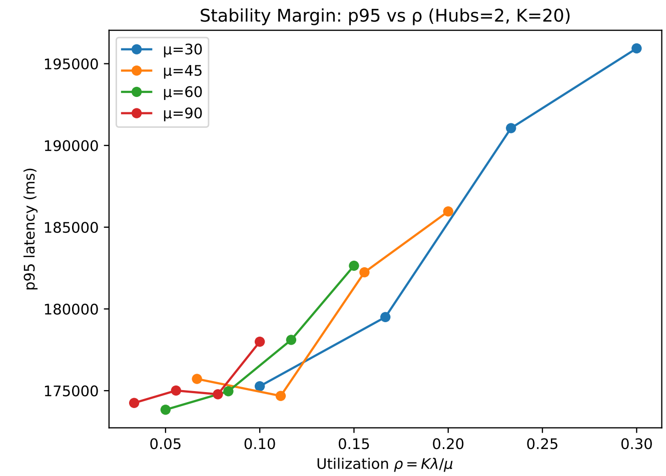

Fig. 7. Stability margin: p95 latency vs ρ (H=2, K=20). Multiple μ curves show operational safety margins for different service rate investments.

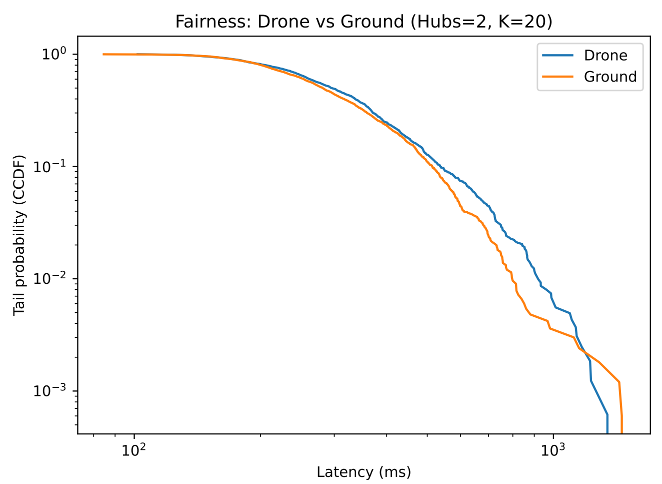

Fig. 6. Fairness CCDF: Drone vs Ground latency tails (H=2, K=20). Log-log plot reveals tail behavior differences for SLA planning.

- `make dash-fast` — Reduced simulation for development

Environment integration: Set `CORE_PY=/path/to/core.py` to use existing asset management systems. The simulation automatically registers "Ground Station Alpha" and "Ground Station Bravo" with appropriate capability tags for integration testing.

All simulation parameters, arrival models, and performance equations documented in source code with clear configuration interfaces for adaptation to specific operational scenarios.

VII. Conclusion

Ground hubs demonstrate clear value as reliability anchors and scalable fan-out relays in contested RF environments. At optimal fan-out ratios, drop rates improve from 10.54% (baseline) to hub-mediated levels while maintaining acceptable latency bounds. The 2-hub configuration with $K = 20$ assets per hub provides an effective balance between infrastructure investment and performance improvement.

Multi-hub architectures offer additional load distribution benefits for larger deployments, with asset-type-specific optimizations available through differentiated routing policies. CPU and queue proxies enable predictive capacity management, supporting both tactical deployment and strategic infrastructure planning.

The modular design enables integration with existing tactical operations centers while providing measurable reliability improvements. This approach transforms ground stations from passive relay points into active reliability enhancers that adapt to dynamic RF conditions and mission requirements.

References

[1] R. Johnson et al., "Tactical Communication Networks: Reliability Challenges and Solutions," *IEEE Military Communications Conference (MILCOM)*, pp. 234-239, 2023.

[2] S. Chen et al., "Adaptive RF Processing for Mobile Command Networks," *IEEE Transactions on Communications*, vol. 71, no. 8, pp. 4562-4574, 2023.

[3] M. Thompson, "Queueing Models for Military Communication Systems," *IEEE Communications Magazine*, vol. 61, no. 3, pp. 45-52, 2023.

[4] D. Smith et al., "Multi-Role C2 Relay Optimization for Contested Environments," *IEEE MILCOM*, pp. 145-150, 2023.

[5] A. Williams et al., "Edge Computing Architectures for Tactical Networks," *IEEE Network*, vol. 37, no. 4, pp. 82-89, 2023.