Frozen Radio Silicon Need Not Freeze Access: Concurrent Satellite-IoT Uplink on Deployed Radios

Rohith Reddy Vennam, William Hunter, Dinesh Bharadia
University of California San Diego, La Jolla, CA, USA
{rvennam, wshunter, dineshb}@ucsd.edu

Abstract

Satellite IoT faces an energy–scale–compatibility trilemma: nodes must conserve energy, networks must absorb many sporadic transmitters, and deployed radios are difficult to replace. Each choice gives up one: LoRa limits concurrency, NB-IoT adds overhead, LR-FHSS needs new node hardware. We argue the trilemma is partly architectural, not fundamental: a radio's advertised packet modes need not bound what its silicon can emit, and an upgradeable gateway absorbs separation. CodeHop gives each transmission its own signature: a node draws a code and hop pattern from shared pools and sends raw symbols as cyclic shifts of that code, so each is a hypothesis the gateway tests by correlation. Emission and recovery are tested separately: two unsynchronized SX1276-class nodes emit the structure from firmware alone, and a Gold-127 simulation recovers 93.6% of raw symbols from 8 fully overlapping equal-power signals at −6 dB each, given distinct codes and exact timing and no packets or FEC. CodeHop opens a broader direction: evolve access in firmware and gateway software, leaving the node's radio silicon frozen.

CCS Concepts

• **Networks** → **Wireless access networks**; • **Computer systems organization** → **Embedded systems**.

Keywords

Satellite IoT, concurrent access, legacy radio silicon, frequency hopping, multiuser reception

ACM Reference Format:
Rohith Reddy Vennam, William Hunter, Dinesh Bharadia. 2026. Frozen Radio Silicon Need Not Freeze Access: Concurrent Satellite-IoT Uplink on Deployed Radios. In *The 25th ACM Workshop on Hot Topics in Networks (HotNets '26), November 16–17, 2026, Salt Lake City, UT, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3847290.3859959

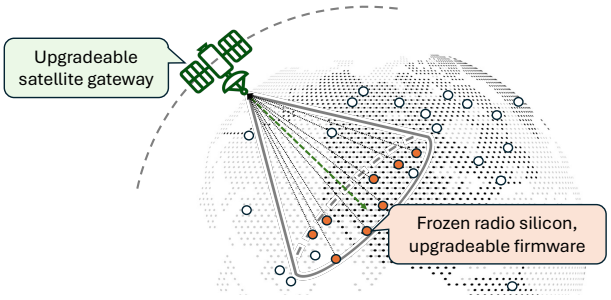


Figure 1: Many nodes share one short satellite pass. Orange nodes inside the cone actively transmit, while white ones are either out of view or inactive. Each node's radio is fixed, but its firmware and the gateway can change: that gap is what CodeHop uses.

1 Introduction

Satellite IoT collects sensor messages beyond terrestrial coverage over a ground-to-satellite uplink. Direct-to-satellite LoRa deployments receive these messages in orbit and relay them through ground stations [7, 8, 21], serving applications such as smart agriculture, logistics, remote utilities, wildlife monitoring, defense, and environmental sensing.

Energy, scale and compatibility form a trilemma. Nodes must conserve energy, the network must absorb many sporadic transmitters, and deployed radios are hard to replace. A single node sends one or two packets during a contact, so no one node strains the link; the fleet does, because many sporadic, unsynchronized, low-power transmitters share a large satellite footprint with timing and carrier uncertainty (Figure 1). What limits a deployment is how many successful packets a satellite collects in a pass. That leaves a fleet with a false choice: preserve its hardware and retain legacy access, or replace its radios to adopt a clean-slate PHY.

Most satellite IoT deployments use LoRa, LR-FHSS or NB-IoT, and each gives up one of these demands. LoRa preserves compatibility with the deployed nodes and their energy efficiency, but cannot handle a high number of concurrent transmissions. Some previous works show that overlapping LoRa packets can be separated [6, 19], but independent evaluations show that low-SNR detection failures can limit the scalability gains from concurrent reception and reduce throughput [15].

This work is licensed under a Creative Commons Attribution 4.0 International License.
HotNets '26, Salt Lake City, UT, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3029-0/2026/11
https://doi.org/10.1145/3847290.3859959

LR-FHSS supports many concurrent transmissions and is energy efficient, but requires new node PHY support [3, 4]. NB-IoT NTN also supports many concurrent transmissions, adding timing, synchronization and resource control, but through a different node PHY and stack [13] and with measured node energy above LoRaWAN's [22]. These are credible designs for supported or new nodes. The missing point lies between gateway-only salvage of a legacy packet and replacement of the node radio.

Frozen radio silicon need not freeze access. That middle ground exists because the upgrade cycle is asymmetric: node radios can remain in the field for years, while a satellite or ground gateway can repeatedly replace its signal-processing software. Yet many physical-layer upgrades require a new modulation engine, coding pipeline, synchronization procedure, or protocol stack at the node. We argue that this trilemma is *partly architectural rather than fundamental*: present designs often place the decisive access machinery on the side that cannot evolve. This claim does not remove link-budget or multiple-access limits. It asks how much searchable structure simple firmware can synthesize from controls the old silicon already exposes.

CodeHop explores that middle ground. Nodes and the gateway share finite code and hopping-pattern pools. At each opportunity a node independently draws an ephemeral code/hop signature; cyclic shifts of its code carry raw symbols, and firmware emits short constant-envelope blocks while retuning across the selected hop pattern. Each transmission therefore arrives with a separation hypothesis, which a common chirp family does not provide, and the gateway must discover the realized choices, search timing and carrier offset, correlate code/shift hypotheses, and apply successive interference cancellation (SIC) when a peel is reliable. Two requirements follow: the waveform must be reachable from firmware-visible controls alone, and the gateway must recover transmissions whose codes and arrival times it does not know. The individual ingredients are familiar; the contribution is their compatibility-constrained composition and control split.

Initial evidence addresses each requirement separately, and supports layers of this architecture rather than an end-to-end PHY. For reachability, two unsynchronized SX1276-class nodes sharing no clock or trigger emit short coded blocks on a 1 ms hopping schedule across eight 125 kHz channels; short hop periods and narrow channels jointly force a 31-chip code on that bench, 6.1 dB below the 127-chip point the simulations use, so the measured and simulated operating points are deliberately distinct. For separation, a controlled, fully overlapping Gold-127 simulation recovers 93.6% of raw symbols at $K = 8$, 88.4% at $K = 10$, and 76.7% with correlation-only decoding under the same conditions, with equal received power, −6 dB per node, known distinct active codes and timing, Gaussian per-node residual carrier offset ($\sigma = 197$ Hz), and no packet or FEC. For acquisition, a separate 12-block experiment searches timing and CFO over a known active set, reaching the same criterion at $K = 6$; the two must not be combined.

We make three contributions: a compatibility-constrained design lens for concurrent access on frozen radio silicon; CodeHop, a structured-random-access waveform skeleton built from firmware-reachable code and retuning primitives; and scoped component evidence with explicit falsifiers for the missing discovery, packet, and coding layers. CodeHop is not a finished satellite-IoT PHY; it is a concrete test of how far access can evolve through node firmware and gateway software before new radio silicon becomes unavoidable.

2 Why Does Frozen Silicon Freeze Access?

Satellite-IoT access is normally designed for a new node; our question is what remains possible when the node's radio silicon is already deployed. Here "frozen" does not mean immutable node software. Host firmware may change, but the radio's modulation primitives, interfaces, and regulatory envelope do not. The relevant comparison is therefore not universal sensitivity, airtime, or capacity. It is the upgrade path available to a fleet whose gateway can change but whose radios cannot acquire an arbitrary new baseband. This distinction separates three compatibility levels that are often collapsed. A gateway-only receiver can decode a legacy emission without touching a node. Updated node firmware can exercise controls that the packaged modem mode did not expose. New radio silicon can add a genuinely new baseband. CodeHop targets the middle level: it does not promise that an untouched legacy packet becomes a CodeHop packet, and it does not require replacing the transceiver. That boundary is operationally meaningful for a fleet that can receive firmware updates but cannot change its radio silicon.

LoRa anchors the processing-gain comparison, not a sensitivity claim. At 125 kHz, LoRa SF7/SF8 provide nominal spreading ratios of 21.07/24.08 dB; Gold-127/Kasami-255 nearly match at 21.04/24.07 dB. Direct-to-satellite modeling selects SF7/SF8 during favorable portions of a modeled LEO pass under its assumptions [2]. Thus CodeHop inhabits satellite-relevant processing-gain regimes, but not proven sensitivity or link closure: duration, modulation, coding, receiver loss, and interference differ.

The remaining problem is what happens when several such signals arrive together. Standard LoRa exposes no independently selected code signature for general same-channel, same-SF multi-packet separation. Commodity reception can capture a packet under such overlap, but does not expose all equal-power, fully overlapping packets as separate streams.

Research receivers such as Choir and CIC recover important classes of LoRa collisions using custom gateway processing [6, 19]. Their success rules out "collisions are inherently undecodable" as our gap. Their demonstrated reach, however, is bounded by link quality: an independent reimplementation of four such receivers reports gains near 0 dB SNR, but finds that around −15 dB most fail at *preamble detection* rather than at separation, and the one receiver still detecting packets delivers lower throughput than standard LoRa [15]. Direct-to-satellite links sit at that low end [21, 23]. This sharpens our question: can updated firmware on the same radio silicon emit deliberately searchable code/hop signatures, allowing an evolving gateway to separate simultaneous nodes? Whether the resulting packet system preserves LoRa-class node energy remains open until acquisition, packet/FEC overhead, link closure, and retransmissions are measured.

Figure 2 contrasts collision granularity: LoRa overlaps span packets, LR-FHSS collisions span fragments, and short CodeHop blocks can avoid overlap for some offsets while full overlaps remain code-structured.

LR-FHSS and NB-IoT NTN require node-side PHY support. LR-FHSS uses narrow hopping fragments, coding, and repeated headers to buy sensitivity and erasure resilience [3, 4]. Working direct-to-satellite receivers show that LEO Doppler requires explicit search and tracking, not that LR-FHSS fails [11]. Semtech enables it in software on supported SX1261 and SX1262 parts [18], but its standardized transmission format is absent from the targeted SX127x modem modes [17]. NB-IoT NTN takes a coordinated cellular path: its terminal participates in timing, positioning, synchronization, and resource control unavailable to a LoRa-class modem [1, 13], and terrestrial measurements put its node energy above LoRaWAN's on the same hardware [22], with the dominant cost in connected-state paging and repeated control messages rather than payload size [14, 24]. Small-data optimizations make generic signaling-overhead attacks unfair; the durable point is that a different cellular PHY and stack cannot be installed by upgrading an SX127x gateway.

These foils also locate what CodeHop does *not* optimize. LR-FHSS trades narrow coded fragments and airtime for sensitivity and erasure resilience; CodeHop explores shorter, channel-width code blocks that bound within-block phase evolution and expose structured overlaps to a compute-rich gateway. This is a different compatibility-constrained operating point, not an optimality result. ChirpTransformer also changes encoding on COTS LoRa radios [12]; CodeHop instead composes FSK/MSK spreading-code blocks, cyclic-shift data and code-aware separation. CodeHop asks which new transmission structures remain expressible after the silicon choice has already been made; we do not claim a matched sensitivity, goodput, or capacity win. Those comparisons require a completed packet and link, which our present evidence does not supply [5].

3 CodeHop Design

The argument so far has been architectural: node radios are fixed while gateway software can change, so the question is what access structure firmware can still synthesize. We answer it concretely here, in three parts. First, what a node emits: a code, a cyclic shift carrying the symbol, a hop placing it in frequency, and the firmware-visible controls that make all three reachable on an SX127x. Second, what the gateway must do: the search and cancellation that turn the mixture back into symbols. Third, the tradeoffs that bind the two, from pool size against search cost to code length against carrier tolerance.

Both sides share the same finite code and hop pools, and that shared contract is what the rest of this section builds on. A node draws an ephemeral signature $s_u = (c_u, h_u)$ from those pools, a code c_u paired with a hop pattern h_u, maps each raw symbol to a cyclic shift of c_u, and emits the resulting blocks along h_u. The gateway knows the pools, but not how many nodes are active or which signatures they drew. This is structured random access, in the spirit of massive uncoordinated multiple access [16] and satellite random-access schemes that resolve collisions via replicas and SIC [5]. Random choice avoids stable per-node assignments, but two nodes drawing the same signature are not separable by code; only their arrival times could distinguish them.

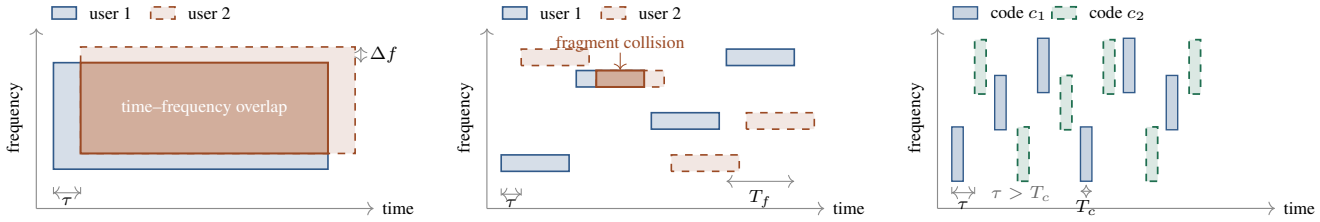
3.1 What the node emits

A code supplies correlation structure; its cyclic shift carries data. For a selected length-N code $c_u[n] \in \{-1, +1\}$, a raw symbol a_u selects a cyclic shift:

$$x_u[n] = c_u[(n - a_u) \bmod N]. \tag{1}$$

This cyclic code-shift keying (CCSK) lets the receiver test possible shifts by correlation: the matching code and shift add coherently, while other codes contribute cross-correlation interference. Distinct codes thus provide different separation hypotheses for overlapping signals, rather than perfect orthogonality.

Modulation and hopping turn code blocks into emissions. Firmware maps the chips to a minimum-shift-keying (MSK) frequency trajectory at modulation index 0.5 and retunes between blocks. Hopping places each block in frequency and opens a path to diversity and cross-hop coding. Distinct codes, distinct hop patterns and relative delay can each help separate transmissions. Our all-user evidence covers distinct codes, which separate even at full overlap; hop-only separation is untested and delay is scored victim-only. CCSK/ MSK



(a) LoRa: overlap spans the packet **(b) LR-FHSS: overlap spans a fragment** **(c) CodeHop: overlap can be avoided**

Figure 2: Three collision granularities (schematic, not to scale). Small time and frequency offsets leave most of two LoRa packets overlapping; the illustrated LR-FHSS offset confines overlap to one hopping fragment; CodeHop's shorter blocks admit offsets with no block overlap ($\tau > T_c$). Panel (c) illustrates one such offset, not the full-overlap SIC regime evaluated in Section 4.

hopping is established in other links [10]. Figure 3 shows the resulting node and gateway chains.

The SX127x exposes a modulation path beyond its packet engine: its frequency-shift-keying (FSK) path has bitrate, deviation and shaping controls that can realize MSK and GMSK [17], and its raw-chip, packet-handling and carrier controls let firmware emit a code block, retune and repeat. It cannot synthesize arbitrary I/Q samples. Our implementation uses fixed code and hop choices; random signature sampling and the joint-acquisition pilot are outside the hardware study.

Applying the design to another radio means checking controls, not advertised features: firmware must reach modulation data, retuning and packet handling at the needed timing and bandwidth. A closed packet interface can block that access even on a radio that documents FSK and hopping. Spectral containment, PA behavior and certification also require radio-specific validation.

3.2 What the gateway must do

As illustrated in Figure 3, the gateway includes five stages. A software-defined front end first captures the band as I/Q samples and channelizes it. Detection comes next, doing the job packet detection does in a conventional receiver, but with no single common preamble to lock onto: it must search the shared pools for which code and hop signatures are active. Acquisition then estimates arrival time, carrier offset and channel for each candidate. Correlation against cyclic shifts of the hypothesized code reads the raw symbols, and ordered cancellation removes what has been read. The chain emits raw shifts; packet and FEC recovery sits downstream and is not composed here.

For an active set $\mathcal{A}$, we model the captured mixture as

$$y(t) = \sum_{u \in \mathcal{A}} \alpha_u x(c_u, h_u, a_u; t - \tau_u) e^{j(2\pi \nu_u t + \phi_u)} + w(t). \quad (2)$$

Here α_u and ϕ_u describe received amplitude and phase, τ_u arrival time, ν_u frequency offset, and $w(t)$ noise. Every per-user parameter here is unknown, as is $\mathcal{A}$ itself and its size.

The unknowns are also coupled. A correlation peak only sharpens once timing and carrier offset are close enough for the block to add coherently, yet those estimates presuppose a code hypothesis and the hypothesis is confirmed by the peak. The five stages are therefore responsibilities rather than a single forward pass: the gateway searches code, hop, timing and frequency together and revisits earlier decisions, and the size of that product is what it pays. Section 4 measures the timing and frequency search with the active set supplied; discovering code and hop signatures blindly remains open.

A confident shift decision opens the way to successive interference cancellation (SIC): estimated amplitude, phase, timing and frequency permit waveform reconstruction and subtraction, and testing the residual can expose peaks that the removed user masked. A wrong shift instead subtracts a full-power wrong waveform and can increase interference, so confidence and stopping rules are essential.

Iteration is what the retained samples buy. The same capture can be re-scored under new hypotheses, and re-scored again after each subtraction; a packet radio emits a demodulated result and discards the evidence behind it. This is where CodeHop's complexity lives, by design: hypothesis search and waveform reconstruction are expensive, but they run where the processing is free to change, not where the radio's controls are fixed. Ordered SIC is evaluated separately from acquisition with timing supplied; the two stages have not been composed.

3.3 Tradeoffs that bind them

Pool size trades selection collisions against search cost. For a joint pool of M signatures and K simultaneously active nodes, independent uniform draws repeat a signature with

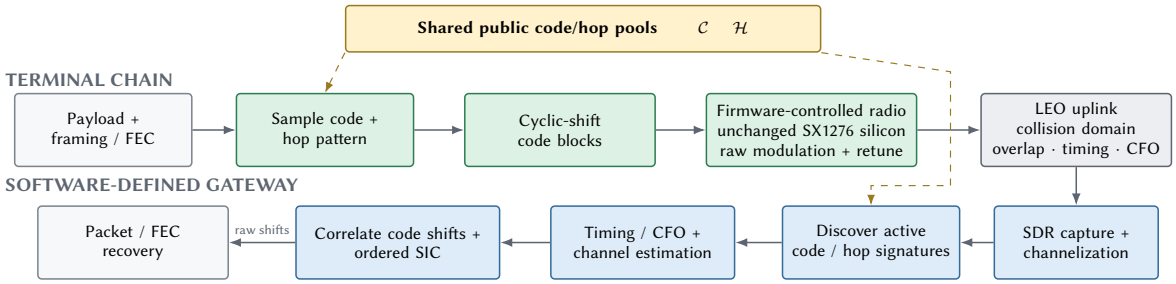


Figure 3: Proposed node and gateway chains. Firmware constructs code/shift blocks and controls the radio; gateway software discovers, acquires and separates signals. Hardware validates emissions; separate simulations study receiver components. Blind discovery and packet/FEC composition remain open.

probability

$$p_{\text{coll}} = 1 - \prod_{i=0}^{K-1} \left(1 - \frac{i}{M}\right), \tag{3}$$

so a larger pool makes repeats rarer but enlarges the space the gateway must search. Node identity, retry behavior and false-alarm control over the timing–CFO–code–hop search space remain open.

Short blocks bound phase curvature, not frequency uncertainty. Bulk Doppler requires search or coarse correction; Doppler rate curves phase within a block; residual CFO rotates phase throughout coherent integration. A residual offset of $1/(4T)$ rotates phase by a quarter cycle over duration T: about 197 Hz for a 1.27 ms Gold-127 block and 97.7 Hz for a 2.56 ms Kasami-255 block. These engineering scales are not estimator guarantees. A long physical fragment in another PHY need not be one untracked coherent observation; LR-FHSS tracking must be included in any comparison.

Delay tolerance depends on the whole transmission train. Arrival differences change which blocks overlap, but a delay larger than one block may move it into another block of a transmission train. A useful signature family therefore needs appropriate cross-correlation across joint delay and Doppler uncertainty, not merely a long self-repetition period. No delay-derived capacity gain is established here.

Code length sets a joint tradeoff among integration time, signature pool and frequency tolerance. Gold-127 at roughly 100 kchip/s integrates for 1.27 ms and offers a 129-member structured family; we test this point for data-bearing CCSK with ordered SIC and, separately, for joint timing and CFO acquisition. Large-Kasami-255 at 99.6 kchip/s integrates for 2.56 ms and offers 4,101 usable code signatures in our audit [25], adding roughly 3 dB of nominal processing gain but tightening CFO tolerance and changing duration. Its present multiuser evidence is signature-only.

Acquiring cyclic-shift data requires a timing reference because both arrival offset and data shifts move a correlation peak. The tested acquisition construction uses four pilot blocks with public cover shifts to resolve the mother code's repeated timing ambiguity, followed by eight data blocks. Four of twelve blocks are pilots, so acquisition overhead must enter an eventual efficiency comparison. The active set remains supplied. Code length, alphabet, signature pool, acquisition overhead and spectral containment must therefore be selected together; nominal processing gain alone does not identify a winner.

Table 1: Hardware emission test: unmodified SX1276, fixed code and hop.

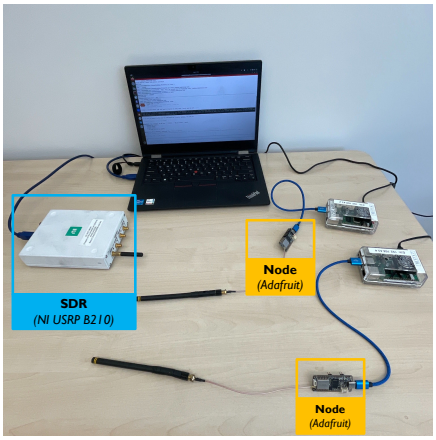
Node	RP2040 + SX1276 ($20 Feather)
Path	FSK/MSK controls, not the LoRa packet engine
Code	31-chip Gold + 1 guard chip, 62.5 kchip/s
Symbol	one of eight cyclic shifts
Slot	1 ms: 512 µs emission, ≈490 µs control budget
Band	eight 125 kHz channels at 915 MHz

4 Initial Validations

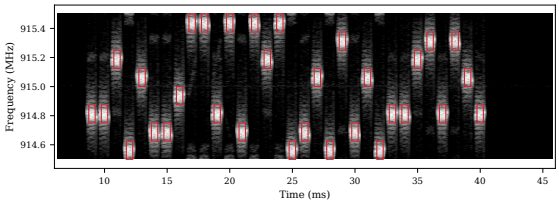
We test two prerequisites for CodeHop: commodity radio emission and concurrent signal recovery. Here K counts simulated signals, not admitted nodes or delivered packets. Separation scores raw cyclic-shift symbols across all users; acquisition transmits all K but acquires and scores $\min(K, 8)$ per trial; signature diagnostics score unmodulated signatures. None of these establishes blind random access, packet goodput or energy per delivered message. K_{90} is the last tested load meeting 90% raw-symbol recovery in the acquisition study, or the median per-seed 90% crossing in the diagnostic; brackets are seed ranges, not confidence intervals.

4.1 Can commodity radios emit the structure?

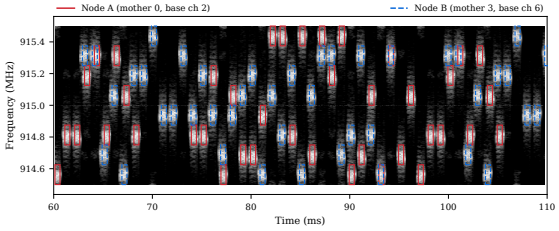
The bench asks whether unmodified SX1276 silicon can hold a coded hopping schedule, alone and then concurrently. Table 1 gives the operating point; the two nodes run from



(a) Measurement setup.



(b) One node: complete 32-slot burst.



(c) Two nodes: concurrent emissions.

Figure 4: Hardware emission feasibility. (a) Two independently clocked commodity nodes and a USRP recorder. (b) Two preamble and 30 hopping slots; 512 µs emissions follow a 1 ms grid. (c) Concurrent nodes use different base channels. Boxes follow known hop schedules and attribute received energy.

separate hosts sharing no clock, trigger or frequency reference, and a USRP B210 observes without generating or synchronizing anything (Figure 4a). The scored observable is emitted energy in the slots a node's known schedule predicts, never a decoded payload.

Scoring slides each node's 32-slot template over the energy map and compares its cells against the all-channel mean; a slot below its channel noise floor counts as absent. All 32 slots clear the noise floor in every one of 23 bursts of the single-node capture (Figure 4b), at a received-energy period of 999.92 µs (s.d. 0.06 µs) against a nominal 1000 µs.

Both nodes' known patterns remain attributable during concurrent activity. With a 5 ms inter-burst gap each node is inside a burst 86% of the time, at about 44% RF duty cycle because each emission fills 512 µs of its 1 ms slot, so each node is detected while the other is transmitting (Figure 4c). Over 3 s each yields its expected 82 bursts at median pattern contrast 4.90 and 4.80 against a 1.5 threshold, and no burst is lost to the other's presence. Period estimates drift to 999.66 and 999.36 µs (s.d. 0.41 and 1.00 µs), further from nominal than the single-node case.

Attribution works because both schedules are known in advance and the nodes sit on different base channels, so each template predicts a distinct set of cells. Schedule feasibility does not establish spectral compliance. Averaged over 400 hops, bandwidth is 39 kHz at −3 dB, but unshaped MSK sidelobes extend the −30 dB width to 164 kHz, beyond one 125 kHz allocation. Gaussian shaping is available in the radio but is not used here.

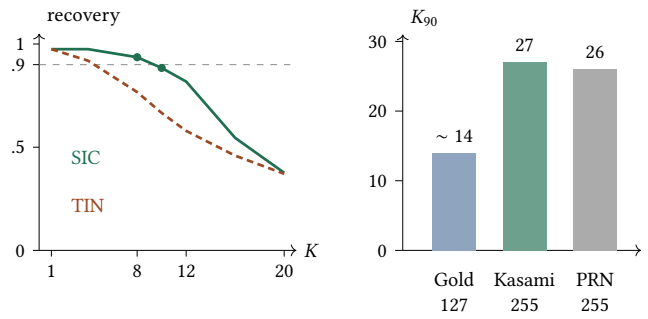
The hardware and simulation operating points cannot be equated. A 128-chip emission at 62.5 kchip/s needs 2.05 ms, so the 1 ms slot forces $N = 31$ and about 6.1 dB less nominal processing gain than the simulated $N = 127$. The bench therefore shows that commodity silicon can emit the structure, not that it can emit the simulated waveform.

4.2 Can the structure support concurrent recovery?

These are simulations of the gateway, not bench captures. All three transmit at equal received power, −6 dB per user, carry no packets or FEC, and are given the active set; they differ in what else the receiver is handed. Separation also supplies exact timing, so the receiver only has to read each block's shift. Acquisition supplies neither timing nor carrier offset and must find both. The signature diagnostic carries no data at all and scores presence only.

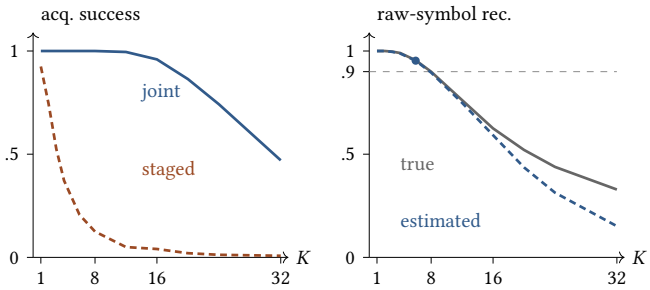
With full overlap and Gaussian per-user residual CFO ($\sigma = 196.85$ Hz), ordered SIC reaches 93.6% median raw-symbol recovery at $K = 8$ ([92.8, 95.9]), falling below 90% by $K = 10$ (Figure 5a). The matched correlation-only control treats interference as noise (TIN), omitting cancellation; it clears the criterion only through $K = 4$ and reaches 76.7% at $K = 8$, a 16.9 percentage-point gap, because subtracting reliably decoded blocks lowers remaining interference. The benefit requires correct decisions. At $K = 20$, the arms reach 37.6% and 37.1%; the experiment does not distinguish failed decoding from wrong subtraction. This internal comparison excludes LoRa and LR-FHSS.

Acquisition then removes the timing and carrier genies: the continuous 12-block Gold construction searches arrival time over ±334 chips and CFO over ±10 kHz. Figure 6a runs



(a) Raw-symbol recovery (b) Signature only

Figure 5: K fully overlapping signals at −6 dB/user. (a) SIC versus interference-as-noise (TIN), given codes/timing. (b) Signature-only K_{90}: load achieving 90% recovery. No packets/FEC.



(a) Acquisition success (b) Raw-symbol recovery

Figure 6: Acquisition with K known-active signals, up to eight scored. (a) Joint versus staged search: timing error <0.5 chip, frequency <197 Hz. (b) Recovery from estimated versus true parameters.

that search two ways on the same mixture: the staged reference fixes timing non-coherently before estimating carrier offset, while the joint arm scores lag and offset on one coherent surface, searching more hypotheses rather than fewer. It is a coherent-search reference, not a proposed receiver. The search interval is a particular prior, not coverage of every satellite Doppler and oscillator offset.

Finding the parameters is not the same as reading the data (Figure 6). At $K = 6$, estimated and true parameters yield 95.3% and 95.5% recovery, so estimation costs almost nothing there. At $K = 8$, median acquisition success remains 100%, yet raw-symbol recovery falls to 89.5%, and even the true-parameter control reaches only 89.96%: what limits recovery at that load is interference, not parameter error. This receiver uses no SIC, so combining its estimates with reliable cancellation remains open.

Signature diagnostics show longer-code headroom without a measured Kasami-family advantage. Under unmodulated TIN (Figure 5b), Gold-127 reaches roughly 14 signatures; Kasami-255 has median per-seed $K_{90} = 27$ [25, 28], versus 26 [25, 28] for same-length PRN. Data-bearing recovery at $N = 255$ is unmeasured. Separately, partial overlap scores one victim rather than all users, establishing no delay-derived supported-user or capacity gain.

Together these establish that the signature structure supports concurrent raw-symbol recovery at the loads tested, and that a receiver can find its own timing and carrier when it knows who is active. Discovering the active set, and composing acquisition with cancellation, remain unjoined.

5 Conclusions and Future Work

CodeHop asks how far satellite access can evolve when the node's radio is frozen. The answer so far is partial but concrete: firmware alone drives two unsynchronized SX1276-class nodes to emit the structure; a gateway holding the samples recovers over 90% of raw symbols from eight fully overlapping transmissions given codes and timing; and a separate receiver finds timing and carrier frequency offsets for six known-active signals. Three questions decide whether these components become a system.

Can acquisition work without a known active set? Near-far powers make that harder still. The coherent joint search held up with load where a staged one collapsed, but with the active set supplied; whether that margin survives without it is the open question, and if it does not the control split loses its point.

Does raw-symbol recovery translate into delivered packets? Acquisition must compose with cancellation, framing, cross-hop FEC and retries, and node energy must count failed attempts. Two costs are visible: a third of the tested blocks are pilots, and cancellation needs a confidence rule it lacks, its advantage gone by $K = 20$. If those exceed what concurrency returns, the gain never reaches users.

Can CodeHop extend to other radios? Extension needs a capability model: which controls, at what timing and bandwidth, yield acceptable emissions. We validate one radio family and establish no count of compatible deployed devices. TI's CC1101 documents serial data modes and frequency hopping [20], yet its asynchronous serial mode excludes MSK. Without node source or build access another control path is needed, and if emissions are immutable, receiver-only interpretation is its own opportunity, as ZigZag shows [9].

Acknowledgments

This research was partially supported by the National Science Foundation under grant 2232481.

References

[1] 3GPP. 2021. *Study on Narrow-Band Internet of Things (NB-IoT) / Enhanced Machine Type Communication (eMTC) Support for Non-Terrestrial Networks (NTN).* Technical Report 36.763. 3rd Generation Partnership Project. Version 17.0.0 (Release 17).

[2] Guido Alvarez, Juan A. Fraire, Khaled Abdelfadeel Hassan, Sandra Cespedes, and Dirk Pesch. 2022. Uplink Transmission Policies for LoRa-Based Direct-to-Satellite IoT. *IEEE Access* 10 (2022), 72687–72701. doi:10.1109/ACCESS.2022.3189647

[3] Guillem Boquet, Pere Tuset-Peiró, Ferran Adelantado, Thomas Watteyne, and Xavier Vilajosana. 2021. LR-FHSS: Overview and Performance Analysis. *IEEE Communications Magazine* 59, 3 (2021), 30–36. doi:10.1109/MCOM.001.2000627

[4] Jumana Bukhari and Zhenghao Zhang. 2024. Understanding Long Range-Frequency Hopping Spread Spectrum (LR-FHSS) with Real-World Packet Traces. *ACM Transactions on Sensor Networks* 20, 6, Article 117 (2024), 30 pages. doi:10.1145/3694971

[5] Oscar del Río Herrero and Riccardo De Gaudenzi. 2012. High Efficiency Satellite Multiple Access Scheme for Machine-to-Machine Communications. *IEEE Trans. Aerospace Electron. Systems* 48, 4 (2012), 2961–2989. doi:10.1109/TAES.2012.6324671

[6] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of the ACM SIGCOMM 2017 Conference*. Association for Computing Machinery, New York, NY, USA, 13 pages. doi:10.1145/3098822.3098845

[7] European Space Agency. n.d.. Internet of Things Technology Demonstration. ARTES project description. https://resilience.esa.int/archives/projects/internet-things-technology-demonstration

[8] FOSSA Systems. 2026. Frequently Asked Questions. https://fossa.systems/frequently-asked-questions/. Commercial LoRa-based satellite IoT constellation; accessed 2026-10-22.

[9] Shyamnath Gollakota and Dina Katabi. 2008. ZigZag Decoding: Combating Hidden Terminals in Wireless Networks. In *Proceedings of the ACM SIGCOMM 2008 Conference*. Association for Computing Machinery, New York, NY, USA, 12 pages. https://people.csail.mit.edu/gshyam/Papers/ZigZag.pdf

[10] Zhanjun Huang, Jinkai Li, and An Zhang. 2024. A Tactical Data Link Model Based on CCSK-FMSK. *Physical Communication* 64 (2024), 102345. doi:10.1016/j.phycom.2024.102345

[11] Sooyeob Jung, Seongah Jeong, Jinkyu Kang, Gyeongrae Im, Sangjae Lee, Mi-Kyung Oh, Joon Gyu Ryu, and Joonhyuk Kang. 2025. LR-FHSS Transceiver for Direct-to-Satellite IoT Communications: Design, Implementation, and Verification. *IEEE Open Journal of Vehicular Technology* 6 (2025), 1920–1942. doi:10.1109/OJVT.2025.3585160

[12] Chenning Li, Yidong Ren, Shuai Tong, Shakhrul Iman Siam, Mi Zhang, Yunhao Liu, Jiliang Wang, and Zhichao Cao. 2024. ChirpTransformer: Versatile LoRa Encoding for Low-power Wide-area IoT. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. Association for Computing Machinery, New York, NY, USA, 13 pages. doi:10.1145/3643832.3661861

[13] Olof Liberg, Stefan Eriksson Löwenmark, Sebastian Euler, Björn Hofström, Talha Khan, Xingqin Lin, and Jonas Sedin. 2020. Narrowband Internet of Things for Non-Terrestrial Networks. *IEEE Communications Standards Magazine* 4, 4 (2020), 49–55. doi:10.1109/MCOMSTD.001.2000004

[14] Foivos Michelinakis, Anas Saeed Al-Selwi, Martina Capuzzo, Andrea Zanella, Kashif Mahmood, and Ahmed Elmokashfi. 2021. Dissecting Energy Consumption of NB-IoT Devices Empirically. *IEEE Internet of Things Journal* 8, 2 (2021), 1224–1242.

[15] Manan Mishra, Daniel Koch, Muhammad Osama Shahid, Bhuvana Krishnaswamy, Krishna Chintalapudi, and Suman Banerjee. 2023. OpenLoRa: Validating LoRa Implementations through an Extensible and Open-sourced Framework. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1165–1183.

[16] Yury Polyanskiy. 2017. A Perspective on Massive Random-Access. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*. 2523–2527. doi:10.1109/ISIT.2017.8006984

[17] Semtech Corporation. 2015. *SX1276/77/78/79 – 137 MHz to 1020 MHz Low Power Long Range Transceiver.* Technical Report Rev. 4. Semtech Corporation.

[18] Semtech Corporation. 2022. Semtech Announces New Tool Suite Enabling Dense Deployments and Satellite Connectivity for LoRa. https://www.semtech.com/company/press/semtech-announces-new-tool-suite-enabling-dense-deployments-and-satellite-connectivity-for-lora

[19] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent Interference Cancellation: Decoding Multi-Packet Collisions in LoRa. In *Proceedings of the ACM SIGCOMM 2021 Conference*. Association for Computing Machinery, New York, NY, USA, 503–515. doi:10.1145/3452296.3472931

[20] Texas Instruments. 2013. *CC1101: Low-Power Sub-1 GHz RF Transceiver.* Texas Instruments. https://www.ti.com/lit/ds/symlink/cc1101.pdf Sections 27–28; accessed October 21, 2026.

[21] TinyGS. 2026. TinyGS – Satellites. https://app.tinygs.com/satellites. Community open-source ground-station network; accessed 2026-07-17.

[22] Stephen Ugwuanyi, Greig Paul, and James Irvine. 2021. Survey of IoT for developing countries: Performance analysis of LoRaWAN and cellular NB-IoT networks. *Electronics* 10, 18 (2021), 2224.

[23] Rohith Reddy Vennam, Maiyun Zhang, Raghav Subbaraman, Deepak Vasisht, and Dinesh Bharadia. 2025. FSMA: Gateway-Controlled Access for Scalable LoRa Non-Terrestrial Networks. *arXiv preprint arXiv:2510.02800* (2025).

[24] Deliang Yang, Xianghui Zhang, Xuan Huang, Liqian Shen, Jun Huang, Xiangmao Chang, and Guoliang Xing. 2020. Understanding Power Consumption of NB-IoT in the Wild: Tool and Large-scale Measurement. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20)*. 732–744.

[25] Xiangyong Zeng, John Q. Liu, and Lei Hu. 2007. Generalized Kasami Sequences: The Large Set. *IEEE Transactions on Information Theory* 53, 7 (July 2007), 2587–2598. doi:10.1109/TIT.2007.899528