algorithms

MDPI

Article

Asymmetric Interference Perception Method in Short-Range Wireless Systems

Volodymyr Sokolov * and Pavlo Skladannyi

Department of Information and Cyber Security, Borys Grinchenko Kyiv Metropolitan University, 04212 Kyiv, Ukraine; p.skladannyi@kubg.edu.ua
* Correspondence: v.sokolov@kubg.edu.ua

Abstract

In the 2.4 GHz band, the channel of a short-range link is usually selected at one end; in Bluetooth Low Energy, a peripheral's channel classification is advisory. This paper examines, for one instrumented case, whether the two ends of a link perceive interference similarly. A purpose-built protocol on two nRF52840 devices reported each end's noise floor once per adaptation window. On a 4.2 m indoor link with one swept narrowband interferer, 656 time-paired windows were obtained. In the 614 quiet windows, the transmitter read a standing +6.30 dB more than the receiver, whereas a co-location control gave a mean offset of 0.66 dB. In the 42 windows within fifteen interference episodes, the difference had a standard deviation of 19.2 dB, and the two ends classified the channel differently in 38. Co-located devices also differed by a median of 9 dB under a directed source, so these differences cannot be attributed wholly to position. With two simultaneously sampled receive chains 60 cm apart, the levels received from a moving source were uncorrelated. A negotiation in which the receiver selects among the transmitter's candidate channels completed 17 of 17 changes on a quiet bench and 43 of 44 under interference; in all 43, however, it selected the channel that transmitter-only selection would have used. Endpoint perception therefore differed substantially in this configuration, whereas an advantage of the method over simpler alternatives was not demonstrated.

Keywords: asymmetric interference; interference perception; channel selection; Bluetooth Low Energy; adaptive frequency hopping; spectrum sensing; software-defined radio; cross-technology interference; wireless testbed

Check for updates

Academic Editor: Carolina Del-Valle-Soto

Received: 10 September 2026
Revised: 6 October 2026
Accepted: 9 October 2026
Published: 14 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The 2.4 GHz Industrial, Scientific and Medical (ISM) band is shared by technologies whose transmit powers differ by one to two orders of magnitude. A Wi-Fi access point radiating 20 dBm and a Bluetooth Low Energy (BLE) sensor radiating 0 dBm occupy the same spectrum under rules that were never designed to arbitrate between them. The resulting coexistence failures are not symmetric: the high-power device frequently does not detect the low-power one at all, transmits over it, and never learns that it did so [1–3].

Most mechanisms that respond to this situation form their view of channel quality at one end of the link, with some exceptions. IEEE 802.11h gives the detecting station binding authority for one specific hazard; time-slotted channel hopping (TSCH) lets each end keep its own per-neighbour blacklist; and cooperative spectrum sensing has long addressed disagreement between distributed sensors (Section 2). A short-range link in which the far end's measurement determines the choice is, however, unusual.

BLE is the extreme case of the vulnerable side. It performs no clear channel assessment before transmitting, so it cannot defer to an occupied channel [3,4]. Its adaptive frequency hopping (AFH) mechanism does maintain a channel map, but three properties of that mechanism limit how well it can respond to noise. First, the classification that populates the map is inferred from packet errors accumulated over many connection events rather than measured directly, so it is both slow and confounded: a channel that fails because of fading and a channel that fails because of a Wi-Fi beacon look identical. Second, the specification prescribes no classification algorithm at all, so behaviour varies between implementations [5,6].

Third, and centrally for this paper, the decision is one-sided. Until Core 5.3, the map was decided by the central alone. Core 5.3 added the Channel Classification Enhancement, in which a peripheral may classify channels itself and report them to the central through the `LL_CHANNEL_REPORTING_IND` and `LL_CHANNEL_STATUS_IND` procedures, so that they "can be considered when updating the central device's channel map" [7]. The report is therefore advisory. The central retains the decision, and no mechanism makes the peripheral's assessment binding. Because no classification algorithm is specified for either end, nothing constrains how far the two classifications may diverge.

These mechanisms rest on an assumption that is rarely stated: that channel quality is a property of the link and can therefore be established at either end. Interference, however, is local. A narrowband emitter three metres from one node and eight metres from the other raises the two noise floors by different amounts, and the difference need not average out over the time in which a channel decision is taken. If the two ends disagree about which channels are usable, a channel chosen by one end may be unusable at the other.

1.1. Problem Statement

The problem addressed here is that interference perception in a short-range wireless link may be asymmetric, and that prevailing channel-selection mechanisms represent this possibility weakly or not at all. Mechanisms in which the non-deciding end can express an opinion do exist: BLE Core 5.3 channel reporting [7], IEEE 802.11h dynamic frequency selection (DFS) [8], and per-neighbour blacklisting in IEEE 802.15.4e TSCH [9]. In the first, the opinion is advisory; in the second, it is binding only for one specific hazard; in the third, it is applied to a hopping sequence rather than to a negotiated single channel. What is missing is a measurement of how much the two ends actually disagree, together with a mechanism in which the receiver's assessment determines the choice.

1.2. Objective and Contributions

The objective of this work is to define, implement and experimentally characterise a method of *asymmetric interference perception*: a scheme in which both ends of a link measure the spectrum independently, the transmitter proposes, the receiver chooses among the proposals against its own measurements, and the resulting agreement, or its absence, governs whether the session continues.

The contributions are as follows:

1. A time-paired measurement of interference–perception asymmetry between the two ends of a 4.2 m link under a controlled narrowband interferer, comprising 656 adaptation windows on 26 channels, of which 42 fall within fifteen interference episodes, with autocorrelation-aware intervals and with the window-level description separated from the episode-level inference (Section 4.2). The result applies to one geometry, one pair of devices and one interferer type; it demonstrates that the asymmetry is large enough to matter in that configuration and does not characterise short-range links in general.

2. A co-location control that separates receiver calibration from the other causes of the difference and identifies its own limit (Section 4.3). In a quiet band, the same two devices 10 cm apart differ by a mean of 0.66 dB, against 6.30 dB across the link; under a directed source, they differ by a median of 9 dB. The standing difference is therefore attributed to the two ends' positions or operating roles rather than to receiver calibration, whereas the differences observed during interference probably overstate spatial decorrelation.
3. A two-antenna measurement with a software-defined radio (SDR), whose two receive chains share one clock and are sampled simultaneously, in which a transmitter was moved along a 20° arc at 3.5 m from a 60 cm baseline (Section 4.4). It shows how strongly a separation of a few wavelengths decouples two identical receivers under a directed source, and it examines whether the wider baseline improves direction finding.
4. A channel-selection method in which the receiver's measurement, rather than the transmitter's, determines which channel is used: a ranked, cryptographically sealed candidate list is offered by the transmitter on dedicated control channels, and the receiver selects from it against its own noise floor and returns that measurement, with a bounded termination rule when no exchange is possible (Sections 3.8 and 4.5). Its negotiations under interference are compared with transmitter-only selection on the recorded data, and the decision rules of transmitter-only, receiver-only and intersection-based selection are compared on the paired windows (Section 4.6).
5. A purpose-built experimental protocol and testbed, not constrained by any standard, in which every adaptation loop is observable and controllable, and in which an SDR independently demodulates and decodes every frame the nodes exchange (Section 3).
6. A per-stage processing-cost characterisation showing that, across the six configurations of this ladder, the authenticated-encryption cost is constant in payload length, while the forward-error-correction cost is predicted by correction work rather than by payload length (Section 4.7).

Section 2 reviews the literature on asymmetric interference and its mitigation. Section 3 describes the testbed, the protocol and the method. Section 4 presents the measurements. Section 5 discusses the interpretation of the results, their limitations and their generalisation. Section 6 concludes.

2. Related Work

2.1. The Mechanism of Asymmetric Interference

Asymmetric interference arises when heterogeneous technologies with very different transmit powers coexist so that high-power devices corrupt low-power traffic without detecting it [1–3]. BLE and Zigbee are the canonical victims, sharing spectrum with Wi-Fi transmitters operating at 10–100× higher power [2,3].

The mechanism has a spatial structure. In asymmetric regions the Zigbee signal is too weak to trigger Wi-Fi's clear channel assessment, so Wi-Fi treats the channel as idle and transmits, corrupting any bit of a Zigbee packet with equal probability [2,10]. In symmetric regions, Zigbee activity can trigger Wi-Fi backoff, and the residual corruption concentrates in the packet header [2,10]. BLE is more exposed than Zigbee because it does not employ clear channel assessment at all [3].

2.2. Detection and Perception

Several schemes detect interference from the low-power device's own perspective. An interference awareness scheme for BLE monitors link-layer connection events, achieving supervision timeout and packet loss rates below 3% in chamber tests and below 1% in office environments [5]. A passive TinyML system on ESP32 microcontrollers extracts received-

signal-strength distributions and activity indicators from BLE scan windows, detecting interference with sub-millisecond latency without modifying the BLE protocol [11].

From the high-power side, an enhanced clear channel assessment mechanism uses signal identification so that Wi-Fi treats the channel as busy when a Zigbee signal is identified even below the energy threshold [1]. Commodity Wi-Fi cards can classify Zigbee, microwave and LTE-U interference from receiver error traces, where burst length and inter-burst timing encode the interfering technology's signature, with over 95% classification accuracy [12].

Common to all of these is that perception is established at one node and then treated as a property of the channel. Core 5.3 reporting [7] is the exception in intent, since it combines classifications from both ends, but it neither compares them quantitatively nor makes the disagreement binding. None of this work pairs simultaneous measurements at the two ends of a link and asks by how much they differ.

2.3. Mitigation Strategies

BuzzBuzz adds header and payload redundancy through multi-headers and Reed–Solomon coding, improving Zigbee delivery by 70% and reducing retransmissions threefold [2]. B-hop adjusts BLE frequency hopping in the time domain to avoid predicted Wi-Fi beacons, the first standard-compliant time-domain approach to BLE–Wi-Fi cross-technology interference [3]. Smoggy-Link maintains a link model translating interference fingerprints into spatiotemporal link states for adaptive Zigbee link selection [13]. CoSense and ZiSense detect Zigbee signals from correlation and time-domain features, reducing false wake-ups and saving up to 63% of energy under dynamic interference [10,14].

2.4. Spatial and Directional Asymmetry

Asymmetry also appears spatially. Targeting the source rather than the destination of a link produces different throughput degradation, because the source receives control information while the destination receives the payload [15]. Pulse-modulated noise degrades a link more than stationary additive white Gaussian noise of the same average power at both endpoints, which challenges test methods that use the latter as a benchmark [15]. For BLE phase-based ranging, Wi-Fi interference increases random distance error more than threefold at short range while becoming marginal beyond 10 m [16].

Spatial asymmetry has a well-understood physical basis. Indoors, the received field is a superposition of many reflected components, and the small-scale fading it produces decorrelates over distances of the order of half a wavelength [17], which is 6 cm at 2.4 GHz. Two receivers a few wavelengths apart may therefore see uncorrelated levels from the same source. The same property limits two-antenna (interferometric) direction finding: the phase difference between two antennas a distance d apart determines the bearing only within an ambiguity interval set by d/λ, and multipath biases it [18]. The plane-wave model on which that relation rests is conventionally taken to hold beyond the Fraunhofer distance $2d^2/\lambda$ [19]. Section 4.4 uses these relations to interpret a two-antenna measurement.

2.5. Standardised Mechanisms with Far-End Input

Four standardised or widely deployed mechanisms bear directly on the claim of this paper and must be distinguished from it.

BLE Core 5.3 channel classification reporting. As described in Section 1.1, a peripheral may report its own classification to the central [7]. This is the closest prior art to the method proposed here, and it differs in exactly one respect: the report is advisory. The central may act on it or ignore it, the specification requires nothing, and there is no protocol state in which the peripheral's assessment prevents a channel from being used. The

present work lets that assessment select among the transmitter's candidates and records, for 43 negotiations, whether it changed the outcome (it did not; Section 4.6).

IEEE 802.11h dynamic frequency selection. In DFS, a station that detects radar must report it, and the access point is then obliged to vacate the channel [8]. This is a genuine binding veto held by the detecting side, and it is deployed at scale. It is narrower than the mechanism studied here in two ways: it is defined for one specific hazard rather than for channel quality in general, and it is a veto on the channel in use rather than a selection among proposed alternatives.

IEEE 802.15.4e TSCH blacklisting. TSCH networks maintain per-link channel blacklists and adapt the hopping sequence to them [9,20], which is architecturally the closest analogue to per-end perception. The difference is structural: TSCH spreads each link over a hopping sequence, so a channel rejected by one end is diluted rather than avoided, whereas the protocol studied here occupies one channel at a time and therefore has to agree on which.

LoRaWAN adaptive data rate. Adaptive data rate (ADR) is a further example of a one-sided decision: the network server sets data rate and transmit power from uplink measurements the end device cannot contest [21].

Cooperative spectrum sensing in cognitive radio is the body of work that comes closest to asking the question posed here, since it fuses observations from several sensing nodes and must therefore confront the possibility that they disagree [22,23]. That literature is concerned with fusing many observations of one primary user into a better estimate. Here, the concern is different: with only two observers, the disagreement itself is the quantity of interest, and averaging removes it.

2.6. Channel Selection and Frequency Hopping

Recent work covers the channel-selection problem from several directions. Eltholth compares random, chaotic, adaptive and optimised chaotic frequency hopping for Wi-Fi/BLE coexistence in the 2.4 GHz band, reporting that the optimised chaotic scheme balances interference reduction against connection success [24]. Lee et al. revisit AFH map prediction with machine-learning classifiers, treating the map as something to be forecast rather than measured [6]. Yang et al. propose diffusion-based frequency hopping for collision mitigation in dense Bluetooth deployments [25]. Kim enables WLAN and WPAN coexistence through cross-technology communication, so that the two standards can address each other directly [26]. On the detection side, Del-Valle-Soto et al. develop low-cost jamming detection from performance metrics in cluster-based wireless sensor networks [27], and Cortés-Leal et al. extend that analysis to collaborative and cooperative schemes in industrial networks [28].

All of these works treat the channel decision as single-valued. The map is predicted, optimised or diffused, but it is one map, and the question of whether the two ends of a link would populate it identically is not raised. That question is the subject of Section 4.2.

2.7. Antecedents of the Present Testbed

This study continues a line of experimental work on the resilience of short-range wireless systems to interference, and three specific results from it shape the design used here.

The jamming susceptibility of BLE beacons [29], ZigBee networks [30] and Wi-Fi [31] was characterised by moving a narrowband source through the band and recording where each technology failed; the observation that the failure threshold depended on where the source stood, and not only on its power, motivated the paired measurement reported here. The jump-stay strategy of [32] showed that an interferer sweeping a band defeats a victim whose evasion is periodic, which is why the candidate ordering in Section 3.8 is keyed rather than cyclic. The repeater study of [33] measured how strongly received power

varies with position in an ordinary room, giving the prior expectation that two endpoints four metres apart would not see the same noise floor.

These results motivated the hypothesis, tested here, that the two ends of one link may perceive different noise floors. None of them measured both ends of a link at the same instant.

2.8. Gap Analysis

Table 1 places the present work against the closest prior approaches. The distinguishing property is not that the receiver measures the spectrum, since several schemes have it do so, but that its measurement is binding on the channel decision. The final column gives the cost of acting on a change: blind hopping involves no decision and no air-time cost, a unilateral AFH map update is a single control procedure applied at a future connection event, and the negotiation proposed here is a bounded exchange that completed in 0.95 ± 0.004 s (current firmware, $n = 12$) on a quiet bench. The comparison is qualitative, because the entries are structural costs from specifications set beside one measurement from one testbed; it indicates an order of magnitude, not a ranking.

Table 1. Where channel-quality perception is formed and who decides.

Approach	Perception at	Decision by	Is the Far End's Assessment Binding?	Cost of Acting on a Change
BLE AFH, Core 4.0–5.2	central, from packet errors	central	no mechanism exists	one control procedure, applied at a future connection event
BLE AFH, Core 5.3+ [7]	both; peripheral may report	central	no; advisory only	as above, plus a reporting procedure the central may ignore
IEEE 802.11h DFS [8]	detecting station	AP, on report	yes, but only for radar	channel switch announcement; vacate within 10 s
802.15.4e TSCH [9]	per link, per neighbour	local blacklist	partly; dilutes rather than avoids	none on the air; the sequence absorbs it
LoRaWAN ADR [21]	network server, uplink	server	no	downlink MAC command, one round trip at duty-cycle rates
B-hop [3]	central	central	no; time-domain avoidance	none; avoidance is scheduled, not negotiated
IAS for BLE [5]	both, from link events	central	no	as BLE AFH
AFH map prediction [6]	central	central	no	as BLE AFH
Chaotic FH [24]	none; blind sequence	sequence	no	none; nothing is decided
E-CCA/ZShark [1]	interferer	interferer	n/a [1]	n/a
Smoggy-Link [13]	sender	sender	no	sender-side, no agreement
This work	both, by energy detection	jointly, receiver selects	**yes, over the offered set**	0.95 s measured, quiet bench [2]

[1] E-CCA protects the low-power victim from the interferer's side; the victim's own receiver plays no part, so the column does not apply. Rows are classified from the mechanism each cited work describes, not from any claim those works make about one another. [2] The final column is qualitative, and the entries are not directly comparable: only the last is a measurement from this testbed, while the others are the mechanism's structural cost as specified. It indicates the order of magnitude of each cost; it does not rank the approaches.

3. Materials and Methods

3.1. Rationale for a Purpose-Built Protocol

The method under study cannot be expressed inside BLE. A standard-compliant implementation cannot make the peripheral's assessment binding on the central, cannot add dedicated control channels, and cannot expose the per-window internal state (noise floor, symbol error rate, discriminator output) that the measurement requires. Working inside the standard would have restricted the experiment to what the standard already permits, which is precisely the limitation under examination.

A complete link protocol was therefore implemented from the radio registers upwards, with no Bluetooth stack, so that every mechanism is both observable and controllable. The cost of this choice is that the resulting system is not interoperable with anything, and its absolute throughput is well below BLE's on the same silicon (Section 4.8). The benefit is that every adaptation loop can be instrumented, frozen, forced or replaced, and that an independent receiver can reconstruct every frame from the air.

3.2. Hardware Testbed

The testbed (Figure 1) consists of two nodes: an independent observer and a controlled interferer:

- **Node A**, an Adafruit Feather nRF52840 Express, acting as transmitter.
- **Node B**, a nice!nano v2 (nRF52840), acting as receiver.
- **Observer**: a Nuand bladeRF 2.0 micro xA4 SDR, receive-only, demodulating and decoding every frame independently of both nodes.
- **Interferer**: a narrowband noise generator, manually retuned across the band and stepped in 3 dB increments.

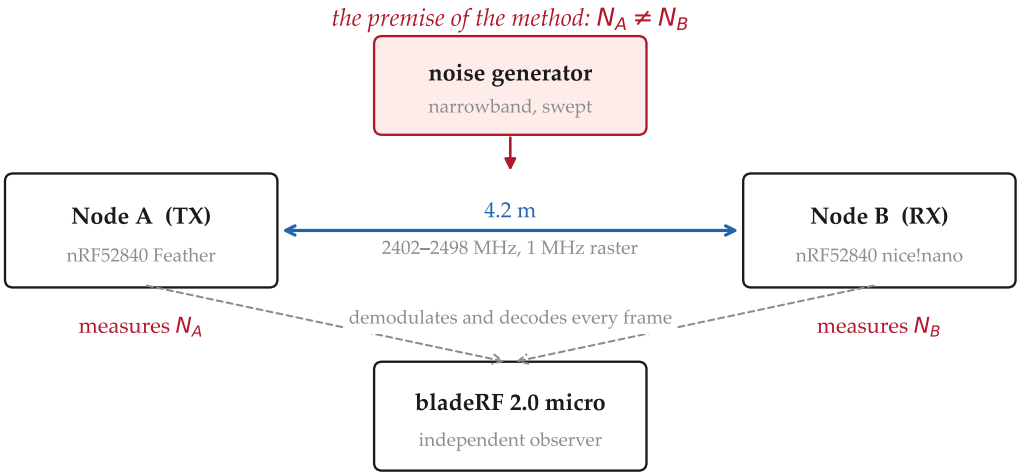


Figure 1. The testbed. Both nodes run the same energy detector and report a noise floor for the channel in use once per adaptation window; the SDR reconstructs every frame from the air, so no claim about what was transmitted rests on the two devices under test agreeing with each other.

The two nodes were separated by $d = 4.2$ m in an ordinary indoor office environment, at a transmit power of 0 dBm, with printed antennas and no ground plane. The free-space path loss at this separation,

$$L_{\text{fs}} = 20 \log_{10} d_{\text{km}} + 20 \log_{10} f_{\text{MHz}} + 32.44, \tag{1}$$

evaluates to 52.6 dB at 2420 MHz, so a 0 dBm transmitter should be received at about −52.6 dBm. The measured received level was −58 to −62 dBm, 6–9 dB below free space,

which is consistent with the antennas and the environment and provides a sanity check on the link budget.

3.3. Software

The node firmware is written for Zephyr RTOS under the nRF Connect SDK and drives the `RADIO` peripheral directly. Cryptographic operations use the PSA API on the CryptoCell 310 accelerator. The host software is Python 3.14.5: a serial telemetry collector, a protocol mirror that reproduces the on-air format byte for byte, a Reed–Solomon implementation, and an SDR receive chain comprising a fused mix-and-decimate channeliser, an energy gate with a quantile noise floor, GFSK and Coded-PHY demodulators, and the same frame decoder the nodes use. Because the SDR chain shares no code with the firmware and is driven only by the air interface, agreement between the two is evidence about the air rather than about the implementation.

3.4. Channel Plan

The band is divided as shown in Figure 2. Channel numbers are $c = f_{\mathrm{MHz}} - 2400$ in the radio's default frequency map. The channel plan is outlined as follows:

- **Data**: 2402–2498 MHz on a 1 MHz raster, 97 channels.
- **Control**: 2401, 2450 and 2499 MHz, carrying channel-switch negotiation only.
- **Open**: 2499 MHz, with 2401 MHz as a fallback, used for pairing before a key exists.

Ninety-seven data channels rather than the forty of BLE is a deliberate choice: a single narrowband interferer was repeatedly observed to raise two of four widely spaced channels above threshold simultaneously, which leaves a four-channel plan with no usable alternative. The three control channels are placed at the edges and the centre of the band, 98 MHz apart, so that an interferer must be wideband to prevent the two ends from agreeing where to meet.

This plan is an experimental abstraction and is not proposed as a channel plan for BLE. A fine raster maximises the number of distinguishable measurement points the asymmetry experiment can draw on, and a channel that is 1 MHz wide in a plan and 1.34 MHz wide on the air is acceptable for that purpose in a way it would not be in a standard. Nothing in the method depends on the number: the negotiation carries channel indices, and a 40-channel or 16-channel plan would work unchanged. Where the comparison with BLE is drawn in this paper, it is a comparison of *mechanism* (who perceives, who decides) and not of channelisation.

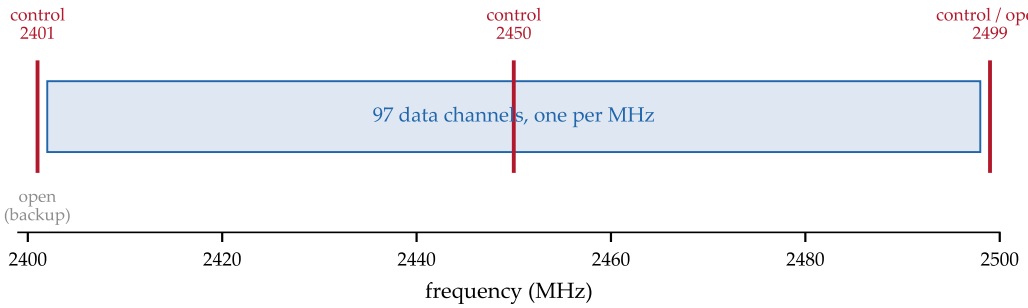


Figure 2. Channel plan. Control channels never carry data, so the negotiation that follows a failure does not depend on the channel that failed.

3.5. Frame Types

Four frame types share one 14-byte header, transmitted in the clear and covered as associated data by the authenticated encryption (Table 2). The header carries the segment number, a transmit counter that also supplies the nonce for the authenticated encryption

with associated data (AEAD), the mode index, the channel in use and an announced next channel, and a cumulative acknowledgement base with a 32-bit selective-repeat bitmap.

Table 2. Frame types. τ is the authentication tag length in bytes.

Type	Body (B)	Protection	Channel	Purpose
DATA	L_{seg}	AES-128-CCM*, τ	data	payload segment, carries feedback
ACK	8	AES-128-CCM*, τ	data	cumulative base plus selective bitmap
CTRL	16	AES-128-CCM*, τ	control	channel-switch request, acknowledgement or refusal
RDV	8–72	none (public exchange)	open	pairing: public keys, commitment, nonces, check values

Rendezvous frames carry no authenticated encryption because the exchange is what produces the key; they are whitened with an initialisation vector derived from the channel number alone for the same reason. Everything else is sealed under the link key, so a channel switch cannot be steered by a party that has not paired.

The nonce. Because CCM* fails catastrophically if a nonce repeats under one key, the construction is given here in full rather than left to the implementation. The 13-byte nonce is

$$\text{nonce} = \underbrace{d}_{1} \parallel \underbrace{\text{id}}_{4} \parallel \underbrace{e}_{2} \parallel \underbrace{v}_{4} \parallel \underbrace{q}_{2}, \tag{2}$$

with $d \in \{0,1\}$ the direction, id a 32-bit device identifier, e the 16-bit epoch, v the 32-bit transmit counter and q the 16-bit segment number; the header transmits the low 16 bits of v, and each end reconstructs the rest from its own count.

Uniqueness within a session follows from three properties, each of which is a requirement on the implementation rather than an observation about it. The two directions never collide because d differs. The two nodes never collide because id differs. And within one direction, v is strictly monotonic: it is incremented on every transmission, never reset while a session is running, and advanced by a guard of 4096 on entering the resynchronisation mode and 65,536 on leaving it, so that frames in flight across the transition cannot reuse a value. The epoch e is redundant with respect to uniqueness, since it is also an input to the HKDF-SHA256 derivation of the traffic key, and therefore, every epoch is a different key; it is included so that a frame carries its own key identifier. At 32 bits and one frame per 10 ms slot, v would take about 500 days of continuous transmission to wrap, and every mode or channel change rekeys long before that.

Two cases fall outside this argument. A node restarted without pairing resumes at epoch 1 with $v = 0$ under the same built-in key and will then reissue nonces it has already used; the built-in key is a test key, present so that the SDR can decode a link that has not paired, and pairing, which derives a fresh K_{master} per session, makes the construction sound. In addition, id is at present a constant per role rather than a device identity, which is adequate for two nodes but must be replaced first if more than two nodes are used (Section 5.5).

3.6. Perception, Scoring and Discrimination

Each node measures the energy on channel j with the radio's received-signal-strength indicator (RSSI) while not transmitting, producing $N_j[k]$ in dBm. The quantity is called the noise floor throughout this paper, and the term is operational rather than physical: N_j is the RSSI-derived energy estimate produced by this measurement procedure (one sample per visit, on a receiver configured for the mode in use, with no calibration against a reference source) and not a calibrated measurement of noise power. The nRF52840's RSSI has 1 dB resolution and a specified accuracy of ± 2 dB over an input range of -90 to -20 dBm [34].

The quiet floors in this study (about -100 dBm) lie below that range, which bears on how the differences in Section 4.2 should be read and is taken up there.

An energy detector also cannot say *what* raised the energy. Elevated N_j may be a hostile emitter, another legitimate user of the band, a harmonic, or the node's own transmitter leaking between slots. The paper therefore distinguishes three things that are easily conflated: an *energy elevation*, which is the quantity measured; *interference*, which is an inference about its cause; and *channel unusability*, which is a decision. The threshold N^* turns the first into the third directly, and the discriminator of Equation (8) is the only place where the second is even attempted, from the joint behaviour of energy, symbol errors and delivery.

Because 97 channels cannot be swept every superframe, the sweep is a rolling slice of eight channels, so each channel is refreshed roughly every twelve superframes; the channel in use is additionally measured in every superframe. A first-order filter smooths the measurement,

$$\bar{N}_j[k] = \alpha\, \bar{N}_j[k-1] + (1-\alpha)\, N_j[k], \qquad \alpha = 0.7. \tag{3}$$

The value $\alpha = 0.7$ was determined empirically. The filter is updated at each measurement, so its time constant is $-1/\ln\alpha$ updates, i.e., 2.8 updates at $\alpha = 0.7$ and 32.8 at the initial value $\alpha = 0.97$, which suits a slowly varying ambient floor. The channel in use is measured every superframe of 64 transmissions (1.24 s in the run of Section 4.2, which used a 20 ms slot), giving time constants of about 3.5 s and 41 s, respectively; the other channels are measured about every twelve superframes, so their time constants are about twelve times longer. Since the interferer in these experiments remained on a frequency for only a few seconds at a time, at $\alpha = 0.97$, the score lagged the band by tens of seconds or more (Section 5.4). The channel loop decides once per superframe, so the filter must settle within a few updates, and 0.7 is the largest value that does so. Slower interferers would justify a larger value. The blacklist and the discriminator use the raw measurement, so the filter affects only the score.

A channel score combines noise, delivery and a blacklist,

$$\Phi_j = w_n \frac{N_{\max} - \bar{N}_j}{N_{\max} - N_{\min}} + w_p\,(1 - \mathrm{PER}_j) - w_b\, \mathbb{1}\!\left[t < t_j^{\mathrm{ban}}\right], \tag{4}$$

with $w_n = 0.6$, $w_p = 0.4$, $w_b = 1.0$, $N_{\max} = -30$ dBm and $N_{\min} = -100$ dBm, and $\bar{N}_j$ clamped to that range. A channel whose *measured* noise exceeds the interference threshold $N^* = -80$ dBm is blacklisted for t^{ban}, set to 15 s for the reported runs (the compile-time default is 60 s; the value used is recorded in every log). The blacklist test uses the raw measurement rather than (3): it is a hazard test asking whether the channel is unusable now, and on the average, a channel stays blacklisted long after the interferer has left it.

Link quality is tracked by the residual symbol error rate seen by the decoder,

$$p_s = \frac{\sum_i e_i}{\sum_i n_i}, \tag{5}$$

where e_i is the number of symbols Reed–Solomon corrected in codeword i and n_i its length, and by a signal-to-interference proxy

$$\gamma = \bar{S} - \bar{N}, \tag{6}$$

with $\bar{S}$ being the smoothed received level of valid frames. Transmit power follows an integral law, clamped to the hardware grid and moved at most one step per eight frames,

$$P_{\text{tx}} \leftarrow P_{\text{tx}} + \beta(\gamma^* - \gamma), \qquad \beta = 0.5. \tag{7}$$

A discriminator D names the reason a window degraded, which allows the three adaptation loops to respond differently to causes that look alike in a packet-error count:

$$D = \begin{cases} \text{ATTACK} & \text{if } n_{\text{mac}} > 0, \\ \text{INTERFERENCE} & \text{if } q > T_{\text{lost}} \text{ and } N > N^*, \\ \text{DESYNC} & \text{if } q > T_{\text{lost}}, \\ \text{INTERFERENCE} & \text{if } \left(p_s > p_{\uparrow}^* \text{ or } \Pi\right) \text{ and } N > N^*, \\ \text{FADING} & \text{if } p_s > p_{\uparrow}^* \text{ or } \Pi, \\ \text{MARGIN} & \text{if } \bar{S} \geq S_{\text{low}} \text{ and } p_s < p_{\uparrow}^*/4, \\ \text{NONE} & \text{otherwise,} \end{cases} \tag{8}$$

where q is the time since the last valid frame, T_{lost} the silence threshold, and Π is the packet-error alarm

$$\Pi \;=\; (a \geq 8) \wedge (4b > a), \tag{9}$$

with a the transmissions attempted in the window and b those that failed. The minimum-sample gate $a \geq 8$ matters: without it, a window truncated by a channel change would raise the alarm on one or two frames.

A frame that passes its cyclic redundancy check but fails its authentication tag is classified ATTACK. This freezes the mode ladder (the up and down run-length counters are cleared) and excludes the window from triggering a channel change, so that an adversary cannot steer either loop by forging traffic. The power loop is not frozen, since it is driven by the peer's authenticated feedback rather than by the discriminator. After $\text{ATT}_{\text{thr}} = 16$ consecutive tag failures, the node increments its own epoch and rekeys unilaterally, on the assumption that a sustained run of failures is more likely to be lost epoch agreement than a live adversary; a false trigger costs one additional rekey.

3.7. Mode Ladder and Switching Mechanics

Six modes trade throughput for robustness (Table 3). Each fixes a physical layer, a Reed–Solomon correction power t, an interleaver depth D, an in-frame repetition factor R and a segment size L_{seg}. The air-frame length follows as

$$n_{\text{air}} = R \cdot D \cdot (k + 2t), \qquad k = \left\lceil \frac{H + L_{\text{seg}} + \tau}{D} \right\rceil, \tag{10}$$

with $H = 14$ the header length. Codewords are shortened Reed–Solomon over GF(256), so parity remains $2t$ symbols, and each mode keeps its nominal correction power.

L5 carries no payload; it is a resynchronisation mode whose beacon is authenticated with a message authentication code keyed on the master key rather than the epoch key so that it remains verifiable at an end whose epoch has diverged.

Three adaptation loops run at different rates (Figure 3): power per slot, mode per 32-frame window, channel per superframe. The mode moves up when p_s exceeds the mode's threshold $p_{\uparrow}^*$ for $W_{\uparrow} = 2$ consecutive windows and down after $W_{\downarrow} = 8$ windows below $p_{\uparrow}^*/3$; the asymmetry is deliberate, since degrading quickly and recovering slowly is the safer error. Every mode or channel change increments an epoch, and the epoch

re-derives the traffic key through HKDF-SHA256, so a change of operating point is also a rekey.

Table 3. The mode ladder. t is the Reed–Solomon correction power, D the interleaver depth, R the repetition factor, τ the tag length, γ_{req} the lower edge of the operating window.

Mode	PHY	t	D	R	L_{seg} (B)	τ (B)	n_{air} (B)	Frame (µs)	γ_{req} (dB)
L0	GFSK 2 Mbit/s	0	1	1	224	8	246	1028	22
L1	GFSK 1 Mbit/s	8	2	1	128	8	182	1536	16
L2	GFSK 1 Mbit/s	16	4	1	64	8	216	1800	10
L3	Coded $S = 2$	16	4	1	64	12	220	3982	5
L4	Coded $S = 8$	16	2	2	32	16	252	16,848	0
L5	Coded $S = 8$	16	1	3	24	8	234	15,696	—

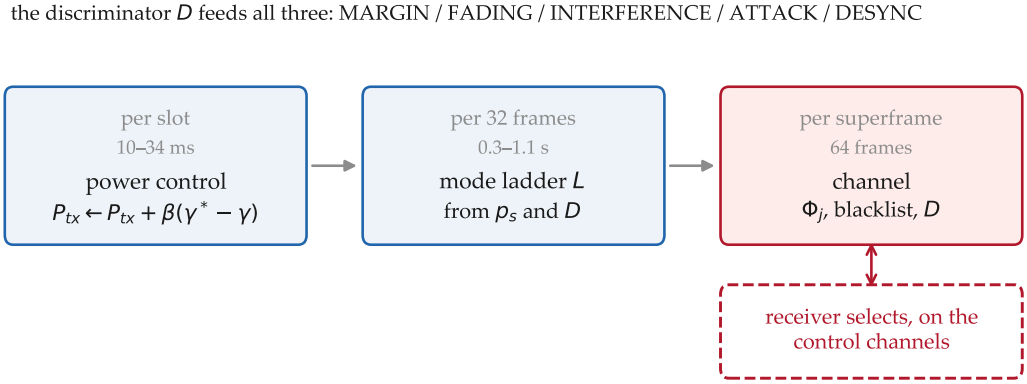


Figure 3. The three adaptation loops. The solid boxes are conventional; the dashed box is the addition made by this work, and it is attached to the channel loop because that is the only one of the three whose decision the far end can contest. The discriminator output D is shared by all three, which lets the channel loop respond to interference while the mode loop responds to fading.

3.8. The Asymmetric Interference Perception Method

The method replaces a unilateral channel decision with a two-sided one. It has four parts:

(i) Independent perception. Both ends run (3) and (4) over the same 97 channels. Neither end's estimate is communicated as ground truth; each is authoritative only for the node that made it.

(ii) A ranked proposal, not a decision. When the discriminator returns INTERFERENCE or DESYNC, the transmitter forms a shortlist of the $m = 6$ best-scoring channels that are not blacklisted,

$$\mathcal{C} = \underset{|\mathcal{C}|=m,\, c \notin \mathcal{B}}{\arg\max} \ \Phi_c, \tag{11}$$

and rotates the entry selected by AES-CMAC under the hop key to the head of the list. The list is sealed under the link key and sent as a CTRL frame. Only the transmitter initiates a negotiation, and its discriminator uses its own noise reading. Interference local to the receiver therefore reaches the decision indirectly: acknowledgement loss raises the packet-error alarm, which is classified as FADING and answered with more power while the transmitter's own floor is low, and a negotiation starts only when the link falls silent for longer than T_{lost} (DESYNC). The receiver, for its part, moves to the control channels after T_{quiet} of silence (part (iv)).

Choice of the shortlist length. The value $m = 6$ follows from two considerations. First, the probability that every candidate is rejected is not the binding constraint. In the paired data of Section 4.2, the receiver classified a channel as interfered in 4 of the 618 windows in which the transmitter classified it as clean, so the estimated probability that a channel

acceptable to the transmitter is rejected by the receiver is $q_{\text{rej}} = 4/618 \approx 0.0065$. This estimate refers to the channel in use rather than to candidates; in the 43 negotiations under interference, the first candidate was always accepted (Section 4.6). The probability that all m candidates are rejected lies between q_{rej}^m (independent rejections: 6.5×10^{-3}, 2.7×10^{-7} and 7.4×10^{-14} for $m = 1, 3$ and 6) and q_{rej} (fully correlated rejections, e.g., under a wideband interferer). Even in the latter case it is small, and because of the fallback in part (iii), a total rejection does not cause a failed negotiation. A list with $m > 1$ is required for a different reason: with $m = 1$, the receiver can only accept or reject, and the mechanism reduces to unilateral selection, whereas with $m > 1$, the receiver can choose among channels that are acceptable to the transmitter (Equation (12)).

Second, the frame budget fixes the exact value. The 16-byte CTRL body contains the version, kind, candidate count and chosen channel (four bytes), the epoch and the transmit counter at which the change takes effect (two bytes each), the responder's noise reading and the triggering cause (one byte each), and six one-byte candidates. A longer list would require a longer air frame on the control channels, where the exchange must survive a degraded link, and would lengthen every repetition within the bounded sweep of part (iv). Convergence time is set by the dwell relation (13), not by m. Any value between three and eight would serve; six is the largest value that fits the existing frame.

Security limitations. The rotation is a single rotation, not a permutation, and it protects only against an adversary who can observe transmission timing but cannot decrypt the exchange: without K_{master}, the sealed list cannot be read at all, and with K_{master}, which unauthenticated pairing does not protect, the hop key and the list are both available. In addition, to tolerate epoch divergence, the receiver attempts decryption against the current epoch and four epochs on either side, so a validly sealed control frame can be replayed within that window. An attacker can thereby force a renegotiation or return the pair to a previously agreed channel, but it cannot introduce a channel of their choice. Three known remedies, none implemented here, would close this exposure in increasing order of cost: (1) rejecting any CTRL frame whose `at_txc` field does not exceed the highest value already accepted under the current key; (2) binding a session nonce from the pairing exchange into the key derivation, which also covers node restarts; and (3) a challenge-response in the request and reply, against an adversary who delays rather than replays frames. The security claims of this paper are accordingly limited to the integrity and confidentiality of the control exchange *within a trusted pairing model*.

(iii) The receiver decides among them. The receiver replies with the first candidate that is also acceptable where it is standing,

$$c^* = \min\{i : N_B(c_i) \le N^*\}, \tag{12}$$

falling back to $\arg\min_i N_B(c_i)$ when no candidate satisfies the threshold, and returns its own measured $N_B(c^*)$ with the reply. The transmitter commits by repeating the request with c^* filled in; both then move. Because the reply is armed before the request arrives (the radio's interrupt service routine turns the transceiver around autonomously), the answer reflects the previous request, which costs one repetition and is accounted for in the timing below.

The fallback bounds what the mechanism guarantees. The receiver always returns a channel: if none of the six candidates is below N^* at its position, it returns the quietest of them rather than refusing. This is deliberate, because a refusal would leave the pair on the channel that is already failing, and the transmitter would have no better information with which to construct a second list. The receiver therefore does not hold a veto in the sense of IEEE 802.11h [8], where a detecting station can compel a channel to be abandoned; it holds

the narrower power to determine *which* of the transmitter's candidates is used. The only outright refusal in the design is termination of the session, described below.

(iv) Rendezvous and termination. The exchange runs on the three control channels in turn (Figure 4). The transmitter holds each for $T_{\text{tx}} = 400$ ms, repeating every 40 ms; the receiver rotates the same three with a dwell of $T_{\text{rx}} = 120$ ms, going there after $T_{\text{quiet}} = 2$ s of silence. Overlap is guaranteed by construction whenever

$$T_{\text{tx}} > n_c \cdot T_{\text{rx}}, \tag{13}$$

with $n_c = 3$ control channels. The receiver visits a given channel for T_{rx} in every $n_c T_{\text{rx}}$, so a transmitter dwell overlaps its visits to that channel for a contiguous interval of at least $(T_{\text{tx}} - (n_c - 1)T_{\text{rx}})/2 = 80$ ms, which contains at least two request repetitions at the 40 ms period. When the link is still alive, the transmitter first announces its intent in-band, by setting the header's next-channel field to a reserved value, so the receiver goes to the control channels immediately instead of waiting out the silence timer. If no reply arrives on any of the three channels over six sweeps, the session terminates: the alternative is to transmit on a channel the receiver was never told about and to interpret the resulting silence as a link.

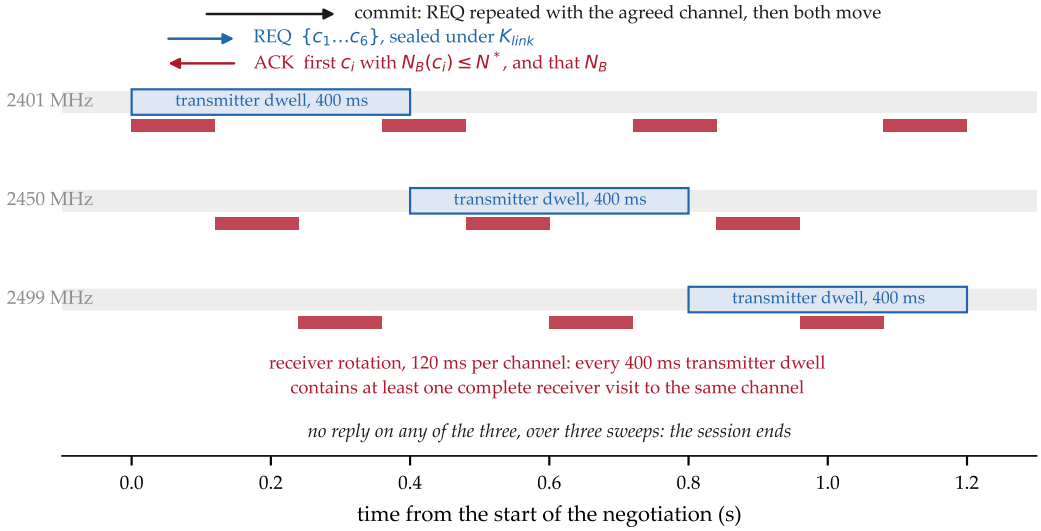


Figure 4. The negotiation, drawn to the implemented dwell times. The dwell relation (13) makes rendezvous deterministic without the two ends being able to agree in advance where to meet.

3.9. Measurement Methodology

Three data sources were used: the two nodes and the SDR, which is independent of both. The nodes report internal state as one JSON object per line over USB: per-window noise floor, smoothed signal, symbol error rate, discriminator output, counters. The SDR reconstructs frames from the air and applies the same decoder. The host records everything with a wall-clock timestamp, which permits the two nodes' windows to be paired.

For the asymmetry measurement, a window from node A was paired with the temporally nearest window from node B, subject to both reporting the same channel and to a maximum separation of 0.5 s. The achieved alignment is much tighter than this tolerance. The distribution of $|t_A - t_B|$ is bimodal: 51% of pairs are aligned to within 1 ms, the median is 0.78 ms, and the remainder cluster near 207 ms, half an adaptation window, with a maximum of 418 ms. The 0.5 s tolerance is therefore a cut-off that admits the second cluster; it was chosen because it is shorter than the several seconds for which the interferer held a frequency and level. Its effect is examined in Section 4.2 by recomputing all statistics at 1, 10 and 100 ms.

The two nodes free-run and their window boundaries are not aligned, so 13% of pairs share a node-B window with a neighbouring pair (never more than two). This is a mild additional dependence on top of the temporal autocorrelation treated below.

Statistical treatment. The paired windows are consecutive observations roughly 0.41 s apart within a single 546 s run, and they are not independent. The difference series has a lag-1 autocorrelation of 0.51 and an integrated autocorrelation time of $\tau_{\mathrm{ac}} = 3.49$ windows (used in Equation (15)). The dependence is physical rather than an artefact: the filter of Equation (3) smooths each node's reading over several windows, and the interferer occupied one frequency and level for several seconds at a time. Every interval reported below is therefore a moving-block bootstrap [35] over 20,000 resamples with a block length of 28 windows ($\approx 8\tau_{\mathrm{ac}}$), which preserves the correlation structure; point estimates are additionally checked by thinning the record to one window per second ($n = 237$), which moves the 3 dB exceedance from 83.8% to 83.1%, the 6 dB exceedance from 57.6% to 56.1% and the mean from $+7.36$ to $+6.99$ dB.

Where the unit of analysis matters, as it does for the direction of the asymmetry, results are reported over *episodes* rather than over windows. An episode is a contiguous run of windows in which at least one end classifies the channel as interfered against N^{*}, runs separated by fewer than three quiet windows being merged; consecutive windows within one episode are one event observed repeatedly, not repeated events.

The co-location control. A difference between two receivers standing in two places is not evidence about the two places until the difference between the two *instruments* is known. Both ends were therefore also measured side by side: the boards were placed 10 cm apart in the same orientation, both links were stopped so that neither transmitted (at that spacing a live link would put roughly -20 dBm into the peer's front end and the measurement would be of the peer rather than of the room), and both nodes swept all 97 channels once a second for ten minutes. Each channel is thus measured by both receivers in nominally the same field, and the two sweeps align to a median of 2 ms.

The band was not left quiet throughout. A narrowband generator ran from the second to the fourth minute and a wideband generator from the sixth to the eighth, with the first, third and fifth two-minute intervals quiet. The schedule was recorded by hand rather than inferred from the data, which avoids selecting the reference samples on the basis of the data they are meant to characterise. Only the quiet intervals bound the instrument offset; the illuminated intervals address a different question (Section 4.3).

Ten centimetres is about 0.8 wavelengths at 2.44 GHz, so the two antennas do not sample the same point; any residual field difference is included in the instrument offset, which is therefore an upper bound. The control also differs from the link measurement in operating state: during the link measurement node A transmitted and node B received in the configured mode, whereas in the control, neither transmitted and both swept the band. The control therefore isolates receiver calibration, not the combined effect of position and operating role (Section 4.3).

Signal-to-noise estimator. The signal-to-noise ratio used for the capacity bound is the mean of $S - N$ over every adaptation window the receiver reported while operating in that mode, which is between 136 and 458 windows depending on the mode. Its standard deviation is reported alongside, because the spread is not negligible, and a capacity computed from a single sample would not be reproducible.

Processing times were measured on the target with a 1 MHz timer around each stage, with the radio stopped, reporting count, sum and sum of squares so that means and standard deviations are computed exactly on the host without floating-point arithmetic in the firmware. Goodput was measured as unique application segments delivered and acknowledged over a 12 s window after waiting for the link to reach steady state.

3.10. Two-Antenna Measurement with a Moving Transmitter

The co-location control compares two different devices, so it cannot separate the receiver response from position. An additional experiment was therefore performed with two receivers that are identical by construction: the two receive chains of the bladeRF, which share one local oscillator and one sample clock and are sampled simultaneously. Two parallel dipoles were placed $d = 0.60$ m apart, i.e., 4.98λ at the measurement frequency of 2490 MHz ($\lambda = 0.1204$ m). The SDR was tuned to 2487 MHz at 10 MS/s with a fixed gain of 24 dB. A node transmitted at 0 dBm in mode L2 on channel 90 with a 10 ms slot, and the other node, placed beside the SDR, acknowledged its frames. In this experiment the nice!nano was the transmitter and the Feather the receiver, i.e., the roles of Section 3 were exchanged; the transmitter ran a firmware variant that starts transmitting at power-up and was powered through a long USB cable.

The transmitter was placed 3.5 m from the midpoint of the baseline at a bearing of -10° from broadside. The planned schedule, cued to the operator by recorded speech, was a stationary interval of 15 s, motion by hand along the arc from -10° to $+10^\circ$ and back for 60 s and a final stationary interval of 10 s; the times actually observed are given in Section 4.4. The operator and the transmitter's cable moved near the feeders during the run. Each capture was mixed to baseband and decimated by a ten-sample boxcar filter, whose nulls fall on the 1 MHz channel raster. A burst was declared where the combined power exceeded the median by 10 dB, and the central 80% of each burst was used to compute the two levels s_0 and s_1, the phase difference $\varphi = \arg \sum x_0 x_1^*$ and the coherence $|\sum x_0 x_1^*| / \sqrt{\sum |x_0|^2 \sum |x_1|^2}$. DATA frames (1.84 ms) from the moving transmitter and ACK frames (0.57 ms) from the fixed node were separated by duration and level so that the fixed node provides a simultaneous reference for the stability of the instrument phase.

For a source in the far field at bearing θ from broadside, the phase difference is

$$\varphi = \frac{2\pi d}{\lambda} \sin\theta + \varphi_0, \tag{14}$$

where φ_0 is a constant instrument offset [18]. The sensitivity at broadside is $2\pi d/\lambda = 31.3$ rad per unit $\sin\theta$, i.e., 0.546 rad per degree, six times the value for a 10 cm spacing. The phase is unambiguous only for $|\sin\theta| < \lambda/(2d)$, i.e., within $\pm 5.8^\circ$, and the $\pm 10^\circ$ arc corresponds to 1.73 phase cycles. The Fraunhofer distance of the pair, $2d^2/\lambda$, is 5.98 m [19], so the 3.5 m range does not satisfy the formal far-field criterion. Evaluated directly, however, the difference between the exact path difference and its plane-wave approximation over the arc is at most 0.019 rad at 3.5 m, which is negligible against the measured phase noise (Section 4.4).

Before the motion run, each dipole was removed in turn for 10 s while the transmitter was stationary. Removing either dipole changed the level of the corresponding chain by less than 4 dB, so a substantial part of the signal is received by the feeder cables and connectors in addition to the dipoles. A preliminary run, in which the source was a node a few centimetres from the SDR, showed changes below 1.5 dB, i.e., pickup that bypasses the antennas almost entirely. Both observations bear on the interpretation of the SDR measurements and are taken up in Sections 4.3 and 4.4.

3.11. Decision-Level Replay of Selection Strategies

A live comparison of channel-selection strategies, i.e., with each strategy run on the link under a repeated interference schedule and its packet error rate and outage measured, was not performed. Two analyses of the recorded data provide a partial comparison. The first operates on the 656 paired windows: for each strategy, it counts the windows in which the deciding end would keep a channel that the other end classifies as interfered against N^*,

and the windows in which the strategy would leave the channel. Transmitter-only selection decides on N_A, receiver-only selection on N_B, and intersection-based selection requires both $N_A \leq N^*$ and $N_B \leq N^*$. This replay uses the classifications only; it does not reproduce the indirect feedback through packet loss, by which interference at the receiver also degrades the transmitter's view of the link. The second operates on the 44 negotiations initiated during the same run: it reports their completion and duration and, for each completed negotiation, compares the channel selected by the receiver with the transmitter's first candidate, which is the channel transmitter-only selection would have imposed. Random selection and a full receiver-only or intersection selection over all 97 channels cannot be replayed, because the receiver's complete channel map was not logged.

4. Results

4.1. Correctness of the Implementation

The implementation was verified before the measurements were taken. Reed–Solomon encoding and decoding passed 1200 randomised cases per node across all six geometries with mixed errors and erasures and no failures. The host and firmware frame chains agree byte for byte. The SDR receiver recovered 11 of 11 mode and direction combinations from synthetic signals at 25 dB and at 12 dB signal-to-noise ratio. On air, all six modes carried traffic in both directions and were decoded independently by the SDR with zero residual payload bit errors and zero corrected symbols on a clean link. The pairing key derivation reproduces the published Bluetooth test vectors for the f_4, f_5 and f_6 functions with no failures on either node; the expected values were additionally cross-checked against an independent open-source implementation of the same specification, which stores them in the opposite byte order, and agree after reversal.

4.2. Asymmetry of Interference Perception

The central measurement pairs 656 adaptation windows in which both nodes reported the same channel within a median of 1 ms of each other, over a 546 s run during which a narrowband generator was moved across the band forty times and the pair was made to leave the occupied channel each time. The pair did so through 43 negotiated channel changes (Section 4.5), which spread the record across 26 of the 97 data channels. Figure 5 shows the result and the paragraphs below give the numbers.

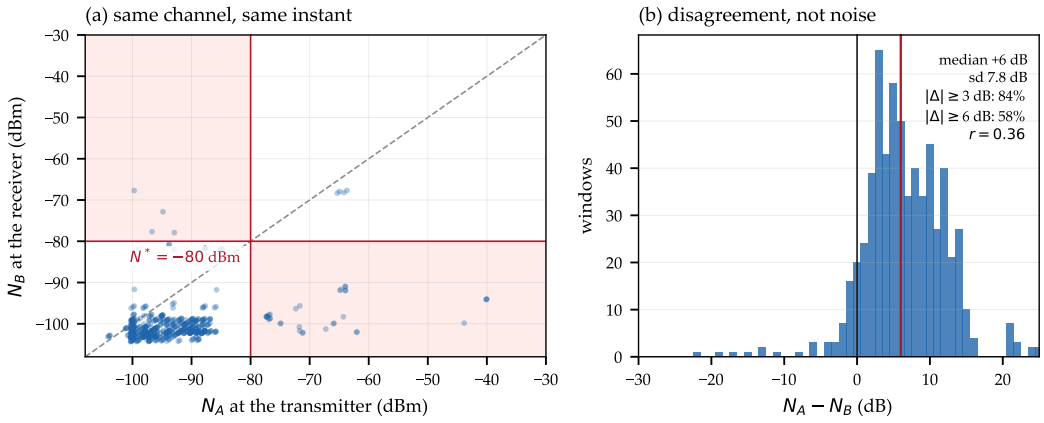


Figure 5. Interference perception at the two ends of a 4.2 m link. (**a**) Every paired window; points off the diagonal are disagreement, and points in the shaded quadrants are disagreement about the decision itself. (**b**) The distribution of $\Delta = N_A - N_B$: centred away from zero, with tails reaching tens of decibels.

Across those 656 pairs, the difference $\Delta = N_A - N_B$ has a mean of +7.36 dB (95% confidence interval (CI) 6.35–8.66), a median of +6.00 dB and a standard deviation of

7.79 dB, ranging from −32 to +56 dB. The 50th, 75th, 90th and 95th percentiles of $|\Delta|$ are 6, 11, 14 and 21 dB, and $P(|\Delta| \geq 3\ \text{dB})$ is 83.8% (79.0–89.3), $P(|\Delta| \geq 6\ \text{dB})$ is 57.6% (49.4–68.0) and $P(|\Delta| \geq 10\ \text{dB})$ is 32.9%. The two readings correlate at $r = 0.361$. Successive windows are not independent: the integrated autocorrelation time of Δ is $\tau_{\text{ac}} = 3.49$ windows, giving an effective sample size

$$n_{\text{eff}} = \frac{n}{\tau_{\text{ac}}} = \frac{656}{3.49} \approx 188, \tag{15}$$

so every interval quoted here is a moving-block bootstrap rather than a normal-theory interval.

Classified against $N^* = -80$ dBm, the two ends agree that the channel is clean in 614 windows and that it is interfered in 4; the transmitter alone calls it interfered in 34 and the receiver alone in 4. The disagreement rate is therefore 5.79% (38 of 656, CI 2.90–9.91).

Magnitude. The intervals above are wide because the effective sample size is 188, not 656. The two readings correlate at only $r = 0.361$, so $r^2 \approx 0.13$: about an eighth of the variance in one end's measurement is explained by a linear relationship with the other's. The residual is not necessarily local interference, since it also contains measurement noise, temporal misalignment, quantisation and device-specific effects. Moreover, most of the exceedance reflects the standing offset described next: in the 614 quiet windows alone, $P(|\Delta| \geq 3\ \text{dB})$ is 82.7% (Table 4). The interference-specific result is the behaviour of the 42 windows within episodes, described below.

Table 4. The co-location control: the same two boards 10 cm apart, neither transmitting, sweeping the band once a second for ten minutes. The generators ran on the stated schedule; only the quiet rows bound the instrument offset.

Regime	Samples (n)	Median Δ (dB)	Mean Δ (dB)	sd (dB)	$P(\lvert\Delta\rvert \geq 3)$ (%)
Quiet, channel dark [1]	34,076	+1	+0.66	2.48	21.4
Source on, this channel dark	15,775	−1	—	5.07	—
Narrowband source on this channel	981	+9	—	9.44	—
Wideband source on this channel	6524	−7	—	23.80	—
For comparison, the same statistic across the 4.2 m link (Section 4.2)					
Quiet windows, 4.2 m separation	614	+6	+6.30	4.77	82.7

[1] The instrument bound. The other rows are not instrument offsets and must not be read as such; see the text.

Cause: the standing offset. The mean difference of +7.36 dB could be produced by a persistent offset between the two receivers without any environmental asymmetry. Restricting to the 614 windows in which *neither* end saw interference, the mean Δ is +6.30 dB (CI 5.28–7.58) with a standard deviation of 4.77 dB; that is, node A reads on average 6.3 dB higher than node B.

The co-location control (Section 4.3) indicates that this offset does not originate in receiver calibration. With the two boards 10 cm apart in the same field and neither transmitting, the offset is +0.66 dB (CI 0.61–0.71), with a standard deviation of 2.48 dB over 34,076 paired measurements, against +6.30 dB and 4.77 dB at 4.2 m. Removing the measured per-channel offset from the interference data leaves the result substantially intact: $P(|\Delta| \geq 3\ \text{dB})$ falls from 83.8% to 81.1%, $P(|\Delta| \geq 6\ \text{dB})$ from 57.6% to 50.6%, $P(|\Delta| \geq 10\ \text{dB})$ from 32.9% to 22.0%, and the classification disagreement rate does not move at all. The standing offset seen across the link is therefore attributed to the two ends' positions or operating roles rather than to receiver calibration; exchanging the two nodes would separate position from role (Section 5.6). The quiet floors (about −100 dBm) lie below the range over which the RSSI accuracy is specified (Section 3.6); since the control measured the two receivers at the same levels, this affects the absolute values but not the comparison between them.

The dispersion leads to the same conclusion. Inside the fifteen interference episodes, the standard deviation of Δ is 19.2 dB, and *every* one of the 42 windows differs by at least 3 dB; outside them it is 4.77 dB. Co-located, in a quiet field, it is 2.48 dB. A constant offset between two receivers cannot produce a fourfold change in spread; the interferer decouples the two measurements rather than shifting them apart by a fixed amount. How much of this decoupling is due to position is limited by the co-location control (see **Limit of the co-location control** below).

Direction: not established. The mean difference is $+7.36$ dB rather than zero, and the transmitter alone declared the channel interfered in 34 windows against the receiver's 4. Treated as 38 independent discordant observations, this would be strongly significant, but they are not independent: they fall within the fifteen episodes defined in Section 3.9, of which twelve are transmitter-side and three receiver-side. An exact binomial test on episodes gives $p = 0.035$; however, the result depends on how an episode boundary is defined. Separating episodes at one quiet window gives sixteen episodes and $p = 0.021$; at five or more windows, fourteen episodes and $p = 0.057$; and at twenty windows, eleven episodes and $p = 0.065$. Because the conclusion changes with an analysis choice that has no physical justification, the direction of the asymmetry is not regarded as established by this dataset, and no conclusion in this paper depends on it.

Consequence. A difference matters to a protocol only if it changes the decision. Table 5 gives, for each candidate threshold x, both the fraction of windows in which $|\Delta| \geq x$ and the quantity relevant to the protocol, namely, the fraction of *those* windows in which the two ends classified the channel differently against N^*. The conditional probability rises monotonically, from 6.0% at 1 dB to 17.6% at 10 dB. Below 3 dB, it is exactly zero, and every window in which the two ends classified the channel differently had $|\Delta| \geq 15$ dB. The monotonic rise follows from this, since all 38 disagreements fall in every row of the table. The 3 dB threshold used to describe the differences is therefore conservative with respect to the decision; it is, however, comparable to the co-located instrument spread (21.4% of co-located quiet samples differ by at least 3 dB; Table 4).

Table 5. How often a difference of a given size changes the decision. The second column is the fraction of all 656 windows in which the two ends' readings differ by at least x; the fourth is the fraction of those windows in which they also classified the channel differently against $N^* = -80$ dBm.

Threshold x (dB)	$P(\lvert\Delta\rvert \geq x)$ (%)	Windows (n)	$P(\text{Differ} \mid \lvert\Delta\rvert \geq x)$ (%)
1	97.0	636	6.0
2	90.9	596	6.4
3	83.8	550	6.9
5	66.5	436	8.7
6	57.6	378	10.1
10	32.9	216	17.6

The rate at which the two ends disagree outright is imprecisely determined. Only 42 of 656 windows have either end above N^*, in fifteen contiguous episodes; the rest of the run is agreed quiet. At the window level, 38 of 656 windows (5.79%, bootstrap interval 2.90–9.91%) were classified differently, but the experimental unit behind that number is the episode, of which there are fifteen, and each contained at least one disagreeing window. The appropriate statement is therefore the compound one: *of 656 paired windows, 42 fell within fifteen interference episodes; 38 of them showed a classification disagreement, and all fifteen episodes contained at least one.* The percentage is a description of this record; narrowing the interval requires more independent episodes, not a longer record. For completeness, raw agreement over the windows is $p_o = 94.2\%$, and Cohen's κ is 0.157 [36]. The difference between them reflects the known behaviour of κ when 94% of observations fall in one cell;

the prevalence-adjusted bias-adjusted kappa is PABAK $= 2p_o - 1 = 0.884$ [37]. Table 5 is more informative for the present purpose.

Robustness. Table 6 varies the two analysis choices, the pairing tolerance and the bootstrap block length. Tightening the pairing tolerance from 500 ms to 1 ms discards half the pairs and moves $P(|\Delta| \geq 3$ dB) from 83.8% to 81.0%, $P(|\Delta| \geq 6$ dB) from 57.6% to 55.4% and the disagreement rate from 5.79% to 3.92%; the mean difference falls from +7.36 to +6.91 dB. The exceedance statistics are therefore not an artefact of loosely aligned pairs. The disagreement statistic is more sensitive: at 1 ms, 13 disagreeing windows remain among 16 interfered windows in seven episodes, each of which still contains a disagreement. Disagreements are thus more frequent among pairs separated by about half an adaptation window, but they do not depend on them. Varying the block length over 14, 28, 56 and 84 windows (from $4\tau_{ac}$ to $24\tau_{ac}$) leaves every point estimate unchanged by construction and widens the intervals by one to two percentage points at each end.

Table 6. Robustness of the headline estimates to the two analysis choices. Upper block: the maximum time separation at which two windows are treated as simultaneous. Lower block: the moving-block bootstrap block length, over 20,000 resamples; point estimates are unaffected by construction and only the intervals are shown.

Variant	Pairs (n)	Mean Δ (dB)	$P(\|\Delta\| \geq 3)$ (%)	$P(\|\Delta\| \geq 6)$ (%)	Disagree (%)
Pairing tolerance					
≤1 ms	332	+6.91	81.0	55.4	3.92
≤10 ms	339	+6.98	81.1	56.0	4.13
≤100 ms	451	+6.96	82.3	54.5	5.10
≤500 ms [1]	656	+7.36	83.8	57.6	5.79
Bootstrap block length, 95% interval					
$L = 14$	656	6.36–8.49	79.3–88.7	50.0–66.2	2.90–9.60
$L = 28$ [1]	656	6.35–8.66	79.0–89.3	49.4–68.0	2.90–9.91
$L = 56$	656	6.55–8.91	78.4–90.4	49.1–71.0	2.90–9.91
$L = 84$	656	6.75–9.11	79.6–91.0	49.5–73.2	2.90–9.60

[1] The value used for the results reported elsewhere in this paper.

Limit of the co-location control. The control bounds the offset in a *diffuse* field, not in a directed one (Section 4.3). Ambient energy arrives from all directions, so a difference in antenna pattern between the two devices averages out, and in this regime the co-located devices agree to within a decibel on average. Under a narrowband source on the same channels, the same two devices, still 10 cm apart, differ by a median of 9 dB. Consequently, the standing offset across the link is attributed to the two ends' positions or operating roles, but an unknown part of the larger differences *during* interference may reflect the two receivers responding differently to a directed signal rather than the two positions experiencing different fields. Section 4.4 shows that, with identical receivers, a separation of 60 cm produces differences of the same order as the 9 dB co-located difference, and Section 5.6 describes the controls that would apportion the two contributions.

Table 7 breaks the result down by channel for the ten channels that contributed most windows. The per-channel medians differ by up to 27 dB, and the disagreement rate is highest on the channels the interferer occupied while the link was still trying to use them.

Table 7. Noise floors measured at the two ends, and how often they classified the channel differently, for the ten channels contributing the most paired windows out of the 26 the run occupied.

Channel (MHz)	Windows (n)	Median N_A (dBm)	Median N_B (dBm)	Disagree (%)
ch7 (2407)	77	−92	−101	11.7
ch9 (2409)	43	−97	−99	18.6
ch5 (2405)	38	−98	−100	2.6
ch10 (2410)	30	−97	−102	0.0
ch4 (2404)	28	−97	−102	3.6
ch17 (2417)	22	−95	−103	9.1
ch14 (2414)	11	−97	−103	0.0
ch6 (2406)	10	−92	−99	0.0
ch12 (2412)	10	−88	−96	0.0
ch8 (2408)	7	−64	−92	85.7

4.3. Instrument Offset: Co-Location Control

Table 4 gives the co-location control. Over the three quiet intervals, and counting only channels on which neither receiver saw anything above N^*, the two boards 10 cm apart differ by a mean of +0.66 dB (95% CI 0.61–0.71), with a median of +1 dB and a standard deviation of 2.48 dB, over 34,076 paired measurements. They differ by 3 dB or more in 21.4% of those measurements and by 6 dB or more in 3.8%. Across the band the per-channel median offset runs from −1 to +6 dB with a standard deviation of 1.13 dB between channels and no trend with frequency (+0.006 dB per MHz), so the correction applied in Section 4.2 is per-channel rather than scalar.

Against the standing offset of +6.30 dB (standard deviation 4.77 dB) across the link (Section 4.2), the control indicates that receiver calibration contributes little to the difference in a quiet band, subject to the difference in operating state noted in Section 3.9.

What the control does not cover. The illuminated rows of Table 4 qualify the result above rather than extending it. On the same channels on which the two co-located boards agree to within a decibel when the band is quiet, they differ by a median of 9 dB when the narrowband generator is placed on them, and the difference is systematic: 83% of those samples have $\Delta > 0$, with an interquartile range of +5 to +12 dB. Channel by channel, the same seven channels that show median offsets of $+1, 0, +1, 0, +1, +1, +2$ dB dark show $+9, +10, -7, +16, +5, +13, +8$ dB lit.

Since the two boards are almost at the same position, the likely explanation is that ambient energy arrives from every direction at once, and a generator does not. A difference in antenna pattern, in near-field coupling between two boards less than a wavelength apart, or in front-end compression at high input averages out of a diffuse field and does not average out of a directed one. The two boards are different designs with different antennas, so a pattern difference of several decibels toward one bearing is to be expected.

The *standing* offset across the link is attributed to the two ends' positions or operating roles, because that conclusion rests on the quiet-field comparison, which is the regime the control covers. However, the much larger differences observed *during* interference cannot be attributed wholly to the two ends occupying different places, because two receivers in the same place also differ under a directed source. How the 19.2 dB in-episode dispersion at 4.2 m divides between position and receiver response is not determined by this experiment.

A check with the software-defined radio. The two nodes never sample a channel at the same instant, so a changing field could in principle contribute to their difference. The two receive chains of the SDR share an oscillator, a clock and a capture buffer, so a pair of readings taken from one block is simultaneous by construction. With a dipole on each input and one node used as a source stepped through the transmit-power grid, the difference

between the two chains behaves as Table 8 shows: it is +3.05 dB on the diffuse ambient floor and between −4.6 and −16.6 dB on the node's signal.

Table 8. The two receive chains of one software radio, sampled from the same capture block so that no time separates them, with parallel dipoles and one node as a source stepped through the transmit-power grid on one channel. Level is referred to that channel's own floor with the source off.

Source	Level Above Floor (dB)	$s_0 - s_1$ (dB)
Diffuse ambient, source off	—	+3.05
Node at +8 dBm	34.2	−4.58
Node at +4 dBm	31.1	−7.77
Node at 0 dBm	30.9	−7.72
Node at −4 dBm	23.5	−12.05
Node at −8 dBm	18.7	−16.60
Node at −12 dBm	14.8	−16.61
Node at −16 dBm	11.1	−16.19
Node at −20 dBm	7.3	−15.66
Node at −40 dBm	0.4	−3.24

The bottom row is the source at its lowest setting, which puts almost nothing above the floor; the reading there is ambient again rather than a weak signal.

The interpretation of the directed rows of Table 8 must be qualified. The source node was on the same bench as the SDR, and a check made during the revision (Section 3.10) showed that a node a few centimetres from the SDR is received with less than 1.5 dB change when either dipole is removed. The directed rows may therefore be dominated by pickup that bypasses the antennas, and they are no longer offered as evidence about antenna response. They do show that two simultaneously sampled chains can differ by more than 15 dB for a strong nearby source while agreeing within 0.23 dB on the ambient floor, so that sampling-time differences are not necessary to produce the effect. The measurement that separates position from instrument under a directed source is the two-antenna experiment of Section 4.4.

The ambient row also shows the sensitivity of such measurements to antenna alignment. With the two dipoles at an angle to one another, the ambient offset was +7.26 dB (standard deviation 1.14 dB; per-channel range +5.4 to +7.5 dB). With the dipoles parallel, it was +3.05 dB (standard deviation 0.23 dB; per-channel range +2.0 to +3.2 dB). Thus, 4.2 dB of the offset and four fifths of its scatter were due to orientation. The remaining offset is flat across all 26 channels in the passband, which indicates a constant gain difference between the two receive chains. Antenna alignment at this level is also a potential source of error in the co-location control of Table 4.

4.4. Two Identical Receivers 60 cm Apart Under a Moving Source

The SDR recorded 5586 frames from the transmitter and 1995 acknowledgements from the fixed node over 85 s (Section 3.10). The transmitter left its initial position after 8.7 s, slightly before the cue, and remained in motion until the end of the record; these observed times, rather than the planned schedule, are used below. While the transmitter was stationary, the two chains differed by +3.67 dB with a standard deviation of 1.05 dB, close to the constant gain difference of +3.05 dB between the chains (Section 4.3). Once it moved, the level in each chain varied over a range of 26–30 dB, the two levels became essentially uncorrelated ($r = 0.027$), and their difference had a standard deviation of 7.71 dB, departing from the stationary value by at least 6 dB in 48.8% of frames. Because the two chains share a clock and are sampled simultaneously, these differences arise from the positions and directional responses of the two antenna–feeder assemblies rather than from the receivers; the absence of correlation indicates that the dominant pickup in each chain

is local to its own position. This is consistent with small-scale fading, which decorrelates over approximately half a wavelength [17], a tenth of the baseline. The moving source provides an ensemble over source positions, whereas the link measurement had fixed nodes; the result therefore does not apportion the 19.2 dB in-episode dispersion of the link between position and receiver response. It shows, however, that a separation of 60 cm produces differences of the same order as the 9 dB co-located difference attributed to receiver response in Section 4.3.

For direction finding, the stationary phase noise of 0.120 rad at a signal-to-noise ratio of 22 dB would correspond to a bearing precision of 0.22° for an ideal 60 cm baseline, against 1.31° for a 10 cm baseline with the same phase noise. The effective baseline was not calibrated, however, because the feeders also receive part of the signal, and precision is not accuracy, because the multipath bias is unknown. The fixed node's acknowledgements kept a constant phase (standard deviation 0.059 rad), so the instrument was stable. Although the 3.5 m range is inside the Fraunhofer distance of 5.98 m, the plane-wave error over the arc is only 0.019 rad, so Equation (14) is adequate at this range. The trajectory of the moving transmitter could not be recovered. Even six traverses of the arc during the 75 s of motion would correspond to a mean geometric phase rate below 1 rad/s, whereas the measured rate had a median of 2.7 rad/s and a 90th percentile of 11.9 rad/s, which indicates that multipath dominated the phase. Deep fades also interrupted tracking: only 47.6% of the motion could be followed without a gap, the longest gap-free segment lasted 3.2 s, and with an unambiguous field of only ±5.8°, the phase cannot be unwrapped across such gaps. In this indoor environment, therefore, the wider baseline promises better precision for a stationary source but does not make the bearing of a moving source observable; additional antennas to resolve the ambiguity and averaging over frequency to suppress multipath would be required [18].

4.5. Negotiation Performance

With the method in place, channel changes were forced on an otherwise quiet bench and measured end to end (Table 9). Two timings are reported because they answer different questions: the node-internal interval from the transmitter announcing its intent to both ends having moved, which is the protocol's own cost, and the host-observed interval from the command being issued to both nodes reporting the change, which includes serial latency and is what an application would perceive.

These are best-case figures: on a quiet bench, the candidate list was uncontested, and the control channels were clear. The negotiations that took place under interference during the run of Section 4.2 are reported at the end of this section.

Table 9. Negotiated channel changes on a quiet bench, twelve consecutive changes with the raw log archived. The earlier column is retained because the firmware changed between the two runs: the exchange now sweeps six times rather than three and the receiver goes to the control channels after two seconds of silence rather than six, both of which were needed before a jammed link could negotiate at all.

Quantity	This Firmware	Before the Repair
Changes attempted	12	5
Changes completed	12	5
Sessions terminated for want of agreement	0	0
Node-internal, announce to moved, mean ± sd (s)	0.950 ± 0.004	0.57 ± 0.22
Node-internal, minimum/maximum (s)	0.95/0.96	0.40/0.81
Host-observed, command to both moved (s)	1.04 ± 0.01	0.78
Agreed on the first control channel tried	0 of 12	3 of 5
Candidates offered per negotiation	6	6
Full 97-channel energy sweep (s)	1.0	1.0

The current firmware is slower and more consistent than the earlier one: 0.95 ± 0.004 s against 0.57 ± 0.22 s. In all twelve negotiations, the exchange was not answered on the first control channel (2401 MHz) and was agreed on the second (2450 MHz), so each negotiation incurred one additional dwell of $T_{\text{tx}} = 400$ ms. This is a consequence of the unrandomised rendezvous phase (Section 5.6): the receiver's rotation and the transmitter's dwell start from the same index, so the control channel on which agreement occurs is determined by the two timers. Randomising the receiver's starting index would reduce the mean by avoiding the systematic extra dwell, at the cost of consistency.

On the quiet bench, all 17 attempted changes completed, and no session was terminated for want of agreement. The spread of the earlier timing (0.40 to 0.81 s, bimodal) reflects the phase between the transmitter's 400 ms dwell and the receiver's 120 ms rotation at the start of the negotiation, not measurement error.

Negotiations under interference. During the run of Section 4.2, the current firmware initiated 44 negotiations, 42 for INTERFERENCE and 2 for DESYNC. Of these, 43 ended in a channel change. For the 42 with complete records, the time from initiation to the change had a median of 0.70 s, a mean of 0.90 ± 0.84 s and a range of 0.16–5.02 s; agreement was reached on the first control channel tried in 13 cases, on the second in 15, on the third in 6 and in a later sweep in 8. The remaining negotiation received no reply in six sweeps (18 control-channel dwells) and terminated the session, after which the test harness restarted the link. The termination rule therefore operated as designed, but the negotiation is not guaranteed to complete under interference: the dwell relation (13) guarantees overlap in time, not reception.

4.6. Comparison with Simpler Selection Strategies

In the window-level analysis (Section 3.11), transmitter-only selection would have kept a channel that the receiver classified as interfered in 4 of the 618 windows in which the transmitter classified it as clean (0.65%), and receiver-only selection would have kept a channel that the transmitter classified as interfered in 34 of 648 windows (5.25%). Intersection-based selection avoids both errors by construction, at the cost of leaving the channel in 42 windows rather than 38 or 8 and requiring both channel maps to be exchanged. The two errors do not have equal consequences: for data flowing from the transmitter, interference at the receiver corrupts the data, whereas interference at the transmitter corrupts only the acknowledgements. The proposed method combines these rules: the decision to leave a channel is taken by the transmitter (Section 3.8), so at the window level, it behaves as transmitter-only selection, whereas the choice of the new channel is an intersection restricted to the shortlist. In the negotiation analysis, the receiver accepted the transmitter's first candidate in all 43 completed negotiations, with N_B between -104 and -96 dBm at the selected channel, so the proposed method selected exactly the channel that transmitter-only selection would have used. The candidates, chosen away from the interferer, were clean at both ends. The receiver's choice can change the outcome only when the receiver observes interference on a channel that the transmitter considers clean, which occurred in 0.65% of windows. An advantage of the method over transmitter-only selection is therefore not demonstrated by these data; demonstrating it requires an outcome-level comparison with interference local to the receiver (Section 5.3).

4.7. Per-Stage Processing Cost

Table 10 and Figure 6 report the time taken by each stage of encoding and decoding a frame, measured on the target over 500 iterations per mode with the radio stopped. Because each mode fixes its own segment size, the six modes are also six packet lengths spanning 224 bytes at L0 down to 24 at L5.

Table 10. Per-stage processing time, mean ± sample standard deviation over 500 iterations per mode.

Mode	PHY	L_{seg} (B)	Seal (µs)	FEC Enc. (µs)	FEC Dec. (µs)	Open (µs)	Enc. + Dec. (µs)
L0	GFSK 2M	224	310.7 ± 6.6	128.5 ± 8.8	140.2 ± 0.4	314.7 ± 8.0	897
L1	GFSK 1M	128	301.7 ± 4.7	1831.5 ± 21.4	1442.4 ± 15.8	307.2 ± 9.8	3887
L2	GFSK 1M	64	297.0 ± 9.6	5090.8 ± 31.5	3318.9 ± 23.1	300.3 ± 0.5	9011
L3	Coded $S = 2$	64	297.7 ± 9.6	5157.5 ± 30.5	3380.3 ± 23.6	303.1 ± 10.0	9143
L4	Coded $S = 8$	32	295.3 ± 7.5	2863.8 ± 24.8	1946.4 ± 19.6	300.9 ± 8.8	5410
L5	Coded $S = 8$	24	264.0 ± 5.5	1717.5 ± 19.9	1209.4 ± 15.4	266.5 ± 9.4	3461

The last column is measured independently, around the whole chain and is not the sum of the four preceding columns; it additionally contains the header pack and the buffer handling, which is why it exceeds that sum by 3–5 µs.

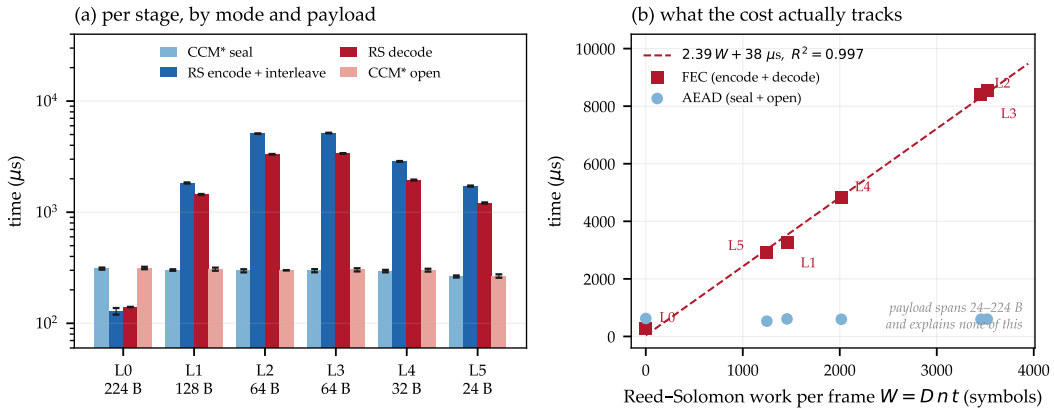


Figure 6. Processing cost per stage. (**a**) Absolute times by mode (logarithmic axis). (**b**) The same data against payload size: the authenticated encryption is flat, while the error-correction cost is ordered by t and not by length.

Two results are of interest beyond the protocol.

The authenticated-encryption cost is constant in payload length. Seal and open take 264–315 µs across a tenfold range of payload, because AES-128-CCM* on the hardware accelerator is dominated by the cost of the call into the cryptographic driver rather than by the number of blocks processed. The 24-byte and 224-byte cases differ by less than 18%.

The forward-error-correction cost tracks correction work, not payload length. Mode L0, with $t = 0$, spends 128 µs encoding the largest packet in the ladder; modes L2 and L3, with $t = 16$ over four codewords, spend 5.1 ms on a quarter of the payload. The ordering across the six modes follows t and the number of codewords rather than L_{seg}, and a single-parameter fit against the product $W = D \cdot n \cdot t$ (interleaver depth times codeword length times correction power, i.e., the number of syndrome-symbol operations the decoder performs) accounts for the six points to $R^2 = 0.997$ (Figure 6b).

This statement has limits, because the six modes are six points rather than a factorial design: the ladder deliberately pairs short payloads with strong coding, so payload length and correction power are confounded. Only L2 and L3 differ mainly in one factor (the physical layer; their tag lengths are 8 and 12 bytes), and their processing times differ by 1.5%. The data therefore support the narrower claim that, *on this implementation and hardware, across these six configurations, processing time is predicted well by correction work and poorly by payload length*. The 40-fold difference between L0 and L2 should not be read as an effect of t alone, since the two modes differ in payload, codeword count, interleaver depth and physical layer; a controlled benchmark varying L_{seg} and t independently would be required to separate them. Even the narrower claim implies that a scheduler on this platform that keys its cost model to packet length would be inaccurate.

Expressing processing as a fraction of the frame's own air time,

$$\eta = \frac{T_{\text{enc}} + T_{\text{dec}}}{T_{\text{air}}}, \tag{16}$$

gives $\eta = 0.87$ at L0, 2.53 at L1, 5.01 at L2, 2.30 at L3, 0.32 at L4 and 0.22 at L5. At L2 the pair spends five times as long on the codec as on the air. For this reason, the slot is derived from the mode rather than from the air time; the modes that are slow on the air leave the most processing headroom.

4.8. Goodput and the Shannon Bound

Table 11 reports application goodput per mode at 4.2 m against two reference rates. The first is the one-segment-per-slot ceiling, which is what this protocol could deliver if nothing were ever lost or retransmitted; it is the relevant efficiency metric. The second is the Shannon capacity [38] of the same channel at the signal-to-noise ratio the receiver itself reported, averaged over every adaptation window in that mode (Section 3.9),

$$C = B \log_2\left(1 + \frac{S}{N}\right), \tag{17}$$

with B the occupied bandwidth from Carson's rule applied to the radio's deviations: 1.34 MHz for 1 Mbit/s GFSK and both Coded modes ($\pm$170 kHz deviation at 500 kHz), 2.64 MHz for 2 Mbit/s ($\pm$320 kHz at 1 MHz). For L0, the measured mean signal-to-noise ratio is 38.77 dB, a power ratio of 7534, $\log_2(1 + 7534) = 12.879$, and $C = 2.64 \times 10^6 \times 12.879 = 34.0$ Mbit/s.

Two cautions apply. First, Carson's rule is a closed-form estimate that ignores the Gaussian pulse shaping, the transmit filter and the actual spectral mask, so B is nominal; the SDR could measure the occupied bandwidth directly, but this was not done. Second, C is an upper bound over all modulation and coding schemes. Because the radio fixes the physical layer at 1 or 2 Mbit/s GFSK and the slot schedule is deliberately sparse, the ratio of goodput to C mainly measures the distance between a 2 Mbit/s radio and the capacity of its channel, not the performance of the protocol. The efficiency claim of this paper is therefore made against the one-segment-per-slot ceiling.

Table 11. Goodput at 4.2 m, 0 dBm transmit power, five repetitions of 12 s per mode. L5 is listed for completeness: it is the resynchronisation mode and carries no application payload, so it has no goodput to report.

Mode	PHY	L_{seg} (B)	Goodput (bit/s)	sd (bit/s)	cv (%)	SNR (dB) (n)	C (Mbit/s)	Goodput/C (%)
L0	GFSK 2M	224	172,950	6853	4.0	38.8 ± 3.5 (430)	34.00	0.51
L1	GFSK 1M	128	101,295	706	0.7	39.2 ± 8.9 (451)	17.44	0.58
L2	GFSK 1M	64	50,437	695	1.4	38.1 ± 4.2 (458)	16.94	0.30
L3	Coded $S = 2$	64	38,932	100	0.3	33.5 ± 6.9 (350)	14.92	0.26
L4	Coded $S = 8$	32	7449	558	7.5	36.8 ± 2.3 (136)	16.40	0.05
L5	Coded $S = 8$	24	*carries no application payload; see below*					

L5 appears in Table 11 and Figure 7 without a value because it carries no application data by design: it re-establishes synchronisation after the link has failed, and its body is an authenticated beacon rather than an encrypted segment.

The measured goodput of the five data modes is 0.05–0.58% of capacity. The gap decomposes into three factors, only the last of which is a property of the protocol rather than of the platform. The modulation is fixed at 1 or 2 Mbit/s where capacity would permit tens of megabits: at L0, that is $34.0/2.0 = 17.0$. The slot is fixed and mostly idle: a 1028 µs

frame in a 10 ms slot is a factor of 9.73. Framing overhead consumes 9% of the air frame at L0 (efficiency $224/246 = 0.911$) and 70% at L2. Their combination predicts

$$\hat{R} = \frac{C}{17.0 \times 9.73} \times 0.911 = \frac{34.0 \times 10^6}{165.4} \times 0.911 = 187 \text{ kbit/s}, \tag{18}$$

against 173 kbit/s measured, or 92% of the prediction. The decomposition shows that, at L0, the gap of more than two orders of magnitude is accounted for by the fixed modulation rate, the slot duty cycle and the framing, in that order of size. Against the one-segment-per-slot ceiling of $224 \times 8/10$ ms $= 179.2$ kbit/s, the link achieves 96.5%.

Goodput is linear in payload above 64 bytes, at 772–791 bit/s per payload byte across a fourfold range, because in that range one segment is delivered per slot regardless of its size (Figure 7). Below it, the relationship breaks: L4's 32-byte segment achieves 233 bit/s per byte, because its frame is 16.8 ms of air in a 33.8 ms slot. The comparison between L2 and L3 is the cleanest available measure of what the Coded physical layer costs when it is not needed: identical payload, similar framing, 23% less goodput for 2.2 times the air time (3982 µs against 1800 µs, Table 3).

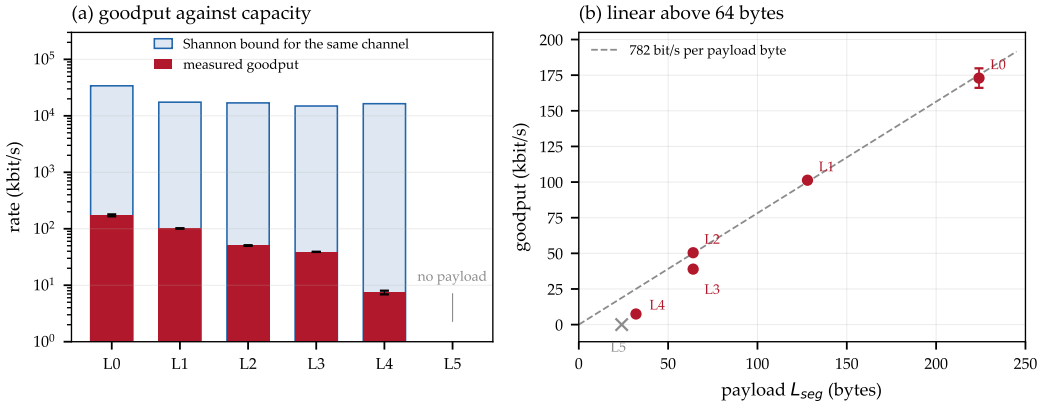


Figure 7. (**a**) Measured goodput against the Shannon bound for the same channel and measured signal-to-noise ratio, logarithmic axis. (**b**) Goodput against payload size, showing the linear regime above 64 bytes.

5. Discussion

5.1. Interpretation of the Asymmetry Measurement

The measurement in Section 4.2 concerns a premise rather than a protocol. Any scheme that lets one end classify channels on behalf of both assumes that a measurement at one position predicts the other. On the 4.2 m link, this assumption failed in two distinct ways. In the 614 quiet windows, the transmitter read a standing +6.30 dB more than the receiver, which the co-location control does not attribute to receiver calibration. In the 42 windows within the fifteen interference episodes, the difference had a standard deviation of 19.2 dB, and the two ends classified the channel differently in 38 windows; no classification disagreement followed a difference smaller than 15 dB (Table 5).

The scope of this result is limited. What was demonstrated is that *under one indoor geometry, with one pair of nRF52840 devices at one separation and one swept narrowband interferer, endpoint noise-floor measurements differed substantially and sometimes produced different channel classifications, and the standing part of that difference is not attributable to receiver calibration*. Whether this is characteristic of short-range links in general would require several separations, exchanged node positions, several antenna orientations, interferer positions and interferer types, and many more independent episodes than the fifteen recorded here (Section 5.6).

Four qualifications apply. First, raw agreement is high (94.2% of windows are classified identically) because 94% of the record is quiet; the disagreements are concentrated in the periods of interference, which is when a channel decision is taken, and they rest on fifteen episodes. Second, the standing difference is attributed to the two ends' positions or operating roles; the co-location control excludes receiver calibration but not the difference in operating state, which an exchange of the two nodes would test. Third, the direction of the asymmetry is not established, because its significance depends on the segmentation of episodes. Fourth, the co-location control bounds the instrument offset in a diffuse field only; two devices 10 cm apart differ by a median of 9 dB under a directed source, so the in-episode differences at 4.2 m probably overstate spatial decorrelation. The two-antenna measurement of Section 4.4 shows that, with identical receivers, a separation of 60 cm produces differences of the same order as that 9 dB (standard deviation 7.7 dB for a moving source). How the in-episode dispersion of the link divides between position and receiver response remains undetermined.

In this configuration, therefore, one end's measurement did not predict the other's, whether because of where it stood or of how it responded to a directed source. Because the direction of the asymmetry is not known in advance, a method that does not depend on it is preferable: if the deciding end is the more exposed one, a unilateral scheme abandons channels unnecessarily; if it is the less exposed one, it selects channels that its peer cannot use. The analysis of Section 4.6 shows, however, that in the recorded scenario the proposed method made the same choices as transmitter-only selection, so its benefit remains to be demonstrated.

5.2. Negotiation Cost

On a quiet bench, a negotiated change cost 0.95 ± 0.004 s of node-internal time with the current firmware (0.57 ± 0.22 s with the earlier one), most of which is the bounded control-channel sweep; this is a best-case figure for this firmware. Under interference, 43 of 44 negotiations completed, with a median of 0.70 s and a maximum of 5.02 s, and one terminated the session. Because the receiver returns its own noise reading with the agreement, the transmitter also obtains information about the far end at every negotiation. Because the exchange either concludes or terminates, there is no state in which one end transmits on a channel that the other has not accepted.

5.3. Untested Aspects of the Method

An outcome-level comparison. The analysis of Section 4.6 compares the decisions of the strategies on recorded data, but it does not measure their consequences for packet error rate, goodput or outage duration, and it cannot evaluate random selection or full-band receiver-only and intersection selection, because the receiver's channel map was not logged. The testbed supports a live comparison, since the channel decision is reachable from the host on either node. Table 12 specifies it. Each strategy would be run on the same link under the same interference schedule, alternating in short blocks, with interference placed near the transmitter and near the receiver in turn. The prediction implied by Section 4.2 is that transmitter-only selection should show longer outages than negotiated selection when interference is local to the receiver and that the strategies should be indistinguishable when the band is quiet or when the candidates are clean at both ends, as in the recorded run. Because only the transmitter initiates a negotiation (Section 3.8), the same experiment would also show how quickly receiver-local interference leads to a channel change.

A stress test of the negotiation. The 17 changes on a quiet bench and the 44 negotiations under one narrowband interferer show that the mechanism is correctly implemented and usually completes, but they do not establish its reliability. The cases that violate the

assumptions of the dwell relation were not tested systematically: occupied control channels, frames lost during the exchange, clock drift between the two ends, a node restarting during a negotiation, and replayed control frames (Section 3.8). A proper evaluation would run hundreds of negotiations with randomised starting phase and candidate lists across a range of interference states, and report the failure rate and the distribution of convergence times.

Table 12. Specification of a live comparison of selection strategies and the extent to which each arm was evaluated in this work.

Strategy	Information Used	Channel-Selection Rule	Evaluated Here
Random	none	uniform over non-blacklisted channels	no
Transmitter-only	N_A	the transmitter's choice, announced and imposed (BLE-like)	decisions only
Receiver-only	N_B	the receiver's choice, requested and granted unconditionally	decisions only
Intersection	N_A, N_B	best channel with $\max(N_A, N_B) \leq N^*$; requires both maps	decisions only
Proposed	N_A list, N_B choice	receiver selects among the transmitter's candidates (Equation (12))	negotiations measured

Metrics for each arm: packet error rate, goodput, outage duration, recovery time after a channel is lost, channel changes per minute, failed negotiations, and energy per delivered byte.

5.4. Implementation Experience

Several implementation defects encountered during this work had one property in common: none of them produced an error, and each was detected only as a disagreement between the three independent accounts of the same interval (the transmitter's, the receiver's and the SDR's). They were as follows. (1) A noise filter with $\alpha = 0.97$ (a time constant of about 33 updates, i.e., tens of seconds or more) kept reporting a channel as it had been long before, so the link failed while the channel still appeared clean; the discriminator and the blacklist now use the raw measurement. (2) An override that converted total acknowledgement loss into FADING was guarded against INTERFERENCE but not DESYNC and rewrote the cause in 1184 of 1184 such windows, so that no channel change was ever triggered. (3) Widening the plan from 4 to 97 channels exposed costs that were linear in the channel count: a full energy sweep took 6.0 s until the waits for radio events were bounded by the clock, after which it took 1.0 s. (4) Because EasyDMA is single-buffered and the radio turns around autonomously, every acknowledgement reflects the receive window one frame earlier; this caused exactly two transmissions per delivered segment on a clean link, and granting the most recent segment one slot of grace doubled the goodput of every mode. (5) Two recovery mechanisms depended on the agreement they were meant to restore: the resynchronisation beacon was keyed on the epoch, and the epoch exit rule led the two ends to different epochs. (6) The cryptographic accelerator's DMA cannot read flash memory, so an operation on a constant array failed while the same operation on a stack buffer succeeded.

5.5. Generalisation Beyond a Pair

The method as implemented is pairwise, and the obstacles to generalisation are specific. The AEAD nonce contains a per-role constant identifier, so three nodes sharing a key would reuse nonces; a real node identity must be introduced first. The 14-byte header has no address field, and the hardware auto-turnaround admits exactly one responder, so multiple responders require scheduled reply slots. Per-peer state is small (an ARQ window, a few

filters and counters, about 150–200 bytes), and the 97-channel spectrum arrays are shared across peers.

For a star topology, the existing 64-transmission superframe can be divided into per-node slots. Processing, not spectrum, is then the binding constraint (Section 4.7): encoding plus decoding takes 9.0 ms at L2 against 0.9 ms at L0, so a coordinator can serve about ten times as many nodes at L0 as at L2. For a mesh, the header is readable but not verifiable by a relay, which suggests separate network and application keys, and channel agility conflicts with relaying: if each link selects its own channel, a single-radio relay must time-multiplex, and if the network shares a channel, the receiver's choice must become a quorum. With N receivers, each rejecting a channel proposed as clean with probability p, the rejections are not expected to be independent, because the interference that causes them is spatially structured and nearby receivers should reject together. Assuming that the acceptances are positively associated, the probability of unanimous acceptance satisfies

$$P_{\text{unanimous}} \geq (1-p)^N, \tag{19}$$

so the independence expression is a lower bound on the probability of unanimity, i.e., an upper bound on the difficulty, and not an estimate. The assumption is plausible but not guaranteed: a single localised interferer among dispersed receivers need not produce positive association. For illustration only, with $p = q_{\text{rej}} = 0.0065$ (Section 3.8), the bound is 0.85 for 25 receivers and 0.72 for 50. Since q_{rej} comes from one link, these figures indicate only that a quorum rule with explicit handling of dissenting nodes may be required. The three control channels, being broadcast-shaped and 98 MHz apart, generalise without change. The design has no clear channel assessment, which is adequate for a scheduled pair but not for a mesh with hidden terminals, and its maximum transmission unit varies with the mode (224 to 32 bytes, and zero during resynchronisation), which a network layer would have to accommodate. Conversely, the per-link quantities γ, p_s and the discriminator's cause codes could serve directly as a routing metric.

5.6. Limitations

The experiment is narrow. The asymmetry measurement comes from a single indoor link at one separation, with two devices and one interferer, and 26 of the 97 data channels contributed paired windows. A campaign supporting a general claim would include several separations spanning at least an order of magnitude, the two nodes physically exchanged with the geometry held fixed, several antenna orientations, interferer positions and powers, more than one interferer waveform (e.g., Wi-Fi beacons and microwave ovens), and enough independent episodes for episode-level statistics to have power.

The statistical basis is limited. The 656 paired windows are consecutive samples from one 546 s run, with an effective sample size near 188; the intervals account for this and are robust to the block length (Table 6). The classification disagreements arise from fifteen episodes, so the number of independent events behind that statistic is fifteen, not 656. For this reason, the disagreement rate is reported as a description of the record, the direction of the asymmetry is reported as not established, and the block bootstrap is not taken to compensate for the shortage, since it preserves dependence within one realisation but cannot create further realisations.

The calibration covers a diffuse field only. The co-location control justifies attributing the standing offset to position, but under a directed source, the same two boards differ by a median of 9 dB, so an unknown part of the in-episode dispersion is receiver and antenna response, and the standing difference may reflect the operating roles of the two nodes as well as their positions. Three further controls would separate the two, in increasing order of effort: repeating the link and co-location runs with the two boards exchanged in

place, which separates device and role from position; rotating each board through several orientations under a fixed source; and feeding both receivers from one source through a splitter. The SDR measurements do not replace these controls. The directed rows of Table 8 may be dominated by pickup that bypasses the antennas, and the two-antenna measurement of Section 4.4 shows the size of position effects for identical receivers but not the receiver response of the nRF52840 boards.

The two-antenna measurement is a single run. It was made at one range, with manual motion along one arc, roles exchanged relative to the main experiment, and a substantial part of the signal received by the feeders. Its conclusions (spatial decorrelation at 60 cm, a predicted stationary precision of 0.22° for an uncalibrated baseline, and the impossibility of tracking a moving source with two antennas in this environment) are therefore indicative.

The evaluation lacks an outcome-level baseline comparison and a stress test (Section 5.3); the negotiation cost of 0.95 s is a best-case figure obtained on a quiet bench.

The security model is a bench model. Pairing uses an unauthenticated key agreement, so an attacker present at pairing can become the peer; control frames are accepted within an epoch window of ± 4 without a freshness check, so a captured frame can be replayed within that window; and the AEAD nonce contains a per-role constant rather than a device identity, which is sound for two nodes but not for three. Remedies are known but were not implemented or evaluated (Sections 3.8 and 5.5). The work therefore offers integrity and confidentiality of the channel-selection exchange *within a trusted pairing model*, and the word "sealed" is used in that restricted sense.

Two implementation limitations are known. The rendezvous phase is not randomised, which causes the systematic extra dwell observed in Table 9. During the first twelve or so superframes of a session, the rolling scan has not covered the whole band, so early decisions are made on partial information, which the telemetry does not flag.

6. Conclusions

This paper examined, for one instrumented configuration, whether the two ends of a short-range wireless link perceive interference similarly. On a 4.2 m indoor link under one swept narrowband interferer, 656 time-paired adaptation windows were obtained. In the 614 quiet windows, the transmitter read a standing $+6.30$ dB more than the receiver, against a mean offset of 0.66 dB for the same devices side by side. In the 42 windows within fifteen interference episodes, the difference had a standard deviation of 19.2 dB, and 38 windows, distributed across all fifteen episodes, were classified differently against the interference threshold; every such disagreement followed a difference of at least 15 dB.

These results carry three qualifications. They describe one run, in which the classification result rests on fifteen episodes rather than on 656 trials. The direction of the asymmetry is not established, because its significance depends on how episodes are segmented. The co-location control attributes the standing difference to the two ends' positions or operating roles rather than to receiver calibration, but it covers a diffuse field only; under a directed source, co-located devices differ by a median of 9 dB, so the larger differences during interference probably overstate spatial decorrelation. An additional measurement with two identical, simultaneously sampled receivers 60 cm apart showed that such a separation decorrelated the levels received from a moving source (standard deviation of the difference 7.7 dB), so position effects of this size are physically plausible, although their share in the link measurement remains undetermined. The same measurement indicated a stationary bearing precision of about 0.22° for an ideal 60 cm baseline, but indoor multipath prevented the trajectory of a moving source from being recovered with two antennas.

The proposed method treats each end's measurement as authoritative only for that end: the transmitter offers a ranked, sealed shortlist on dedicated control channels, the

receiver selects the candidate that is acceptable at its position and returns its measurement, and the session terminates rather than continuing on a channel that was not accepted. Its guarantee lies between the advisory reporting of Bluetooth Core 5.3 and the radar veto of IEEE 802.11h [8]. Implemented on a purpose-built protocol over 97 one-megahertz channels, the mechanism completed 17 of 17 channel changes on a quiet bench, in 0.95 ± 0.004 s with the current firmware, and 43 of 44 under interference, with a median of 0.70 s. In all 43, however, the receiver selected the transmitter's first candidate, because the candidates, chosen away from the interferer, were clean at both ends. No advantage over simpler strategies is therefore claimed; the contribution is a quantification of the premise together with an instrumented realisation of one way to act on it.

A subsidiary result is that, across the six configurations measured, authenticated-encryption cost is constant in payload length and error-correction cost is predicted by correction work rather than by payload length on this platform; a controlled sweep would be needed to generalise this statement, because the six modes confound the two factors.

Future research. Four experiments would strengthen the present results, in order of priority. The first is an outcome-level comparison of the strategies in Table 12, with interference local to each end in turn, since it is the only experiment that can show whether the method improves on transmitter-only selection. The second is an extension of the co-location control, with the boards exchanged in place, rotated under a fixed source and fed through a splitter. The third is a replication campaign across separations, orientations, interferer positions and waveforms, with enough independent episodes for episode-level statistics. The fourth is several hundred negotiations under interference, with randomised starting phase, occupied control channels and induced frame loss. The security additions of Section 3.8 should be evaluated in the same round. The extension to more than two nodes requires a real node identity in the nonce, an address field, per-node keys and scheduled reply slots and, for a shared channel, a quorum rule informed by the disagreement statistic reported here.

Author Contributions: Conceptualisation, methodology, software, validation, formal analysis, investigation, data curation, V.S.; writing—original draft preparation, writing—review and editing, and visualisation, P.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The firmware, the host software, the measurement scripts and the raw telemetry logs from which every figure and table in this article was produced are openly available at https://github.com/volodymyr-sokolov/Asymmetric-Interference-Perception (accessed on 7 October 2026). The repository README gives the bill of materials, the toolchain versions, the flashing procedure and the exact command line behind each table and figure.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ADR	Adaptive Data Rate
AEAD	Authenticated Encryption with Associated Data
AES	Advanced Encryption Standard
AFH	Adaptive Frequency Hopping
ARQ	Automatic Repeat Request
BLE	Bluetooth Low Energy

CCA	Clear Channel Assessment
CCM	Counter with CBC-MAC
CI	Confidence Interval
CMAC	Cipher-based Message Authentication Code
DFS	Dynamic Frequency Selection
FEC	Forward Error Correction
GFSK	Gaussian Frequency-Shift Keying
HKDF	HMAC-based Key Derivation Function
IAS	Interference Awareness Scheme
ISM	Industrial, Scientific and Medical
PABAK	Prevalence-Adjusted Bias-Adjusted Kappa
PER	Packet Error Rate
PHY	Physical Layer
RSSI	Received Signal Strength Indicator
SDR	Software-Defined Radio
SNR	Signal-to-Noise Ratio
TSCH	Time-Slotted Channel Hopping

References

1. Yao, J.; Lou, W.; Xie, R.; Jiao, X.; Wu, K. Mitigating Cross-Technology Interference Through Fast Signal Identification. *IEEE Trans. Veh. Technol.* **2023**, *72*, 2521–2534. [CrossRef]
2. Liang, C.-J.M.; Priyantha, N.B.; Liu, J.; Terzis, A. Surviving Wi-Fi Interference in Low Power ZigBee Networks. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys), Zurich, Switzerland, 3–5 November 2010; pp. 309–322. [CrossRef]
3. Park, W.; Eo, S.; Joo, C.; Bahk, S. B-hop: Time-Domain Adjustment of BLE Frequency Hopping Against Wi-Fi Beacon Interference. *IEEE Internet Things J.* **2024**, *11*, 12853–12863. [CrossRef]
4. Bluetooth SIG. *Bluetooth Core Specification*, Version 5.4; Bluetooth Special Interest Group: Kirkland, WA, USA, 2023. Available online: https://www.bluetooth.com/specifications/specs/core-specification-5-4/ (accessed on 8 September 2026).
5. Pang, B.; T'Jonck, K.; Claeys, T.; Pissoort, D.; Hallez, H.; Boydens, J. Bluetooth Low Energy Interference Awareness Scheme and Improved Channel Selection Algorithm for Connection Robustness. *Sensors* **2021**, *21*, 2257. [CrossRef] [PubMed]
6. Lee, J.; Park, C.; Roh, H. Revisiting Adaptive Frequency Hopping Map Prediction in Bluetooth with Machine Learning Classifiers. *Energies* **2021**, *14*, 928. [CrossRef]
7. Woolley, M. *Bluetooth Core Specification Version 5.3 Feature Enhancements*, v1.0.2; Bluetooth Special Interest Group: Kirkland, WA, USA, 2021. Available online: https://www.bluetooth.com/bluetooth-resources/bluetooth-core-specification-version-5-3-feature-enhancements/ (accessed on 8 September 2026).
8. *IEEE Std 802.11h-2003*; Spectrum and Transmit Power Management Extensions in the 5 GHz Band in Europe. IEEE: New York, NY, USA, 2003. [CrossRef]
9. Cena, G.; Scanzio, S.; Valenzano, A. Ultra-Low Power Wireless Sensor Networks Based on Time Slotted Channel Hopping with Probabilistic Blacklisting. *Electronics* **2022**, *11*, 304. [CrossRef]
10. Lim, S.; Seo, S. Towards Interference-Aware ZigBee Transmissions in Heterogeneous Wireless Networks. *Comput. Mater. Contin.* **2020**, *66*, 751–765. [CrossRef]
11. Jozwik, C.; Zhu, R. Passive On-Device BLE Interference Detection Using TinyML. In *Proceedings of the 2026 IEEE 2nd International Conference on Secure IoT, Assured and Trusted Computing (SATC)*; IEEE: Piscataway, NJ, USA, 2026; pp. 1–5. [CrossRef]
12. Croce, D.; Garlisi, D.; Giuliano, F.; Inzerillo, N.; Tinnirello, I. Learning From Errors: Detecting Cross-Technology Interference in WiFi Networks. *IEEE Trans. Cogn. Commun. Netw.* **2018**, *4*, 347–356. [CrossRef]
13. Jin, M.; He, Y.; Zheng, X.; Fang, D.; Xu, D.; Xing, T.; Chen, X. Exploiting Interference Fingerprints for Predictable Wireless Concurrency. *IEEE Trans. Mob. Comput.* **2021**, *20*, 2354–2366. [CrossRef]
14. Zheng, X.; Cao, Z.; Wang, J.; He, Y.; Liu, Y. Interference Resilient Duty Cycling for Sensor Networks Under Co-Existing Environments. *IEEE Trans. Commun.* **2017**, *65*, 2971–2984. [CrossRef]
15. Pirrone, M.; Bernhardt, J.; Wunderlich, A. Assessing Directional Time-Dependent Interference Vulnerabilities in Closed-Box Wireless Systems. *IEEE Trans. Electromagn. Compat.* **2025**, *67*, 644–657. [CrossRef] [PubMed]
16. Kravets, I.; Kotliar, N.; Karpin, O.; Luchechko, A. BLE Phase-Based Ranging: Accuracy and Capability Under Strong Wi-Fi Interference. *Int. J. Electron. Telecommun.* **2025**, *71*, 1–7. [CrossRef]
17. Rappaport, T.S. *Wireless Communications: Principles and Practice*, 2nd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.

18. Krim, H.; Viberg, M. Two Decades of Array Signal Processing Research: The Parametric Approach. *IEEE Signal Process. Mag.* **1996**, *13*, 67–94. [CrossRef]
19. Balanis, C.A. *Antenna Theory: Analysis and Design*, 4th ed.; Wiley: Hoboken, NJ, USA, 2016.
20. Urke, A.R.; Kure, Ø.; Øvsthus, K. A Survey of 802.15.4 TSCH Schedulers for a Standardized Industrial Internet of Things. *Sensors* **2022**, *22*, 15. [CrossRef] [PubMed]
21. Kufakunesu, R.; Hancke, G.P.; Abu-Mahfouz, A.M. A Survey on Adaptive Data Rate Optimization in LoRaWAN: Recent Solutions and Major Challenges. *Sensors* **2020**, *20*, 5044. [CrossRef] [PubMed]
22. Akyildiz, I.F.; Lo, B.F.; Balakrishnan, R. Cooperative Spectrum Sensing in Cognitive Radio Networks: A Survey. *Phys. Commun.* **2011**, *4*, 40–62. [CrossRef]
23. Trigka, M.; Dritsas, E. An Efficient Distributed Approach for Cooperative Spectrum Sensing in Varying Interests Cognitive Radio Networks. *Sensors* **2022**, *22*, 6692. [CrossRef] [PubMed]
24. Eltholth, A.A. Improved Spectrum Coexistence in 2.4 GHz ISM Band Using Optimized Chaotic Frequency Hopping for Wi-Fi and Bluetooth Signals. *Sensors* **2023**, *23*, 5183. [CrossRef] [PubMed]
25. Yang, G.; Shin, H.; Kim, H. Diffusion-Based Frequency Hopping for Collision Mitigation in Dense Bluetooth Networks. *Sensors* **2025**, *25*, 5893. [CrossRef] [PubMed]
26. Kim, S. Enabling WLAN and WPAN Coexistence via Cross-Technology Communication. *Sensors* **2022**, *22*, 707. [CrossRef] [PubMed]
27. Del-Valle-Soto, C.; Mex-Perera, C.; Nolazco-Flores, J.A.; Rodríguez, A.; Rosas-Caro, J.C.; Martínez-Herrera, A.F. A Low-Cost Jamming Detection Approach Using Performance Metrics in Cluster-Based Wireless Sensor Networks. *Sensors* **2021**, *21*, 1179. [CrossRef] [PubMed]
28. Cortés-Leal, A.; Del-Valle-Soto, C.; Cardenas, C.; Valdivia, L.J.; Del Puerto-Flores, J.A. Performance Metric Analysis for a Jamming Detection Mechanism under Collaborative and Cooperative Schemes in Industrial Wireless Sensor Networks. *Sensors* **2022**, *22*, 178. [CrossRef] [PubMed]
29. Sokolov, V.; Skladannyi, P.; Astapenya, V. Bluetooth Low-Energy Beacon Resistance to Jamming Attack. In Proceedings of the 2023 IEEE 13th International Conference on Electronics and Information Technologies (ELIT), Lviv, Ukraine, 26–28 September 2023; pp. 270–274. [CrossRef]
30. Sokolov, V.; Skladannyi, P.; Korshun, N. ZigBee Network Resistance to Jamming Attacks. In Proceedings of the 2023 IEEE International Conference on Information and Telecommunication Technologies and Radio Electronics (UkrMiCo), Kyiv, Ukraine, 13–18 November 2023; pp. 161–165. [CrossRef]
31. Sokolov, V.; Skladannyi, P.; Astapenya, V. Wi-Fi Interference Resistance to Jamming Attack. In Proceedings of the 2023 IEEE 5th International Conference on Advanced Information and Communication Technologies (AICT), Lviv, Ukraine, 21–25 November 2023; pp. 1–4. [CrossRef]
32. Sokolov, V.; Skladannyi, P.; Platonenko, A. Jump-Stay Jamming Attack on Wi-Fi Systems. In Proceedings of the 2023 IEEE 18th International Conference on Computer Science and Information Technologies (CSIT), Lviv, Ukraine, 19–21 October 2023; pp. 1–5. [CrossRef]
33. Sokolov, V.; Skladannyi, P.; Mazur, N. Wi-Fi Repeater Influence on Wireless Access. In Proceedings of the 2023 IEEE 5th International Conference on Advanced Information and Communication Technologies (AICT), Lviv, Ukraine, 21–25 November 2023; pp. 33–36. [CrossRef]
34. Nordic Semiconductor. *nRF52840 Product Specification*, v1.11; Nordic Semiconductor ASA: Trondheim, Norway, 2023.
35. Künsch, H.R. The Jackknife and the Bootstrap for General Stationary Observations. *Ann. Stat.* **1989**, *17*, 1217–1241. [CrossRef]
36. Cohen, J. A Coefficient of Agreement for Nominal Scales. *Educ. Psychol. Meas.* **1960**, *20*, 37–46. [CrossRef]
37. Byrt, T.; Bishop, J.; Carlin, J.B. Bias, Prevalence and Kappa. *J. Clin. Epidemiol.* **1993**, *46*, 423–429. [CrossRef] [PubMed]
38. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.