STADIUM TECH REPORT

Hyper-Directional Wi-Fi for the Stadium Bowl

Two Approaches to Long-Throw, Focused-Beam Coverage of a Stadium Bowl

Wesley Terry and Bill Anderson
July 2026

2026

Introduction

This briefing reflects AmpThink's opinion, formed over fifteen years of designing, building, and operating **high-density Wi-Fi networks** across more than one hundred venues. That experience spans the full range of fan-facing approaches — overhead, back-to-front, infill, proximate (under-seat and handrail), and hyper-directional. Our early under-seat and handrail work was foundational, helping define how proximate solutions are deployed. Hyper-directional builds on those learnings and is now part of how we build networks. This briefing sets out, from that vantage, how the two hyper-directional families compare. The premise of the approach is straightforward: instead of blanketing a bowl with closely spaced access points and fighting the interference that results, a smaller number of antennas each throw a tight, long-reach beam at a single seating block, so the same channels are re-used more aggressively across the bowl. [1][2]

The original concepts — derived from the work Panasonic and Everest brought to market in 2017 and since carried forward by Cisco — are well understood, and in our experience the advantages are real and measurable: longer throw, tighter co-channel isolation, fewer mounts, and lower installed cost than proximate under-seat or handrail designs. [18][19] That established approach is a powered, software-controlled antenna that shapes and steers its own beam.

There is now a **second approach**. A passive radio-frequency lens (most prominently from MatSing) achieves the same long-throw, focused-beam coverage by optical rather than electronic means — forming many fixed beams from a single, unpowered lens. Its economics, failure modes, and installation profile differ from the electronic array. [6]

This briefing presents the two approaches side by side — the established active-electronic array and the newer passive lens — and sets out the differences that bear on a deployment: how each controls its beam, how each is installed and powered, how each scales, and how each fails. We do not regard either family as inherently superior; in our view they are two sides of the same coin, and the right choice is venue-specific.

The decisive factors are usually the venue's commercial objectives — capital and operating cost, sponsorship and brand visibility, and aesthetic and ownership priorities — constrained by available mounting structures, which differ markedly between the two platforms. The intent here is to inform that judgment; the purchasing decision, and the criteria that govern it, rest with the buyer.

Disclosure (conflict of interest). AmpThink is a commercial systems integrator that designs, builds, and operates Wi-Fi networks in large public venues. We deploy hyper-directional systems from the vendors discussed here. The authors are AmpThink personnel. AmpThink therefore has a direct commercial interest in these technologies and the conclusions this briefing reaches. This is AmpThink's analysis and opinion, drawn from its work in venues — not independent, third-party research — and some of the supporting figures come from vendor and trade-press sources. The document should be read with that interest in mind.

How to read this briefing. This paper serves a wide readership, and not every section is for every reader. Those concerned with why hyper-directional antennas help — RF and wireless designers — will find the argument in the Capacity Problem (§2), Why Not Simply More Access Points (§3), and Realized Capacity and the Client Side (§11), with the information-theory basis in Appendix A. Those deciding which approach to deploy — venue owners and operators — can go straight to How the Differences Bear on a Deployment (§12), supported by Two Families (§4) and Objections and Responses (§13). Those concerned with how the systems are built and mounted — constructors and integrators — will want Mechanism of RF Control through Radio Density (§7–§10), and Installation (§9) in particular. The lineage sections (§5–§6) give the provenance and maturity of each approach for any reader who wants it. Sections §1 and §4 are the shared foundation and are worth reading first.

Table of Contents

Table of Contents

What "Hyper-Directional" Means — Concept and Origin

A **hyper-directional antenna** (sometimes written as one word, "hyperdirectional") is an antenna engineered to throw a single narrow, high-gain beam far enough to cover a fixed seating section from afar. These antennas reach roughly three times the distance of a conventional access point — on the order of 175–200 feet — while starving every other direction of energy. The aim is isolation, not range: the main lobe puts energy on the target section, while deep side-lobe and back-lobe suppression keeps that energy out of every other section — on both transmit and receive. Beamwidth sets where a radio serves; the side-lobe envelope sets how well it isolates. Therefore, if each radio covers only its own block of seats and is deaf to its neighbors, many radios re-use the same channels in one bowl without colliding. The term "hyper-directional" is engineering shorthand, not a standards term; "directional" antennas are a century old, and "hyper-" signals the unusually tight beam and long throw of these stadium designs. [1][2]

Origin of the term. The product and the phrase have separate origins. As a product category, the hyper-directional stadium antenna originates with **Everest Networks** and its parent at the time, **Panasonic**, which in 2017 shipped the first antenna to couple a highly directional array with RF filtering for long-throw Wi-Fi. [2][19] The phrase itself — "hyper-directional antenna" — was coined and championed by AmpThink to describe this long-reach, RF-filtered, focused-beam class, and it was popularized by the trade publication *Stadium Tech Report* (editors Paul Kapustka and Pat Coyle), whose reporting carried it into common industry use. [1][2]

The Capacity Problem

A stadium network's task is to deliver **capacity** — usable aggregate throughput — to tens of thousands of densely packed users at once. Aggregate capacity reduces to three quantities multiplied together — the same three the wireless literature uses, named here as that literature names them:

$$\textbf{venue capacity} = B \times \eta \times \rho$$

- Where **B** is the usable **bandwidth** — the spectrum, in hertz,
- **η** (eta) is the **spectral efficiency** — the data each hertz carries, in bits per second per hertz,
- and **ρ** (rho) is the **spatial re-use factor** — how many times that spectrum runs at once in different parts of the venue.

Their product is aggregate throughput in bits per second; the per-unit-area form of the same product is what the literature calls *area spectral efficiency*. Because the terms multiply, the smallest caps the result and a gain in any individual parameter scales the whole. The decomposition is an effort to convey the levers that drive performance, not a claim that the terms are independent: they are coupled, most importantly through signal-to-noise ratio. Packing cells closer to raise ρ normally raises co-channel interference, which lowers SINR and therefore η — so in conventional designs ρ and η trade against each other rather than multiplying freely, and the true joint quantity is the interference-aware area spectral efficiency the cited literature derives, not a free product of three independent numbers. The value of writing capacity this way is that it shows where the headroom lies; the value of a directional antenna, as §2.3 develops, is that it raises ρ without paying the usual η penalty — which is the specific reason the two terms can move together here rather than against each other.

The three terms, and where they come from

Bandwidth, B — the spectrum. The raw resource: the total width of radio spectrum, in hertz, that a regulator releases for use. Wi-Fi draws on 2.4 GHz (effectively three non-overlapping channels), 5 GHz, and the 6 GHz band that Wi-Fi 6E opened with 1,200 MHz of new unlicensed spectrum. Slicing B into channels of width W gives the familiar channel count, $N = B / W$; the older “number of channels × data per channel” picture is just B counted in integer slices, and it collapses back to $B \times \eta$ once the channel width cancels. For a venue designer B is essentially fixed — you use the spectrum you are allowed — so it is a multiplier you cannot grow at will.

Spectral efficiency, η — bits per hertz. How much data each hertz of spectrum carries, in bits per second per hertz. Its ceiling is the Shannon–Hartley theorem, $\eta = \log_2(1 + SNR)$: on an additive-white-Gaussian-noise channel, per-link efficiency is fixed by signal-to-noise ratio and rises only logarithmically with it.[21] Realized efficiency in a bowl, though, tracks SINR rather than SNR alone — the co-channel interference the rest of the bowl contributes joins noise in the denominator — and is delivered through modulation and coding (MCS) and, per spatial stream, MIMO, net of MAC and PHY overhead, retries, contention, and client mix. The per-link coding gap to that bound is already small — modern radios sit within a few decibels of it (Appendix A)[22] — so the remaining efficiency headroom is not in the coding but in raising SINR, the lever §2.3 develops. This is the term information theory governs.

Spatial re-use, ρ — how many times a channel is used in a defined space. The number of times the same channel or channels can be used simultaneously in different parts of the bowl without the copies interfering — the spatial dimension of capacity, rooted in the cellular concept (its derivation and sources are in §2.2). Unlike B and η, ρ is capped by neither regulation nor information theory; it is set by how the network is physically built, and it is where nearly all the headroom lives. Its limit is **co-channel interference:** two radios on the same channel that can hear each other must take turns, so the spectrum is not truly re-used between them.

Why re-use (ρ) is the dominant lever

The history of wireless is, overwhelmingly, the history of raising ρ. Martin Cooper’s observation — “Cooper’s Law” — is that the number of simultaneous wireless conversations possible in a given area has roughly doubled every thirty months since Marconi, about a trillion-fold in all. Decomposed onto the three terms, more spectrum (B) accounts for only about a 25-fold gain and better modulation and coding (η) another ≈25-fold, while the remaining gains — the vast majority — came from spatial re-use (ρ): dividing space into ever-smaller cells so the same frequencies are used repetitively.[22] One peer-reviewed accounting attributes roughly a 2,700-fold capacity gain between 1950 and 2000 to spatial re-use alone.[23]

The cellular concept — cover an area with many small cells and re-use each channel in non-adjacent cells — dates to Bell Labs and was formalized by MacDonald in 1979; capacity scales with how finely space is divided.[24] Alouini and Goldsmith later quantified this as area spectral efficiency — data rate per unit bandwidth per unit area — and showed analytically how sharply capacity rises as cell size and re-use distance shrink: the per-area form of $\eta \times \rho$.[25]

But the same literature that credits re-use also bounds it. Because Wi-Fi shares the medium by carrier sense and contention, packing co-channel cells closer enlarges each radio's contention domain and multiplies hidden- and exposed-node conflicts and overlapping-BSS activity. Beyond a point aggregate capacity saturates rather than scaling with cell count. [29][30][42] Area spectral efficiency keeps rising only where the interference between cells falls as fast as the cells shrink — so the binding question is not how small a cell can be made, but how well it can be isolated from its neighbors, which is where a directional pattern earns its place (developed in §3).

The practical route to a higher ρ is spatial isolation: shrink the footprint of each radio's energy so the same channel can live again a short distance away. Isolation is not beamwidth alone. Main-lobe beamwidth and gain define where a radio serves, while side-lobe level, front-to-back ratio, cross-polarization discrimination, mounting geometry, and multipath determine how well it isolates. A narrow 3-dB beam alone does not guarantee a re-usable cell. [34] Sectorized and directional antennas have always been a means to this end; the hyper-directional antenna is that idea pushed to its limit — a beam so tight it covers a single block of seats. In the terms of the equation, a hyper-directional antenna is fundamentally a ρ-multiplier — though, as the next subsection shows, that is not the whole of it.

Power: the trade a directional antenna refuses

The ρ-multiplier description is incomplete in a way that bears on a second piece of conventional wisdom. The standard remedy for co-channel interference is to turn power *down*: a lower-power access point covers a smaller cell, the cells overlap less, and the interference falls. In the terms of the equation this looks like buying re-use (ρ) — and it does. But it is not free. Power sets signal-to-noise ratio, and SNR sets spectral efficiency — the rate a link can sustain (§2.1; Appendix A). Turning power down to shrink the cell pulls SNR down with it, forcing the radios to a lower data rate; the throughput lost on every link is material. The conventional fix trades η for ρ.

Cisco's John M. (Matt) Swartz — Cisco's distinguished engineer and a named inventor on the company's steerable-array patent [8] — frames the dilemma as a choice between "co-channel and no-channel": reduce power far enough to suppress the interference and you have throttled the very links you were trying to protect.

A hyper-directional antenna refuses the trade. It breaks the coupling between power and interference that forces the choice. Because the beam is confined to one section, main-lobe gain raises the desired signal ("S") delivered to that section, while side-lobe and back-lobe suppression holds the off-axis interference ("I")I it contributes to every other section low; both move SINR the same way — and with it the number of sections that can re-use a channel at once. High power no longer means high interference. The antenna can both run hot and stay isolated. With isolation, the equation moves the terms together. The re-use of (ρ) rises because adjacent sections no longer hear each other. Spectral efficiency (η) rises because a high-gain, high-power beam puts a strong, high-SINR signal on its target rather than the weak one a power-reduced omnidirectional cell would. Where the conventional remedy spends η to buy ρ, the directional antenna buys both. [32] That is the nuance a pure "ρ-multiplier" account misses, and it is part of why measured gains can exceed what re-use alone would predict. The next section asks why this is preferable to deploying more conventional access points; the remainder then examines the two families of antenna that pursue it.

Why Hyper-Directional, Not Simply More Access Points

Before comparing the two hyper-directional approaches, it is worth stating why the category exists at all. The alternative is not "do nothing". Instead, the alternative is conventional dense designs that preceded it: omnidirectional access points distributed through the bowl using proximate strategies; under-seat or handrail placements. The capacity model of §2 explains the difference. Adding conventional access points does little for the spectrum and spectral-efficiency terms (B and η). Spectrum is fixed and modern radios already run near the Shannon bound — so proximate designs push on re-use (ρ) by shrinking cells. Cells shrunk with omnidirectional radios still interfere. Every added co-channel access point is also a new source the others must hear and defer to. Capacity and contention rise together and the gain saturates. [29][30][42] A hyper-directional antenna raises re-use (ρ) a different way — by physical isolation rather than proximity. Hyper-directional antennas allow re-use in adjacent sections without the two copies hearing each other. It lifts the term with the most headroom without the interference penalty that simply adding radios incurs.

The operational case is documented. Proximate under-seat and handrail designs carry real, recurring costs: core-drilling and conduit, environmental-penetration, trip-hazard risk, and ongoing cleaning and maintenance around thousands of in-crowd enclosures. AmpThink's estimation of those costs were documented in an article published in Stadium Tech Report. Our own review of the cost of the under-seat systems we have deployed finds that the combination of maintenance costs and attributable lost revenues can reach into the millions of dollars per annum. The cost of proximate designs are most significant in arenas with retractable seating. [2][19] Overhead hyper-directional placement transfers much of the cost from an under-seat deployment. It re-uses existing roof structure, stays out of the fans' sightline, and is ideally serviced from catwalks. [18] Leveraging access points with multiple radios in a single package, device count falls accordingly. Overhead directional designs can often lower the access-point count of an equivalent under-seat plan — lowering both capital and operating cost. [10] These are the measurable advantages referred to in the introduction, and they anchor the emerging shift from "under the seat" to top-down. [19][1]

Spectrum and protocol are multiplicative with the antenna, not a substitute for it. A fair objection is that recent Wi-Fi gains — 6 GHz spectrum, OFDMA, BSS coloring — deliver capacity on any access point, making antenna gains marginal. The capacity model shows why this is a false choice: those features raise B (more spectrum) and η (and ease contention), while the antenna raises spatial re-use, ρ. They multiply rather than replace one another — protocol features schedule and de-conflict transmissions in time and frequency, while the antenna provides the physical isolation that lets the same channel live again a section away, which adaptive carrier-sense can approximate but not manufacture. [29][32] The strongest designs use both.

Hyper-directional is a bowl technology, not a whole-venue one. Its long, narrow beams are built to cover tiered seating from a distance; concourses, clubs, suites, and back-of-house are still served by conventional access points. A complete venue design is therefore a hybrid. Hyper-directional antennas covering the seating bowl, conventional coverage everywhere else. Stating this is not a concession; it is the deployment model, and it bounds where the comparison in this briefing applies.

Two Families of Hyper-Directional Antenna

Both families exist to raise ρ by the same logic: confine each radio's energy to a tight cone aimed at one block of seats so that co-channel radios are effectively invisible to one another. They reach that isolation by two completely different mechanisms:

- **Active electronic beam control (RF filtering / side-lobe suppression).** A powered, directional patch-array access point selects among preset beam states and aggressively suppresses its side lobes electronically, through switching and tapered feed networks. This is the Cisco approach (the "Marlin" 9130-STA/9104 and its Wi-Fi 7 successor, the CW9179F) and the Everest / Fortinet approach. [1][7]
- **Passive optical beam control (RF-lens refraction).** A passive spherical dielectric lens bends — refracts — radio waves the way a glass lens bends light, forming many independent pencil beams from one mounting point. This is the MatSing approach. [9][12]

Both families deliver roughly three times the reach of a conventional stadium access point — on the order of 175–200 feet — and both cut device count and avoid in-bowl construction effort. They differ in almost everything else. [1][2]

Attribute	Active electronic (Cisco; Everest / Fortinet)	Passive RF-lens (Extreme Networks + MatSing)
RF method	Directional patch array with switched, tapered feed networks; electronic beam shaping across preset states with side-lobe suppression ("RF filtering").	Gradient-index dielectric (Luneburg) lens; beams formed by refraction — no active electronics in the antenna.
Where the beam is aimed	By the array geometry, amplitude tapering, and switching among preset states — set in the electronics.	By the physical position of each feed on the lens surface — set mechanically. (Minimizing side lobes via antenna design)
Radios	A full radio set inside every access point.	16 access points co-located with the lens, in two side cars of eight; the lens itself is passive, each AP driving one beam.
Form factor	One focused-beam AP per section, rigging-mounted.	One large spherical lens plus co-located AP side cars, on a single high mount.
Reach	~200 ft (~3× a traditional AP).	~175–240 ft from a central mount.
Maturity (Wi-Fi)	Cisco CW9179F shipping at scale; Everest now Fortinet FortiAP-HD / Wi-Fi 7 narrow-directional.	Cellular-proven for years; purpose-built Wi-Fi 5/6Ghz lens introduced in 2026 in partnership with Extreme Networks new AP5022FX access point model.

Lineage of the Active Approach: Everest → Fortinet, and Cisco

Everest: Montreal patent to Fortinet

The active hyper-directional category was effectively created by Everest Networks. The company began life under a different name — **Tembo Systems, Inc.** — whose foundational filing, "Wireless system with configurable radio and antenna resources" (priority October 2013, inventor Arbinder Singh Pabla), was later assigned to Tembo and then renamed to Everest Networks in 2018. [3][5]

Crucially for understanding the technology, Everest's core engineering IP is anchored in Montreal. Its key stadium patent — US 10,575,188, "Overhead wireless access point deployments for arenas, stadiums and other venues" (issued 25 February 2020) — names three Montreal-area inventors: Samer Abielmona (Montréal), Van Hoang Nguyen (Senneville), and **Jean-François Frigon** (Brossard). [10] Frigon is a professor and Chair of the Department of Electrical Engineering at **Polytechnique Montréal**, with a long research record in MIMO beamforming and dynamically reconfigurable antenna radiation patterns — the academic foundation directly aligned with the directional, software-controlled antenna systems Everest commercialized and Fortinet now ships. He is also a named inventor on Everest's multi-port / column-sharing beamforming patents. [14][15]

Everest's product reached the market through Panasonic. In April 2017, Panasonic Enterprise Solutions launched the Everest solution as, by independent (Miercom) testing, the first true four-radio 802.11ac Wave 2 access point built for ultra-high-density venues — pairing those radios with patented, software-controlled directional antennas in an IP67 enclosure and claiming roughly one-third the hardware of legacy systems. [16] The first units advertised ~175 ft reach with four antennas (three 5 GHz, one 2.4 GHz) in one enclosure. [2]

Everest won a marquee NFL reference at Lincoln Financial Field (Philadelphia Eagles) in 2017, followed by MetLife Stadium and the Florida Panthers' arena. After a Panasonic reorganization it became a standalone company, but it never built out the enterprise-management features large buyers expect and faded from public view. [2][15]

In May 2025, Fortinet acquired Everest, folding its patented directional technology into the FortiAP-HD line and marketing it as "SpotLight" directional coverage; Everest has since filed Wi-Fi 7 narrow-directional access points (e.g., the AP750NRE) with the FCC. [7][17]

Cisco: Enterprise Ready take on Everest

A popular account holds that Everest's Eagles win prompted Cisco to build a rival. AmpThink's understanding is different: unable to scale commercially, Everest sought a buyer, and Cisco was among those that looked. That part is not independently documented and is offered as AmpThink's read, not established fact. What is verifiable is the patent position, and it is the substantive point.

Everest's patents do not claim the two things that make the product work — the RF filtering and the beam directionality it produces. What they claimed was the surrounding architecture: multi-radio resource provisioning and management (US 2015/0110212, originally filed by Tembo Systems under Arbinder Singh Pabla), [3] overhead access-point deployment for arenas and stadiums (US 10,575,188), [10] and multi-port, multi-beam antenna structures that share radio ports across columns (for example US 11,621,755). [15] With directionality and RF filtering unclaimed, a competitor could implement its own side-lobe suppression without colliding with that portfolio — the verifiable basis for building rather than buying.

Cisco's commercial success with the product is material. Cisco's self-reported unit sales exceed 23,000 stadium-antenna access points — the public cumulative figure, released after model-level data had been provided but without breaking the total out by model, so it cannot be attributed to any single product. [19] Its first "Marlin" 9104 antenna shipped in 2021, followed by the Wi-Fi 7 CW9179F, which adds a 6 GHz radio (1,200 MHz of new unlicensed spectrum) and runs indoors or outdoors under FCC Part 15 Subpart E in the US. [1] Cisco's own mechanism is protected by US 11,024,961, "Electronically steerable antenna array" (granted June 2021), whose named inventors include **John M. Swartz**, Fred J. Anderson, Paul J. Stager and others. [8]

Lineage of the Passive Approach: Extreme Networks and Matsing

MatSing was founded in 2005 in Singapore by Dr. Serguei Matitsine together with his sons Leo and Michael Matytsine — the name combining "materials" and "Singapore." The family's breakthrough was a new lightweight dielectric metamaterial that made large Luneburg lenses manufacturable; the company is now headquartered in Irvine, California, with offices remaining in Singapore. [40][11]

Its lens patents — for example US 2017/0040687, "Lens based antenna for super high-capacity wireless communications systems" (applicant MatSing; inventors Anthony DeMarco, Serguei Matitsine, Leonid Matytsine) — describe stacked cylindrical lens discs combined with feed elements to produce multi-beam sector coverage from a single passive structure. [13] The technology proved itself first in cellular networks (carriers including AT&T, Verizon, and T-Mobile, and Facebook's rural "SuperCell" concept) before being adapted for Wi-Fi. [9]

At MWC Barcelona 2026, MatSing introduced the MS-16.16W45, its first purpose-built 5/6Ghz Wi-Fi lens antenna — sixteen independent beams, 4×4 MIMO each, serving thousands of users from one mount. Each lens forms about sixteen directional sectors, so the number of mounts follows the bowl's geometry rather than a per-section antenna count — far fewer mounts than an active design, each distributed for clear line of sight. The platform is proven at venue scale in cellular service: MatSing's neutral-host cellular DAS at Allegiant Stadium uses about sixty lens antennas (a cellular, distributed-antenna-system deployment for the carriers — cited here only as evidence the lens hardware works at stadium scale, not as a Wi-Fi antenna-count comparison). Two constraints apply: it is new at Wi-Fi scale, and it depends on suitable high mounts with clear lines of sight. [6][39]

Because the lens is passive, the Wi-Fi itself comes from conventional access points feeding it. MatSing has partnered with Extreme Networks to provide the wi-fi radio infrastructure to power the lens antenna solution. In the integration developed with Extreme Networks, sixteen AP5022FXs are co-located with the lens in two side cars each with eight APs. The kit - as supplied - powers each AP over a discrete PoE cable; the lens supplies the beam-forming physics while the side cars supply the radios. AmpThink would add a co-located switch to consolidate the cluster onto a single fiber pair with remote DC power (§9). [20][41]

Mechanism of RF Control

Active electronic control — Everest and Cisco / Fortinet

In the active family, beam control lives in the antenna's electronics and geometry, and **Everest got there first**. In 2017 it shipped a software-controlled, highly directional antenna system that paired multiple simultaneous radios in one enclosure — the first hyper-directional Wi-Fi product, four years before Cisco's Marlin. [2][4]

The shared principle is straightforward: the access point carries a narrow-beam array and, alongside it, a wider-beam array. Switches route the radios to whichever pattern is needed, and the narrow beam is sharpened by feeding the outer radiating elements with deliberately reduced power so that side lobes are deeply suppressed.

Cisco entered the category four years later with a clean-sheet design of its own. Its mechanism patent describes the method precisely: unequal power dividers and attenuators taper the array so that side lobes sit roughly 30 dB below the main lobe — low enough that two of these access points mounted within about fifty feet of each other cannot hear one another's transmissions. [8] Cisco's team likens the filtering to "earmuffs" on the antenna, cutting the sideways leakage so the radio can concentrate power in the one direction that matters and reach farther. [1]

Passive optical control — Extreme with MatSing

MatSing removes the electronics from the antenna entirely. Its core element is a **Luneburg lens**: a sphere built from many concentric layers of dielectric metamaterial whose refractive index varies with radius. A single radiating element placed against the surface of the sphere emits energy that the lens focuses into a tight, high-gain beam exiting the opposite side. Because the sphere is symmetric, many feeds can be arranged around it, each producing its own independent beam aimed wherever that feed sits — all from one passive lens. Each feed is excited by its own access point, mounted alongside the lens rather than at a distant head-end (see §9–10). [12][13]

The underlying idea is old: Rudolf Luneburg described the lens in 1944, built on Maxwell's fish-eye lens of 1854, but the dielectric materials were historically too heavy for anything but shipborne radar. MatSing's contribution was a lightweight patented metamaterial that makes large lenses practical. [11]

Summary

The control philosophies are therefore opposites. In the active family, RF control is electronic and software-selectable — the array switches among preset beam states (narrow and wide patterns, and discrete steering offsets) formed by switching and tapered feed networks, rather than the arbitrary, continuous beamforming of a fully phased array. In the lens family, RF control is static and physical — fixed by the optics of the lens and the placement of each feed. The lens cannot re-steer a beam in software, but it also has nothing in the radiating element to power, tune, or fail.

Radiation Patterns and Isolation

The radiation pattern — main-lobe width together with side-lobe and front-to-back suppression — is the clearest expression of each mechanism. It is the whole pattern, not beamwidth alone, that determines isolation: a narrow main lobe with high side lobes is not a re-usable cell. The active arrays produce a narrow electronically defined sector with deeply suppressed side lobes. The ~30 dB taper in Cisco's design is what lets densely packed APs share channels without colliding. [8] The lens produces a set of fixed pencil/fan beams, each typically very narrow in azimuth, physically separated from its neighbors by the geometry of the sphere so that beam-to-beam coupling is inherently low. [12]

	Active Electronic Array	Passive RF-lens
Main Lobe	Narrow, electronically tapered sector aimed at one seating block.	Narrow fixed pencil/fan beam per feed; many beams per lens.
Side Lobes	Aggressively suppressed (~−30 dB) by unequal power weighting.	Inherently low — beams are physically separated by the lens.
Steerability	Beam can be re-shaped / re-pointed in software and switched narrow↔wide.	Fixed at install by feed position; re-aim means moving a feed.
Isolation Goal	Co-located APs (within ~50 ft) do not detect each other.	Adjacent beams stay orthogonal by optics, not by power control.

Practically, both achieve the low co-channel contention a bowl demands, but the active array buys it with electronics it can later re-tune, while the lens buys it with geometry that is locked in once the lens is hung.

Beamwidth by band and product

Published beamwidth-by-band figures are uneven across the field. Cisco documents them precisely for both generations of its stadium antenna. Everest/Fortinet publish no beamwidth in degrees in any source — press, vendor materials, or third-party releases; for Everest, the figures below were extracted directly from the radiation-pattern plots in an FCC certification filing (method noted beneath the table). Cisco's values are 3-dB azimuth × elevation, and several of its beams are software-switchable between narrow and wide. [26][27] Extreme's beams are also 3-dB azimuth × elevation and software-switchable between narrow and wide patterns leveraging the built-in antenna products in the AP5050D and AP5060D.

Product / antenna	2.4 GHz	5 GHz	6 GHz
Cisco C-ANT9104 (with Cat 9130 AP; Wi-Fi 6)	70° × 70° (wide only)	25° × 25° narrow 80° × 25° wide	— not supported

Product / antenna	2.4 GHz	5 GHz	6 GHz
Cisco CW9179F (Wi-Fi 7)	70° × 70°	35° × 35° (both slots)	35° × 35° narrow 70° × 35° wide
Extreme Networks AP5050D (Wi-Fi 6E)	70° × 70°	30° × 70° narrow 70° × 70° wide	30° × 70° narrow 70° × 70° wide
Extreme Networks AP5060D (Wi-FI 7)	70° × 70°	20° × 70° narrow 70° × 70° wide	20° × 70° narrow 70° × 70° wide
Everest → Fortinet (AP750NRE, Wi-Fi 7)	not published	≈24–35° ‡ (read from FCC plots)	not published
Extreme AP5022FX (Wi-Fi 7) with MatSing MS-16.16W45 (5/6GHz Wi-Fi lens)	— not covered	16 beams† 11° × 11° narrow	16 beams † 11° × 11° narrow

Notes. Cisco 9104: the two 5 GHz beams also steer 0°/10°/20° off-center; CW9179F: in wide mode the two 5 GHz beams splay 15° apart, and the 6 GHz slot runs 7 dBi wide / 12 dBi narrow. Cisco C9130AXE AP End of sale as of Dec 31,2026 ‡ Everest/Fortinet: no beamwidth in degrees is published in press, in Panasonic/Everest/Fortinet materials, or by third parties — all describe the antenna only as "highly directional" or "narrow-directional." The figures here were extracted from the radiation-pattern plots in the FCC certification of the current Wi-Fi 7 narrow-directional unit (FCC ID 2AGMR-AP750NRE): reading the half-power (−3 dB) points of the main lobe across the two principal-plane cuts gives ≈26°×31° for the U-NII-3 (upper 5 GHz) panel — peak gain 16.47 dBi — and ≈24°×35° for the U-NII-1/2a (lower 5 GHz) panel — peak 16.84 dBi — i.e., a focused beam of order 25–35° in 5 GHz, the low and high bands closely matched. The 2.4 GHz and 6 GHz panels are not in this filing. These values are read off plots rather than tabulated; but an FCC certification exhibit is a regulatory filing, not marketing material, and the plotted patterns are corroborated by the filing's own published gain figures (16.47 / 16.84 dBi) — so they are approximate — with an estimated read precision of about ±3° — yet rest on a more authoritative source than a vendor datasheet. † MatSing MS-16.16W45 operates 5.125–7.125 GHz as one continuous band group of sixteen 4×4-MIMO beams, so it spans upper-5 GHz and 6 GHz but not 2.4 GHz; Horizontal and vertical Beamwidth (8dB level) is 11°, horizontal and vertical Beamwidth (3dB level) is 7°.[6][7][28][55]

Two things stand out. The active antennas trade beamwidth for band: their tightest focus is in the mid-bands, while 2.4 GHz stays deliberately wide (70°) because it is a coverage, not a capacity, band — and the older 9104 actually held a narrower 5 GHz beam (25°) than its Wi-Fi 7 successor (35°), trading some focus for the simpler, more robust pattern set Cisco adopted in the CW9179F. Across the active products the 5 GHz focus converges on a similar window — Cisco's 25–35°, Extreme's 20-30° and Everest's ≈24–35° read from its FCC plots — roughly the angular width of one seating block. The lens, by contrast, applies the same narrow beam across its whole operating range, because the focusing is optical and essentially frequency-flat — a structural difference the band-by-band view makes visible.

Installation

The two families place very different demands on the building. The active approach scatters many self-contained access points around the canopy, roofline, or rigging to cover individual seating sections, each leveraging a single Ethernet cable (data + power), and comparatively light and independently serviceable. The lens approach concentrates coverage onto far fewer high mounts (catwalk, roof steel, scoreboard), where each large spherical lens covers a section of the bowl. Each lens antenna forms sixteen beams — equivalent sixteen directional APs — so the count follows the bowl's geometry. As an example, where the active design places one antenna per section, the lens design places roughly one mount per eight or sixteen sections, far fewer in total but each heavy and dependent on a clear, high sightline. The exact number is set by how the bowl divides into sections and by where each device can be physically mounted, not by a fixed antenna-for-antenna ratio.

Unlike many distributed-antenna deployments, the access points feeding the MatSing lens are not remote. The architecture co-locates the sixteen APs with the lens; in Extreme's architecture they are split across two side cars of eight, kept nearby such that each AP reaches its assigned feed ensuring low RF loss. Adding a co-located switch — which would let each cluster be fed by a single fiber pair with remote DC power rather than running a separate cable to every AP and fan — is a logical refinement that AmpThink sees as a potential efficiency gain, though it is not part of the kit as supplied. [20]

Installation Factor	Active electronic	Passive RF-lens
Mount points	Many — distributed around the bowl perimeter / canopy.	Far fewer — each lens mount covers about sixteen sections.
Per-unit weight / rigging	Lighter individual APs; standard rigging mounts.	Higher weight lenses; structural mounting and load review required.
Cabling	Power + backhaul to every AP location.	Consolidated cabling to each lens mounting. 16 PoE cables for APs.
Line-of-sight dependence	Per-section LOS, shorter throws, more forgiving.	LOS from lens mount to 16 seating sections served. Lens provides interference isolation between coverage areas.
Serviceability	Swap a single AP without affecting others.	Lens is passive, its two side cars (16 APs plus their fans) must be serviced — at dozens of locations.

The trade is between hundreds of simpler installs and dozens of more involved ones: the active family multiplies the number of mounting, power, and backhaul endpoints, while the lens family collapses the number of mounts — makes each position a somewhat more structurally demanding module that carries APs, side car chassis, and their cabling.

Radio Density

"Radio density" diverges because of where the radios physically live. In the active family, every access point carries its own full radio complement (Cisco's CW9179F, for example, integrates four client radios spanning 2.4 GHz, dual 5 GHz, and 6 GHz, plus IoT and GNSS), so the total radio count scales directly with the number of APs hung in the bowl. In the lens family the radios are not removed and not centralized — they are clustered. Sixteen access points sit alongside each passive lens — two side cars of eight APs in Extreme's kit — one AP per beam. The total radio count is comparable to the active family — sixteen APs per lens across the bowl's lens mounts — but the radios are grouped into per-lens clusters rather than spread across many independently sited APs. As supplied each AP is cabled individually; consolidating the cluster behind a co-located switch and shared backhaul is a refinement AmpThink would add (§9). [1][6]

	Active electronic array	Passive RF-lens
Radio location	Inside each AP, distributed through the bowl.	16 APs co-located with each lens, in two side cars of eight.
Per-unit radios	~4 client radios per AP (CW9179F: 2.4 / dual-5 / 6 GHz + IoT, GNSS).	16 APs (AP5022FX- 5/6GHz) each per lens, one per beam, each 4×4 MIMO
Scaling	Radio count rises with the number of per-section APs.	Comparable radio count, clustered into per-lens side cars.
Failure domain	One AP down = one section affected.	One AP down = one beam (APs are hot-swappable in the sidecar) one side car down = eight beams; loss of the mount = antenna coverage area.

Realized Capacity and the Client Side of the Link

Sections 2 through 6 establish the antenna's contribution to capacity: a hyper-directional pattern multiplies spatial re-use (ρ) by isolating each cell from its neighbors, and both families do this well. That gain is real. What turns it into delivered capacity is the other half of the link — the client side. A packed bowl is tens of thousands of independent, omnidirectional devices, and Wi-Fi is half-duplex and contention-based, so every device, access point and handset alike, must find the channel clear before it sends. [29] Characterizing that client side is part of the antenna evaluation, not separate from it. This section sets out what the antenna delivers there, what the radios and the venue deliver, and how the two families differ once both are counted.

Two isolations realize the gain

Spatial re-use is realized through two isolations acting together. The antenna secures the first; the crowd and the protocol govern the second; the delivered re-use follows both.

- **AP-side isolation** — co-channel access points do not hear one another. This is what the antenna delivers, through the ≈30 dB side-lobe suppression of §7–8. [8]
- **Client-side isolation** — co-channel clients, and the access points receiving them, do not collide or force one another to defer. This is set by client transmit behavior, the RF environment, and the MAC layer. [30][42]

Because both must hold, delivered re-use tracks whichever is tighter. The antenna can raise AP-side isolation almost without limit, so in a dense bowl the client side is the one worth characterizing — and it is the subject of the rest of this section. [31]

The client side in a packed bowl

Begin with what the client side is on game day. A full bowl concentrates tens of thousands of handsets in a small volume, each transmitting on its own schedule through a small, near-omnidirectional antenna, and the traffic is upload-heavy as the crowd posts and streams. In that setting the uplink — client to access point — is the binding direction. Measurements in a full stadium show devices running at or near their maximum permitted transmit power yet reaching only modest modulation, while the downlink, carrying the infrastructure's power advantage, has headroom to spare. [36][37][43]

Several things in this picture belong to the client and the channel rather than the antenna. Each device sets its own transmit power and runs its own clear-channel assessment, choosing to send or defer on what it hears — at thresholds only a few decibels above the noise floor. [31] Its omnidirectional uplink raises the noise floor. Multipath in a reflective bowl carries energy off-beam no matter how cleanly the main lobe is formed. [34] Together these set the client-side isolation, and they tighten as device count climbs through hidden- and exposed-node contention and overlapping-BSS activity. [29][30]

Directionality helps both directions

The directional antenna helps this client link in both directions, and it helps the weaker one most. By reciprocity an antenna's receive pattern equals its transmit pattern, so the gain that concentrates the downlink also concentrates the faint uplink it hears and rejects off-axis client signals by the margin it suppresses its own side lobes. [35][33] The infrastructure receiver compounds the effect. Lower noise figures and multiple combined receive chains — receive diversity and maximal-ratio combining — let an access point recover a usable signal from a client at, or beyond, the range at which that client could cleanly hear the access point. This specifically addresses the problem that power-limited uplinks in a stadium presents. [38][37] Both effects land on the receive side and both families of hyper-directional antennas are reciprocal, so both inherit them. Reciprocity lives in the antenna and the channel, not the whole link. Transmit powers and noise figures still differ at each end — so it narrows the asymmetry rather than erasing it, and only within the beam.

Directionality also shapes a client-owned behavior — roaming — favorably in a bowl. A seated spectator is stationary, so most roaming is unnecessary, and a sharp cell edge keeps the serving antenna decisively stronger than its neighbors across the seats it covers, so a device locks on and stays. When a hand-off is genuinely warranted, the clean signal differential between adjacent cells is the kind of input a client's roaming algorithm needs to choose well, where soft, overlapping coverage would leave it lingering. The roaming decision remains the client's with the antenna significantly reducing ambiguity. Though client scan policy, body loss, aiming error, and multipath can still blur or move the cell edge.

This feature is however conditional and depends upon a consistent overhead layout — our recommended best practice.

Where each contributor is controlled

Each contributor has a clear owner and a known lever:

Interference contributor	Primarily controlled by	Lever, and where it lives
AP-to-AP co-channel leakage	Antenna	Beam shaping and side-lobe suppression — the subject of §7–8.
Off-axis client uplink into a non-target mount	Antenna (receive) + MAC	Reciprocal receive directionality; OBSS/PD with BSS color.
Client deferral; hidden / exposed nodes	MAC / protocol	BSS coloring, adaptive carrier-sense threshold, RTS/CTS.
Uncoordinated client uplink timing	MAC / protocol	OFDMA with trigger-based, AP-scheduled uplink.
Excess client uplink power	MAC / client	Transmit-power control.
Reflective multipath	RF environment / siting	Mount geometry, line of sight.
Reflective multipath	Regulatory / design	6 GHz channel plan and width.

The levers that remain after the antenna live in the radios behind it and in the venue design — scheduled uplink and adaptive carrier sense in Wi-Fi 6/7, transmit-power control, line-of-sight siting, and 6 GHz spectrum.[29][32] Both families rely on these same mechanisms, so the client-side budget is largely common-mode between them: it moves the result for either approach together, and does not, on its own, decide active versus lens.

How it enters the evaluation

Two differentiators survive this common-mode factor: the active array can adapt — re-steering beams and retuning carrier sense, coloring, power, and scheduling per radio in response to a moving crowd or a localized hot spot — whereas the lens's beams are fixed at install and its only client-side levers are in the access points feeding it; and the quality of those mitigations rides on the chosen AP platform, which differs between the two approaches.

How the Differences Bear on Deployment

Both families deliver hyper-directional coverage. The differences that matter in practice track the building, the operating model, and tolerance for newer technology.

Conditions that suit the active electronic array (Cisco; Everest / Fortinet)

- The venue is a retrofit or lacks a few strong central overhead mounts with clear line of sight to all seats.
- Per-section flexibility matters — beams can be re-shaped and re-pointed in software, and a single failed AP affects only one section.
- Integrated enterprise management and security are priorities (Cisco's management stack; Fortinet's security fabric), or shipping-today Wi-Fi 7 at scale is required — Cisco reports its deploying tens of thousands of stadium-antennas (vendor-furnished total; model mix not broken out).
- Distributed redundancy and lighter, independently serviceable units are preferred over a few heavy central installs.

Conditions that suit the passive RF-lens (Extreme with MatSing)

- The structure offers a few strong, high, clear-line-of-sight mounts (catwalk, roof, scoreboard) that can carry heavy lenses.
- Minimizing the number of mounts is the priority — dozens of high positions rather than hundreds of distributed APs — and the lens itself is passive, even though each mount still carries side cars of APs (and their fans) to power and service.
- Coverage of the floor/pitch and premium seating from above is needed, and consolidating many per-section antennas into far fewer high lens mounts — about sixteen beams each — is attractive.
- Integrated enterprise management and security is served through Extreme Networks cloud or on-premise network management platforms.
- Adoption of Wi-Fi technology whose beams are fixed at install and not subject to re-steering after installation. Wi-Fi 7 is only available using Extreme Networks AP5022FX radio platform.

Summary. Active electronic arrays distribute intelligence and risk across many reconfigurable radios and tolerate awkward geometry; the passive lens concentrates coverage into far fewer — though more structurally demanding — mounts, each a lens plus its side cars of APs, and depends on the sightlines to support them. Neither is the better technology in the abstract. In our view the two are sides of the same coin, and the choice follows the venue — its commercial objectives (cost, sponsorship and brand visibility, aesthetic and ownership priorities) and the mounting structures it can actually offer — far more than any intrinsic superiority of one family over the other.

Limitations. One real constraint sits on each side, and we leave both standing. The active array multiplies powered radios — broadly one focused-beam access point per section, each needing power and a backhaul drop — so device count, installed cost, and management surface scale with the size of the bowl. The passive lens concentrates coverage onto a few heavy mounts that demand clear, sustained lines of sight — so a venue without suitable high structure cannot use it at all. Neither point is argued away here; each weigh on the choice.

Objections and Responses

The strongest case against hyper-directional Wi-Fi comes less from a competing hyper-directional vendor than from one with no such product. The primary objections, stated plainly and answered on the merits:

"The technology is unproven." The maturity is not uniform, and conflating the two approaches is an error. The active electronic approach is field-proven at scale — by Cisco's own (vendor-furnished) figures, hundreds of 9104 units shipped in the first year and the stadium-antenna line now totals tens of thousands of units, though Cisco did not break the total out by model; its Wi-Fi 7 CW9179F is shipping, and Everest's antennas have been in NFL venues since 2017. The passive lens is new at Wi-Fi scale but is cellular proven over years (§4, §6). The category is not unproven; one implementation is newer than the other, and the briefing says so. [1][19]

"6 GHz, OFDMA, and BSS coloring deliver the gains on any access point; the antenna is marginal." Multiplicative, not substitutes (§2–§3). The protocol and spectrum features raise B and η (and ease contention); the antenna raises spatial re-use, ρ, the term with the most headroom, through physical isolation that adaptive carrier-sense can approximate but not manufacture. The strongest designs use both. [29][32]

"Clients transmit at low power through poor antennas, so the link is uplink-limited, and the antenna only helps the downlink." Reciprocity places the antenna's gain on the uplink-receive path, and the access point's superior receiver compounds it, so the imbalance narrows rather than persists (§11.3). Stadium traffic is upload-heavy — which is exactly where the reciprocal receive gain is felt. [35][36]

"Narrow beams leave coverage gaps and break roaming." The seating bowl is covered by design and the rest of the venue by conventional access points, in a hybrid the briefing states outright (§3, §11.3). On roaming the premise is backwards for a bowl: spectators are stationary, so roaming should be minimized, and the hard cell edge help keep a device locked to its section, and when a hand-off is genuinely needed it gives the client's roaming algorithm a large, clean signal differential to choose well. Soft, overlapping coverage is what breaks roaming; a uniform directional design is the fix (§11.3).

"The beamwidths are read off a graph, and the capacity claim is unproven." The Everest beamwidths come from an FCC certification — a regulatory filing, more authoritative than a marketing datasheet — and are corroborated by that filing's own published gain figures (§8.1). The capacity argument rests on the peer-reviewed re-use literature and on documented deployments (§2), with its evidentiary limits stated rather than hidden (§11). [28][23]

Appendix A — The Shannon Limit and the Spectral-Efficiency Ceiling

This appendix expands on a claim in §2: that the spectral-efficiency term (η) is largely fixed because modern radios already operate within a few decibels of the Shannon bound. The supporting detail is collected in the End Notes of Appendix B.

The bound

A channel's spectral efficiency — the most data it can carry per hertz with arbitrarily low error — is fixed by the Shannon–Hartley theorem: $\boldsymbol{\eta = \log_2(1 + S/N)}$ bits per second per hertz, where S/N is the linear signal-to-noise ratio. (Equivalently, a channel of width W carries at most $C = W \cdot \log_2(1 + S/N)$ bits per second; η is that rate per hertz. W denotes a single channel here, distinct from the model's total bandwidth B.)[44] It is not an engineering target but an information-theoretic ceiling: no modulation or coding scheme, however clever, can exceed it.[45]

Why there is little headroom

Capacity rises only *logarithmically* with signal-to-noise ratio. At the high SNR a stadium radio enjoys over the block it serves, each further doubling of received power — an added 3 dB — buys only about one more bit per second per hertz, while the cost of that power, and the interference it creates elsewhere, is real.[46] Capacity rises linearly with bandwidth B, which is why more spectrum and wider channels help — but B is set by regulation, not by the radio. The first lever, η, is near its logarithmic ceiling; the second, B, is fixed by regulation — so both per-link terms are close to exhausted.

How close "close" is

The distance between a real system and the bound is the gap to capacity: the extra signal-to-noise ratio, in decibels, a practical link needs to hit a target rate compared with the theoretical minimum.[47] Two advances have nearly closed it. Capacity-approaching error-correction codes — turbo codes and low-density parity-check (LDPC) codes — operate within a fraction of a decibel of the bound in additive-noise channels, and both Wi-Fi 6 and Wi-Fi 7 use LDPC.[48] High-order modulation supplies the constellation density to exploit that coding: Wi-Fi 6 added 1024-QAM and Wi-Fi 7 added 4096-QAM — ten and twelve bits per symbol — usable only at the high SNR these directional links provide.[49]

What remains of the gap

The residual few decibels are structural, not slack a vendor can reclaim: a roughly 1.5 dB "shaping" penalty for using uniform QAM rather than a Gaussian-distributed input, plus OFDM cyclic-prefix, pilot, and preamble overhead and ordinary implementation losses — error-vector magnitude, phase noise, quantization.[50] These trim a few decibels off the ideal but do not move the ceiling.

How long radios have been there — and where the gains came from

The gap did not close overnight. Shannon set the bound in 1948, but for four decades practical codes sat several decibels short of it — Hamming codes, convolutional codes with Viterbi decoding, Reed–Solomon, and the concatenated schemes used in deep-space and early digital links all left visible headroom.[51] The leap came in 1993 with turbo codes, which reached within about a decibel of capacity, and was consolidated as LDPC codes were refined to within hundredths of a decibel and written into the mainstream standards — DVB-S2, then Wi-Fi from 802.11n, then 4G and 5G.[52] For practical purposes the per-link coding gap has been essentially closed since the mid-2000s. In that sense radios have been at the Shannon limit for roughly two decades.

Which raises the obvious question. If the spectral-efficiency term (η) has sat near its ceiling for twenty years, how did wireless throughput rise so dramatically over the same period — Wi-Fi from tens of megabits to multiple gigabits, cellular from 2G to 5G? The answer is the premise of this briefing: almost none of that growth came from η. It came from the other two terms.

Decomposed over the long run, the gains line up lopsidedly. Martin Cooper's accounting of the roughly trillion-fold rise in wireless capacity since 1900 assigns a factor of about 25 to more spectrum (B) and about 25 to better modulation and coding (η) — the whole of the per-link budget — while the remaining factor of roughly a thousand came from spatial re-use (ρ): dividing space into ever-smaller, more numerous cells so the same frequencies are used over and over.[53] A peer-reviewed accounting of 1950–2000 reaches the same shape, crediting spatial re-use with a ~2,700-fold gain against comparatively small contributions from spectrum and efficiency (§2.2).

Even the gains that look internal to the link were largely spatial. Multiple-antenna (MIMO) systems, introduced to Wi-Fi and cellular in the 2000s, raise a link's rate not by beating the per-stream Shannon limit but by running several spatial streams in parallel — capacity scales with the number of antennas, a spatial-multiplexing gain layered on top of the single-channel bound, not a violation of it.[54] The pattern is consistent: for half a century the way to move wireless forward has been to add spatial dimensions and re-use, not to extract more from a single channel that was already near its limit.

The consequence

History and theory converge here. With the spectral-efficiency term η pressed against its theoretical limit and the bandwidth B fixed by spectrum regulation, the capacity equation of §2 has meaningful headroom in only one term — spatial re-use, ρ. That is the analytical reason a hyper-directional antenna, which multiplies ρ, addresses the part of the problem still open rather than chasing gains where almost none remain.

Appendix B — End Notes

All technical and historical claims above are drawn from the sources and notes below.

[1] Coyle, P. "Cisco improves 'Marlin' hyper-directional antenna line with Wi-Fi 7 (CW9179F)." *Stadium Tech Report*, 10 June 2025. stadiumtechreport.com. https://www.stadiumtechreport.com/

[2] Coyle, P. "Hyper-directional gets hyper-competitive." *Stadium Tech Report*, 19 May 2025. stadiumtechreport.com. https://stadiumtechreport.com/feature/hyper-directional-gets-hyper-competitive/

[3] U.S. Patent Application US 2015/0110212 A1, "Wireless system with configurable radio and antenna resources." Inventor: Arbinder Singh Pabla; orig. assignee Tembo Systems, Inc.; later Everest Networks, Inc. (priority 20 Oct 2013). https://patents.google.com/patent/US20150110212A1

[4] Panasonic Enterprise Solutions. "Panasonic Introduces Transformational Ultra-High Density Wi-Fi Solution for Sports and Entertainment." PR Newswire, 20 April 2017. https://www.prnewswire.com/news-releases/panasonic-introduces-transformational-ultra-high-density-wi-fi-solution-for-sports-and-entertainment-300441748.html

[5] Google Patents / USPTO assignment record, US 2015/0110212 A1 — assignment to Tembo Systems, Inc. (28 Feb 2018) and change of name to Everest Networks, Inc. (18 April 2018). https://patents.google.com/patent/US20150110212A1

[6] MatSing MS-16.16W45 Wi-Fi 6E lens antenna press materials — 16 beams, 4×4 MIMO each, 5.125–7.125 GHz; introduced at MWC Barcelona 2026. https://www.matsing.com/news/press-releases/matsing-to-debut-new-high-capacity-wifi-6e-lens-antenna-at-mwc-barcelona/

[7] Fortinet. "FortiAP High Density Series" data sheet (Everest acquisition; "SpotLight" directional coverage). fortinet.com; and Tech Field Day, Mobility Field Day 13, 8 May 2025. https://www.fortinet.com/products/wireless-access-points/high-density

[8] U.S. Patent 11,024,961 B2, "Electronically steerable antenna array." Assignee: Cisco Technology, Inc.; inventors incl. Fred J. Anderson, Paul J. Stager, John M. Swartz et al.; granted 1 June 2021 (−30 dB side-lobe taper; ~50 ft AP isolation). https://patents.google.com/patent/US11024961B2

[9] Saunders, M. "MatSing supplies lens antenna tech for Biden inauguration." *Fierce Network*, 20 Jan 2021 (MatSing Luneburg-lens deployments in cellular and large-venue settings). https://www.fierce-network.com/wireless/matsing-supplies-lens-antenna-tech-for-biden-inauguration

[10] U.S. Patent 10,575,188 B2, "Overhead wireless access point deployments for arenas, stadiums and other venues." Assignee: Everest Networks, Inc.; inventors: Samer Abielmona, Van Hoang Nguyen, Jean-François Frigon; issued 25 Feb 2020. https://patents.google.com/patent/US10575188B2

[11] "Mat-Sing: Praise for the Lens Antenna." *EE Times*, 10 Feb 2021 (Luneburg lens, 1944; Maxwell fish-eye, 1854; metamaterial). https://www.eetimes.com/mat-sing-praise-for-the-lens-antenna/

[12] MatSing. "About MatSing — RF Antenna Lens" (lens refraction; multiple independent beams from one lens). matsing.com. https://matsing.com/

[13] U.S. Patent Application US 2017/0040687 A1, "Lens based antenna for super high capacity wireless communications systems." Applicant: MatSing, Inc.; inventors: Anthony DeMarco, Serguei Matitsine, Leonid Matytsine. https://patents.google.com/patent/US20170040687A1

[14] Frigon, J.-F. Faculty and Chair, Department of Electrical Engineering, Polytechnique Montréal; research profile and publications (MIMO beamforming, dynamic radiation-pattern diversity). polymtl.ca; Google Scholar. https://www.polymtl.ca/expertises/en/frigon-jean-francois

[15] Patents assigned to Everest Networks, Inc. (multi-segment / multi-port antenna systems; column-sharing beamforming), inventors incl. Van Hoang Nguyen, Jean-François Frigon, Samer Abielmona, Arbinder Singh Pabla. Justia Patents. https://patents.google.com/patent/US11621755B2

[16] Miercom comparative test commentary in Panasonic/Everest launch materials (first four-radio 802.11ac Wave 2 HD AP), 2017. https://www.prnewswire.com/news-releases/panasonic-introduces-transformational-ultra-high-density-wi-fi-solution-for-sports-and-entertainment-300441748.html

[17] FCC ID filings for Everest Networks AP750NRE (4-radio Narrow-Directional Outdoor Wi-Fi 7 access point) and AP750WRE (omni). fccid.io. https://fcc.report/FCC-ID/2AGMR-AP750NRE

[18] Kapustka, P. "New long-reach 'hyper-directional' wireless antennas provide revenue and operational benefits for stadium networks" (market report). *Stadium Tech Report*, 2025. stadiumtechreport.com. https://www.stadiumtechreport.com/feature/new-long-reach-hyper-directional-wireless-antennas-provide-revenue-and-operational-benefits-for-stadium-networks/

[19] "Cisco's sales success shows rapid market acceptance of the hyper-directional Wi-Fi antenna strategy." *Stadium Tech Report*, 6 March 2026. https://stadiumtechreport.com/feature/ciscos-sales-success-shows-rapid-market-acceptance-of-the-hyper-directional-wi-fi-antenna-strategy/

[20] Corroborated in public reporting by Stadium Tech Report on the " system combines 16 separate Extreme Wi-Fi 7 access points, housed in two hardware chassis, with a single MatSing lens antenna". https://stadiumtechreport.com/news/new-extreme-matsing-platform-brings-competition-to-hyper-directional-stadium-wi-fi/.irectional

[21] Shannon, C. E. "A Mathematical Theory of Communication," *Bell System Technical Journal*, 1948 (Shannon–Hartley channel-capacity limit, $C = B \cdot \log_2(1+SNR)$). https://en.wikipedia.org/wiki/A_Mathematical_Theory_of_Communication

[22] Cooper, M., "Cooper's Law" of spectral efficiency: wireless capacity per unit area has roughly doubled every ~30 months since Marconi (~10^{12}× overall), with spectrum (~25×) and modulation (~25×) minor next to spatial re-use (~10^3×). See ArrayComm/Cooper materials and discussion in J. Andrews et al. https://en.wikipedia.org/wiki/Cooper%27s_law

[23] Andrews, J. G., Zhang, X., Durgin, G., Gupta, A. "Are We Approaching the Fundamental Limits of Wireless Network Densification?" *IEEE Communications Magazine*, Oct. 2016 (~2,700× capacity gain 1950–2000 from spatial re-use; spectrum and efficiency comparatively small). https://arxiv.org/abs/1512.00413

[24] MacDonald, V. H. "The Cellular Concept," *Bell System Technical Journal 58*(1):15–41, 1979 (frequency re-use across cells; capacity scales as space is divided). https://en.wikipedia.org/wiki/Cellular_network

[25] Alouini, M.-S. & Goldsmith, A. J. "Area Spectral Efficiency of Cellular Mobile Radio Systems," *IEEE Transactions on Vehicular Technology* 48(4):1047–1066, 1999 (data rate per unit bandwidth per unit area as a function of re-use distance / cell size). https://doi.org/10.1109/25.775355

[26] Cisco Catalyst 9104 Stadium Antenna (C-ANT9104) Deployment Guide and Installation Guide (with Catalyst 9130 AP): 2.4 GHz 70°×70° (6 dBi); 5 GHz narrow 25°×25° (10 dBi) / wide 80°×25° (7 dBi); 5 GHz beam steering 0°/10°/20°; no 6 GHz. cisco.com. https://www.cisco.com/c/en/us/support/docs/wireless/

aironet-antennas-accessories/221788-catalyst-9104-stadium-antenna-c-ant9104.html

[27] Cisco Wireless 9179F Access Point Deployment Guide (beam-pattern table): 2.4 GHz 70°×70°; 5 GHz (both slots) 35°×35°; 6 GHz wide 70°×35° (7 dBi) / narrow 35°×35° (12 dBi); wide-mode 5 GHz beams splay 15°. cisco.com. https://www.cisco.com/c/en/us/support/docs/wireless/wireless-9179f-access-point/224989-cisco-wireless-9179f-access-point.html

[28] Everest Networks, Inc., FCC ID 2AGMR-AP750NRE ("AP750UNe," 4-radio narrow-directional outdoor Wi-Fi 7 AP), Antenna Specifications test exhibit, 2024: cross-polarized microstrip patch panels, two radios at 2×2:2 MIMO, peak gain ~16.5–16.84 dBi across U-NII-1/2a/3. The exhibit's radiation-pattern plots are not tabulated; the 3-dB beamwidths cited here — ≈26°×31° (U-NII-3) and ≈24°×35° (U-NII-1/2a) — were read from those plots (half-power points relative to the 16.47 / 16.84 dBi peaks). Panels assembled by Innovox, Montréal. fcc.report / fccid.io. https://fcc.report/FCC-ID/2AGMR-AP750NRE

[29] Cisco. "802.11ax: The Sixth Generation of Wi-Fi" (white paper): as client and AP density rise, channel contention and airtime from intra- and overlapping-BSS use degrade throughput (co-channel interference); OFDMA provides contention-free scheduled transmission to multiple clients on both uplink and downlink within a transmit opportunity. cisco.com. https://www.cisco.com/c/en/us/products/collateral/wireless/white-paper-c11-740788.html

[30] Barrachina-Muñoz, S. et al. "Hidden- and exposed-node problems in dense IEEE 802.11 networks": under CSMA/CA these become more frequent in crowded scenarios because carrier-sense range exceeds communication range, and overlapping BSSs enlarge each node's contention domain. arXiv:1607.08138. https://arxiv.org/abs/1607.08138

[31] Coverage and rate analysis of dense Wi-Fi using stochastic geometry: clients transmit at lower power than access points, so uplink viability can be the binding constraint; the clear-channel-assessment threshold is typically ~−82 dBm, so radios defer to signals only a few dB above noise. arXiv:1704.08152. https://arxiv.org/abs/1704.08152

[32] Wilhelmi, F. et al. "Spatial Reuse in IEEE 802.11ax WLANs": BSS coloring with OBSS/PD lets a radio apply a more aggressive carrier-sense threshold to inter-BSS frames (with transmit-power limits) to increase parallel transmissions; 802.11be adds parametrized spatial reuse for scheduled uplink sharing. arXiv:1907.04141. https://arxiv.org/abs/1907.04141

[33] Directional-antenna interference analyses: a directional pattern provides passive interference suppression as a function of off-axis angle, applying on receive as well as transmit (reciprocity); sectorization reduces co-channel interference in multi-client deployments. arXiv:1602.02718; sector-vs-omnidirectional deployment guidance. https://arxiv.org/abs/1602.02718

[34] Experimental characterization of sectorized antennas in dense 802.11 networks (MobiHoc 2009): spatial reuse is constrained by the antenna's receive radiation pattern, and several idealized directional-MAC design assumptions do not hold in dense, reflective (multipath) environments. https://www.researchgate.net/publication/221628556_Experimental_characterization_of_sectorized_antennas_in_dense_80211_wireless_mesh_networks

[35] Antenna reciprocity: an antenna's receive pattern, directivity, and gain are identical to its transmit pattern, by the Lorentz reciprocity theorem — so directional gain applies equally on transmit and receive. Balanis, *Antenna Theory: Analysis and Design*; antenna-theory.com; "Antenna (radio)," Wikipedia. https://en.wikipedia.org/wiki/Antenna_(radio)

[36] "Comprehensive Analysis of Cellular Uplink Performance in a Dense Stadium Deployment" (arXiv:2604.04371): on game day the uplink is fundamentally power-limited — user devices operate near their maximum permitted transmit power yet achieve only low MCS — a structural downlink/uplink link-budget asymmetry independent of congestion. https://arxiv.org/abs/2604.04371

[37] U.S. Patent 9,860,855 B2, "Systems and methods for wireless link balancing in wireless networks": access-point–client links are inherently asymmetric in link budget, and a beyond-range client can remain audible to the access point even where it cannot cleanly receive the access point. https://patents.google.com/patent/US9860855B2

[38] Receive diversity and combining at the infrastructure end (multiple receive chains, maximal-ratio combining, beamforming to raise received SINR), long standard in cellular uplink. "Downlink and Uplink Transmission," ScienceDirect Topics; LTE dual-receive-antenna baseline. https://www.sciencedirect.com/

[39] MatSing neutral-host cellular DAS at Allegiant Stadium (integrated by DAS Group Professionals): about sixty multibeam lens antennas covering the seating bowl and venue for the major carriers, including C-Band 5G. This is a cellular distributed-antenna-system deployment, cited only as evidence the lens platform operates at venue scale — not as a Wi-Fi antenna-count comparison, for which it would not be a valid basis. AFL Wireless / DGP. https://aflwireless.com/matsing-lens-antennas-allegiant-stadium/

[40] Matytsine, L., interview, *Pulse 2.0*, 2026: MatSing founded 2005 in Singapore by the Matitsine family; lightweight metamaterial Luneburg lenses. https://pulse2.com/

[41] Extreme Unveils Industry's First Multi-Beam Wi-Fi Solution, Delivering Unprecedented Capacity and Performance in Large Venues https://investor.extremenetworks.com/news/news-details/2026/Extreme-Unveils-Industrys-First-Multi-Beam-Wi-Fi-Solution-Delivering-Unprecedented-Capacity-and-Performance-in-Large-Venues/default.aspx

[42] Throughput degradation with client count in dense 802.11: aggregate and per-client throughput fall as the number of contending stations rises. arXiv:1406.3147. https://arxiv.org/abs/1406.3147

[43] U.S. Patent 10,021,655 B2, "Systems and methods for wireless link balancing in wireless networks" (related to US 9,860,855): infrastructure-side link-budget and power-balancing methods for asymmetric AP–client links. https://patents.google.com/patent/US10021655B2

[44] The relation $C = W \cdot \log_2(1 + S/N)$ is the Shannon–Hartley theorem for a channel of bandwidth W; dividing by W gives the per-hertz form $\eta = \log_2(1 + S/N)$ used here. C. E. Shannon, "A Mathematical Theory of Communication," Bell System Technical Journal, 1948. S/N is a linear power ratio; in decibels, $SNR(dB) = 10 \cdot \log_{10}(S/N)$.

[45] Capacity is an achievability-and-converse result: rates below C are achievable with arbitrarily low error given sufficient coding length, and rates above C are not achievable at any error rate. It bounds every possible scheme, not merely those tried so far.

[46] At high SNR, $\log_2(1 + S/N) \approx \log_2(S/N)$, so capacity grows by about one bit per second per hertz for each 3 dB (factor-of-two) increase in S/N — sharply diminishing returns from raising power. Capacity is, by contrast, linear in bandwidth B, the motivation for 6 GHz and wider channels.

[47] The "gap to capacity" (or SNR gap, Γ) is the additional signal-to-noise ratio, in dB, a real modem requires to reach a given spectral efficiency relative to the Shannon minimum; it bundles coding sub-optimality, modulation and shaping loss, and implementation margin.

[48] Turbo codes (Berrou, Glavieux & Thitimajshima, 1993) first demonstrated operation within roughly 0.5–1 dB

of capacity; carefully designed irregular LDPC codes (Gallager, 1962; Chung, Forney, Richardson & Urbanke, 2001) have been shown within about 0.0045 dB of the additive-white-Gaussian-noise bound. IEEE 802.11ax (Wi-Fi 6) and 802.11be (Wi-Fi 7) both specify LDPC coding.

[49] 802.11ax introduced 1024-QAM (10 bits per symbol) and 802.11be introduced 4096-QAM (12 bits per symbol). Each step up the constellation requires markedly higher SNR and lower error-vector magnitude to demodulate — precisely what a long-throw, high-gain directional link delivers to its target section.

[50] Even with ideal coding, uniform QAM sits about 1.53 dB short of a Gaussian-distributed input (the asymptotic "shaping gap"); constellation shaping recovers part of it. Further losses come from OFDM overhead (cyclic prefix, pilot subcarriers, preambles) and radio implementation (transmit EVM, phase noise, ADC/DAC quantization). Together these account for the "few decibels" of real-world margin referenced in §2.

[51] Pre-turbo error correction advanced steadily but left headroom: Hamming codes (1950), convolutional codes (1955) with Viterbi decoding (1967), Reed–Solomon codes (1960), and Forney's concatenated codes (1966). Even the strong concatenated codes used on deep-space and broadcast links remained roughly 2–3 dB from the Shannon bound into the 1980s.

[52] Turbo codes (Berrou et al., 1993) were the inflection. LDPC codes (Gallager, 1962; rediscovered by MacKay & Neal, mid-1990s) entered deployed standards over the following decade — DVB-S2 (2005), Wi-Fi from 802.11n (2009), and 5G NR data channels (2017) by which point near-capacity coding was the norm in shipping radios.

[53] Martin Cooper's decomposition of the ~10^{12} rise in wireless capacity since ~1900: roughly 25× from additional spectrum, ~25× from improved modulation and coding, and the dominant ~10^{3}× from spatial re-use (smaller, more numerous cells). Consistent with J. Andrews et al. (2016), which credits spatial re-use with ~2,700× over 1950–2000. See also refs [22], [23].

[54] MIMO spatial multiplexing (Foschini, 1996; Telatar, 1999): in a rich-scattering channel a link's capacity grows roughly in proportion to the number of antenna streams, min (N_tx, N_rx), because independent data streams share the band over distinct spatial paths. This adds spatial dimensions to the link; it does not raise the per-stream Shannon limit.

[55] MatSing MS-16.16W45 datasheet - https://www.matsing.com/wp-content/uploads/Product_Documents/MS-16.16W45/MS-16.16W45_Datasheet_V05312026.pdf