AMATEUR RADIO · AUGUST 16, 2026

SDR Receiver and Digital Decoding Practice, August 2026: Hardware, Architecture, and the Theory Coming into Practice

State of play across SDR in August 2026: direct-sampling HF receivers are now mainstream from $30 to $10k (RTL-SDR V4L, RX-888, Aurora AU-520); GNU Radio 4.0 reached RC1; neural receivers had their first real 5G deployments via Sionna; OTFS/AFDM, FTN, PAC codes and diffusion decoders are the not-yet-mainstream frontier.

11 min read · 28 sources

6G waveforms | GNU Radio 4.0 | SDR | channel coding | digital decoding | neural receivers | receiver design | software defined radio

CONTENTS

The 2026 receiver market: what changed and what to buy 2
The transceiver frontier: direct-sampling SDRs arrive 4
Receiver architecture: the design rules that still hold 5
The digital decoding chain in practice 6
Software platform: GNU Radio 4.0 goes release-candidate 7
Coding theory reaching practice: 6G channel coding and neural decoding 9
Neural receivers: the first real deployments 10
Waveforms and signalling not yet mainstream 11
Bottom line 12
Sources 13

Software-defined radio in mid-2026 is at an inflection point. Direct RF sampling has moved down-market to every price tier — from the $30 dongle to the $10,000 flagship transceiver — the dominant open-source framework, GNU Radio, is on the verge of a ground-up 4.0 release, and for the first time neural-network receivers trained in simulation have been deployed and demonstrated inside real 5G NR networks running commercial equipment. This document covers the current receiver hardware, the design rules that still govern good SDR receiver and decoding chains, and the mathematical techniques — neural decoders, PAC codes, OTFS/AFDM, faster-than-Nyquist signalling — that are genuinely new and not yet widely adapted.

103 dB	16-bit @ 245.76 MSPS	RC1	+0.4 dB	Real-time
Airspy HF+ Discovery in-band DR (measured)	FlexRadio Aurora AU-520 ADC	GNU Radio 4.0 status (22 Mar 2026)	GNN decoder gain vs weighted BP, BCH(63,45)	Neural receiver in 5G NR (Sionna RK demo)

- 22 Mar 2026 — **GNU Radio 4.0 RC1**
  API locked; compile-time composition, SimdFFT, WASM, SoapySDR
- 15 May 2026 — **FlexRadio Aurora at Hamvention**
  DX Engineering named US distributor of the Aurora 500W SDR line
- 21 May 2026 — **GR4 stewardship settled**
  Community release line continues from MIT core; GSI/FAIR goes separate way
- ~May 2026 — **RTL-SDR Blog V4 end-of-line**
  R828D chip stockpile exhausted; counterfeits found in remaining reels
- 11 Aug 2026 — **RTL-SDR V4L (Lite) on sale**
  R828S tuner, diplexed front end; Osmocom driver support lands days later
- 16 Aug 2026 — **Status snapshot**
  No 6G waveform standardised; neural receivers demoed live, not shipped

The 2026 receiver market: what changed and what to buy

The most consequential hardware news of 2026 is at the budget end of the market. The RTL-SDR Blog V4 was made end-of-line in May 2026 because the Rafael R828D tuner chip it depends on is out of production and the remaining third-party stockpile turned out to be counterfeit — the last reels had the R828D logo lasered onto a different, cheaper chip of the same package size. Its replacement, the RTL-SDR Blog V4L (“Lite”), went on sale on 11 August 2026 using the R828S variant, which has one fewer RF input than the R828D. That forces a diplexed instead of triplexed front-end filter, hence the “Lite” designation. The trade is a small one: slightly better sensitivity on the front end because there is less filtering loss, against weaker rejection of strong off-band signals — the two-tone test results on the V4L are identical to the V3[3][4]. The V4L keeps the V4’s HF upconverter architecture, 1 PPM TCXO, 4.5 V bias-T and LED, and improved PCB layout; the official Osmocom driver support landed within days of launch, with the rtlsdrblog driver fork and a Windows ExtIO available earlier. The chip stockpile is expected to last roughly a year, and the design is itself “limited edition” — a V5 is in early development but not expected to appear before 2027[3][4].

RECEIVER	STREET PRICE (2026)	ADC	INSTANTANEOUS BANDWIDTH	ARCHITECTURE	BEST FOR
RTL-SDR Blog V4L	~$30	8-bit	2.4 MHz	Upconverter + direct sampling	Entry, VHF/UHF monitoring
Airspy HF+ Discovery	$169	18-bit effective	768 kHz	Direct conversion	Weak-signal HF, MW DX
SDRplay RSP1B	~$143	14-bit	10 MHz	Single tuner, preselection	General coverage, first upgrade
SDRplay RSPdx-R2	$219	14-bit	10 MHz	Tuner + direct sampling	One-box 1 kHz–2 GHz
RX-888 MkII	~$150	16-bit @ 130 MSPS	64 MHz	True direct sampling	Whole-HF-band recording
Elad FDM-S3	$689	16-bit	24 MHz	Direct sampling	Wideband HF recording

Prices reflect July 2026 dealer checks and move around[1][2][5].

For HF weak-signal work the consensus pick remains the Airspy HF+ Discovery at $169: an 18-bit effective ADC (oversampled from a 16-bit converter) with a measured 103 dB in-band dynamic range on 40 metres during a contest, and field reports of decoding FT8 at −24 dB SNR that were completely invisible on an RTL-SDR running the same antenna. Its limits are real: only 768 kHz of instantaneous bandwidth and no coverage between 31 and 60 MHz[1][2]. The SDRplay RSPdx-R2 at $219 is the all-rounder — continuous coverage from 1 kHz to 2 GHz with 10 MHz of visible spectrum, enough to see the entire 20-metre band at once — but it sits roughly 6–8 dB behind the Airspy on dynamic range in crowded HF conditions. The RSP1B at about $143 is the strongest value general-coverage receiver: 14-bit, 10 MHz bandwidth, eleven switched preselection filters, and a bias-T[5]. For operators who want the entire HF spectrum digitised at once, the RX-888 MkII streams 10 kHz to 64 MHz — a 16-bit ADC at 130 MSPS — over USB 3.0 and lets the host PC do all the DSP; the MkII revision added a tunable HF-path attenuator, a −10 to +33 dB variable-gain amplifier, the newer R828D tuner, and external reference-clock selection[6]. Two cautions from the vendor ecosystem: counterfeit RX-888 units are circulating and the vendor publishes guidance on identifying them, and the Chinese “Malachite” screen receivers measure only about 70 dB of dynamic range — RTL-SDR territory in a bigger, more expensive box[1][6].

Instantaneous bandwidth by receiver

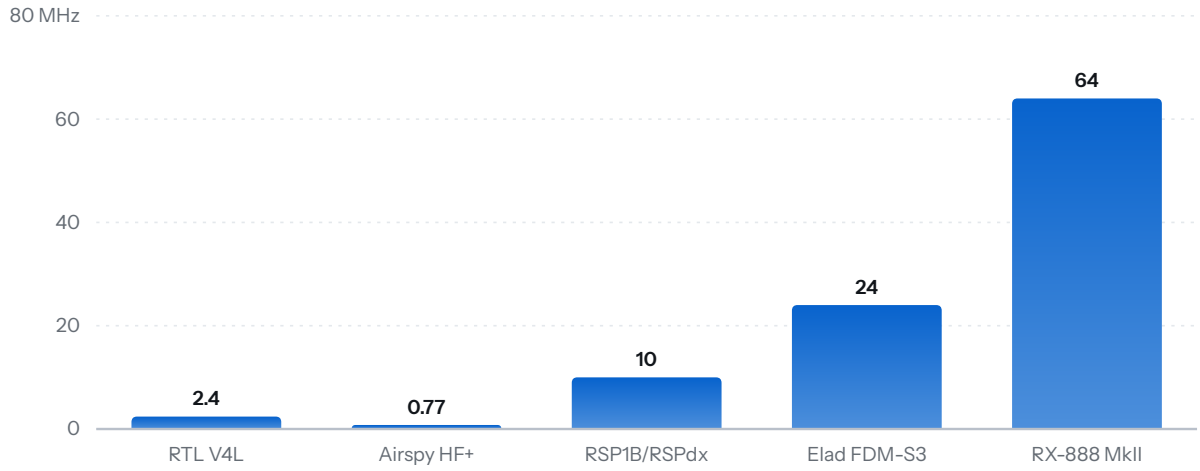

Manufacturer figures, 2026; Airspy figure per HamPost review.

CAUTION

If you buy the new RTL-SDR V4L, check driver status before relying on it: Osmocom support landed only in the days after the 11 August 2026 launch, and the R828S tuner is new to the RTL-SDR ecosystem [3].

The transceiver frontier: direct-sampling SDRs arrive

The top end of the market moved decisively to direct sampling this year, and the flagship story is FlexRadio's Aurora series. Announced through 2025–2026 and given a US distribution deal with DX Engineering at Hamvention on 15 May 2026, the Aurora is a fully integrated 500-watt HF/6-metre SDR transceiver: amplifier, automatic antenna tuner, and wide-range AC power supply all inside one 18 lb chassis, with the 500 W PA running at 80% nominal efficiency via switching amplifier technology rather than conventional LDMOS — which is why the chassis stays small and the heatsink stays cool [8] [9]. The AU-520 pairs that transmitter with two independent Spectral Capture Units, each digitising wide slices of the HF/6 m spectrum: 16-bit ADCs at 245.76 MSPS direct sampling, four slice receivers, four panadapters, 110/115 dB RMDR at 2 kHz spacing, better than 145 dB dynamic range at maximum zoom, seventh-order preselectors, and integrated GNSS providing 0.001 ppm disciplined frequency stability (or a 10 MHz GPSDO option at 5×10^{-13}). It also ships with SmartSignal adaptive predistortion for spectral purity and Controlled Envelope SSB, which the vendor says adds up to 3 dB of effective talk power without distortion. The 500 W amplifier and DPD/envelope-control combination is the quiet headline of 2026: this class of performance used to require a separate amplifier and a rack of gear [8] [9].

Icom's answer, the IC-7760, is a modular direct-RF-sampling design with a LAN-connected remote RF deck, two independent receivers, a 7-inch main touchscreen plus a 2.4-inch secondary display, a digital preselector (DIGI-SEL) that filters before the main receiver chain to protect against overload, and digital pre-distortion on its 200 W transmitter[10]. The architectural takeaway is identical across both vendors: the heterodyne chain's analog mixer and IF stages have been replaced by a high-speed ADC and FPGA, and the remaining analog content is preselectors, LNAs, and the power amplifier. At the open-source end, the Hermes-Lite 2 remains the reference for an affordable, fully open direct-conversion HF transceiver — a broadband-modem-chip ADC/DAC plus FPGA and gigabit Ethernet on one board, entirely open hardware and software, over 500 units built, roughly $300 for the board. It is the community route into the same direct-sampling territory the FlexRadio and Icom boxes occupy, and it is a first-class target for SDR++-class software, including the SDR++ Brown fork with its built-in FT8/DSD decoding and Hermes-Lite 2 transmit support[11][2].

Receiver architecture: the design rules that still hold

Whatever the hardware, the engineering rules for a good SDR receiver have not changed — they have just moved further into the digital domain. A systematic 2026 treatment of SDR architecture frames the design space as a continuum from fully analog (level 0) to the theoretical "ideal software radio" with the ADC at the antenna (level 4); practical receivers live at levels 2–3, digitising after the IF stage or after the LNA[1].

ARCHITECTURE	ADC POSITION	TYPICAL RANGE	ADC SPEED NEEDED	ANALOG COMPLEXITY
Direct RF sampling	After LNA	DC–100 MHz	Very high (>250 MSPS)	Minimal
Single conversion	After one mixer	100 MHz–6 GHz	Medium (50–250 MSPS)	Moderate
Dual conversion	After two mixers	1–40 GHz	Low (10–50 MSPS)	High
Direct conversion (zero-IF)	At baseband I/Q	All	Low–medium	Moderate

Direct RF sampling is now preferred at HF because a modern 14–16 bit ADC at 125+ MSPS digitises the entire 1.8–30 MHz spectrum in a single Nyquist zone; the superheterodyne architecture survives at VHF/UHF where wide coverage matters, and zero-IF survives where integration density matters[1]. NI's guidance for RF engineers puts the trend in context: 12-bit ADCs went from 250 MS/s in 2005 to 6.4 GS/s by 2018, which is what made direct RF digitisation of L- and S-band practical — and it simplifies phase-coherent arrays because there are no LO chains left to synchronise, only clocks[12].

The quantitative core that every designer still uses begins with noise. The LNA dominates total noise figure through the Friis cascade because the first stage's contribution is undivided by any preceding gain. At HF, atmospheric and man-made noise exceed thermal noise by 20–40 dB, so linearity matters more than a sub-1 dB noise figure: a 3 dB NF LNA with +30 dBm OIP3 is preferable to a 0.5 dB NF LNA with +15 dBm OIP3 in an HF application. Dynamic range is budgeted through two equations: spurious-free dynamic range, $\mathrm{SFDR} = (2/3)(\mathrm{IIP3} - \text{noise floor})$, and blocking dynamic range, BDR =

P1dB − MDS, where MDS = −174 + NF + 10log10(BW) + required SNR. The worked example in the reference (NF 10 dB, IIP3 +10 dBm, 2.4 kHz SSB) yields SFDR ≈ 93.5 dB and BDR ≈ 120 dB[1]. ADC physics then sets the ceiling: SNR = 6.02N + 1.76 dB for an ideal N-bit converter, and the effective number of bits (ENOB) is the number that matters in practice — a 14-bit converter typically delivers about 12 ENOB, roughly 74 dB of SINAD and 84 dB of dynamic range[1].

ADC BITS	THEORETICAL SNR	TYPICAL ENOB	TYPICAL SINAD	DYNAMIC RANGE
8	49.9 dB	7.0	43.9 dB	48 dB
12	74.0 dB	10.5	65.0 dB	72 dB
14	86.0 dB	12.0	74.0 dB	84 dB
16	98.1 dB	13.5	83.0 dB	96 dB

Aperture jitter is the hidden ceiling: SNR_jitter = −20log10(2π·f_in·t_j), so at 150 MHz input, 500 fs of RMS jitter caps SNR at about 52 dB — below what an 8-bit ADC already achieves. This is why high-end receivers spend real money on OCXO or GPSDO clocking, exactly as the Aurora's GNSS-disciplined reference demonstrates[1][8]. Oversampling, meanwhile, is free dynamic range: every factor-of-four oversample adds roughly one bit, and sampling 125 MSPS with a 3 kHz SSB channel gives M = 20,833, i.e. +43.2 dB of processing gain or about 7.2 extra bits after decimation — the reason a 14-bit ADC at 125 MSPS can exceed 120 dB effective range in a narrow channel[1]. The digital down-conversion chain that exploits this is the familiar NCO–CIC–FIR cascade: a 32-bit NCO phase accumulator gives about 0.03 Hz resolution at 125 MSPS; the CIC decimator uses no multipliers but grows bits alarmingly fast — a fifth-order, 128× CIC grows 35 bits, so a 14-bit ADC input needs roughly 49-bit integrators — and is followed by a sinc-compensation FIR, a half-band filter, and a final channel-selection FIR. For many simultaneous channels, a polyphase filter bank plus FFT channeliser is $O(K \log_2 K)$ per sample against $O(K^2)$ for per-channel DDCs, which is the technique behind "see the whole band at once" receivers. Zero-IF receivers additionally need DC-offset removal (a slow IIR high-pass) and adaptive 2×2 I/Q imbalance correction, which takes image rejection from roughly 25–40 dB uncorrected to better than 60 dB corrected. Inside the box, FPGAs absorb everything that runs at sample rate — ADC interface, DDC, channelisation, timing — while CPUs and DSPs run the variable-rate work: FEC, protocol stacks, spectrum sensing[1][12].

The digital decoding chain in practice

Everything downstream of the DDC is synchronisation, and the four-stage structure is the backbone of every digital demodulator in production. Coarse frequency acquisition uses FFT peak detection and converges in a single block. Fine frequency and phase tracking uses a second-order type-II PLL, converging over tens of symbols to better than 1% of the symbol rate. Symbol timing uses a Gardner timing-error detector or Mueller–Müller, converging over 50–200 symbols to within a few percent of

the symbol period. Frame synchronisation correlates against a known preamble and locks in one frame [1].

Timing recovery in practice means the Gardner algorithm running in a closed loop. The receiver samples at twice the symbol rate and forms the error metric $m[n] = (x[n] - x[n-2])\cdot x[n-1]$: the sign of the metric tells you whether you are sampling early or late, and its magnitude scales with the offset. A worked, code-level demonstration running over a real soundcard channel shows the loop centring the eye diagram and re-converging even when the stream randomly skips samples, which is exactly the failure mode a drifting clock or buffer underrun produces [26]. Loop bandwidth is the acquisition-versus-noise trade: typical values run 0.01–0.05 of the symbol rate, and the loop filter is a proportional-integral controller in the classic sense. Above timing sits the rest of the chain — carrier recovery with a Costas loop for PSK or a PLL for QAM, adaptive equalisation (LMS-based, or blind equalisation such as CMA for unknown constellations) for dispersive channels, and frame sync via preamble correlation. Computational cost scales steeply with constellation: BPSK/QPSK demodulation runs about 30 MFLOPS/ksps, 16-QAM with an equaliser about 50, and OFDM about 100 [1].

FUNCTION	ALGORITHM	CONVERGENCE	ACCURACY
Coarse frequency	FFT peak detection	1 block	±1 sub-carrier
Fine frequency	2nd-order type-II PLL	10–100 symbols	<1% of symbol rate
Symbol timing	Gardner TED / Mueller–Müller	50–200 symbols	<5% of symbol period
Frame sync	Preamble correlation	1 frame	exact

The 2025–2026 software side of this chain deserves equal attention, because decoding is increasingly a built-in feature of the receiver software rather than a separate audio pipe. The SDR++ Brown fork decodes FT8/FT4 across the whole SDR passband — the decoder is extracted from MSHV — with PSKreporter reporting, plus DSD for DMR/P25/NXDN digital voice and direct KiwiSDR connectivity. WSJT-X 3.0.x added optional parallel multithreaded FT8 decoding, an automated release pipeline, and native TCI rig control; WSJT-Z merged the JTDX iterative decoder into the WSJT-X 3.0 baseline with a 1–12 thread decode control. The practical effect for a monitoring or contest station is more decodes per period on ordinary hardware, with the decoder living inside the SDR application itself [2] [15] [28]. On the deep-signal-analysis side, a February 2026 Nature study of OneWeb Ku-band signals demonstrates the modern workflow: blind signal parameter estimation and demodulation of a commercial constellation from raw SDR captures — a reminder that the Gardner/PLL/equaliser stack scales from soundcards to LEO satellite downlinks [search: Nature OneWeb].

Software platform: GNU Radio 4.0 goes release-candidate

GNU Radio 4.0 — a clean-slate re-architecture with roots in the DARPA SDR 4.0 program and the GSI/FAIR accelerator facility — reached RC1 on 22 March 2026, with the project stating that the core architecture, execution model, and API are stable and no longer subject to major breaking

changes[5][6]. Three capabilities define the release. First, modular schedulers decouple the signal-processing graph from the execution strategy, so a developer can select or write a scheduler matched to the workload — low-latency streaming, high-throughput batch, or heterogeneous CPU/GPU/accelerator targets — and custom buffer types and scheduling domains control data movement per device[5][6]. Second, compile-time composition fuses block chains into zero-overhead pipelines: the published single-threaded benchmark on one million float samples shows a src→mult→div→add→sink chain at 87 M ops/s in the runtime, 187 M ops/s merged, and 25.4 G ops/s when the whole graph is compiled constexpr — with the merged and constexpr modes eliminating intermediate buffers entirely, and feedback-heavy structures like IIR filters and PLLs gaining orders of magnitude. SimdFFT makes SIMD a first-class concern across the signal chain[5].

GNU Radio 4: 1M-sample pipeline throughput

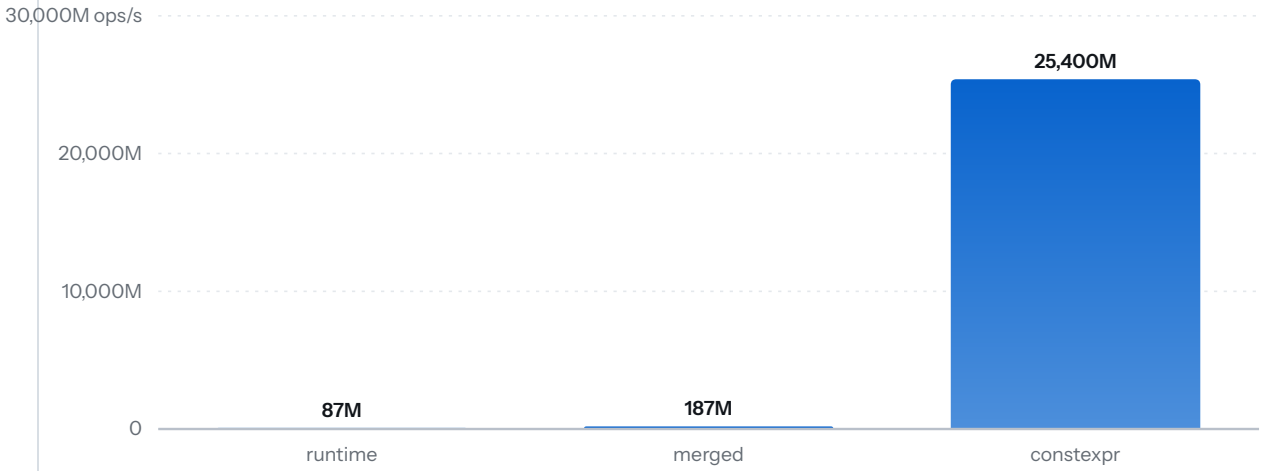

GR4 RC1 announcement benchmarks, x86-64, single-threaded; constexpr reaches tens of GS/s.

A block in GR4 is this minimal — the whole framework handles scheduling, buffering, and SIMD width:

```cpp
template<typename T>
struct MultiplyConst : gr::Block<MultiplyConst<T>> {
    PortIn<T> in;
    PortOut<T> out;
    T k = static_cast<T>(1);
    GR_MAKE_REFLECTABLE(MultiplyConst, in, out, k);
    [[nodiscard]] constexpr T processOne(T x) const noexcept {
        return x * k;
    }
};
```

The third defining capability is the AI-ready data model. GR4 replaces the old polymorphic message type with a tensor-aware, metadata-rich model — signal data becomes multidimensional tensors with

shape, type, and context such as time and frequency, natively consumable by PyTorch, TensorFlow, and ONNX, which the project explicitly positions as the foundation for AI-enabled SDR. The core is a header-only C++23 library: a full signal-processing chain compiles into a 2–3 MB standalone executable with no GNU Radio runtime dependency, multi-block chains run with under 10 µs latency, and early SYCL/CUDA prototypes demonstrate cross-device scheduling. WebAssembly support runs DSP pipelines in the browser, demonstrated by the OpenDigitizer example, and built-in SoapySDR abstraction covers the hardware zoo. A plugin system with runtime reflection makes graphs self-describing — block metadata, ports, and parameters are introspectable, which is the hook for remote control middleware such as OpenCMW and for AI-driven workflow tooling [5] [6].

Governance was settled in May 2026. GSI/FAIR, which led the 4.0 development, proposed a separate GSI-centred direction with LGPLv3 licensing; the GNU Radio Board declined and on 21 May 2026 confirmed that the 4.0 line continues as an officially community-maintained release from the MIT-licensed core, with ported 3.x blocks retaining GPLv3 and the GNU Radio Studio GUI remaining GPLv3. The practical message for practitioners: build new receivers against GR4 now — the API is locked, the core is permissively licensed, and the performance gains are measured, not promised [7].

Coding theory reaching practice: 6G channel coding and neural decoding

Channel coding is where the “new mathematics not yet widely adapted” question has its most concrete answer. The Open6GHub white paper “Towards a Unified Coding Scheme for 6G” — from the University of Stuttgart, RPTU Kaiserslautern, and KIT, published March 2026 — argues that 5G’s split between LDPC codes for the data channel and polar codes with successive-cancellation-list decoding for control should give way to a single, adaptable “unified” code family covering all 6G use cases, because coding gain alone is no longer the figure of merit. The new KPIs it targets are sub-millisecond decoding latency, ten-fold energy efficiency, ten-fold cost efficiency, ten-year battery life, and a bit error rate down to 10^{-12} with no error floor for Ethernet compatibility — an envelope that is frankly hostile to compute-heavy neural decoders [14].

The strongest coding-theory candidate remains the polar-code family, proven capacity-achieving under successive cancellation and already deployed in 5G control channels. The 2024–2026 literature adds a clear trajectory. PAC codes — polarization-adjusted convolutional codes — add a rate-one convolutional pre-transform before the polar transform and deliver near-dispersion-bound performance at short block lengths; 2025 work produced Monte-Carlo rate-profile construction valid for any code length and rate, and 2026 results give generalized fast list decoding that tames the large list sizes PAC codes otherwise demand [search: ScienceDirect 2025; IEEE WCL 2026; arXiv 2501.18080]. In parallel, the deep-learning-for-channel-coding literature has matured into a full taxonomy: learned belief propagation (weighted BP, neural BP with per-iteration weights), transformer-based decoders, and generative approaches, with polar-code construction itself now done by learned frozen-bit selection [15].

The most concrete neural-decoding results come from NVIDIA's GNN decoder, reproducible in the Sionna simulator. By letting a graph neural network learn a generalised message-passing algorithm over the code's factor graph, the decoder beats weighted belief propagation by about 0.4 dB on the (63,45) BCH code with only 9,640 trainable parameters, is competitive with conventional BP on 5G NR LDPC codes, and needs significantly fewer decoding iterations — while scaling to arbitrary block lengths, which most neural decoders cannot [16]. A May 2026 paper from Sun Yat-sen University goes further in the same direction, integrating belief propagation into a variational diffusion decoder and claiming state-of-the-art BER at drastically reduced model size and compute, explicitly because earlier neural decoders were too large for the millisecond-latency, power-constrained reality of 5G-class receivers [17].

KEY FINDING

None of these neural decoders is in deployed silicon. What changed in 2025–2026 is the direction of travel: the field pivoted from "bigger model, better BER" to "model-based, low-complexity, hardware-realistic" — deep unfolding, GNNs with thousands of weights, and BP-integrated diffusion decoders are all attempts to fit neural decoding inside the 6G KPI envelope [14] [15] [16] [17].

Neural receivers: the first real deployments

The most significant "not yet widely adapted but now demonstrably real" development is the neural receiver — a trainable network that replaces the channel-estimation, equalisation, and demapping blocks of a conventional receiver — deployed in a live network. The vehicle is NVIDIA's Sionna ecosystem. Sionna 2.0.1 is a differentiable, GPU-accelerated open-source library combining a ray tracer for radio propagation, a link-level PHY simulator, and system-level abstraction, in which every signal-processing block can be replaced by a trainable network and gradients flow end-to-end through the whole transceiver [18] [19]. On top of it, the Sionna Research Kit (v1.3.0) couples an NVIDIA DGX Spark with OpenAirInterface and Ettus USRPs into a real-time, software-defined 5G RAN — the documentation claims you can connect commercial 5G equipment to a network running your own transceiver algorithms in an afternoon, with tutorials covering GPU-accelerated LDPC decoding, a neural demapper, and a full 5G NR PUSCH neural receiver [18]. The proof point was presented at IEEE ICMLCN 2025: a real-time neural receiver, trained in Sionna and run with TensorRT inference, deployed in a 5G NR network using commercial user equipment [20]. The same pattern is visible on NI USRP hardware, where a DeepRx-style neural receiver runs inside the OpenAirInterface 5G stack coordinated by an O-RAN near-real-time RIC [search: NI 5G/6G prototyping]. The DeepRx lineage itself matters: the original DeepRx work demonstrated that a learned receiver could beat a conventional 5G NR receiver chain on block error rate with lower complexity, and the 2025–2026 deployments are the first time that idea has run in real time on real hardware with real UEs.

What makes this more than a demo is the standardisation track. A June 2026 paper by Xingqin Lin frames the 3GPP perspective on an AI-native 6G air interface: standardisation should not prescribe model architectures, training methods, or weights, but should define the protocol framework for

configuring, validating, activating, monitoring, and safely reverting AI-enabled functions — with a neural-receiver-assisted reference-signal adaptation as the concrete case study[21]. In other words, the industry is already designing the plumbing that will let trained receivers ship in 6G, building on the stepping stones already in 3GPP Release 18/19 AI/ML work for NR (CSI feedback, beam management, positioning). The model-based middle path — deep unfolding, which turns iterative algorithms such as ISTA, ADMM, and belief propagation into structured neural layers with interpretability preserved — is the subject of a comprehensive 2025 survey covering detection, channel estimation, beamforming, decoding, sensing, and security, and is widely regarded as the complexity-bounded route to deployment[22].

Waveforms and signalling not yet mainstream

Four mathematical directions are active in research and conspicuously absent from consumer gear. Orthogonal Time Frequency Space modulation moves signalling into the delay-Doppler domain, where high-mobility channels look quasi-static; it dramatically improves BER in high-Doppler scenarios such as LEO satellites and high-speed rail, but requires 2D FFT processing and a new modem architecture. Its challenger, AFDM (affine frequency division multiplexing), uses chirp pre- and post-rotation around a standard FFT, delivering delay-Doppler benefits with near-OFDM hardware reuse — and 2026 LEO channel-model studies (TDL-B, TDL-C) show AFDM outperforming OTFS, making it the emerging favourite for 6G standardisation. Per 3GPP status as of mid-2026, no 6G waveform is standardised yet; R20 is the study phase, and the working bet is that OFDM stays for mid-band eMBB while AFDM or OTFS is added for non-terrestrial and high-mobility verticals[23]. Industry analysts caution that whether any of the new waveforms fixes energy efficiency remains unproven.

	OTFS	vs	AFDM
Domain	Delay-Doppler	·	Chirp-rotated FFT
Doppler resilience	Excellent	·	Excellent
Hardware reuse vs OFDM	Partial (2D FFT, new modem)	·	Near-OFDM (standard FFT core)
2026 LEO studies (TDL-B/C)	Baseline	·	Outperforms OTFS
3GPP status	R20 study candidate	·	Emerging favourite

Per 5Gto6G waveform comparison, 2026; no 6G waveform standardised yet.

Faster-than-Nyquist signalling packs symbols at rates above the Nyquist limit, accepting deliberately introduced intersymbol interference and paying for it with sequence-estimation receivers. The theory is two decades mature — the Mazo limit and the Anderson–Rusek survey — and DVB-S2/S2X satellite studies have repeatedly shown spectral-efficiency gains. The hardware record is thin but real:

an FPGA FTN transceiver exists, and a GNU Radio implementation from Cal State Fresno and the University of Ottawa built a working BPSK FTN link on SDR hardware, reaching BER 10^{-3} at a packing factor $\tau = 0.8$ with minimum-Euclidean-distance estimation and $\tau = 0.9$ with a probabilistic data association detector [24]. It remains a niche technique precisely because the receiver must undo the deliberate ISI [24].

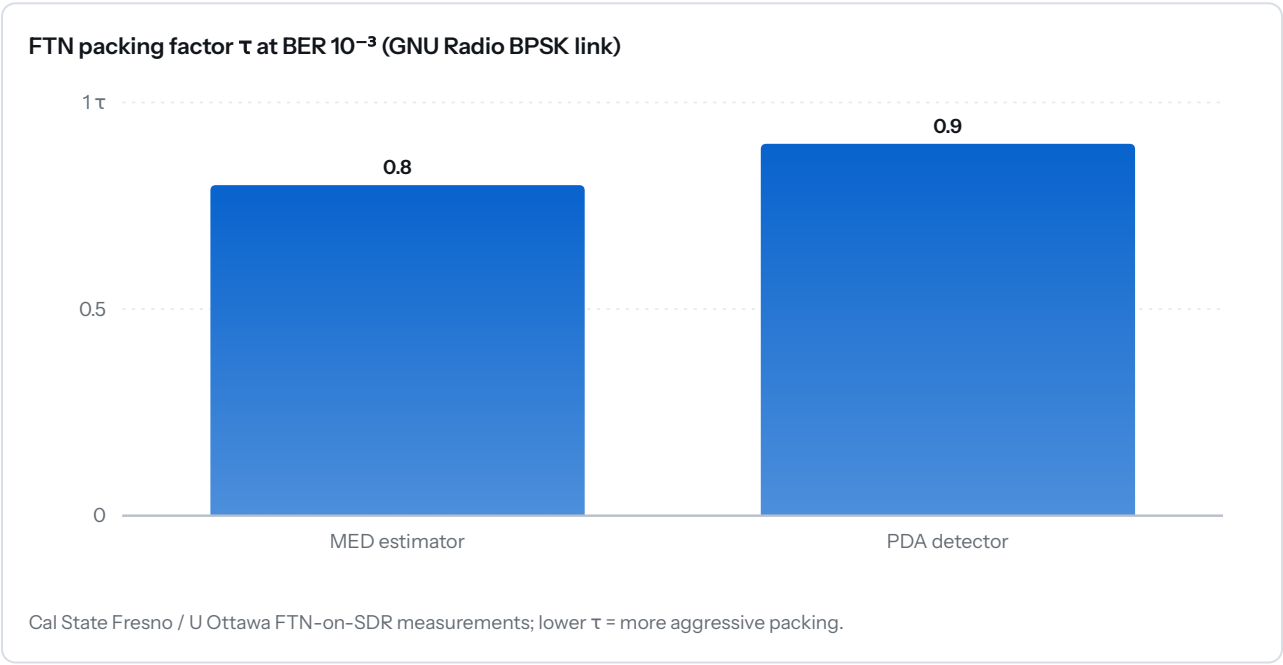

Cal State Fresno / U Ottawa FTN-on-SDR measurements; lower τ = more aggressive packing.

Compressed and sub-Nyquist spectrum sensing is the third thread: the Modulated Wideband Converter demonstrated the first wideband sub-Nyquist receiver sampling multiband signals at rates proportional to actual occupied bandwidth, and compressive-sensing research including sparse FFTs and deep-learning-assisted reconstruction continues — but no commercial SDR ships it, and the closest mainstream relatives are the sparse-signal-processing and AI-denoising features appearing in 2026 weak-signal software [search: Weizmann MWC]. Finally, the AI-native 6G literature includes semantic communication, where the air interface carries task-relevant meaning rather than raw bits [search: Frontiers 2025]. It is the furthest from practice of the four — no reference hardware, no standardisation vehicle — and should be treated as vision rather than roadmap.

Bottom line

For buyers in August 2026: the entry point is the RTL-SDR V4L at about $30, but its chip stockpile runs roughly a year and driver support just landed — treat it as a stopgap until the V5. For serious HF weak-signal work the Airspy HF+ Discovery at $169 remains the price-performance king; the SDRplay RSPdx-R2 at $219 is the one-box 1 kHz–2 GHz option; the RX-888 MkII gives whole-band HF recording for about $150. For transceiving, the Hermes-Lite 2 (about $300, fully open source) is the entry, and FlexRadio's Aurora AU-520 and Icom's IC-7760 are the 2026 flagships — all three prove that direct RF sampling is now the mainstream HF architecture.

For designers: the unchanged core is LNA-first noise budgeting, SFDR/BDR budgeting, aperture-jitter control, oversampling for processing gain, the NCO–CIC–FIR DDC, polyphase channelisation, and the Gardner/Mueller–Müller plus PLL synchronisation chain. Build new software on GNU Radio 4.0 — RC1, API stable, MIT core, AI-ready tensors. Anything touching 6G-class decoding should be designed against the KPI envelope of sub-millisecond latency, 10^{-12} BER, and ten-fold energy efficiency, which is exactly the constraint pushing neural decoders toward deep unfolding, GNNs, and BP-integrated diffusion models rather than giant black boxes.

The watchlist: GNU Radio 4.0.0 GA; the RTL-SDR V5 in 2027; neural receivers moving from Sionna Research Kit demonstrations into 3GPP Release 20+ AI/ML frameworks; the 6G waveform decision, with AFDM currently ahead of OTFS in LEO studies; and whether PAC codes or a unified polar-family code wins the 6G coding slot. Two cautions throughout: your local noise floor, not the receiver, usually decides what you can hear — a $32 dongle on a good antenna in a quiet location outperforms a $700 receiver in a noisy one — and vendor AI claims should be checked against published measurements, because in this field the gap between preprint and product is still measured in years.

SOURCES

1 Software-Defined Radio System Architecture | The Space-RF Journal
https://space-rf.org/articles/software-defined-radio-architecture/

2 Best SDR Receivers for Listening and Decoding (2026) | HamPost
https://www.hampost.com/guides/best-sdr-receivers-ham-radio

3 RTL-SDR Blog V4L (Lite) Now Available for Purchase! | rtl-sdr.com
https://www.rtl-sdr.com/rtl-sdr-blog-v4l-lite-now-available-for-purchase/

4 RTL-SDR Blog V4 End Of Line | rtl-sdr.com
https://www.rtl-sdr.com/rtl-sdr-blog-v4-end-of-line/

5 First Release Candidate for GR4 | GNU Radio
https://www.gnuradio.org/news/2026-03-22-gr4-release-candidate-1/

6 How GR4 Can Transform Your SDR Workflows | GNU Radio
https://www.gnuradio.org/news/2025-12-17-gr4-transform-sdr-workflows/

7 GNU Radio 4: Community Stewardship and the Road to 4.0 | GNU Radio
https://www.gnuradio.org/news/2026-05-21-gr4-community-stewardship/

8 FlexRadio Aurora AU-520 Dual-SCU 500W Integrated SDR Transceiver
https://www.flexradio.com/products/aurora-au-520-signature-series/

9 FlexRadio Announces DX Engineering as an Additional Distributor of Its New Aurora Series Transceivers | OnAllBands
https://www.onallbands.com/flexradio-announces-dx-engineering-as-an-additional-distributor-of-its-new-aurora-series-transceivers/

10 Icom IC-7760 Review and Complete Operating Guide | HamShack Reviews
https://hamshackreviews.com/icom-ic-7760-hf-50-mhz-transceiver/

11 RSP1B | SDRplay
https://www.sdrplay.com/rsp1b/

12 RX-888 | Products
https://www.rx-888.com/rx/

13 Advantages of Direct RF Sampling Architectures | NI
https://www.ni.com/en/solutions/aerospace-defense/electromagnetic-spectrum-operations/advantages-of-direct-rf-sampling-architectures.html

14 Towards a Unified Coding Scheme for 6G (Open6GHub White Paper)
https://arxiv.org/pdf/2603.23123

15 Recent Advances in Deep Learning for Channel Coding: A Survey
https://arxiv.org/html/2406.19664v1

16 NVlabs/gnn-decoder: Graph Neural Networks for Channel Decoding
https://github.com/NVlabs/gnn-decoder

17 Variational Diffusion Channel Decoder
https://arxiv.org/html/2605.18902

18 Sionna Research Kit (RK) documentation
https://nvlabs.github.io/sionna/rk/index.html

19 Sionna Documentation
https://nvlabs.github.io/sionna/index.html

20 Sionna Research Kit: A GPU-Accelerated Research Platform for AI-RAN
https://arxiv.org/abs/2505.15848

21 Toward AI-Native 6G Air Interface: A 3GPP Perspective on Protocol Framework
https://arxiv.org/abs/2606.27466

22 Comprehensive Review of Deep Unfolding Techniques for Next-Generation Wireless Communication Systems
https://arxiv.org/html/2502.05952v1

23 6G Waveform Comparison: OFDM vs OTFS vs AFDM | 5Gto6G
https://5gto6g.com/tools/waveform-comparison

24 Implementation of Faster-than-Nyquist Signaling using Software Defined Radios (GNU Radio Conference paper)
https://events.gnuradio.org/event/18/contributions/265/attachments/64/160/Implementation_of_FTN_Signaling_using_SDR_paper.pdf

25 Symbol Timing Recovery (Gardner algorithm) | DSPIllustrations
https://dspillustrations.com/pages/pages/courses/course_singlecarrier/html/05%20-%20Symbol%20Timing%20Recovery.html

26 RTL-SDR vs Airspy vs SDRplay: Best SDR Dongle 2026 Ranked | HamDeals
https://hamdeals.com/blog/sdr-dongles-ranked/

27 Hermes-Lite | Open Source HF SDR Transceiver
http://hermeslite.com/

28 Weak-Signal Receiver Software for Amateur Radio: Advances Through August 2026 (knowledge base entry)
https://intel.innerpulse.net/doc/weak-signal-receiver-software-for-amateur-radio-advances-through-august-2026

Researched and written with deepseek/deepseek-v4-flash-0731, from the sources listed above. Figures are as reported by those sources. Exported November 29, 2026 from https://intel.innerpulse.net/doc/sdr-receiver-and-digital-decoding-practice-august-2026-hardware-architecture, where the current version lives.