5G: Millimeter Wave

Instructor: Dr. Deepak Vasisht
Scribe: Vishakh Suresh Babu
CS 598 WSI: Lecture 19
October 30, 2024

Overview of Topics Covered

- mmWave Characteristics
- Phased Array
- Beam Alignment
- Agile-Link

1 mmWave Characteristics

The spectrum we use consists of various frequency bands:

- Wi-Fi primarily operates on two frequency ranges: 2.4 GHz and 5 GHz
- In cellular, there are three kinds of signals:
  - low band (below 1 GHz)
  - mid band (1-7 GHz)
  - high band or mmWave (above 7 GHz)

The bands for cellular mmWave typically range from 20 to 30 GHz, while Wi-Fi mmWave (802.11ad) operates at around 60 GHz.

1.1 Attenuation

A key challenge in mmWave systems is signal attenuation. Radio waves traveling through the air are absorbed or scattered by O_2, H_2O molecules, etc. Attenuation varies significantly with frequency, as illustrated in Figure 1.

The attenuation at the Wi-Fi mmWave band is particularly high. Even at cellular mmWave frequencies, signal attenuation is significantly greater than in networks operating below 10 GHz. In general, as we move to higher frequencies, attenuation increases. This results in a significantly weaker signal (for the same

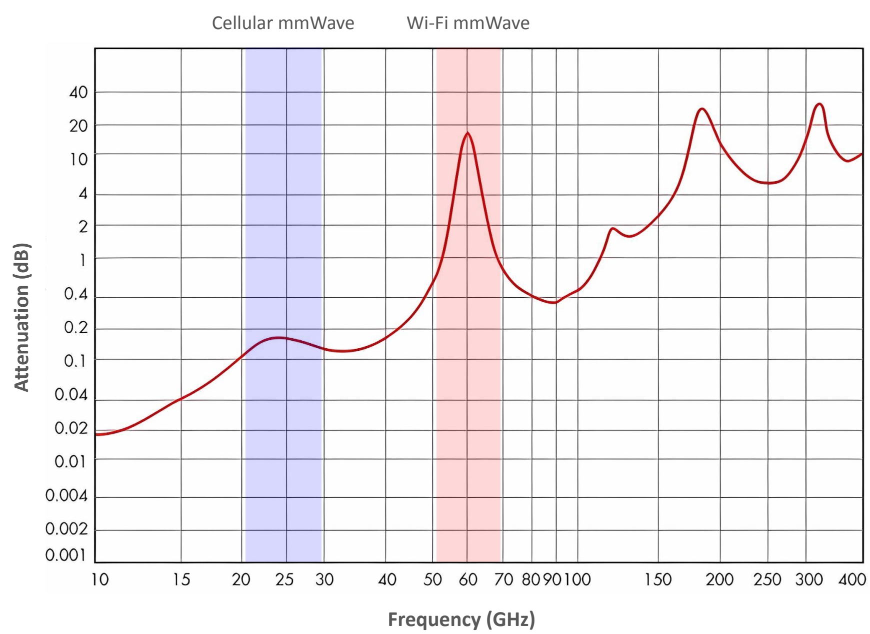

Figure 1: Signal attenuation versus log frequency. Absorption by O_2 and H_2O molecules is responsible for the peaks.

transmit power and distance), which makes implementing mmWave systems much more challenging.

If we have a transmitter antenna and a receiver antenna, with the transmitter sending x, the receiver receives $y = hx + n$. This model still holds for mmWave communications, but identifying y at a single antenna can be difficult due to the weak signal. To address this, we need to find ways to enhance the signal strength, such as by combining signals from multiple antennas or using highly directional antennas.

1.2 Beams

While transmitting more power could help compensate for attenuation, the available power is limited due to the devices' power constraints. Instead of distributing power across all directions, it's more efficient to concentrate it into $5°$ or $10°$ beams.

There are several methods for creating these beams. The simplest approach involves replacing omnidirectional antennas (which broadcast signals in all directions) with directional antennas. By physically adjusting the antennas so that the transmit and receive beams point toward each other, we can achieve strong signal quality.

However, this method isn't ideal for devices that move frequently. If a connection is established and the device moves, the slow mechanical adjustment of the antennas could cause the connection to drop. This approach is better suited for point-to-point links, such as those used in rural areas to extend connectivity from a home to a farm. However, a more mobile-friendly solution would be needed for the farm itself.

This is one way to create beams. Phased arrays present another option.

2 Phased Arrays

Phased arrays can generate and steer beams electronically. To change the beam direction and align the beams, we adjust the phase of the signals at each antenna element, without any need for mechanical movement.

Note: An array of antennas can be used to identify the direction from which a signal is coming. In this context, similar mathematical principles can be applied to steer a beam.

2.1 Recap: Antenna Arrays

In an antenna array, there are multiple antennas, separated by a distance x, as depicted in Figure 2. The signal arrives from a direction θ. When the signal travels towards the array, it covers an additional distance (relative to antenna 0) of $ixcos\theta$ to reach the i^{th} antenna.

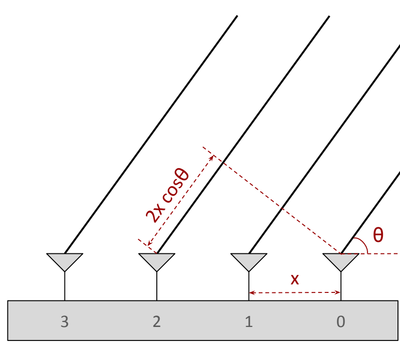

Figure 2: Array with 4 antennas

The channel at the i^{th} antenna can be expressed in terms of x, θ, and the wavelength λ as:

$$h_i = h_0\, e^{-j\frac{2\pi}{\lambda}ix\cos\theta}$$

Define $P(\theta')$ which is the power of the signal coming from direction θ' as:

$$P(\theta') = \left| \sum_i h_i\, e^{j\frac{2\pi}{\lambda}ix\cos\theta'} \right|$$

- When $\theta' = \theta$, all the complex vectors are aligned, and $P(\theta')$ attains its maximum value.
- When $\theta' \neq \theta$, each term has a phase difference that depends on i, leading to a lower value of $P(\theta')$.

This process is called *receive beamforming*. Essentially, we are trying to determine the direction from which the signal is arriving. We consider a particular direction θ' and assess the power of the signal received from that direction.

2.2 Transmit Beamforming

Consider an antenna array with four antennas, as shown in Figure 3. Suppose we want to transmit a signal α, in the direction defined by angle θ'.

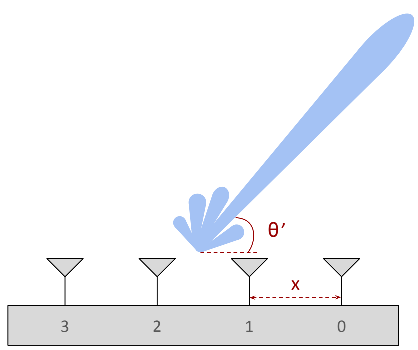

Figure 3: Transmit beamforming

Rather than sending α at each antenna, we introduce a phase shift of $\frac{2\pi}{\lambda}ix\cos\theta'$ for the signal transmitted from the i^{th} antenna. In other words, the signal transmitted by the i^{th} antenna is given by:

$$\alpha\, e^{j\frac{2\pi}{\lambda}ix\cos\theta'}$$

The idea behind this approach is that if a receiver is located in the direction θ', the signals from some of the antennas will travel a longer distance. We can compensate for this extra distance by pre-multiplying the signals with a phase shift at the transmit antennas. As a result, all the signals are aligned in the direction θ', resulting in a much stronger signal. In other directions, the signals from different antennas cancel each other out when received, leading to much weaker signals.

Essentially, while each antenna emits signals in all directions, the array as a whole focuses its transmission in the direction θ'. In effect, this achieves the same result as using a directional antenna, but without actually requiring a physical directional antenna.

2.3 Drawback of Antenna Arrays

Consider a system where the transmitter has 4, 8, 16, 32, or more antennas separated by $\lambda/2$ (which, for mmWave, is 0.5 cm or less), and the receiver also has multiple antennas. This configuration presents a practical challenge. Specifically, it requires the use of multiple (1) downconversion modules, (2) analog-to-digital converters (ADCs), and (3) receive chains on the receiver side, as illustrated in Figure 4. This configuration is costly due to the extra components needed, and it is both power-hungry and computationally expensive.

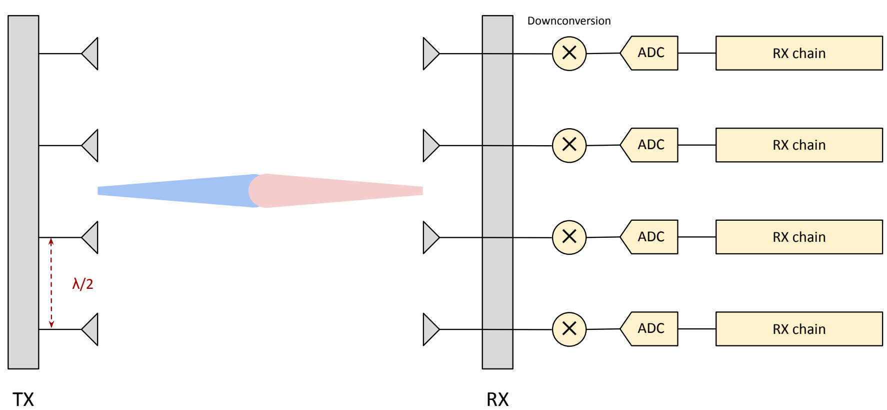

Figure 4: Antenna array with four decoding chains

Instead, we use phased arrays, which, in principle, are similar to traditional antenna arrays.

2.4 Antenna Array vs Phased Array

The main difference between an antenna array and a phased array is shown in the Figure 5.

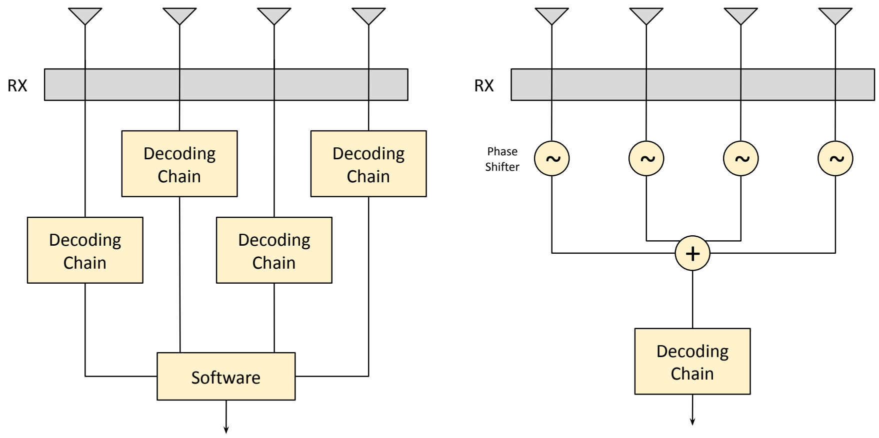

(a) Antenna array **(b)** Phased array

Figure 5: Main difference between an antenna array and a phased array

Consider an antenna array with four antennas (Figure 5a). At the end of each antenna, there is a complete decoding chain. The outputs of these decoding chains are then combined in software, with phase shifts applied to each output, as described below:

$$\left| \sum_i h_i \, e^{ji\phi} \right|$$

Instead of performing this operation in software, a phased array applies the phase shifts in hardware. A phase shifter (a hardware component) is used to apply a different phase shift to the signal at each antenna (as shown in Figure 5b). The signals are then combined and passed through a single decoding chain. Effectively, this process is similar to looking at the signal coming from a specific direction.

The pros and cons of antenna arrays and phased arrays are listed in Table 1.

3 Beam Alignment

For communication to occur, the beams of both the transmitter and receiver must be well aligned. There are several techniques for aligning the beams.

Antenna array	Phased array
Needs multiple decoding chains.	Only requires a single decoder, making it more cost-effective and power-efficient.
The scanning process is computational, enabling us to explore multiple directions simultaneously in software after the signal is received.	Can only beamform in one direction at a time because the phase shifts are applied (in hardware) directly to the antenna elements. Once the signal is combined and the phase shift is set, the array can only focus on that specific direction. This means we are constrained by the hardware and must perform a physical scan for each new direction.
Each antenna receives a weak signal, resulting in a low SNR. During the decoding stages at each antenna, we work with these weak signals, which increases the likelihood of errors. Only at the final stage do we obtain a strong signal with a better SNR.	SNR is high once the correct direction is identified, meaning that each decoding step is performed at a high SNR

Table 1: Pros and cons of an antenna array and a phased array

3.1 Exhaustive Search

The simplest approach to finding the optimal alignment is to scan all possible directions. Suppose there are phased arrays at both the transmitter and receiver sides, as shown in Figure 6. In this case, we can generate beams in all possible directions and perform a scan. To determine the best alignment, the receiver has to repeat the scanning process for each direction of the transmitter's beam. In other words, we need to calculate the SNR for each possible pair of beam directions.

In general, if there are N possible beam directions on each side (for example, if the beam width is 10°, then $N = 18$), the complexity of this approach is $O(N^2)$.

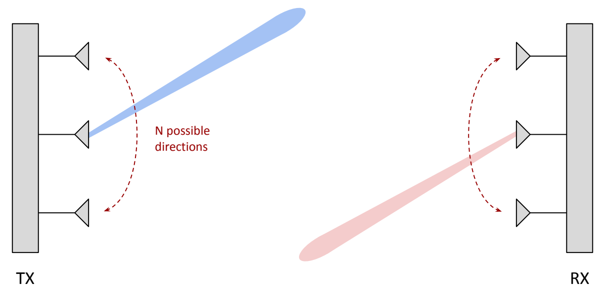

Figure 6: Exhaustive search

The key challenges with this approach are:

- *High setup time*: The initial alignment requires $O(N^2)$ operations, which can be time-consuming.
- *Slow response to mobility*: Even if we perform an exhaustive scan and find the correct pair, we would have to redo the entire scanning process if the position of the transmitter or receiver changes, or if there are any obstacles. This rescanning is triggered only when the connection weakens; in other words, it is reactive.

  For example, suppose we establish a beam as shown in Subfigure 7a, then a person moves into the path and blocks the signal. Assuming there's a wall available to reflect the signal, we can redo the scanning process and find the new alignment shown in Subfigure 7b. This alignment wasn't optimal initially — it only became optimal after the direct path was blocked.

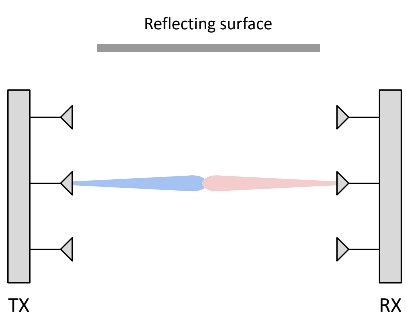

(a) Original beam

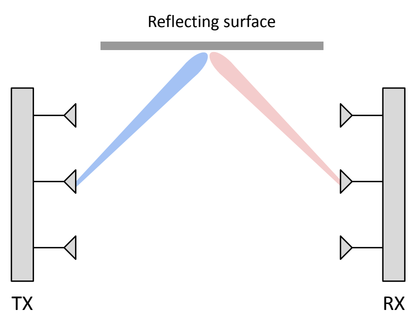

(b) New alignment

Figure 7: Old and new alignments for a situation where the original path was obstructed by an obstacle.

There is a lot of research in the fast beam realignment space. For instance, we can track the top three directions and try the second or third best options if the first one fails. Another approach is to perform intermittent scans to keep the beams updated.

3.2 Alignment in 802.11ad

At the beginning, the antennas on one side — let's say the receiver — do not perform beamforming. Instead, they point in all possible directions. Meanwhile, the transmitter scans through all possible directions one at a time (Subfigure 8a). The transmitter evaluates how the received power varies with direction and selects the beam with the highest power.

Next time, the roles are reversed as shown in Subfigure 8b. The transmitter becomes omnidirectional, and the receiver scans all directions to identify the best beam. The transmitter and receiver then use the two beams selected for communication.

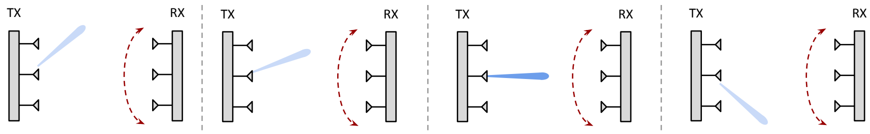

(a) Transmitter scans

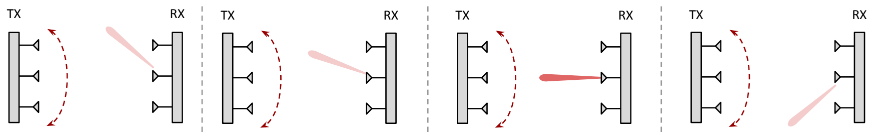

(b) Receiver scans

Figure 8: O(N) beam alignment

With this approach, only 2N scans are required. This effectively reduces the complexity to $O(N)$, which is an order of magnitude faster than an exhaustive search.

This approach does result in some loss of range. To achieve maximum range, the beams should ideally be pointed directly at each other. However, since one of the devices — either the transmitter or the receiver — is omnidirectional during the scanning process, the energy is distributed over a wide area. As a result, the effective range is lower when compared to the previous approach, where both devices were beamforming in the optimal direction.

4 Agile-Link

Overall Idea

- Agile-Link aims to identify the optimal beam alignment without scanning the entire space.
- The *key question* it raises is whether this can be achieved in logarithmic number of measurements.
- The *main intuition* is that multipath is sparse.

4.1 Multi-Arm Beams

The core concept behind Agile-Link is the use of multi-arm beams. Instead of relying on beams that point in a single direction, Agile-Link deploys beams that can simultaneously point in multiple directions (let's say, M directions), as illustrated in Figure 9. By using these multi-arm beams, we can quickly determine if the signal is coming from one of the M directions, though we won't be able to pinpoint the exact direction in the first attempt. In the next step, we deploy a different multi-arm beam to scan a new set of directions.

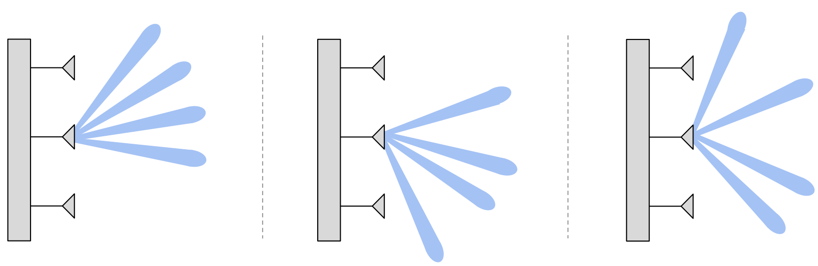

Figure 9: Multi-arm beam scan

For example, consider a scenario where the signal is coming from 60° (ground truth).

- In the first round, we create a multi-arm beam and check if the signal is coming from 45°, 75°, or 105°. We'll get a low SNR, because none of these directions match the actual signal.
- In the next round, we test 20°, 60°, and 100°. This time, we get a high SNR, indicating that the signal must be coming from one of these directions, but we still don't know exactly which one.

- In the third round, we test 45°, 60°, and 75°, and once again we observe a high SNR. Now we know that the signal must be coming from one of these three angles.

By looking at the overlap between the sets of tested directions, we can conclude with high confidence that the signal is coming from 60°.

This approach works only because multipath is sparse. Let K be the maximum number of paths the signal might take. If there are a large number of directions the signal could be coming from (i.e., K is large), we will need to try several different multi-arm beams to detect all the directions. However, if K is small, we can identify all the directions in $O(K \log N)$ time.

4.2 Creating Multi-Arm beams

The idea here is that, instead of configuring all antennas to point in the same direction, we organize the antennas into groups. For each group, we set the phase shifters of the antennas in that subset to direct the beam toward one of the directions we are interested in. This allows the array to simultaneously generate multiple beams, each pointing in a different direction.

For example, consider an array with 12 antennas as shown in Figure 10. We want to create a beam with 3 arms. We split the antennas into three groups: antennas 0–3 are directed in one direction, antennas 4–7 in a second direction, and antennas 8–11 in a third direction. By adjusting the phase shifters for each antenna subset in this manner, the array generates three separate beams, each oriented in a different direction.

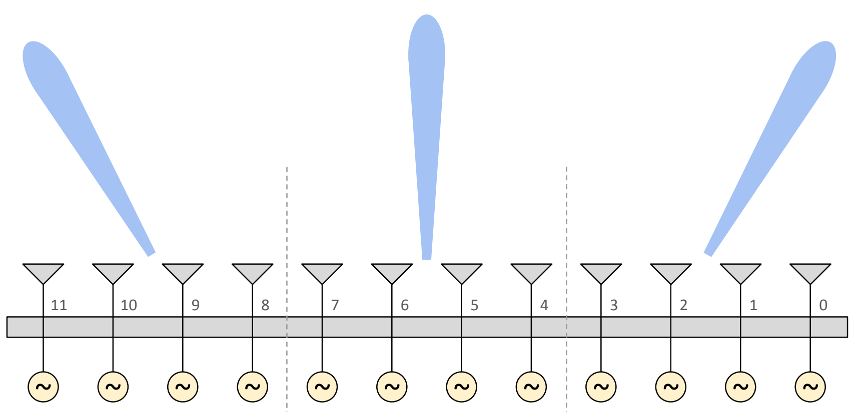

Figure 10: Creating a beam with 3 arms

Note: The original single beam would have been stronger and narrower. As a result, we sacrifice some range, but in return, we gain the ability to scan multiple directions simultaneously.

4.3 Evaluation[1]

4.3.1 Agile-Link Coverage

Figure 11 shows the relationship between distance and the SNR (at the receiver). When the receiver is within 10 meters of the transmitter, the SNR is high enough to support higher modulation schemes like 256-QAM (8 bits/symbol). As the distance increases, the SNR drops, requiring the use of lower modulation schemes. For example, at around 90 meters, only 16-QAM (4 bits/symbol) can be supported.

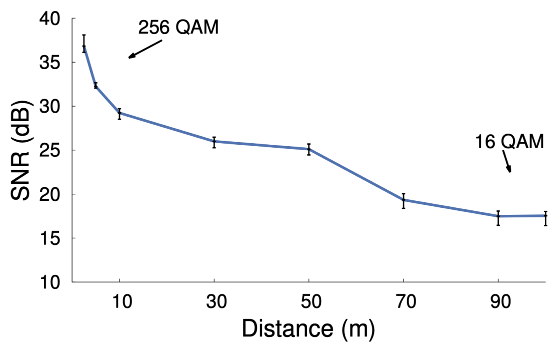

Figure 11: SNR vs distance

4.3.2 Beam Accuracy

Figure 13 illustrates the reduction in the number of measurements required for Agile-Link compared to other techniques. An exhaustive search has a complexity of O(N^2), while 802.11ad operates with a complexity of O(N). In contrast, Agile-Link has a complexity of O(K log N), where K represents the maximum number of paths. This makes Agile-Link significantly faster than both of the other techniques, whether for establishing a connection, maintaining a connection during movement, or other related tasks.

[1]Figures taken from Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast Millimeter-Wave Beam Alignment. In Proceedings of the 2018 ACM SIGCOMM Conference (SIGCOMM '18), Association for Computing Machinery, New York, NY, USA, pp. 432–445.

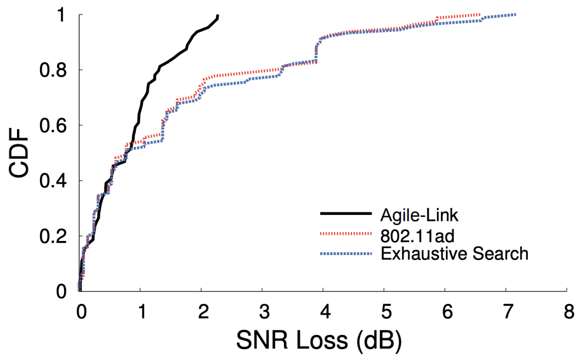

Figure 12: CDF of the SNR loss

4.3.3 Beam Alignment Latency

Figure 13 illustrates the reduction in the number of measurements required for Agile-Link compared to other techniques. An exhaustive search has a complexity of $O(N^2)$, while 802.11ad operates with a complexity of $O(N)$. In contrast, Agile-Link has a complexity of $O(K \log N)$, where K represents the maximum number of paths. This makes Agile-Link significantly faster than both of the other techniques, whether for establishing a connection, maintaining a connection during movement, or other related tasks. This advantage becomes even more significant as the number of antennas increases.

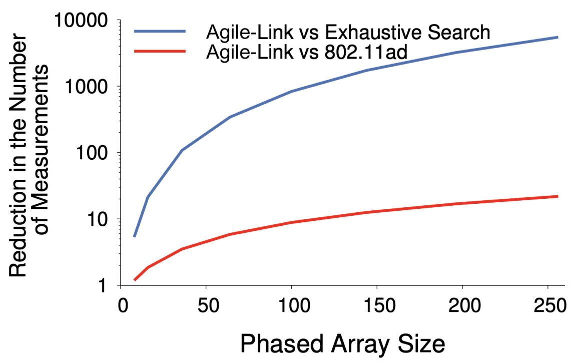

Figure 13: Reduction in the number of measurements