

Resilience Nodes: A Pre-Positioned Mesh MAN with LEO Backhaul for Disaster-Resilient Public-Safety Communications

Joshua Omorodion*, Jinwei Liu*

*Department of Computer and Information Sciences
Florida A&M University, Tallahassee, FL 32307, USA
Email: {joshua1.omorodion, jinwei.liu}@famu.edu

Abstract—Natural disasters can destroy cellular, fiber, and power infrastructure simultaneously, precisely when connectivity is most critical for residents and responders. While low Earth orbit (LEO) satellite systems can restore backhaul, isolated hotspots do not address mixed-device access, local service continuity, or metro-scale coordination. This paper proposes Resilience Nodes, a passive, fault-tolerant metro-area network (MAN) architecture that pre-positions low-cost emergency network units at critical facilities. Nodes remain dormant and activate during disasters to provide Minimum Viable Connectivity (MVC): emergency alerts, SOS messaging, shelter information, and responder coordination. The architecture employs point-to-point directional antennas forming a mesh backbone across 1–3 km urban links, with selected gateway nodes providing LEO satellite backhaul. We present an implementation specification with a per-node bill of materials, a 512 Wh LiFePO₄ energy budget, and a token-bucket MVC scheduler. We evaluate feasibility through (i) an RF link budget showing 27 dB fade margin at 2 km, (ii) an analytical worst-case latency bound of approximately 100 ms for P1 alerts under 1000-user saturation, with the discrete-event simulation measuring P1 p99 of 28 ms in practice, (iii) an energy budget deriving 48 to 72 hour autonomy from the LiFePO₄ pack under graded duty cycling, and (iv) a county-scale cost comparison showing that a 25-node deployment at \$15K–25K is less than a single Cell-on-Wheels deployable.

Index Terms—non-terrestrial networks, public safety, wireless mesh networks, disaster communications, LEO satellite backhaul, metro-area networks, fault tolerance, solar-powered networks

I. INTRODUCTION

Hurricanes, floods, and wildfires damage cell towers, overload radio systems, sever fiber, and interrupt power [1]. The FCC reported that Hurricane Maria destroyed 95.6% of Puerto Rico’s cell sites [2], while Hurricane Katrina disabled over 2,000 cell towers across the Gulf Coast [3]. When infrastructure fails, residents need emergency alerts, responders need coordination channels, and shelters need shared connectivity for diverse devices.

Recent LEO satellite deployments demonstrate that non-terrestrial backhaul can restore connectivity after disasters. Florida deployed Starlink terminals during Hurricane Ian to support hospitals, shelters, and emergency operations [28]. FirstNet provides deployable public-safety assets [4], and T-Mobile’s satellite service enables emergency texting. However, these solutions are either reactive (deployed after damage

occurs), device-specific (requiring supported handsets), or isolated (single hotspots without coordination).

Community mesh networks demonstrate an alternative approach. NYC Mesh operates over 600 nodes with point-to-point links spanning more than 3 km using commodity hardware [5]. Red Hook WiFi, a 35-node solar-powered mesh in Brooklyn, experienced a tenfold usage increase during Hurricane Sandy when other infrastructure failed [26]. The disaster.radio project has shown that solar-powered mesh nodes can operate autonomously on minimal power budgets [6]. These precedents validate that low-cost, self-powered mesh infrastructure can serve communities during crises.

This paper proposes **Resilience Nodes**, combining the pre-positioned reliability of public-safety infrastructure with the low-cost scalability of community mesh networks. Resilience Nodes differ from Red Hook WiFi and similar community meshes in three specific ways. First, deployments are pre-positioned for disaster activation rather than retrofitted from existing community networks. Second, the MVC priority scheduler (Section III-D) enforces graceful service degradation under saturation, which community meshes do not. Third, LEO satellite backhaul is a designed component of the architecture, not an improvised addition after infrastructure loss. To address these gaps, we make the following contributions.

- We introduce a Minimum Viable Connectivity (MVC) model with six priority classes for disaster public-safety services, enabling graceful degradation as network resources become constrained.
- Resilience Nodes provides a passive, fault-tolerant metro-area network architecture using point-to-point directional links for the mesh backbone and LEO satellite connectivity for gateway backhaul.
- We present a COTS-based implementation specification with a per-node bill of materials, an estimated cost of \$200–550 per node, and an analytical energy budget for solar-powered autonomous operation.
- We evaluate the feasibility of the proposed architecture through RF link-budget analysis, comparison with existing emergency-communication solutions, and validation against deployed community mesh networks.

II. DISASTER COMMUNICATION REQUIREMENTS

We define *Minimum Viable Connectivity* (MVC) as the smallest set of communication services required to protect life and coordinate response during infrastructure failure. MVC comprises six priority classes:

P1, Emergency Alerts: IPAWS [7] and WEA broadcast, requiring no uplink if cached locally. **P2, SOS and Distress:** Text-based help requests and medical emergencies at less than 1 KB per message. **P3, Responder Coordination:** Situational awareness, resource requests, and inter-agency messaging. **P4, Shelter and Medical:** Capacity updates, triage status, and pharmacy access. **P5, Family Check-In:** Safe-status messages and reunion registry. **P6, General Information:** Cached news, maps, and restoration timelines.

A. Traffic Characteristics

MVC traffic is predominantly small-message and high-priority, with lower classes tolerating delay. P1 and P2 messages average less than 1 KB and require delivery within 5–15 minutes. P3 and P4 average 1–10 KB with 30-minute tolerance. P5 and P6 are opportunistic. Even at 1,000 concurrent users across 25 nodes, aggregate uplink stays below 500 Kbps, well within a single LEO terminal's 50–200 Mbps downlink capacity [8]. The bottleneck is not bandwidth but rather *coverage and availability*, meaning that users must be able to reach at least one access point.

B. Regulatory and Operational Context

Resilience Nodes operate in the 5 GHz unlicensed Part 15 band and interoperate with licensed public-safety spectrum such as FirstNet Band 14 [4]. CBRS Band 48 [27] offers a regulated alternative where Part 15 congestion is a concern. P2 SOS messages reach the PSAP via gateway backhaul using MVC priority aligned with 3GPP MCPTT [23], NPSTC priority guidance [24], and DHS SAFECOM evolution recommendations [25].

III. SYSTEM MODEL

A. Concepts and Assumptions

A Resilience Node network $G = (V, E)$ consists of nodes $V = \{v_1, \dots, v_n\}$ deployed at critical facilities (hospitals, fire stations, schools, and emergency operations centers) and edges E representing point-to-point wireless links. Each node v_i has attributes: backhaul capability $B_i \in \{0, 1\}$, battery state-of-charge $P_i \in [0, 1]$, link quality vector $\bar{L}_i$, and site priority $S_i \in [0, 1]$.

We assume: (1) facilities are spaced 0.5–3 km apart in urban and suburban areas [5]; (2) line-of-sight or near-line-of-sight exists between roof-mounted antennas; (3) three to five nodes per 25-node network have satellite backhaul, either through existing facility internet or a dedicated terminal; and (4) solar irradiance averages at least 4 kWh/m²/day in the target deployment region [9]. Pre-deployment site surveys resolve marginal links through 1–2 m antenna-height adjustment, consistent with NYC Mesh practice on paths beyond 3 km [5]; and since measured Starlink availability exceeded 99.7% during Hurricane Ian [28], each gateway pairs its satellite uplink with a mesh-relay path so no single backhaul is a point of failure.

B. Problem Statement

Given n nodes with heterogeneous backhaul access, maximize MVC service availability A across the coverage area during infrastructure failure:

$$\max A = \frac{1}{|U|} \sum_{u \in U} \mathbf{1} [\exists v_i \in V : d(u, v_i) \leq r_i \wedge \text{active}(v_i)], \quad (1)$$

where U is the population set, $d(u, v_i)$ is distance from user u to node v_i , and r_i is the effective coverage radius (approximately 100–300 m for Wi-Fi access). The constraint is autonomous operation for at least 72 hours without grid power.

C. Head-Node Election

Nodes with satellite backhaul serve as gateways. Algorithm 1 computes a weighted score H_i over the alive set V , combining attributes $B_i, P_i, \bar{L}_i, C_i, S_i$ under weights $\alpha, \beta, \gamma, \delta, \epsilon$; only satellite-equipped candidates ($B_i=1$) are eligible. It runs in $O(|V|)$ per round, invoked at boot, on heartbeat timeout (Section IV-B), and on detected backhaul loss. A school with existing internet becomes a gateway automatically.

Algorithm 1 Head-Node Election

Input: Alive node set V with attributes $B_i, P_i, \bar{L}_i, C_i, S_i$; heartbeat timeout $\tau = 30$ s; weights $\alpha=0.30, \beta=0.25, \gamma=0.20, \delta=0.10, \epsilon=0.15$

Output: Elected head node v^*

```

1: candidates  $\leftarrow \{v_i \in V : B_i = 1 \wedge \text{alive}(v_i)\}$ 
2: for all  $v_i \in \text{candidates}$  do
3:    $H_i \leftarrow \alpha B_i + \beta P_i + \gamma \bar{L}_i + \delta C_i + \epsilon S_i$ 
4: end for
5:  $v^* \leftarrow \arg \max_{v_i} H_i$ 
6: return  $v^*$ 
```

Here $\bar{L}_i$ is mean link quality to k -nearest neighbors, C_i is hardware capacity, P_i is battery state of charge, and S_i is site priority (EOC/hospital = 1.0, school = 0.6, fire station = 0.4).

D. MVC Priority Scheduler

The six MVC classes are arbitrated by a token-bucket-with-strict-priority scheduler at every Resilience Node uplink (Algorithm 2). Each class $k \in \{1, \dots, 6\}$ holds a leaky bucket with sustained rate r_k and burst b_k , refilled each quantum by $r_k \cdot \Delta t$; the scheduler serves the lowest-numbered backlogged class with sufficient tokens, idling otherwise. Strict priority renders the network-calculus delay bound (2) closed-form for class k provided $\sum_{j \leq k} r_j < C$. Arbitration cost is $O(K)$ per quantum with $K=6$, negligible against link service time.

Algorithm 2 MVC Strict-Priority Scheduler (per uplink quantum)

Input: Class queues $Q_1, \dots, Q_6$; token buckets tok_k with rate r_k , burst b_k ; link capacity C

Output: Packet scheduled for transmission

```

1: Refill all token buckets:  $tok_k \leftarrow \min(b_k, tok_k + r_k \cdot \Delta t)$ 
2: for  $k = 1$  to 6 do
3:   if  $Q_k \neq \emptyset$  and  $tok_k \geq |Q_k.\text{head}|$  then
4:      $pkt \leftarrow Q_k.\text{dequeue}()$ 
5:      $tok_k \leftarrow tok_k - |pkt|$ 
6:     transmit  $pkt$   $\triangleright$  service time =  $|pkt|/C$ 
7:     return
8:   end if
9: end for
10: idle  $\triangleright$  no class has conforming traffic
```

Applying the network-calculus service-curve framework [29] to a constant-rate server of capacity C shared by leaky-bucket-shaped flows, the worst-case queueing delay for class k satisfies

$$D_k \leq \frac{b_k + \sum_{j < k} (b_j + r_j D_k)}{C - \sum_{j < k} r_j}, \quad (2)$$

which, solved for D_k , yields a closed-form upper bound provided $\sum_{j < k} r_j < C$.

Instantiating (2) at $C=500$ Kbps yields the per-class bounds in Table I, all comfortably inside the MVC targets.

TABLE I
MVC SCHEDULER ALLOCATION AND WORST-CASE DELAY BOUNDS AT $C=500$ KBPS.

k	r_k (Kbps)	b_k (KB)	D_k bound	Target SLA
1	50	4	≈ 100 ms	< 15 min
2	100	8	< 500 ms	< 15 min
3	150	16	< 30 s	< 30 min
4	100	16	< 2 min	< 30 min
5	50	32	opportunistic	best effort
6	50	32	opportunistic	best effort

Under flat best-effort FIFO, P1 packets are dropped at the same rate as the P5 and P6 traffic share once offered load exceeds C , whereas strict-priority arbitration guarantees P1 delivery as long as r_1 meets the offered P1 load, motivating the design.

IV. SYSTEM DESIGN

A. Layered Architecture

Fig. 1 illustrates the Resilience Node architecture and its external interfaces. Each node implements five functional layers:

Layer 1 (Access): Open emergency Wi-Fi (802.11n/ac, 2.4 and 5 GHz) with captive portal for MVC services.

Layer 2 (Edge Services): Cached IPAWS alerts, offline maps, shelter data, message queue, and family check-in forms. This layer operates independently of backhaul.

Layer 3 (Mesh Backbone): Point-to-point 5 GHz directional links (16–23 dBi). Routing uses B.A.T.M.A.N. Advanced [11] at Layer 2 over the 802.11 MAC, optionally with 802.11s [10] mesh path selection enabled as backup.

Layer 4 (Backhaul): LEO satellite, LTE/5G, fiber, or relay through the mesh to the nearest gateway. Paths are scored by availability and priority.

Layer 5 (Control): Health monitoring, activation engine, head-node election, MVC traffic shaping, and over-the-air firmware updates.

B. Operating Modes and Failover

Three modes govern node behavior: **Dormant** (heartbeats and content updates, < 2 W draw), **Alert** (cache refresh, neighbor scan, backhaul verification, 60–120 s activation), and **Disaster** (all layers active, mesh established, MVC prioritization enforced). Failover relies on signed 10 s heartbeats; three missed beats trigger local re-election. With no reachable backhaul the partition enters *island mode*, queueing messages via delay-tolerant networking [12] until connectivity restores.




Fig. 1. Resilience Node functional architecture and external interfaces. The five-layer stack runs on every node; gateway nodes additionally activate the LEO backhaul shown above. Head-node election (Algorithm 1) runs in the Control layer, and the MVC priority scheduler (Algorithm 2) runs at the egress of the Mesh Backbone layer.

C. Integration with Existing Public-Safety Systems

Resilience Nodes plug into the incident-response stack rather than replace it, following the DHS SAFECOM interoperability model [25], at three points. (i) *Responder radios*: the mesh carries no push-to-talk voice; where First-Net (Band 14 [4]) retains coverage, an EOC-sited gateway lets those handsets reach the shelter and triage data cached at Layer 2 and relays priority messages under the 3GPP MCPTT model [23]. (ii) *The 911 path*: an SOS typed into the open Wi-Fi is converted to NG911 and carried over the ESInet to the PSAP, landing in the same dispatch queue as a cellular 911 text. (iii) *EOC software*: shelter-capacity and resource forms export as structured records into the county WebEOC/CAD dispatch system, giving commanders a live picture of otherwise-dark neighborhoods. Gateway backhaul runs through an encrypted VPN tunnel to the EOC, and backhaul is deliberately multi-band: where the LEO uplink or 5 GHz backbone is obstructed, a node relays over sub-GHz LoRa in the manner of disaster.radio [6], trading throughput for the range and building penetration that keep P1 and P2 traffic flowing within Part 15 limits.

V. IMPLEMENTATION

A. Hardware Platform

All node classes share a base architecture: a router with directional antenna, solar power, battery, and (at gateways) an SBC plus optional satellite terminal. Mesh radios are MikroTik SXTsq 5 ac (\$49–65) or Ubiquiti LiteBeam 5AC Gen2 (\$89–129), running OpenWrt with B.A.T.M.A.N. Advanced (Layer 2) over the 802.11 MAC, optionally with 802.11s path selection. They integrate 16–23 dBi antennas in weatherproof enclosures, achieving 1–3 km point-to-point links at 5 GHz, a range validated operationally by NYC Mesh [5]. A secondary omnidirectional radio provides local Wi-Fi within 100–300 m. Gateway nodes add a Raspberry Pi 5 (\$80) for edge services;

TABLE II
PER-NODE BILL OF MATERIALS (USD).

Component	Model	Cost
<i>Mesh-only node</i>		
Radio + antenna	MikroTik SXTsq 5ac	49–65
Wi-Fi AP	2.4 GHz omni	30
Battery	LiFePO4 40Ah 12.8V	90–130
Solar + MPPT	50W + 10A ctrl	60–105
Enclosure + cable	NEMA 4X, mast, PoE	73–120
Subtotal		302–450
<i>Gateway add-ons</i>		
Compute	RPi 5 + SD	85–110
Battery upgrade	100Ah delta	120–180
Starlink (opt.)	Starlink Mini	599
GW total	no/with sat	507–740 / 1106–1339

mesh-only nodes run services on the router. The complete per-node bill of materials and cost methodology appear in Section V-F.

B. Point-to-Point Mesh Backbone

The backbone uses dedicated 5 GHz point-to-point links between roof-mounted directional antennas, with two to four links per node for redundancy. B.A.T.M.A.N. Advanced [11] at Layer 2 enables transparent bridging and selects optimal paths based on throughput (TQ metric), re-routing around failures within seconds.

C. Autonomous Power System

A 30–100 W monocrystalline panel feeds an MPPT charge controller into a 20–100 Ah LiFePO4 pack (>2,000 cycles, no thermal runaway risk [13]). A 10–50 F supercapacitor preserves the real-time clock and allows orderly state persistence below 5% SoC; the nodes rest in a dormant state until solar restores charge above 15%. The numerical energy budget appears in Section V-G.

D. Duty Cycling

Nodes implement graded duty cycling tied to battery state: full active above 50% SoC, reduced (40 s/min, P1–P3 only) at 20–50%, minimal (30 s per 5 min, P1–P2 only) below 20%, and graceful shutdown below 5%. This mirrors disaster.radio’s solar duty-cycling [6] with a 50–100× larger energy budget.

E. Physical Deployment

Nodes mount on facility rooftops with grounding per NEC Article 810 [22], antenna 3–10 m above roofline. The facility provides mounting, physical security, and (optionally) backhaul. No trenching, tower construction, or spectrum licensing is required under 5 GHz Part 15 [16].

F. Bill of Materials and Cost Model

Table II itemizes the per-node bill of materials. Mesh-only nodes total \$302–450 in component cost (reconciling to the \$200–300 per-node street price after volume procurement), while gateway nodes add \$205–290 in compute and battery upgrades, or \$804–889 including a Starlink terminal at sites lacking existing internet. Table III compares the methodology behind the cost figures for the alternative architectures.

TABLE III
COST AND RESILIENCE COMPARISON FOR COUNTY-SCALE EMERGENCY COVERAGE. SPOF: SINGLE POINT OF FAILURE; NODE-LOSS RESILIENCE IS MEAN GATEWAY REACHABILITY AT 30% NODE LOSS (SECTION VI).

Solution	Unit cost	Sites	Annual	Node-loss resil.
Resilience Nodes	\$600–1K	25	\$500	83% @ 30% loss
COWS trailer	\$50–150K	1	\$15K	SPOF
FirstNet Deployable	>\$200K	1	Bundled	SPOF
Starlink hotspots	\$3K total	5	\$7.2K	Isolated

G. Energy Budget

Each node carries a 40 Ah, 12.8 V LiFePO4 battery, giving a nominal capacity $C_{\text{batt}} = 512$ Wh. With a measured active draw $P_{\text{active}} = 10$ W, autonomy under a duty cycle fraction $d \in (0, 1]$ follows

$$E_{\text{node}} = \frac{C_{\text{batt}}}{P_{\text{active}} \cdot d} \cdot \eta, \quad (3)$$

where $\eta \approx 0.9$ accounts for inverter, MPPT, and supercap losses. Let $E_{\text{full}} = E_{\text{node}}(d=1) \approx 46.1$ h and $E_{67} = E_{\text{node}}(d=0.67) \approx 68.8$ h; these match the 48 h and 72 h autonomy targets within engineering tolerance. Under the below-20% regime ($d \approx 0.10$), endurance extends to several days. Solar replenishment from a 50 W panel at 4 peak-sun-hours delivers 200 Wh/day, exceeding the 161 Wh/day consumed at $d=0.67$, so steady-state operation is feasible under nominal insolation. The 48 h and 72 h autonomy figures are thus derived from physical parameters rather than assumed.

H. Ownership and Governance

Capital cost can draw on county emergency-management budgets or federal EMPG grants, and maintenance is shared with host staff on existing fire-inspection cycles, with quarterly mesh-formation drills coordinated by the county EOC. On private sites such as hospitals or universities, the county EOC owns the nodes under a standard facility-use agreement, leaving the host only rooftop access and physical security. In a declared disaster, spectrum coordination falls to the EOC under the 47 CFR §95.1255 emergency exception, so no new licensing sits on the critical adoption path.

VI. PERFORMANCE EVALUATION

A. RF Link Budget Analysis

We validate backbone feasibility using the Friis free-space path loss model [17]. For a 2 km link at 5.8 GHz with LiteBeam antennas (23 dBi each):

$$\text{FSPL} = 20 \log_{10}(d) + 20 \log_{10}(f) - 147.55. \quad (4)$$

At $d=2000$ m and $f=5.8 \times 10^9$ Hz this equals $66.0 + 195.3 - 147.55 = 113.7$ dB, so $P_{rx} = P_{tx} + G_{tx} + G_{rx} - \text{FSPL} = 23 + 23 + 23 - 113.7 = -44.7$ dBm. Against the LiteBeam 5AC MCS7 (150 Mbps) sensitivity of -72 dBm [14], this yields 27.3 dB of fade margin, covering rain attenuation (approximately 1 dB/km at 5 GHz during heavy rain [15]), foliage loss (6–12 dB), and Fresnel zone obstruction. Under worst-case urban conditions with an additional 15 dB of loss, the link still maintains over 12 dB of margin at reduced data rates (MCS3, 40 Mbps), above MVC bandwidth requirements.




Fig. 2. Discrete-event simulation results: (a) p99 latency per MVC priority class at $1.0 \times C$ saturation, showing strict-priority protection of P1 to P4 inside their SLA targets; (b) gateway reachability under four failure scenarios with 5th to 95th percentile error bars across 100 Monte Carlo trials.

B. Cost Comparison

Table III compares Resilience Node cost against existing solutions for equivalent county-scale coverage. At nominal 25-site totals, Resilience Nodes cost about \$20K, versus \$100K for a single COWS trailer, \$200K for a FirstNet deployable, and \$30K for five Starlink terminals, providing coordinated mesh coverage at a fraction of the cost of single-site or non-coordinated alternatives.

A 25-node Resilience Node network costs \$15K–25K in total capital expenditure with near-zero recurring costs, as there are no spectrum fees, no monthly satellite subscriptions for mesh-only nodes, and minimal maintenance. By contrast, a single Cell-on-Wheels (COWS) trailer costs \$50K–150K [4], covers only one location, and requires transport, setup crew, and generator fuel. FirstNet deployable assets exceed \$200K per unit [4]. Five standalone Starlink terminals (\$3K hardware plus \$7.2K per year in service) provide no mesh coordination, no cached services, and cover only 5 of the 25 needed sites.

C. Discrete-Event Simulation

We validate the architecture using a discrete-event simulation with 20 real Leon County, Florida critical-facility GPS coordinates (see Fig. 3). Link feasibility uses Friis FSPL plus ITU-R P.833 vegetation attenuation [30] (8 dB/km at 45% canopy fraction, representative of Tallahassee’s live-oak canopy [33]) and 4 dB urban clutter [31]. Of 190 candidate links, 39 close with over 10 dB margin, forming a connected mesh graph. The MVC scheduler is modeled as a strict-priority server at $C=500$ Kbps with Poisson arrivals [32]. At saturation ($1.0 \times$ load), P1 p99 latency is 27.7 ms and P3 p99 is 88.1 ms, both orders of magnitude inside their SLA targets. Under $1.5 \times$ overload, P1 p99 remains below 28 ms while P3 degrades to 169 ms, confirming strict priority protects critical classes.

Fig. 2(a) shows measured p99 latency per MVC class at saturation: P1 alerts at 28 ms, P3 responder at 88 ms, with graceful degradation for lower classes, all inside SLA targets. Failure scenarios inject random node outages (100 Monte Carlo trials). At Idalia-scale 30% node loss, mean gateway reachability is 82.8% (p5=64.3%, p95=92.9%). Under Michael-scale 60% loss, coverage averages 66.7% (Fig. 2(b)). The 90% nominal coverage reflects two peripheral nodes whose only feasible links are marginal under the vegetation model; antenna height adjustment or link-specific tree trimming would close them.



Fig. 3. Simulated Resilience Node deployment over Leon County, Florida: 20 real critical facilities (hospitals, EOC, fire stations, schools/shelters, universities) joined by the 39 feasible 5 GHz mesh links. Ringed markers are gateway-capable (satellite-backhaul) sites. Imagery: Esri World Imagery.

D. Sensitivity to Key Assumptions

The two most uncertain inputs are the vegetation model and the offered load. Raising ITU-R P.833 canopy attenuation from 8 to 10 and 12 dB/km reduces feasible links from 39 to 23 and 20 respectively, with peripheral marginal-link nodes partitioning first. On the demand side, sweeping offered load from $0.5 \times$ to $1.8 \times$ leaves P1 p99 flat within 27–28 ms, because strict priority makes P1 latency a function of r_1 and b_1 , not aggregate load; denser populations degrade only best-effort classes. Autonomy scales inversely with active power, spanning 58 h at 8 W to 38 h at 12 W around the nominal 46 h.

E. Disaster Survivability

Because the network must outlast the event that triggers it, survivability is layered. *Physically*, the NEMA 4X enclosure and wind-rated mast follow the engineering that keeps rooftop cellular antennas standing, and roof mounting places electronics above typical flood stage. *Topologically*, the failure sweep (Fig. 2b) is a cascading-loss model, since each lost node also removes its links; connectivity degrades gracefully to 82.8% at 30% loss and 66.7% at 60% loss rather than collapsing, because any surviving backhaul-equipped node can become a gateway. *Temporally*, a node that loses all backhaul enters island mode and keeps serving local alerts and messaging, queueing SOS traffic until a link or the LEO uplink returns. The residual risk is simultaneous loss of every gateway, a limit shared by every terrestrial alternative.

VII. RELATED WORK

A. Emergency Communication Systems

Wang *et al.* [1] survey emergency communication across satellite, ad hoc, cellular, and hybrid architectures, identifying rapid deployability and interoperability as the open challenges; Resilience Nodes target deployability directly, since pre-positioned hardware eliminates convoy-and-tower-erection latency. Debnath *et al.* [18] propose a QoS classification framework for emergency networks with message-priority tiers; the six-class MVC model in Section II instantiates that

framework for disaster-only operation. The FCC Hurricane Maria report [2], Kwasinski's Katrina analysis [3], and a longitudinal cellular-restoration study [34] establish the 24–72 h recovery gap that motivates pre-positioning.

B. Wireless Mesh Networks

Portmann and Pirzada [19] apply WMNs to public-safety crisis management. The IEEE 802.11s standard [10] defines mesh networking at the MAC layer, while B.A.T.M.A.N. [11] provides a proven Layer-2 routing protocol used by Freifunk's 55,000-node network. Unlike these general WMN proposals, Resilience Nodes are purpose-built for passive pre-positioning with disaster-triggered activation.

C. Non-Terrestrial, Solar-Powered, and Community Networks

Lin *et al.* [20] and Saad *et al.* [21] cover 3GPP NTN integration across LEO, MEO, and GEO constellations (Release 17/18); our architecture uses satellite as one backhaul option among many. Community networks such as NYC Mesh [5], guifi.net, and Freifunk demonstrate operational sustainability at scale. Resilience Nodes differ from community meshes in being institution-hosted, passive until disaster, and designed for MVC prioritization rather than general internet access.

VIII. CONCLUSIONS AND FUTURE WORK

This paper proposes Resilience Nodes, a passive, fault-tolerant metro-area network using LEO backhaul for public-safety critical services. The architecture pre-positions low-cost mesh nodes at critical facilities, eliminating the 24–72 h cellular-restoration gap. Contributions include the MVC priority model for graceful service degradation, a COTS implementation specification at \$200–550 per node, and a discrete-event simulation across 20 real Leon County facilities showing 82.8% gateway reachability under Idalia-scale 30% node loss. Key limitations are that 5 GHz line-of-sight may require site surveys, satellite terminals remain the costliest component, and the delay figures are analytical upper bounds. Future work includes ns-3 protocol simulation, terrain-aware propagation modeling, field-pilot deployment at three to five facilities, equity-aware site selection using census-tract-level Social Vulnerability Index (SVI) data, IPAWS integration for automated alert injection, and a WPA3-Enterprise security framework with signed firmware and heartbeat anomaly detection.

ACKNOWLEDGMENT

This research was supported in part by the U.S. NSF under Grant DUE-2400459 and Grant DBI-2417643; in part by the NTIA/CMC Grant 12-09-C13055; in part by the title III grant; and in part by the generous funding from the Cyber Policy Institute at Florida A&M University. Dr. Jinwei Liu is the corresponding author.

REFERENCES

- [1] Q. Wang, Z. Fang, X. Xu, and J. Jiang, "An overview of emergency communication networks," *Remote Sensing*, vol. 15, no. 6, art. 1595, 2023.
- [2] Federal Communications Commission, "In the matter of the Uniendo a Puerto Rico fund and the Connect USVI fund," FCC Report WC Docket No. 18-143, Aug. 2018.
- [3] A. Kwasinski, W. Weaver, P. Chapman, and P. Krein, "Telecommunications power plant damage assessment for Hurricane Katrina, site survey and follow-up results," *IEEE Syst. J.*, vol. 3, no. 3, pp. 277–287, 2009.
- [4] First Responder Network Authority, "FirstNet deployables," 2024. [Online]. Available: <https://firstnet.gov/network/deployables>
- [5] NYC Mesh, "Network nodes and documentation," 2024. [Online]. Available: <https://www.nycmesh.net/map>
- [6] disaster.radio, "A disaster-resilient communications network powered by the sun," 2020. [Online]. Available: <https://disaster.radio>
- [7] Federal Emergency Management Agency, "Integrated Public Alert and Warning System (IPAWS)," 2024. [Online]. Available: <https://www.fema.gov/ipaws>
- [8] A. Michel, P. Marcos, L. Gracia, and A. Feldmann, "A first look at Starlink performance," in *Proc. ACM IMC*, Nice, 2022, pp. 130–136.
- [9] National Laboratory of the Rockies (formerly National Renewable Energy Laboratory), "Global Horizontal Irradiance, National Solar Radiation Database (NSRDB) Physical Solar Model, 1998–2016," 2024. [Online]. Available: <https://nsrdb.nrl.gov/>
- [10] IEEE, "IEEE Standard 802.11s-2011: Mesh Networking," *IEEE Std 802.11*, amendment 10, 2011.
- [11] B.A.T.M.A.N. developers, "B.A.T.M.A.N. Advanced documentation," Open Mesh, 2024. [Online]. Available: <https://www.open-mesh.org/projects/batman-adv/wiki>
- [12] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proc. ACM SIGCOMM*, Karlsruhe, Germany, 2003, pp. 27–34.
- [13] L. Wang, Y. Wang, and C. Liu, "Lithium iron phosphate batteries for stationary energy storage: A comprehensive review," *Journal of Energy Storage*, vol. 50, art. 104222, Jun. 2022.
- [14] Ubiquiti Inc., "LiteBeam 5AC Gen2 datasheet," 2024. [Online]. Available: <https://store.ui.com/us/en/collections/operator-litebeam>
- [15] ITU-R, "Specific attenuation model for rain for use in prediction methods," Recommendation P.838-3, 2005.
- [16] Federal Communications Commission, "Title 47 CFR Part 15: Radio frequency devices," 2024.
- [17] H. T. Friis, "A note on a simple transmission formula," *Proc. IRE*, vol. 34, no. 5, pp. 254–256, May 1946.
- [18] S. Debnath, A. Saha, and D. De, "A comprehensive survey of emergency communication network and management," *Wireless Personal Communications*, vol. 125, pp. 3435–3467, 2022.
- [19] M. Portmann and A. A. Pirzada, "Wireless mesh networks for public safety and crisis management applications," *IEEE Internet Comput.*, vol. 12, no. 1, pp. 18–25, 2008.
- [20] X. Lin, S. Rommer, S. Euler, E. A. Yavuz, and R. S. Karlsson, "5G from space: An overview of 3GPP non-terrestrial networks," *IEEE Commun. Stand. Mag.*, vol. 5, no. 4, pp. 147–153, 2021.
- [21] M. M. Saad, M. T. R. Khan, and D. Kim, "Non-terrestrial networks: An overview of 3GPP Rel-17 and Rel-18," *IEEE Internet Things Mag.*, vol. 7, no. 1, pp. 20–26, 2024.
- [22] National Fire Protection Association, "NFPA 70: National Electrical Code, Article 810, Radio and Television Equipment," 2023.
- [23] 3rd Generation Partnership Project (3GPP), "Mission Critical Push To Talk (MCPTT) over LTE; Stage 1," 3GPP TS 22.179, Rel-17, 2022.
- [24] National Public Safety Telecommunications Council (NPSTC), "Priority and Quality of Service for Public Safety Mobile Broadband," NPSTC Public Safety Communications Report, 2019.
- [25] U.S. Department of Homeland Security, SAFECOM Program, "Public Safety Communications Evolution," DHS SAFECOM White Paper, 2023.
- [26] A. Powell, "WiFi publics: Producing community and technology," *Inf. Commun. Soc.*, vol. 11, no. 8, pp. 1068–1088, 2008.
- [27] Federal Communications Commission, "Citizens Broadband Radio Service," Title 47 CFR Part 96, U.S. Government Publishing Office, 2024.
- [28] J. Pan, J. Zhao, and L. Cai, "Measuring Starlink satellite network performance during natural disasters," in *Proc. ACM IMC*, 2023.
- [29] J.-Y. Le Boudec and P. Thiran, *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*, Lect. Notes Comput. Sci., vol. 2050. Berlin, Germany: Springer, 2001.
- [30] ITU-R, "Recommendation P.833-10: Attenuation in vegetation," ITU, Geneva, 2021.
- [31] COST Action 231, "Digital Mobile Radio Towards Future Generation Systems, Final Report," European Commission, EUR 18957, 1999.
- [32] L. Kleinrock, *Queueing Systems, Vol. 1*. Wiley, 1975.
- [33] U.S. Geological Survey, "NLCD 2021 Tree Canopy Cover," USGS EROS Center, 2023.
- [34] A. Kwasinski, "Effects of hurricanes Harvey, Irma, and Maria on communication networks," in *Proc. IEEE INTELEC*, 2020.