

Case Study: MeshCore in Production — Atlanta Repeater Network, June 2026

Loomwave — loomwave.org

Status: evidence document · **Data:** three PYMC repeater captures, analyzed 2026-06-13
Companions: [meshcore-routing-analysis.md](#) (code-level), [meshcore-network-characterization.md](#) (earlier snapshot), [case-study-meshtastic.md](#) (the sibling study)

MeshCore is the "lessons learned from Meshtastic" generation: source-routed unicast, flood only for discovery, listen-before-talk, a repeater role distinct from clients. This study measures how those choices behave on a real metro network — what they fix, what they merely relocate, and where the next structural wall stands.

1. The datasets

	Pine Mountain	KM4BA	KD4HME top-tower
Window	79.8 h	15.9 h	~16 h (two stints)
Packets	16,796	2,753	148
Vantage	high site, network core	secondary repeater	new edge node

All are PYMC repeaters logging every reception with route class, type, SNR/RSSI, duplicate/drop disposition, LBT outcome, and re-transmission decision — a per-packet view of the repeater's actual decisions, richer than Meshtastic's gateway view. The top-tower capture doubles as an edge-of-network reference: its packets arrive at a median -9 dB SNR, the far tail of the same mesh Pine hears at $+1.5$.

2. What MeshCore demonstrably does well

- **Users actually talk on it.** Group + direct text is **26% of all packets at Pine** (grp_txt 21%, txt_msg 5%) — versus Meshtastic's 1–2%. Per packet heard, a MeshCore channel carries an order of magnitude more human payload. The protocol-overhead tax exists (advert 16%) but it does not own the channel.
- **Repeaters are real and disciplined.** Pine re-transmits 42% of what it hears and declines the rest with explicit, logged reasons ("marked do-not-retransmit", "transport code not allowed to flood", duplicate suppression). Forwarding is a decision, not a reflex — structurally ahead of flood-everything.
- **LBT works as designed.** Pine detected channel-busy on 25% of receptions and deferred accordingly. CSMA is genuinely operating, not decorative.

- **The quiet-network premium is real.** 79.8 h at the network's best vantage totals ~1.66% heard utilization; Pine's own re-transmissions add only 0.64%. A text-first, advert-budgeted network is sustainable at current size — exactly the regime MeshCore was designed for.

3. What the data shows it cannot do

3.1 The address space is already spent

Source addresses on the air are 1-byte hashes. Pine observed **254 of the 256 possible values in 3.3 days**; KM4BA saw 205 in 16 hours. The namespace is saturated today: distinct nodes are already indistinguishable on the air, route caches alias, and every new node arrives pre-collided. No tuning fixes an 8-bit identifier; the wire format itself is end-of-life at metro scale.

3.2 Flood didn't leave; it moved in with the furniture

Route classes at Pine: **87% flood, 12% direct**. The design intent — flood for discovery, source-routed direct for traffic — inverts in practice: group text (the dominant user traffic) floods, discovery floods, and 45–73% of everything heard is a duplicate re-hearing. The duplicate fraction at KM4BA (73%) exceeds anything in our Meshtastic captures. MeshCore solved unicast flooding and then the community settled on group chat as the primary mode.

3.3 Path accumulation runs away

Flooded packets accumulate their path in-frame. The hop distribution is nearly flat out to 10 (Pine: 968 packets at 3 hops, 613 at 10) with a tail to **62 hops** — 62 bytes of header on a ~100-byte frame, packets ping-ponging across the metro long after delivery. One packet was dropped for exceeding MAX_PATH_SIZE=64: the protocol's own bound, reached in the wild. Airtime per delivered byte degrades with every hop appended, and nothing prunes it.

3.4 CSMA's bill arrives as latency, and the hidden node remains

When Pine defers (25% busy rate), the cost is queueing: median added **491 ms, p90 1.1 s, max 2.5 s** per re-transmission; KM4BA medians 1.0 s with a 3.6 s max. That is the CSMA trade — collisions exchanged for delay — and it compounds per hop. And LBT only hears what the transmitter can hear: the Stone Mountain geometry (~92% of a hub's nodes mutually hidden, [stone-mountain-hidden-node-hub.md](#)) defeats it by construction. Our own top-tower node spent its first day decoding the far tail at -9 dB median SNR — the edge of a CSMA network is permanent twilight.

3.5 Brittleness where source routes meet reality

"Direct: no path" — a unicast packet arriving at a repeater with no usable route — occurred 572 times at Pine (28% of all direct traffic). The companion code analysis predicted exactly

this failure class: cached paths go stale, and recovery is a fresh flood. The data confirms the prediction at a 1-in-4 rate.

4. What could be done staying on MeshCore

1. **Widen the address.** A 2-byte `src_hash` would defer saturation by $\sim 250\times$; even this single change requires a coordinated wire-format break — but it is one break, and the project is young enough to take it.
2. **Path compression/pruning** for flood packets (cap accumulated path at ~ 8 , or hash the prefix): directly recovers the runaway-header airtime and bounds frame growth.
3. **Group-text scoping:** channel-local groups with $TTL \leq 3$ would pull the dominant traffic class out of metro-wide flood; this is policy plus modest code.
4. **Advert budget enforcement** (interval floors are already in PYMC config; making them protocol-mandatory caps the 15–16% advert share).
5. **Repeater-set coordination** — even static TDM among the handful of high sites (Pine, Sweat, Stone) would lift the 25% busy-defer rate at the core. The repeater role exists; what's missing is any peer protocol between repeaters.

These are real improvements and the project's pace suggests some will happen. They extend the runway; they do not change the destination.

5. The boundary: what iteration cannot fix

- **No schedule, no hidden-node answer.** LBT is the ceiling of what sender-side politeness can do, and the measured geometry defeats it. Scheduling requires a time-authority role and a registration state machine — a different protocol.
- **Discovery-by-flood scales as flood.** Adverts at 16% of packets at current size grow superlinearly with N ; the address saturation makes the response (more adverts, re-advertisement on collision suspicion) worse exactly when it is least affordable.
- **No capacity adaptation:** one preset per network, the slowest link taxes the whole metro — same wall as Meshtastic (SF7/BW62.5 here; the edge node decodes at the floor while close-in pairs could run an order of magnitude faster).
- **Latency under load is unbounded** by design: CSMA backoff has no admission control above it, so p90 delay grows with utilization with no mechanism to shed or schedule.

MeshCore is the strongest argument that iterating on the flood paradigm produces a better flood — measurably better (26% user payload vs 2%) — and still hits the same four walls: hidden nodes, discovery cost, single modulation, unbounded latency. Those walls are why Loomwave's design decision record commits to scheduled cells, registered membership, coordinator-owned modulation, and 3-byte self-certifying addresses ([design-decision.md](#) , [packet-format.md](#)).

6. Methodology notes

Captures via `tools/capture_pymc.py` against three PYMC repeaters' APIs (metadata only; encrypted payloads never stored — lengths only). Pine/KM4BA airtime computed from frame lengths at SF7/BW62.5 (their device-reported airtime fields were empty in this PYMC version). Utilization figures are per-vantage (what that repeater's radio experienced), not network-wide integrals. `src_hash` saturation counts raw distinct 1-byte values; some inflation from per-boot identity changes is possible but cannot rescue an 8-bit space. The top-tower capture is thin (148 packets) and used only as an edge-SNR reference, not for rate statistics.