

White Paper - MeshCore Duplicate Suppression and Randomized Backoff Windows

Introduction

MeshCore operates as a controlled flood-based mesh network, where each node may receive multiple copies of the same packet from different neighbors. Duplicate suppression and the **Randomized Transmit Backoff Window** (configured with txdelay) work together to maintain network stability, reduce collisions, and ensure efficient propagation of packets. This document explains how duplicates are identified and removed, how the Randomized Backoff Window shapes retransmission timing, and how both mechanisms reinforce each other to stabilize MeshCore networks.

Purpose of Duplicate Suppression

Duplicate suppression prevents redundant packets from overwhelming the network. In a flood topology, every node forwards packets, which naturally produces multiple copies of the same message. Without suppression, these duplicates would consume airtime, increase collision probability, and destabilize routing.

Key objectives include:

- **Airtime preservation** — Reducing unnecessary retransmissions to maintain channel availability.
- **Collision reduction** — Preventing packet storms and correlated retransmissions.
- **Routing stability** — Avoiding propagation of stale or conflicting routing metadata.
- **Energy efficiency** — Minimizing processing load on battery-powered nodes.
- **ACK flood control** — Ensuring acknowledgment packets do not saturate dense networks.

Conditions for Identifying Duplicates

MeshCore classifies a packet as a duplicate when it meets one or more of the following criteria:

- **Matching identity** — Same (originNodeID, packetID) pair within the packet's TTL window.
- **Non-improving hop count** — Hop count equal to or worse than the previously received version.
- **Inferior routing hints** — RSSI, SNR, or forwarder metadata is worse or inconsistent.
- **Already forwarded** — A node forwards a packet only once; later copies are suppressed.
- **ACK saturation** — Once reliability thresholds are met, additional ACKs are discarded.

Internal Duplicate Suppression Flow

MeshCore's suppression engine follows a deterministic sequence:

1. **Packet reception** — A packet arrives from a neighbor.
2. **Identity lookup** — The (origin, packetID) pair is checked against the recent-packet table.
3. **Metadata evaluation** — Hop count, RSSI/SNR, and arrival timestamp are compared.
4. **Decision** — If the packet is new or meaningfully better, metadata is updated; otherwise, it is dropped.
5. **Forwarding rule** — Even if metadata improves, retransmission never occurs more than once.

This ensures predictable behavior and prevents routing loops or oscillation storms.

The Randomized Transmit Backoff Window

The **Randomized Transmit Backoff Window** determines *when* a node retransmits a packet. It introduces a randomized delay proportional to network density and congestion. This spreads retransmissions across time, reducing the likelihood of simultaneous transmissions and collisions.

The backoff window is influenced by:

- **Estimated neighbor count**
- **Observed duplicate density**
- **Collision probability**
- **Link asymmetry indicators**

By adjusting the retransmission timing, the Randomized Backoff Window reduces correlated duplicates and improves the quality of the first-arriving packet.

Interaction Between Duplicate Suppression and the Randomized Backoff Window

Duplicate suppression and the Randomized Backoff Window reinforce each other in several ways:

Improved First Arrival Quality

The first packet to arrive after the Randomized Backoff Window is typically:

- the shortest path,
- the strongest link,
- the least collision-affected.

This version becomes the anchor for routing hints. Later duplicates are compared against it.

Reduced Duplicate Clustering

Without a randomized delay, neighbors retransmit at nearly the same moment, producing correlated collisions and bursts of duplicates. The Randomized Backoff Window spreads these events out, making duplicates easier to evaluate and discard.

Enhanced Metadata Learning

Even discarded duplicates contribute to:

- link-quality estimation,
- asymmetry detection,
- density inference.

The Randomized Backoff Window improves the SNR and timing separation of duplicates, making them more informative.

Feedback Loop Between the Two Systems

MeshCore uses the rate and spacing of duplicates to tune the Randomized Backoff Window. For example:

- **High duplicate density** → widen the backoff window.
- **Sparse duplicates** → narrow the window.
- **Correlated duplicates** → increase jitter.
- **Late-arriving duplicates** → adjust routing-hint TTL and consistency thresholds.

This creates a self-stabilizing loop:

duplicates → density estimate → Randomized Backoff Window → fewer correlated duplicates → more accurate density estimate

Operational Benefits

The combined system provides:

- **Scalability** — Flood networks remain viable even with high node counts.
- **Deterministic behavior** — Predictable airtime usage and routing convergence.
- **Collision mitigation** — Reduced retransmission overlap.
- **Robust ACK handling** — Stable acknowledgment propagation in dense networks.

- **Improved reliability** — Cleaner first arrivals and more accurate routing hints.

Summary

Duplicate suppression and the Randomized Backoff Window form a tightly integrated system that stabilizes MeshCore's flood-based routing architecture. Duplicate suppression ensures only the most relevant packet is processed and forwarded, while the Randomized Backoff Window reduces collision probability and improves the quality of packet arrivals. Together, they enable MeshCore to operate efficiently and reliably across sparse, medium, and dense network environments.