FEDERATED WIRELESS WHITE PAPER

AI and the New Economics of Shared Spectrum

How Physical AI at the Spectrum Coordination Layer Changes the Capital Math for Network Deployment

Field-Validated Results from Production CBRS Networks

Federated Wireless | June 2026

Abstract

CBRS has answered its foundational question: shared spectrum works at production scale. The remaining question is economic. How can networks be deployed, scaled, and monetized efficiently under shared access constraints?

AI in telecom has concentrated on network operations: automation, fault detection, and workflow optimization. A different class of application is emerging at the physics layer. When AI is applied to RF propagation modeling, interference analysis, and spectrum coordination, the effects are not operational but structural, altering the capital economics of network deployment and the accuracy with which operators can predict coverage before a site is built.

First-generation spectrum sharing relied on conservative static coordination models built on worst-case assumptions to protect incumbents and manage network coexistence. High-resolution environment and interference analysis under multi-dimensional constraints was computationally intractable at commercially viable cost and acceptable response times. Coordination systems also operated on scheduled refresh cycles, so power and channel assignments lagged actual RF conditions.

GPU-accelerated computing and machine learning now make it possible to model RF propagation and interference interactions with physics-level accuracy across large deployment areas, at the spatial resolution and computational speed required for production network decisions. Operators can allocate capital, qualify service at customer addresses, and configure spectrum assignments with a degree of certainty that statistical models cannot provide.

This paper describes how Federated Wireless has applied GPU-accelerated ray tracing and machine learning feedback loops to shared spectrum management, and documents the results observed across hundreds of production CBRS deployments:

- **Up to 5X capacity gain in congested environments**
- **90%+ drive test quality prediction reliability**
- **Up to 50% more spectrum to allocate**
- **1,000X faster planning and deployment cycles**

The Problem: A Computational Bottleneck

RF propagation obeys well-understood physical laws. Physical-layer improvements in coding, modulation, and MIMO are generally speaking approaching diminishing returns. A large share of unrealized performance remained on the table for a different reason: the inability to compute environmentally accurate interference interactions at sufficient resolution and speed.

The Pre-AI Coordination Model

Traditional SAS coordination and planning relied on:

- Terrain-based empirical statistical propagation models such as ITM, eHata, COST-231
- Coarse clutter categories (urban, suburban, rural)
- Static exclusion zones and guard bands
- Worst-case interference assumptions
- Scheduled refresh cycles, with coordination decisions updated on slow timescales rather than continuously

These approaches operate at low spatial resolution, on the order of tens to hundreds of meters, and model the propagation environment as statistically homogeneous. They omit actual clutter, spatial structure, directionality, and multipath geometry.

Three consequences followed. Interference was modeled statistically rather than deterministically. Power limits were set conservatively to protect worst-case incumbents. And assignments were stale: spectrum decisions reflected the RF environment as it existed at the last refresh, hours behind conditions on the ground.

The outcome:

- Underutilized spectrum capacity
- Overbuilt networks with excess sites
- High planning uncertainty
- Iterative and costly deploy-and-fix cycles

Why It Was Intractable

The problem is combinatorial. For a network of N transmitters, interference interactions in the worst-case scale roughly as $O(N^2)$. Each interaction depends on 3D geometry (buildings, terrain, foliage), frequency-dependent propagation, antenna patterns, and environmental conditions.

Solving this at per-home resolution, at decision speed, across thousands of sites was computationally infeasible with CPU-based methods.

The Architecture Shift

The breakthrough comes from combining four elements:

- **GPU-accelerated ray tracing.** Deterministic physics-based modeling of reflection, diffraction, and scattering, with per-link, geometry-aware propagation. The modeling stack is built on NVIDIA Sionna RT.
- **High-resolution geospatial data.** LiDAR-derived 3D terrain and building meshes plus land-use, water and vegetation layers with data curated from multiple sources
- **Machine learning feedback loops.** Continuous model calibration from production network telemetry.
- **Distributed GPU compute.** Continuous SAS decisioning at per-site granularity.

Together these elements shift spectrum coordination from conservative rule enforcement to physics-level optimization: interference constraints are evaluated per site, per sector, and per channel rather than applied as regional averages, and coordination decisions reflect current RF conditions rather than the last scheduled refresh cycle.

The Solution: Physical AI at the RF Physics Layer

1. High-Fidelity Propagation Modeling

The Federated Wireless Adaptive Network Planner (ANP) delivers drive test quality coverage prediction reliability by combining:

- Ray-traced propagation with multi-bounce reflection and diffraction
- Dense LiDAR-based 3D environments
- Material-aware attenuation modeling
- Frequency-specific calibration for the 3.5 GHz CBRS band

Technical basis. Validated against production field measurements across hundreds of live commercial deployments, prediction reliability reaches 90%+ drive-test equivalency. Empirical statistical models carry systematic errors that compound across sites; physics-accurate ray tracing eliminates the bias at the source.

The practical effect: reliable per-home serviceability decisions and elimination of costly post-deployment adjustment cycles.

2. AI-Driven PAL-to-PAL De-Confliction: Up to 5X Capacity in Congested Environments

The capacity gain requires no physical-layer innovation. It comes from removing conservatism that statistical coordination models are structurally required to apply.

Traditional coordination applies uniform power limits across large regions, with guard bands sized to protect the worst-case interference relationship in each zone. This is appropriate for models that cannot distinguish actual interference geometry: when the system cannot determine where interference constraints actually bind, it must assume they bind everywhere. The AI approach computes, for each CBSD, the interference impact on all neighbors, evaluates SINR constraints dynamically, and sets power per site and per sector accordingly. Because worst-case interference is spatially rare, enforcing constraints only where they actually bind releases substantial capacity that uniform limits suppress.

Technical basis. AI-driven de-confliction produces up to 20 dB of effective gain in allowable transmit power per site compared to conservative worst-case power allocation. Twenty decibels represents a 100-fold increase in radiated power equivalent. Under log-distance path loss models applicable to the 3.5 GHz CBRS band, that translates to a coverage radius increase on the order of 3X and a coverage area increase on the order of 9X, without changes to antenna configuration or physical infrastructure.

For operators, the practical consequence is reduced site count requirements, larger achievable cell footprints, and improved deployment economics at equivalent coverage targets.

3. Spectrum Allocation Efficiency: Up to 50% More Spectrum to Allocate

Traditional SAS coordination leaves spectrum on the table. Channel assignments are conservative by design: guard bands are sized to worst-case interference scenarios, PAL and GAA tiers are segregated with rigid separation, and channel plans default to the most protective configuration even when actual RF conditions would support tighter packing. The result is that a meaningful portion of allocated spectrum goes unassigned at any given time, not because the spectrum is unavailable, but because static models cannot safely distinguish where it is safe to use it.

Spectrum AI resolves this at the coordination layer. Physics-accurate interference modeling determines, for each deployment, the precise channel separations required to meet SINR targets across all co-channel and adjacent-channel relationships. Guard bands shrink to what physics requires rather than what worst-case statistics demand. PAL-to-PAL and PAL-to-GAA channel planning is optimized across the full 150 MHz CBRS band rather than partitioned conservatively. The system continuously updates channel assignments as the RF environment changes, recovering spectrum that static models would leave fallow. The outcome, validated across live commercial deployments, is up to 50% more spectrum actively allocated and in use compared to static coordination baselines.

Technical basis. The gain reflects deterministic interference analysis replacing statistical exclusion zones. AI-driven channel packing identifies non-interfering channel combinations that conservative models would prohibit. The result is higher spectral efficiency from the same licensed and unlicensed spectrum resource, with no physical-layer changes required. The 50% figure represents the increase in assignable spectrum capacity per deployment area observed in field-validated deployments relative to static coordination baselines.

4. Continuous Learning Loop

Each CBSD contributes telemetry: measured RSSI and SINR, throughput and link adaptation behavior, environmental variation impacts, and interference events. This data feeds model recalibration and bias correction, environment-specific tuning, and temporal adaptation. The Sionna RT model is natively ML tunable and allows such retuning via automation pipelines.

Unlike static planning tools, model accuracy improves with scale of deployment. Each additional site contributes calibration data that reduces prediction bias in comparable propagation environments, extends the model's coverage of building types and terrain configurations, and improves the system's ability to distinguish interference events from environmental variation. Models trained on production data across hundreds of deployments generalize more reliably than those calibrated on synthetic or lab environments.

5. Simulation Speed: 1,000X Faster Planning and Deployment Cycles

GPU acceleration enables full-market simulations in seconds, Monte Carlo scenario exploration, and rapid what-if analysis.

Task	Legacy	AI System
Site planning scenario	Hours	Seconds
Interference evaluation	Batch, offline	Continuous
Optimization loops	Manual	Automated

The practical effect is a shift in how operators approach network design: scenario evaluation that previously required days of engineering work can be completed in minutes, enabling iterative site selection and coverage optimization before capital is committed.

Production Proof Points

The following results are validated across hundreds of production CBRS deployments spanning fixed wireless access, enterprise, and neutral host networks.

Up to 5X capacity gain in congested environments. Derived from dynamic PAL-to-PAL coordination, observed across diverse geographies, with no antenna or physical-layer upgrades required.

90%+ drive test quality prediction reliability. Validated against field measurements. Eliminates re-climb and redesign cycles.

Up to 50% more spectrum to allocate. Derived from AI-driven channel packing and interference-aware spectrum coordination across the full CBRS band, validated across live commercial deployments where static models would leave spectrum unassigned.

1,000X faster planning and deployment cycles. Supports continuous planning workflows and rapid deployment scaling.

Deployment Economics

The core shift is simple: AI lowers the risk and cost threshold in justifying a network build.

Network Build Economics

Physics-accurate world-model-based modeling changes the site count equation. Traditional planning overbuilds to cover worst-case uncertainty: conservative interference constraints or guard bands lower channel usage and drive excess sites, inaccurate coverage predictions miss at the margin, and re-deployment cycles add cost after the fact. AI-driven planning removes each of these drivers. When spectrum coordination is precise, guard bands shrink. When coverage predictions are field-validated, overbuild disappears. When planning cycles run in seconds rather than hours, operators can simulate as many optimization iterations as needed before committing capital.

Field-validated deployments demonstrate up to 50% reduction in required site count to achieve equivalent coverage targets, and up to 40% lower overall deployment costs, compared to networks planned with conventional methods. The mechanism is direct: physics-accurate interference modeling reduces guard band requirements, which reduces site density requirements, which reduces capital expenditure. Improved coverage prediction accuracy eliminates the overbuild margin operators carry to compensate for statistical model uncertainty, and reduces post-deployment remediation costs where predictions would otherwise have missed.

Additional Effects

- Faster time to deploy
- Reduced engineering overhead
- Improved ROI predictability
- Accurate compliance reporting

Newly Viable Use Cases

- Rural broadband, where BEAD economics improve
- Neutral host networks
- Enterprise private 5G
- Copper replacement
- Dense urban infill
- Fixed wireless access wholesale

In each case, the business case depends on reducing the site count and planning cost required to achieve coverage targets and guarantee customer experience reliably. AI-driven spectrum coordination and physics-accurate propagation modeling address both directly.

Why the Advantage Compounds

Federated Wireless operates the largest commercial Spectrum Access System in the United States. The SAS sits at the spectrum coordination layer, the one position in the network stack with system-wide visibility into interference dynamics across operators, devices, and deployments. RAN-layer tools and network design software see a single network's slice of the environment. The coordination layer sees the system.

The data advantage is structural. Most wireless AI systems are trained on simulations, synthetic RF environments, or the telemetry of a single operator's network. Spectrum AI is trained on a decade of real-world spectrum assignments, interference events, and deployment outcomes drawn from the largest commercial CBRS dataset in the United States. It can further process telemetry continuously from thousands of CBSDs across hundreds of live commercial deployments, spanning fixed wireless access, enterprise, and neutral host network configurations. That dataset encompasses propagation observations across diverse building types, terrain configurations, and interference geometries that no single operator's deployment and no synthetic dataset can replicate.

The position is durable by construction. FCC authorization as a licensed SAS administrator, a nationwide ESC sensor network, and a decade of operator integrations represent infrastructure and regulatory standing that cannot be replicated from a higher layer of the network stack. RAN management tools and network planning software operate downstream of spectrum coordination; they see the assignments the SAS produces, cannot holistically see the interference dynamics and coordination decisions that produced them. The dataset that trains Spectrum AI is accessible only at the coordination layer.

The Platform: Adaptive Network Planner and Spectrum AI

Two AI-native products share the common Physical AI modeling core:

Adaptive Network Planner (ANP). For regional broadband and mobile operators. Use cases include address-level service qualification, Coverage IQ intelligence composed of competition and business data, detailed rooftop receiver-site analysis, and Broadband Data Collection (BDC) reporting.

Spectrum AI. For mobile operators and system integrators. Use cases include PAL and GAA management, EIRP optimization, spectrum de-confliction, boosted spectrum use via advanced antenna modeling, channel planning, and lifecycle governance.

Together they cover the full network lifecycle: deployment produces telemetry, telemetry updates the models, improved predictions inform the next deployment. Each cycle compounds the accuracy advantage.

Roadmap: From CBRS to 6 GHz AFC and Beyond

The modeling architecture developed for CBRS is not spectrum-specific in its foundations. GPU-accelerated ray tracing, LiDAR-based 3D environments, and machine learning feedback loops apply to any shared spectrum regime where interference interactions between incumbents and licensees must be computed at high resolution. The propagation modeling stack, the interference optimization framework, and the continuous learning loop carry forward directly. Operator integrations built on the CBRS SAS provide a deployment and data foundation that extends as those operators expand into adjacent bands.

The 6 GHz AFC regime presents a structurally analogous problem: automated frequency coordination to protect fixed satellite and terrestrial microwave incumbents from unlicensed standard power devices. The interference geometry, the incumbent protection requirements, and the device ecosystem differ from CBRS, but the underlying computational challenge, accurate interference prediction across a large number of heterogeneous transmitters in complex propagation environments, is the same class of problem. Regulatory frameworks for additional shared spectrum bands are in development in both domestic and international venues; each creates a deployment context where physics-accurate coordination produces the same category of advantage observed in CBRS.

The near-term capability trajectory moves from prediction to planning automation: AI-generated site configurations, channel plans, and coverage optimization that operators review and approve rather than derive manually. The longer-term direction is toward closed-loop network operations, where spectrum coordination, site configuration, and performance monitoring are continuously optimized against live network state without operator intervention at each decision point. The data and modeling infrastructure required for that capability is being built now, in production, on the CBRS deployment base.

Conclusion

The physics of RF propagation did not change. The ability to compute and optimize at the required resolution finally did.

CBRS proved that shared spectrum works technically and economically. Physical AI unlocks performance that was always present but inaccessible under static, stale coordination models. The result is a structural shift: more efficient networks, lower capital requirements, and expanded deployment viability. The architecture extends beyond CBRS to 6 GHz AFC and future shared spectrum systems.

Operators making deployment decisions today are doing so with fundamentally better tools than existed even a few years ago. The difference is structural, and it compounds.