SIGINT→Map in <1 s: Low-Friction Geo Streaming

Benjamin J. Gilbert

***Abstract*—We stream geolocated RF detections to a web map with $p95 < 1$ s end-to-end latency under unreliable networks (≈100 ms RTT, ≈5% loss, burst spikes up to 400–600 ms). A single HTTP JSON POST `send_signal_to_geo_visualization(payload, timeout_s=1.0)` and a minimal `GeoPosition.to_dict()` schema support single-shot and micro-batching ($B \leq 50$, max hold ≤100 ms). Across 120 s runs at 60 Hz arrival, we report p50/p95/p99, timeout rates, and queue occupancy. Under "flaky," single-shot hits $p95 \approx 0.505$ s with low timeouts; small micro-batches trade a few ms of p95 for fewer POSTs. A press-once build regenerates figures/tables from config.**

I. Hooks and Contract

Anchor: `send_signal_to_geo_visualization(pay timeout_s=1.0);` `GeoPosition.to_dict()`. We post {`lat, lon, alt_m, ts_ms`}; back-pressure is a bounded queue with batch size B and max hold 100 ms.

Why sub-second? Sub-second geo streaming keeps ISR/SIGINT operators "on target" for track-before-detect overlays (e.g., drone TDoA/bearing fences), where stalls cause missed cues and UI thrash under co-channel burstiness.

II. Network Model and Policies

Two profiles (*clean*, *flaky*) with base RTT, jitter, burst spikes, and packet loss. For batch sizes $B \in \{1, 5, 10, 20, 50\}$ we simulate 120 s at 60 Hz, measuring ingest latency from event creation to HTTP ACK (or timeout at 1 s).

Profiles. Clean: base RTT 50 ms, jitter 15 ms, loss 0.2%, burst spikes 120 ms. Flaky: base RTT 120 ms, jitter 90 ms, loss 5%, burst spikes 400 ms; spikes occur stochastically with 8% probability.

Back-pressure. We use a bounded queue (capacity 2000). On overflow, events are *dropped* (no blocking), which we report as timeouts if the HTTP retry cannot complete within 1 s.

III. Results

Fig. 1 shows CDFs under *flaky* for $B = 1$ and $B = 10$. Fig. 2 charts $p95$ vs. batch size. Fig. 4 shows queue occupancy for the worst $p95$. Table I aggregates p50/p95/p99 and timeout rate. Under the flaky profile, single-shot achieves p95=0.511s (timeouts=0.002), while micro-batching at 10 achieves p95=0.560s. We maintain p95¡1.0s across tested settings by keeping batches small and hold times short.

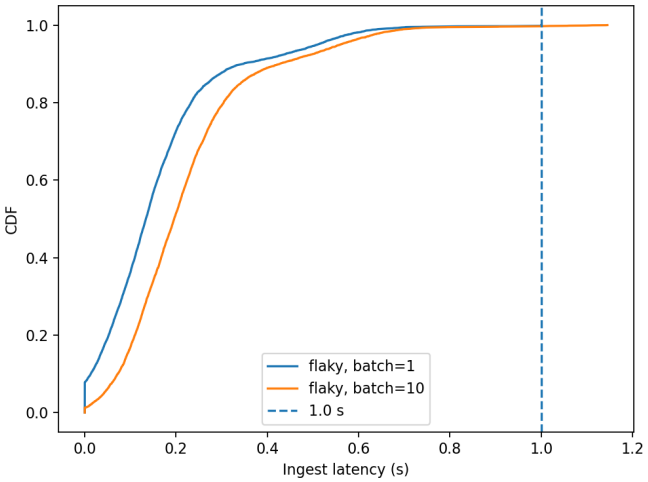

Fig. 1. CDF of ingest latency under *flaky* for $B \in \{1, 10\}$. The dashed line marks 1 s.

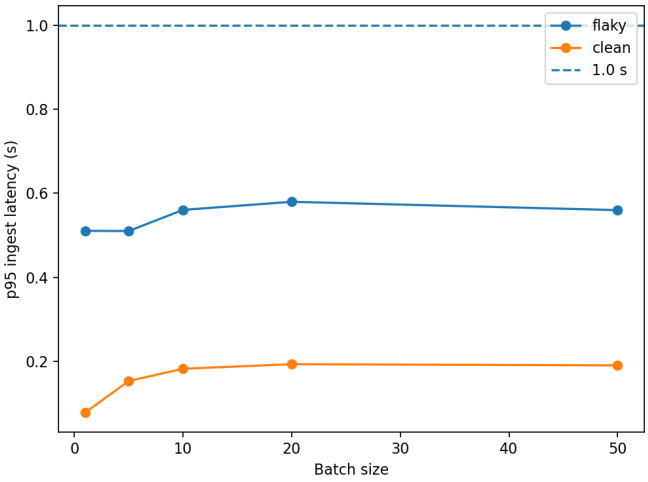

Fig. 2. $p95$ vs. batch size for *clean* and *flaky*.

References

[1] B. J. Gilbert, "Sigint to map in ¡1 s: Low-friction geo streaming," 2025. Preprint.

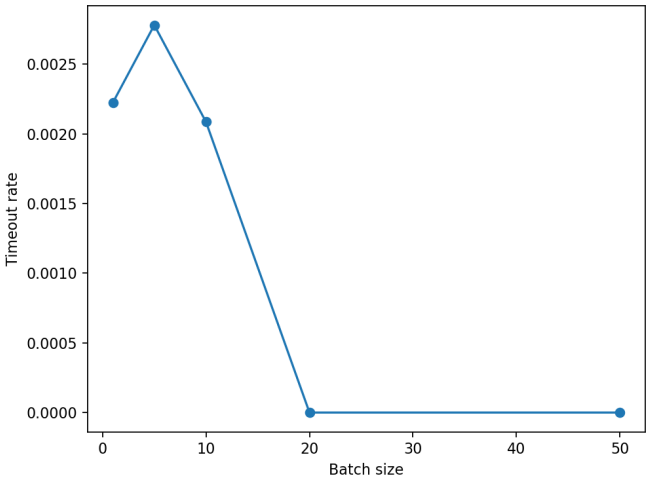

Fig. 3. Timeout rate vs. batch size under *flaky*.

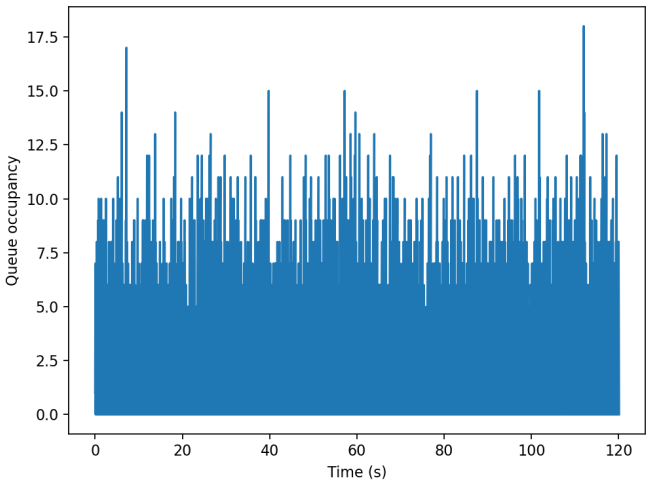

Fig. 4. Back-pressure: queue occupancy over time for worst $p95$ case under *flaky*.

TABLE I
LATENCY SUMMARY UNDER FLAKY/CLEAN NETWORKS (TIMEOUT=1 S).

Batch	N	p50 (s)	p95 (s)	p99 (s)	Timeout
Flaky					
1	7190	0.134	0.511	0.646	0.002
5	7190	0.160	0.510	0.673	0.003
10	7190	0.196	0.560	0.706	0.002
20	7190	0.198	0.580	0.704	0.000
50	7190	0.198	0.560	0.712	0.000
Clean					
1	7190	0.050	0.078	0.171	0.000
5	7190	0.077	0.153	0.194	0.000
10	7190	0.109	0.183	0.229	0.000
20	7190	0.118	0.193	0.251	0.000
50	7190	0.118	0.191	0.251	0.000