Silence and the Jam

A Russian EW team puts a price on knowing the target · ExpressLRS and MILELRS v4 control-link jamming

Treadstone 71 open-source assessment. Source: a Russian-language technical paper and its Telegram teaser from the electronic-warfare group behind the "Тишина" (Silence) reconnaissance product.

TL;DR

A Russian electronic-warfare team ran a bench test to answer one question. How much easier does jamming get when the jammer already knows the target drone's radio settings? Their numbers give a hard answer. Against a standard ExpressLRS control link, full radio reconnaissance stretches jamming range roughly twenty-seven-fold in open air. Against the hardened MILELRS v4 build, the same homework buys about three. Their strongest attack fails against MILELRS entirely. The paper sells their own scanner as the tool that does the homework. For anyone flying drones into that fight, the takeaway is a number and a warning: hide the waveform, or hand the enemy the range.

BLUF

The team has shown, with measured thresholds, a reconnaissance-cued attack against the open-source ExpressLRS link that flies most cheap drones on each side of the war. We assess with high confidence that the work marks a real step in Russian counter-drone jamming, from wideband barrage toward attacks matched to the target's own waveform. The first purpose of the paper is very likely a sales pitch for the "Тишина" scanner to Russian units and buyers. A second purpose, likely, is to rattle Ukrainian operators who trust ExpressLRS. Source reliability rates moderate. A vendor wrote it about a vendor product, so the numbers deserve the weight of a capability claim, not a lab certification.

What the Russians published

The group tested three setups over the air: ExpressLRS at spreading factor 7 and 8, and MILELRS v4 at spreading factor 7. One number carried the whole study. How far does the jammer's power drop before the link still snaps into loss-of-signal and stays there. A jammer that kills the link with less power is the better jammer.

The real contribution is not any waveform. The team ranked jamming by a single axis nobody else names cleanly: how much the attacker has to know about the target first.

Class	What the attacker knows	The attack it unlocks
0	nothing	wideband noise or a sweep across the whole band, power spread thin
1	the hop bandwidth	the same noise squeezed into the target band, denser
2	bandwidth and spreading factor	a chirp matched to the LoRa waveform, riding the receiver's own gain
3	the full hop table	tracking the live channel hop by hop, all power in one place

Each step up costs more intelligence and buys a harder kill. The measured thresholds, in decibels of jammer attenuation the link tolerates before it dies, tell the story straight.

Attacker knowledge	ELRS SF7	MILELRS SF7	Range gain, open air (ELRS SF7)
Class 0, none	12.5	5	baseline
Class 1, hop bandwidth	16	8.5	×1.5
Class 2, + spreading factor	24.75	15	×4.1
Class 3, + hop table	41	not achievable	×27

Walk standard ExpressLRS from a blind barrage to a fully cued synchronous attack and the threshold climbs 28.5 dB. In open air, where every six decibels doubles range, that gain multiplies the kill distance by roughly twenty-seven at the same transmit power. The winning variable is not watts. Homework decides the range.

The chirp move is the clever heart of the paper. LoRa writes each symbol as a sliding tone and reads it back by matching it against a mirror copy. A jammer tuned to the same spreading factor slips into that match and spikes, while plain noise

stays flat and earns nothing from the receiver. Matched to the waveform, the chirp beats the best noise attack by 7.5 to 8.75 dB on ELRS SF7 and by 6.5 dB on MILELRS, with no added power. Shape alone bought the gain.

The team also put down a myth. A reverse chirp, which some jammer builders swear by, points the wrong way against the receiver's forward reference and loses the match. Their bench shows it far weaker, 12 dB against 23.5 for the matched case. Debunking your own community's folklore is not the move of a group faking rigor.

Who: the actor behind "Тишина"

No unit crest sits on the paper. A Proton address and a Telegram handle, nothing more. The shape fits the Russian military-tech volunteer scene that has carried much of Moscow's drone and counter-drone tinkering since 2023. Small teams. Telegram-native. Fast iteration. Selling or donating into the war. The group calls itself, plainly, a builder of radio-electronic reconnaissance, with "Тишина" as its fielded scanner.

Two details argue for a serious shop rather than a weekend project. The paper opens on Aleksandr Svechin, the Soviet strategist Stalin shot in 1938, and on Norbert Wiener, the father of cybernetics. That framing is a choice, aimed at a professional military-technical reader. And the interface screenshot shows a working tool, not a slide. A live spectrum view. A target list. Auto-named tracks for each control channel it catches, with the hop width, spreading factor, and channel count printed beside them.

We assess with moderate confidence that "Тишина" is a real reconnaissance system in at least limited field use. Confidence on who exactly runs it, and at what scale, stays low. A self-published paper and an unverified handle carry an assessment only so far.

So what: a find-fix chain for the radio spectrum

The paper's weight is not one attack. The weight is the loop it proves and advertises. A sensor reads the drone's control settings. A jammer built around those settings hits harder for the same power. Find, fix, finish, run in the RF spectrum, with "Тишина" as the find-and-fix half of the chain.

Three things follow, and they do not all point the same way.

Russian counter-drone jamming is getting smarter. Barrage jamming burns power and floods friendly spectrum. A cued attack kills the one link for a fraction of the cost. A team that measures its thresholds and corrects its own folklore has left the improvising stage behind.

The waveform is the soft spot. Every rung of the advantage rests on what the target lets the enemy learn. Hide the bandwidth and spreading factor and the attacker falls back to plain noise. Keep the hop table out of reach and the 41 dB synchronous attack never gets built.

Hardening works, and the Russians say so themselves, which is the strongest line in the whole document. MILELRS v4 packs 65,536 channels into 20 MHz at about 305 Hz spacing. So dense the grid reads as one smear, with no guard bands to fingerprint. Its hop order runs off a long-period generator that does not repeat inside any window an interceptor gets. The receiver cannot be pinned to a channel, and the sequence cannot be rebuilt. So Class 3 dies on the drawing board. The Russian ceiling against MILELRS sits at 15 dB against 41 for plain ELRS. Reconnaissance still pays, about a threefold range gain, but the top prize stays locked.

Why now

Small FPV drones have become the sharpest cheap weapon of the Russo-Ukrainian war, and the radio link is the thread the whole kill hangs from. ExpressLRS, an open project free for anyone to flash, flies a large slice of those drones on each side. Russia has spent hard on counter-drone jamming, and cued attacks are the natural next move once barrage saturates. The team dates the tests to early summer 2026 in its own teaser. Formal paper plus a Telegram push is the standard Russian mil-tech release, one artifact for the engineers, one for the buyers.

Impacts so far

The paper reframes the drone-versus-jammer fight for anyone reading it closely. Survival against jamming stops being mainly a power contest. What the control waveform leaks about itself, its width, its spreading factor, its hop pattern, now sets the ceiling on how badly a jammer hurts you. MILELRS-style firmware moves from a curiosity to a countermeasure worth fielding. And the marketing bite is real. A tidy Russian paper claiming to lay ExpressLRS bare drops a seed of doubt in every operator who flies it.

Outlook and strategic foresight

Cued jamming will spread across Russian units and volunteer suppliers over the next six to eighteen months. We rate that very likely, since the method is published, the economics favor it, and a productized scanner already exists to feed it.

Machine-assisted signal reading looks likely to shorten the reconnaissance step toward real time, because sorting hop width and spreading factor from a capture is a tractable pattern job, and faster reads mean faster attacks on moving drones.

Allied drone programs will probably push harder toward low-intercept waveforms, near-continuous frequency grids, and cryptographically driven hop orders. MILELRS already shows the fix works, and the Russians signed off on its effect in print.

Whether Russia closes the MILELRS gap is the open question. We give it a roughly even chance inside the window. Chasing a near-continuous grid with wider energy and better estimators runs into a hard information wall the dense grid builds on purpose.

The lesson runs past drones. A waveform whose settings an enemy recovers hands that enemy a jamming edge, and the rule holds for any tactical radio or command net that hops on a rebuildable pattern. NATO signal planners have a live range in front of them. The emissions discipline the FPV fight demands is the same discipline their own datalinks will need against a peer jammer.

What the document proves, and what it only claims

Ranking the evidence keeps the numbers honest.

Item	Standing	Note
Measured RXLOSS thresholds per class	Verified data, in-document	The authors' own bench readings, internal to their rig
The Class 0 to 3 taxonomy	Verified framing	Their analytic model, clean and sound
MILELRS grid: 65,536 channels, ~305 Hz, long-period hop	Vendor description	Technically plausible, not independently confirmed
×27 and ×3.2 range gains	Adversary extrapolation	Physics from a bench delta, never field-measured
"Тишина" reads every parameter a jammer needs	Adversary claim	A sales assertion, unverified
Class 3 impossible against MILELRS	Adversary claim, credible	Follows logically from the grid they describe

Recommendations

The fixes fall into two hands, the units that fly and the enterprise that plans.

For the flyers, the first move is to treat the control-link settings as protected the way any targetable emission is protected. Hop width, spreading factor, and hop order leak a firing solution to a listening enemy. Next, field firmware built on dense or continuous grids with long-period, cryptographic hop generation, since a MILELRS-class design takes the enemy's best attack off the table outright. Build the link to bend, not snap, under matched interference, with error correction and a data-rate fallback instead of a hard cliff into loss-of-signal. And where the platform allows it, carry a second path, an alternate band, a mesh relay, or fiber on the tethered types, so one jammed link does not down the aircraft.

For the enterprise, stand up counter-reconnaissance. The scanner that cues the jammer is an emitter, and an emitter is a target you find and geolocate before it finishes its read. Run the mirror image too. Russian FPV also flies ExpressLRS, so the same class ladder the paper points at Ukrainian drones runs straight back the other way.

Opportunities

The dependency chain the Russians drew works against them. A defender who breaks the reconnaissance step collapses the whole advantage down to plain noise, which makes emissions discipline a cheap countermeasure with heavy payoff. Deception is open ground. A decoy control emitter wastes the enemy's scan and points the jammer at empty air. Push it

further and a fleet-wide move to hardened firmware turns a Ukrainian field fix into an allied baseline that raises the jamming bill for every future opponent.

Gaps and cautions

The study is a relative bench measurement, good to about plus or minus 1.5 dB, run with transmitter, jammer, and receiver crowded together. Real terrain, multipath, and motion never entered the test. The range multipliers are physics drawn from the bench, not field numbers.

Anti-jam features went untested. Error correction, adaptive data rate, and link handoff live in real radios and sit outside the paper. Add them and the thresholds move.

The source is the seller. A capability claim carries advocacy, and the authors admit a naive build performs differently. No outside party has checked the work.

The paper says nothing about detection. The work sharpens the attacker's edge and stays silent on how the target sees the attack coming, or spots the scanner first.

Fielded scale and readiness of "Тишина" stay unconfirmed. The interface looks operational. Open sources fix no numbers behind it.

Treadstone 71 has not independently corroborated the actor, the tool, or the measurements. The assessment rests on the uploaded paper and analyst domain knowledge. Sourcing rates moderate at best, and the sharpest claims, the range gains and the scanner's reach, sit on the vendor's word. The full class-by-class table, the sweep-rate runs, and the reverse-chirp comparison behind the debunked myth sit in the source appendix for anyone auditing the numbers.

Sourcing rated moderate. Vendor-authored source; sharpest figures rest on the adversary's word. Verified data separated from adversary claims within.