WHITE PAPER

Next-Generation FWA Technology Primer

Contents

Abstract

Next-generation fixed wireless access (ngFWA) is an entirely new way of designing and operating large-scale FWA networks. At Tarana, we took a clean sheet approach to designing a platform that offers unprecedented coverage, speed, capacity, non-line-of-sight operation, and reliability in the harsh RF environments common to outdoor deployments. Tarana Gigabit 1 (G1) is the first commercial ngFWA solution and sets a new standard for all other FWA products.

Understanding fully G1's unique capabilities created over the course of the Tarana team's dozen years of wireless technology innovation — which also taps literally decades of team members' wireless R&D efforts prior to Tarana's formation — requires starting with the fundamentals of how radios work, along with comparing and contrasting G1's capabilities as the first instance of ngFWA relative to legacy FWA approaches.

This document provides a grounding in relevant radio principles, next-generation radio design, and the benefits delivered by G1's game-changing capabilities.

Introduction

All wireless transmissions require basic elements, such as radios and antennas, but more is required for optimal performance. Fixed wireless access (FWA) applications present unique challenges along the dimensions of capacity, reliability, non-line-of-sight performance, and interference cancellation that require going beyond simple radio technology.

RF Challenges in Fixed Wireless Access

Independent from the wireless technology applied, the following challenges arise in the fixed wireless RF environment that can degrade radio operation:

- Obstructions
- Interference
- Multipath
- Motion in the channel
- Demand for symmetric link speeds

These challenges affect FWA deployments in unique ways. Until now, two main technologies have been repurposed to deploy FWA: 3GPP (i.e., 4G, 5G) mobile networks and Wi-Fi. But neither technology was designed or optimized for FWA applications and neither address the key FWA challenges.

Mobile networks were designed to perform well in non-line-of-sight (NLoS) and large-scale environments, but they lack end devices powerful enough to model and optimize the RF channel with sufficient precision and control. They also lack the capacity to handle the higher data usage and peak rates typical of broadband subscriber demand. Further, mobile networks are designed for operation in costly licensed spectrum, resulting in added costs necessarily passed on to customers. Finally, while mobile networks might be able to mitigate interference to some degree, they can't actively cancel it, which prevents providers from using unlicensed spectrum for additional cell capacity.

Wi-Fi offers greater capacity, but it is optimized for best-efforts shared access over indoor LANs with mobile clients rather than for large-scale outdoor FWA deployments. Typically, Wi-Fi based FWA products are optimized exclusively for line-of-sight (LoS) operation, which limits their usefulness in the obstructed environments that are prevalent in outdoor wireless broadband applications. Although Wi-Fi is deployed in unlicensed spectrum, it primarily mitigates interference by forcing devices to change channels, which fails when a better channel is unavailable — which is more often than not the case.

It’s possible to address all of these challenges effectively, but doing so requires a fundamentally different approach to wireless system design and architecture, which we’ll explain next.

Terminology

Before beginning a deeper discussion about radios and their design, it’s important to understand how we use illustration and notation vocabulary in this document.

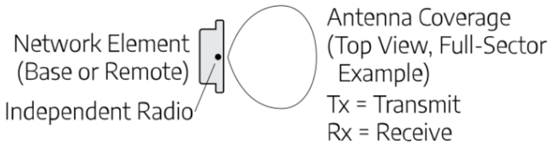

Figure 1: A simple network element with one radio chain

A radio element that concentrates connections at a vertical asset is described as a base, such as radios mounted on a communication tower by a wireless internet service provider.

A radio element that connects end-user locations is described as a remote, such as radios in remote nodes mounted on subscriber homes or buildings.

Logically independent radio chains (for example, transceiver + dedicated antenna structure) are shown as black dots in the element. Figure 1 describes a simple network element with one radio chain.

Antenna patterns are shown projected from a top view, but simplified for clarity to indicate only the relevant highest-gain aperture(s) rather than the complete antenna pattern from a standard 360° polar plot.

“Transmit” and “receive” are abbreviated as Tx and Rx.

Finally, we note the techniques and features unique to Tarana with the burst logo.

Obstructions

Radio networks need to transmit signals successfully from one radio to another, even with obstacles in the way. Obstacles include static objects such as buildings and geological features, and transient obstacles such as moving vehicles and rustling foliage. All of these are common in outdoor radio operation and network plans must take them into account for optimal results. To determine if a wireless link is obstructed, radio network designers consider the state of the Fresnel zone.

The Fresnel zone is an ellipsoid region of space surrounding the link, with the transmitting and receiving antennas at each focus. Because the distance from one focus to a point on the ellipsoid at the other focus is always the same regardless of where the point is on the ellipsoid, the ellipsoid represents the region in which a path is a specified length. In an RF link, if the path is long enough to allow an extra half wavelength of the RF wave, the reflected signal arriving from the edge of the ellipsoid is 180° out of phase with the direct signal and diminishes the total received signal (destructive interference). This ellipsoid is Fresnel Zone 1 (FZ1).

If the ellipsoid is large enough to allow an extra full wavelength in the reflected path, the direct and reflected signals arrive in phase and the total received signal is greater than the direct signal alone (constructive interference). This ellipsoid is Fresnel Zone 2 (FZ2). Fresnel zones are continuously numbered in this way, as long as they are useful to consider in link planning. Odd-numbered Fresnel zones are regions of destructive interference and even-numbered Fresnel zones are regions of constructive interference.

There are many ways in which the signal interacts with the environment such as reflection (bouncing off an obstacle or boundary), refraction (bending through a medium or boundary), and diffraction (bending around an obstacle or boundary). Fresnel zones—specifically FZ1—are also useful in determining the extent to which obstacles might diffract a signal.

The size of the radio aperture also determines the width of the beam, which in turn affects the propagation area of the signal. With all of this in mind, we can describe how a wireless signal propagates and spreads to fill a region of space, rather than existing as a single narrow beam.

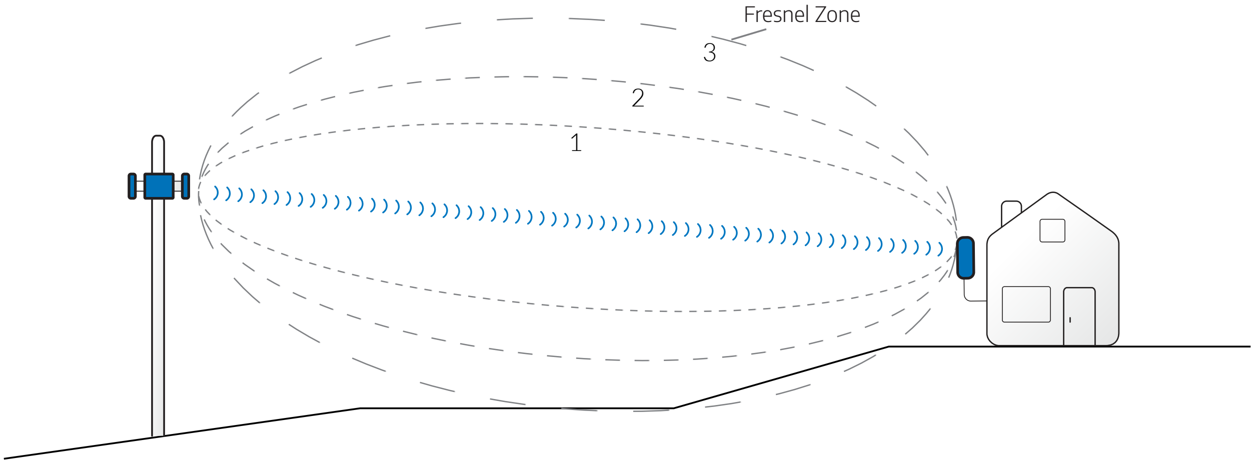

Figure 2: Wireless link with first three Fresnel zones shown (first and second are unobstructed)

A wireless signal with nothing within the first Fresnel zone is conventionally considered unobstructed and the link is referred to as line-of-sight (LoS). It's important to note that a wireless link can be visually line-of-sight, that is, a person standing at the remote node can see the base node, yet the link can be blocked to some degree in the Fresnel zone. Figure 3 illustrates this case on the left side.

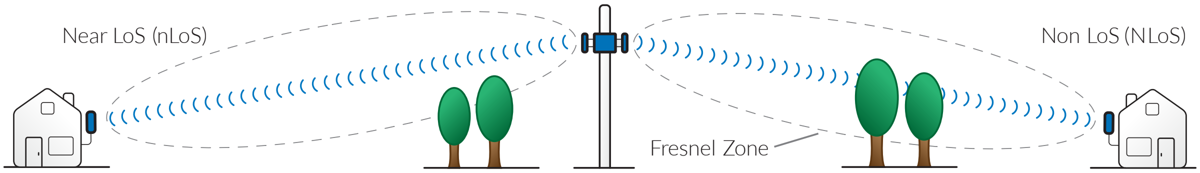

Figure 3: Example of partially obstructed first Fresnel zone (near-line-of-sight) and fully obstructed (non-line-of-sight) links

A wireless link that's partially obstructed is referred to as near-line-of-sight (nLoS). A wireless link whose primary signal path (illustrated by the blue line) is completely blocked is referred to as non-line of sight (NLoS), illustrated on the right side.

FWA deployments encounter many natural obstacles such as buildings and trees in most environments, so they tend to contain a large number of nLoS and NLoS links. An ITU study of links and obstruction over various distances finds that the majority of wireless links in a suburban deployment rapidly become NLoS as link distances increase (Figure 4).

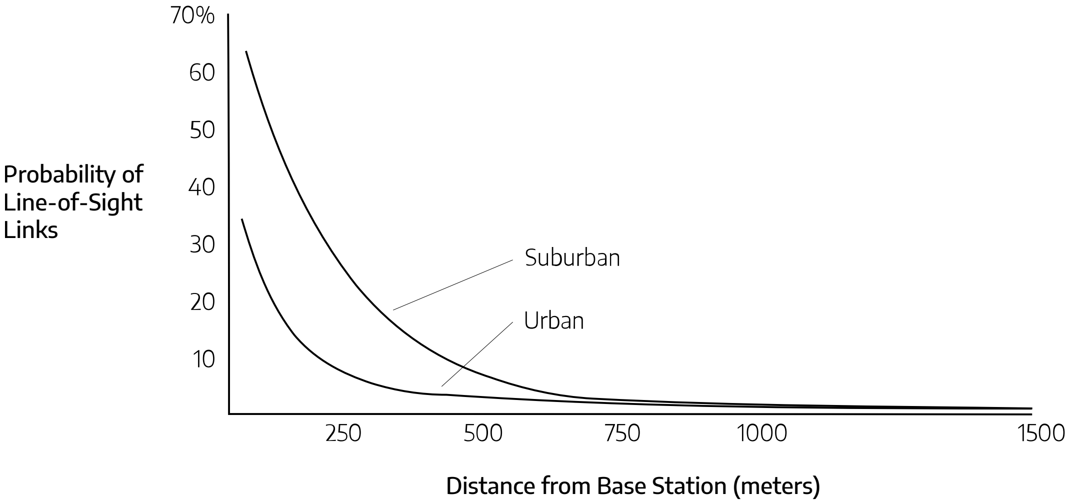

Figure 4: ITU-R survey showing 90% of households are NLoS after 350 meters
(Source: ITU-R M.235, based on extensive field testing)

Interference

Interference occurs when two or more radios transmit on the same or adjacent frequencies in proximity to each other. The result is a receiver that receives multiple unrelated signals and is unable to distinguish between them, leading to degraded performance because the receiver can't reliably decode the data.

Signal-to-interference-and-noise ratio (SINR) is a measure of signal quality that compares the power of the intended signal to unwanted interference and noise. SINR is a key metric for determining wireless link quality and overall performance. Optimizing SINR, either by increasing signal power or decreasing interference and noise, yields the benefits of more base station capacity, higher order modulation, higher peak rates, and greater spectral efficiency. All of these are requirements for FWA, with growing subscriber demand for more bandwidth.

In FWA, there are two common types of interference: self-interference, where an operator's equipment interferes with itself, and external interference, where the interference originates outside the operator's network.

Legacy Attempts to Overcome Obstructions and Interference

Multiple legacy FWA products have attempted to address the challenges of obstructions and interference. One popular technique is static beam switching (Figure 5).

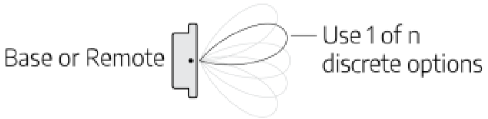

Figure 5: Static beam switching

Static beam switching doesn't model the channel accurately. Instead, it uses a discrete set of selectable signal amplitudes and phases that can shape the directionality of a link (typically by adding or removing physical antenna sub-elements electronically), and then chooses the antenna configuration that best matches the direction of the remote node. This process forms a beam, but the beam isn't optimized and doesn't make adjustments for the RF environment.

Static beamforming is simple, and its simplicity is an advantage because it can be implemented at the physical layer in 3GPP or Wi-Fi, independent of protocol. But its simplicity is a liability because it lacks flexibility and precision. The choices of directionality are limited and unintelligent, it can't create deep nulls to avoid interference, and it can't detect or respond to the nature of the multi-user channel environment.

These limitations are characteristic of legacy FWA. Next-generation FWA designs, such as Tarana G1, use more sophisticated techniques.

Key Principles of Next-Generation FWA Radio Design

Unlike legacy FWA, which attempts to repurpose non-FWA technology and applications such as mobile networks and Wi-Fi, ngFWA is designed specifically to address the challenges of large-scale, outdoor, fixed wireless networks. The ngFWA approach is fundamentally different from what has previously been available to operators. Any effective wireless broadband solution must be able to handle obstructions and interference sources that exist in outdoor environments. Not only can ngFWA meet these challenges, it can extract advantages from otherwise problematic environments, such as in the case of constructive combination of multipath signals.

Next-generation FWA provides greater capacity, link rates, NLoS performance, and reliability by using these technologies:

- Beamforming and RF nulls
- Spatial multiplexing
- NLoS performance
- Interference cancellation, also known as interference immunity
- Adaptation to dynamic motion in the channel
- Symmetric link budgets

This document examines these technologies and discusses how ngFWA uses them to deliver fast, high-capacity, reliable fixed wireless at scale.

Forming Beams and Nulls

Directivity is fundamental to radio design and performance. Beamforming is creating a directed transmission.

There are two approaches to beamforming: static beam switching and active beamforming. Static beam switching allows a system to select from among several static beam positions. Static beam switching stations can sense the RF environment, but can't make fine adjustments to the directionality of the beam. Instead, they must switch to a new preconfigured antenna pattern that's more appropriately directional.

In contrast, active beamforming senses the RF environment and adjusts the amplitude and phase of the independent radio chains to fine tune the beam directivity, enhancing transmission toward the receiver, and diminishing transmission in the direction of other co-network elements.

Static beam switching offers some benefits, but active beamforming offers much more. Using multiple independent radio chains, Tarana radio transmitters create beams (regions of transmission enhancement) and nulls (regions of transmission reduction) that continually adjust to maximize throughput and minimize interference. To do this, the devices adjust the phases of independent radio chains so their signals combine constructively to increase power in directions where it will increase overall network efficiency. At the same time, signals can combine destructively (canceling each other out and reducing power) in directions where the signal is undesired, known as RF nulls or RF nulling.

Greater overall performance gains can be achieved by applying these techniques on both transmitted and received signals to increase power while reducing interference and yielding a higher SINR. Optimizing SINR is crucial to increase capacity, peak rates, and spectral efficiency (Figure 6).

> **NOTE**: Spectral efficiency is the amount of information that can be transmitted over a specific amount of bandwidth. It is used to measure how efficiently a wireless system uses the available spectrum and is calculated by dividing the link capacity in bits per second by the spectrum bandwidth in Hertz. The result is a value measured in bits per second per Hertz (bits/Hz).

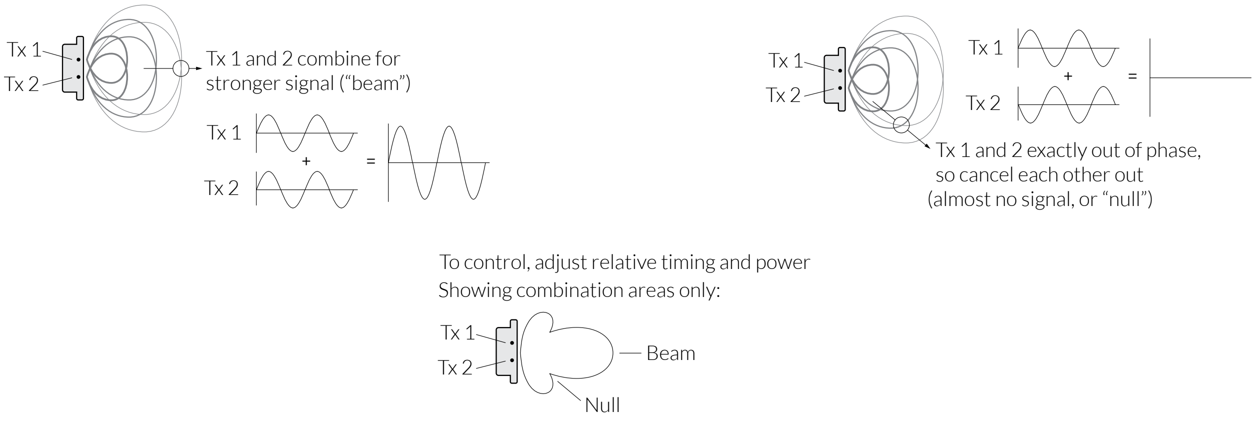

Figure 6: Radios can control the direction and power of the signal through beams and nulls

Beams and nulls can be formed at both ends of a link. However, to do this, radios at both ends must have sufficient radio chains, antennas, and processing power to model the channel accurately and take appropriate action. The high precision of G1's multi-dimensional digital model of the channel, in turn, enables much higher precision in beam directionality and deeper nulls. The more directional the beam, the higher the gain and stronger the signal. The deeper the nulls, the more unwanted interference can be avoided on Tx and rejected on Rx. All of this combines to yield higher SINR.

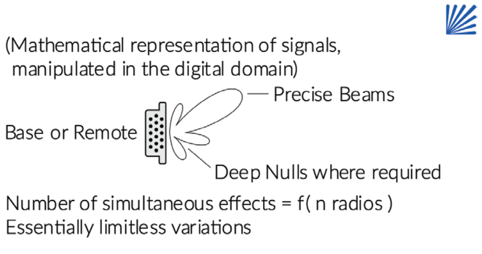

Figure 7: G1's precise digital beams and nulls enable higher signal strength and less interference

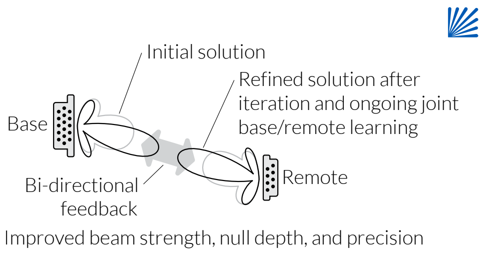

Figure 8: G1 nodes collaborate to quickly find and maintain an optimal solution

To improve our beamforming solutions further, G1 radios coordinate with each other in a learning process. Tarana's patented auto-convergent retrodirective beamforming (ACRB) provides continuous feedback on the quality of the solution to converge and maintain quickly and autonomously on optimal antenna patterns. With ACRB, each beamformer uses the best possible solution to improve RF cell capacity even at cell edges to a degree not possible with any other technology.

MIMO and Spatial Multiplexing

Multiple Inputs, Multiple Outputs (MIMO)

The next tool applied in many radio systems (including G1) is a technique known as MIMO which, as used in the industry, can refer to two different concepts.

MIMO's original definition is the use of different paths in the radio channel between base and remote to send more than one stream of information through the channel at the same time, thereby multiplying the channel capacity. Coding the streams on Tx allows them to be decoded separately on Rx. This is also referred to as spatial multiplexing and relies on the ability of the radios to create multiple streams (also called planes or layers) within the same frequency that can be distinguished from each other.

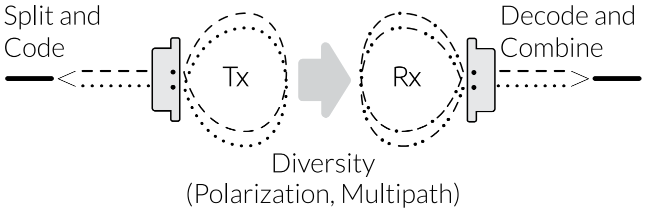

Figure 9: MIMO/spatial multiplexing uses multipath to create multiple unique streams of data

Spatial multiplexing commonly doubles channel capacity by taking advantage of vertical and horizontal antenna polarization to achieve sufficient channel diversity for at least two streams. Multiples of three or four can be achieved in channels with rich multipath diversity, which allows streams to arrive at different times and on different paths. Outside of dense urban areas with many reflections, outdoor systems are typically limited to two streams. Figure 9 illustrates the common scenario of two transmit and receive streams. Figure 10, on the following page, illustrates how multipath in the environment creates additional unique spatial streams.

Massive MIMO architecture is a variation of simple MIMO that involves large numbers of independent radio chains at a base station (as many as 128) to increase spectral efficiency in mobile networks. A mobile remote, like a smartphone, is necessarily limited in terms of radio resources because of tight power, weight, size, and battery-life restrictions. Adding more radio resources at the base is the only practical option for mobile networks in an effort to increase overall performance.

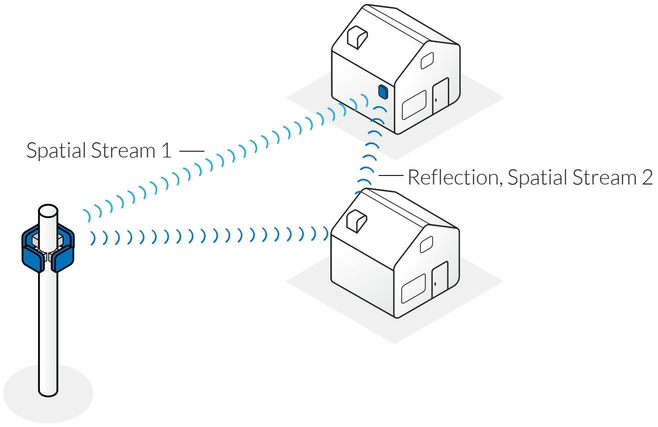

Figure 10: Multiple unique paths create 2 spatial streams

Unfortunately, the massive real-time signal processing load associated with that large number of independent radios requires a trade-off — specifically more coarse modeling of the channel, which means beams and nulls are less precise and less optimized, which results in a suboptimal SINR. Because of this, spectral efficiency gains achieved in practice from massive MIMO have been limited and come at the expense of massive computing costs and physical implementation complexity at the tower.

The massive MIMO approach relies on a powerful, complex base station and a simple, inexpensive remote as shown in Figure 11 and is common to many 3GPP/5G FWA implementations. While massive MIMO can yield real benefits, those benefits are targeted primarily at mobile wireless networks, not fixed networks.

Repurposed 3GPP technology that works well for mobile applications is simply not sufficiently optimized for the task of next-generation FWA and its unique requirements.

Distributed Massive MIMO (DMM)

Tarana G1's ngFWA architecture takes a novel approach to MIMO with an implementation that optimizes spatial multiplexing results through a more nuanced application of massive MIMO: distributed massive MIMO (DMM). The significant difference between our approach and 5G is a more even distribution of radio resources between base and remote endpoints, enabled by looser restrictions on size, weight, and power consumption for the remote node in the ngFWA application. This more even distribution allows G1 to build and continuously update in real time a much more precise model of the channel along all the essential dimensions of space, time, and frequency.

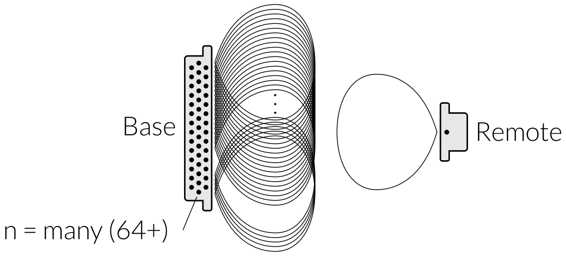

Figure 11: Massive MIMO uses complex base stations and simple CPEs

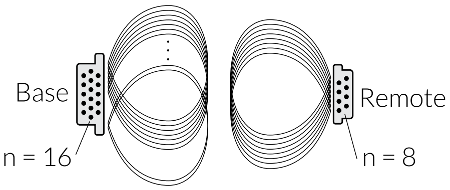

Figure 12: DMM distributes radio resources more evenly across both ends of the link

It's this improved knowledge of RF conditions that gives G1 greater capacity compared to other solutions and, in turn, drives higher spectral efficiency of up to 100 bps/Hz for a single cell site.

Multi-User Distributed Massive MIMO (MU-DMM)

Until now, this document has discussed MIMO in the context of one or more streams of data to a single user (node) scheduled over time, frequency, and space. This is referred to as single-user MIMO (SU-MIMO). Multi-user MIMO (MU-MIMO) also supports multiple spatial streams but does so for multiple remote nodes in the same channel simultaneously. The G1 approach takes this even further by applying MU-MIMO in a distributed massive MIMO implementation that takes advantage of many radio resources at both ends of the link and across multiple users.

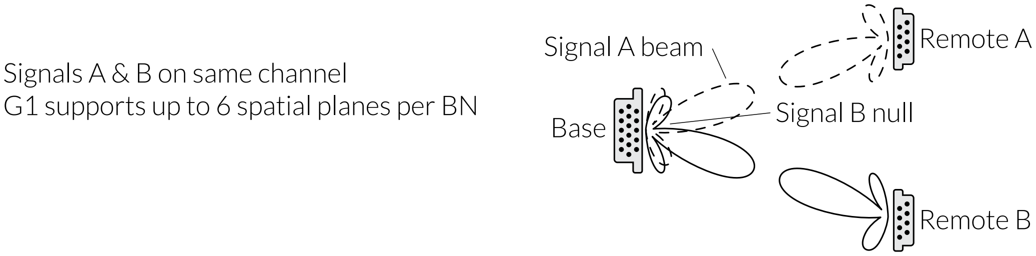

Figure 13: MU-MIMO (a.k.a. spatial planes) allows multiple devices to use the same channel

The high aggregate capacity demands of FWA, with many subscribers downloading gigabytes of data, requires MU-DMM to achieve the kind of gains necessary for efficient operation. G1 supports up to 6 independent spatial planes per base node operating in combination with spatial multiplexing. Spatial planes can be allocated as needed among multiple remote nodes. These resources are allocated as determined by G1's 4D scheduler.

The 4D scheduler operates in the four dimensions of time, frequency, space, and MIMO rank with up to 128 sub-bands, each 625 kHz wide, that can be allocated per remote node across 80 MHz of bandwidth. When taken in conjunction with 6 spatial planes, the scheduler has 768 allocation units it can assign for any given frame at transmission.

> **NOTE**: MIMO rank refers to how many spatial planes are allocated to a device. Usually there are two planes (MIMO rank = 2), but if the conditions are sufficiently adverse, then only one plane is possible (MIMO rank = 1).

Because channel characteristics — such as interference sources and multipath effects — vary by frequency, G1 further enhances precision by calculating transmit and receive beamforming solutions independently for each of the 128 sub-bands.

This level of granularity is possible only with the very precise digital beamforming that's a hallmark of G1. Based on the accuracy of the channel information, beams can be created with greater gain and deeper RF nulls. All of this improves SINR, spectral efficiency, peak rate, and capacity.

Figure 14 illustrates how the 4D scheduler allocates resources on a per-base node basis.

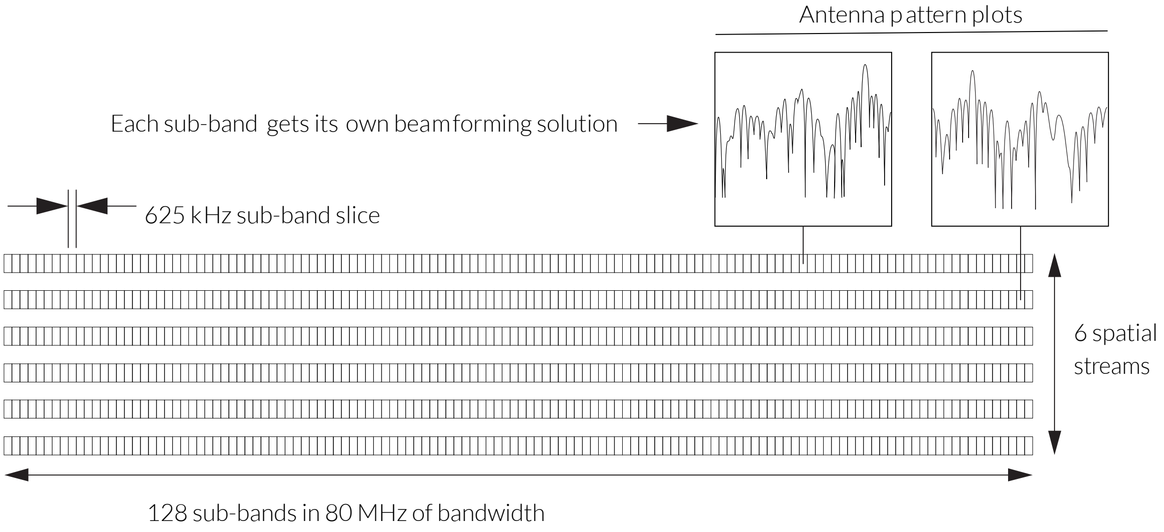

Figure 14: G1's 4D scheduler can allocate up to 768 unique solutions per base node (6 spatial streams x 128 sub-bands)

NLoS Performance

G1's highly precise and continuously perfected beamforming solutions enable several unique capabilities. One of these is essentially perfect multipath integration, which drives our unique level of nLoS and NLoS performance.

As mentioned previously, a challenge of outdoor deployments is the scarcity of pure line-of-sight links. The ability to operate reliably in n/NLoS conditions is, therefore, of critical importance for the feasibility of any FWA solution. It's simply not cost-effective to have base nodes within LoS of every subscriber, given the relative scarcity and nontrivial cost of sufficiently tall vertical assets and their associated backhaul connectivity. NLoS propagation dramatically lowers the range, link speed, and reliability for legacy fixed wireless technologies.

G1 overcomes these obstacles by combining the features described so far with an innovative approach to handling multipath. Multipath, which occurs when signals bounce off or bend around obstacles (i.e., reflection and diffraction) and take multiple paths from the transmitter to the receiver, is usually considered a negative rather than a positive for wireless communications.

G1 can take advantage of a rich diversity of signal paths to find a way for a signal to get around an obstacle, by collecting multipath signals on receive — which have various angles of arrival, amplitude, phases, and delays — and combining them to recreate the original signal. With G1, individual beamforming solutions are computed on a per-sub-carrier basis, yielding perfect NLoS that combines all relevant multipath vectors. This Tarana innovation is called space time frequency adaptive processing (STFAP). STFAP divides the 80 MHz of bandwidth into 6,864 orthogonal frequency-division multiple access (OFDMA) beamlets per layer with their associated N-1 nulls per beamlet.

Unlike other FWA systems where RF energy from multipath is lost, G1 recaptures all available energy, delivering line-of-sight performance even in near- or non-line-of-sight conditions.

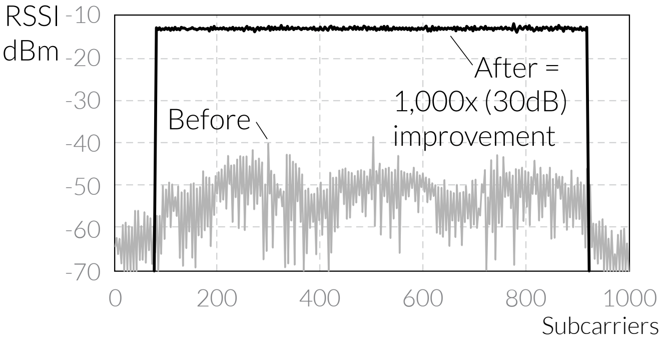

Figure 15: Essentially perfect multipath integration on Rx

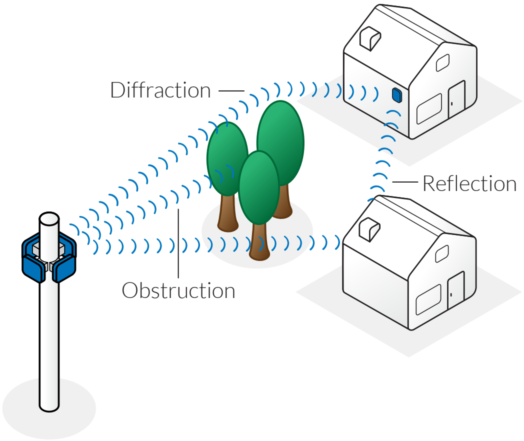

Figure 16: Leveraging multipath to get around obstacles (NLoS)

Interference Cancellation

Another significant aspect of ngFWA is interference cancellation. Interference, even with licensed spectrum, is an expected part of wireless communications. A well-designed ngFWA platform must significantly reduce high levels of interference that would otherwise impair link performance. In contrast, legacy FWA equipment has far fewer tools with which to address this problem, because of the many compromises it must make due to repurposed technology.

There are two types of interference that can affect service: self-interference and external, or out-of-cell, interference.

Self-Interference

Self-interference refers to an operator's network interfering with itself. This typically occurs when multiple sectors operate on the same channel or frequency, either on the same vertical asset, a nearby vertical asset, or both. Before G1, operators would attempt to manage this interference through separate channel use in which, for example, two out of four sectors on a vertical asset (facing in opposite directions) use the same frequency and the two other sectors (facing 90 degrees from the first two) use a second frequency. This configuration is commonly referred to as k=2, where k is the number of channels used for 360° coverage.

G1 radios use a combination of the techniques described previously to create deep RF nulls in the direction of other G1 radios. Figure 17 shows an example of two base nodes on the same channel communicating with multiple remote nodes. By aligning beams more narrowly in constructive (or desired) directions and nulls in destructive (undesired) directions, on both ends of each link, self-interference in a G1 network can be reduced by up to 45 dB.

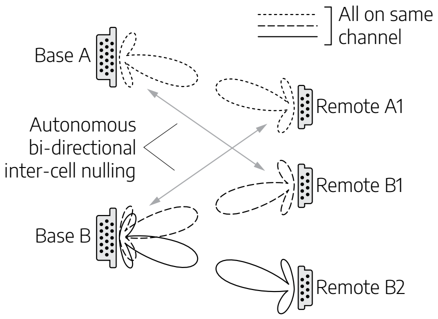

Figure 17: G1 cancels both intra- and inter- cell interference

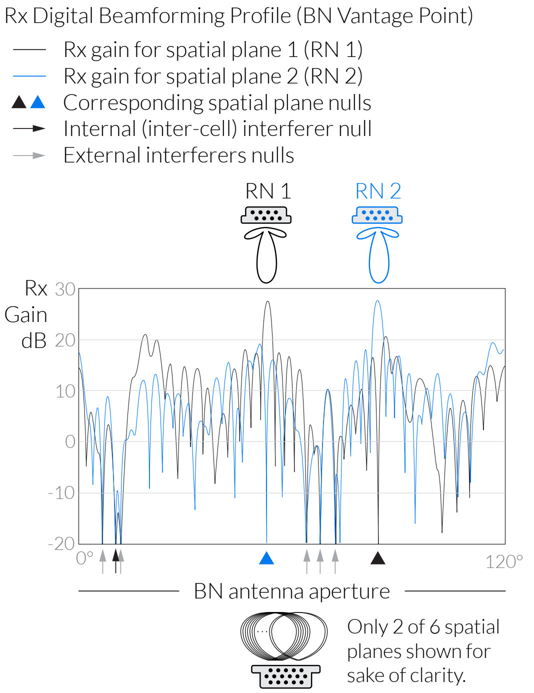

Figure 18: Strong beamforming and deep nulling of undesired signals (internal and external)

This allows G1 to employ universal frequency reuse (k=1) in which all sectors operate on the same frequency without degrading service. This is a hallmark of ngFWA and key to G1's high performance.

Burst Interference

Burst interference originates from outside an operator's network and can occur at any time in the 5 ms frame. It typically arises from other RF sources such as Wi-Fi radios and can't be managed by an administrator or operator with management tools or platforms. This problem is especially prevalent in popular unlicensed spectrum (heavily used in areas where licensed spectrum is unavailable or unaffordable for the fixed broadband use case) in which accessing a clean channel can be nearly impossible.

G1's completely unique asynchronous burst interference cancellation (ABIC) is an industry first and gives operators the ability to deliver licensed-class performance in unlicensed spectrum.

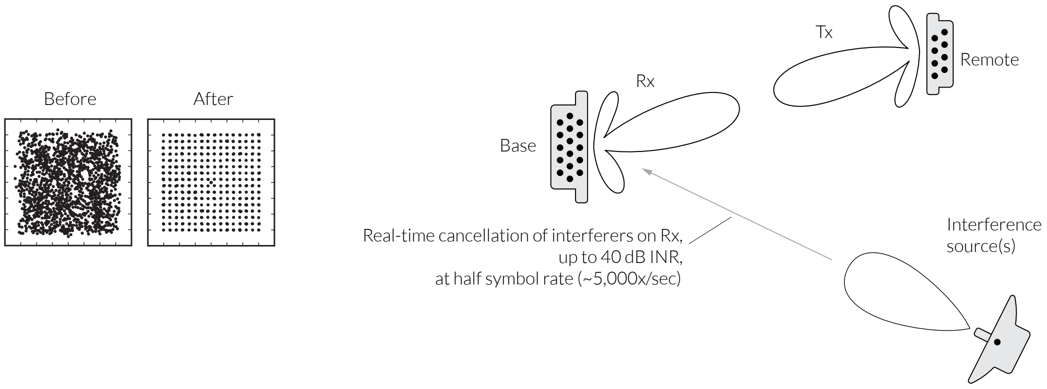

Figure 19: ABIC cancels interference in real-time

The ABIC approach processes combinations of inputs to resolve the G1 signal of interest out of a wavefront scrambled by other interfering signals. Figure 19 illustrates a signal before ABIC, in which the SINR of a signal is 0 dB; that is, the interference is as strong as the desired signal. ABIC effectively cancels out all undesired signals in the channel at the receive end of the link yielding a perfect 256 QAM constellation. ABIC can remove up to 40 dB of interference from sources outside the network, operating on the received signals at both base node and remote node. As with any radio innovation, ABIC has its limits — in particular less (and even no) success at cancelling very strong interference coming from a source aimed directly at (i.e., co-linear with) the boresight of the base node or remote node in question. This high-power, direct-boresight class of interference offers the ABIC algorithms insufficient diversity across Rx radio chains to extract the signal of interest successfully. Alternative mounting locations are required to address these cases.

Dynamic Motion in Channel

Because RF characteristics vary with time as well as distance, given moving obstructions and reflections in the cell, G1's ACRB calculations are performed at half the symbol rate (5,000 times per second) to eliminate the decay in signal quality that would otherwise occur over the duration of the system's 5 ms frame.

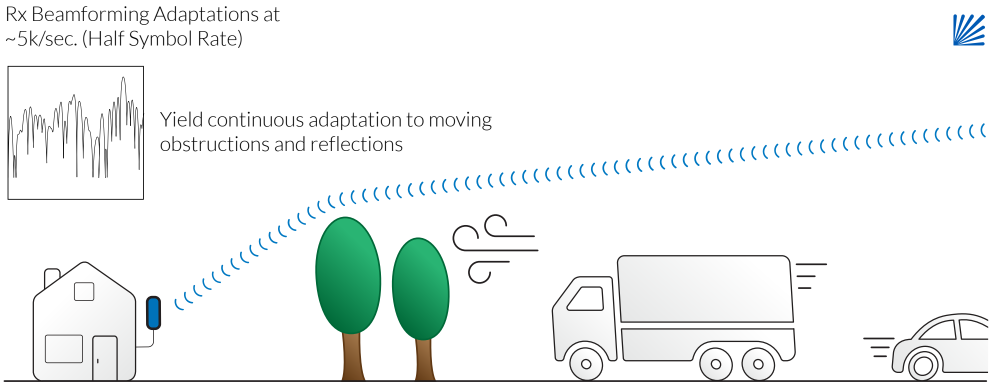

Figure 20: Motion and changing reflections can yield unreliable results unless beamforming is properly adapted

Symmetric Performance

Finally, increasing demand for symmetric performance (e.g., 100 Mbps downlink / 100 Mbps uplink service), raises a challenge for FWA, specifically for uplink. In many wireless technologies, uplink performance trails that of downlink. A primary reason is asymmetric link budgets, in which the uplink is less than the downlink, given the higher number of radio resources on the base node vs. the remote in legacy FWA architectures. While this was not seen as a problem in early FWA deployments, the increasing demand for symmetric performance exposes this weakness in legacy systems.

Summary

Wireless communications technology is relied on by millions of people and billions of devices for everything from casual to mission-critical applications. As we discuss here, there are architecture choices that must be made for any approach that affect radio performance and reliability along the dimensions of cell and network capacity, NLoS performance, resilience to motion and interference in the channel, and symmetrical throughput. Some technologies (such as 5G and Wi-Fi) have been repurposed from their original intent to supply fixed-wireless access with clearly mixed success. A fresh approach is required to take fixed wireless to the next generation.

Tarana has invested $400M in the research and development of real-time digital signal processing, custom silicon, and teraflops of processing power, a custom-designed transceiver (in 3 GHz), and A/D conversion with enough fidelity to match the precision of our calculations in the digital domain. All of this has been done in the pursuit — and now well-proven delivery — of fiber-class reliability and end-game performance in fixed wireless access deployments at large scale.

Interested in learning more about our innovative solutions? Get in touch with us at taranawireless.com/how-to-buy

Tarana's mission is to accelerate the deployment of fast, affordable internet access around the world. Through a decade of R&D and more than $400M of investment, the Tarana team has created a unique next-generation fixed wireless access (ngFWA) technology instantiated in its first commercial platform, Gigabit 1 (G1). It delivers a game-changing advance in broadband economics in both mainstream and underserved markets, using either licensed or unlicensed spectrum. G1 started production in mid-2021 and has been embraced by more than 200 operators in 21 countries and 45 states. Tarana is headquartered in Milpitas, California, with additional research and development in Pune, India.

@taranawireless taranawireless.com

TARANA