

A Millimeter Wave Network for Billions of Things

Mohammad H. Mazaheri
University of Waterloo
mh2mazah@uwaterloo.ca

Ali Abedi
University of Waterloo
ali.abedi@uwaterloo.ca

Soroush Ameli
University of Waterloo
soroush.ameli@uwaterloo.ca

Omid Abari
University of Waterloo
omid.abari@uwaterloo.ca

ABSTRACT

With the advent of the Internet of Things (IoT), billions of new connected devices will come online, placing a huge strain on today's WiFi and cellular spectrum. This problem will be further exacerbated by the fact that many of these IoT devices are low-power devices that use low-rate modulation schemes and therefore do not use the spectrum efficiently. Millimeter wave (mmWave) technology promises to revolutionize wireless networks and solve spectrum shortage problem through the usage of massive chunks of high-frequency spectrum. However, adapting this technology presents challenges. Past work has addressed challenges in using mmWave for emerging applications, such as 5G, virtual reality and data centers, which require multiple-gigabits-per-second links, while having substantial energy and computing power. In contrast, this paper focuses on designing a mmWave network for low-power, low-cost IoT devices. We address the key challenges that prevent existing mmWave technology from being used for such IoT devices. First, current mmWave radios are power hungry and expensive. Second, mmWave radios use directional antennas to search for the best beam alignment. Existing beam searching techniques are complex and require feedback from access points (AP), which makes them unsuitable for low-power, low-cost IoT devices. We present mmX, a novel mmWave network that addresses existing challenges in exploiting mmWave for IoT devices. We implemented mmX and evaluated it empirically.

CCS CONCEPTS

• **Hardware** → **Beamforming**; **Networking hardware**; *Wireless devices*; Radio frequency and wireless circuits; • **Networks** → Wireless access points, base stations and infrastructure; Wireless access networks;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM '19, August 19–23, 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5956-6/19/08...\$15.00

<https://doi.org/10.1145/3341302.3342068>

KEYWORDS

Wireless, Millimeter wave (mmWave), Internet-of-Things (IoT), low-power

ACM Reference Format:

Mohammad H. Mazaheri, Soroush Ameli, Ali Abedi, and Omid Abari. 2019. A Millimeter Wave Network for Billions of Things. In *SIGCOMM '19: 2019 Conference of the ACM Special Interest Group on Data Communication, August 19–23, 2019, Beijing, China*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3341302.3342068>

1 INTRODUCTION

It is anticipated that by the year 2025, 75 billion Internet of Things (IoT) devices will be installed, enabling new applications into every aspect of our daily lives: from smart homes and security cameras to smart cities and autonomous cars [28]. While these prospects sound exciting, the reality is that billions of devices will require wireless connectivity to the internet. In fact, many of these IoT devices will be sensors, such as cameras, which require real-time data streaming. The latest projections predict there will be 45 billion cameras connected by 2022 [30]. Unfortunately, such estimations will place a growing strain on requirements of wireless networks, which cannot be supported by contemporary WiFi and cellular networks. There are two main reasons for this shortage. First, WiFi and cellular bands are already congested, and as such, cannot support additional wireless devices. Second, many of the IoT sensors will be low-power devices, which transmit at rates much lower than channel capacity, and since these devices use omni-directional antennas, they are very inefficient in their use of shared spectrum.

Millimeter wave (mmWave) frequency bands have the potential to address this problem by offering multi-GHz of unlicensed bandwidth, 200x more than the bandwidth allocated to today's WiFi and cellular networks [33]. Spectrum availability at such high frequencies promises to enable higher network throughput than existing wireless networks. Recent studies have explored this technology in enabling high throughput wireless links for emerging applications, including 5G, virtual reality and data centers, which require multiple-gigabits-per-second throughput, while having substantial energy and computing power [3, 4, 31, 49].

In this paper, we focus on using mmWave to enable a wireless network for low-cost, low-power IoT devices. This has two significant advantages. First, it removes the load of low-power and low-cost IoT devices from today's WiFi spectrum. Second, directionality property of mmWave communication allows us to perform a spatial reuse of the spectrum, making the spectrum usage much

more efficient. As a result, a device that uses low rate will not impact others. However, in order to use mmWave for low-cost low-power IoT sensors, there are multiple challenges that need to be addressed. The main challenges include:

High power consumption: Unfortunately, existing mm Wave radios have high power consumption, which makes them unsuitable for low-power IoT sensors. For example, recent mmWave platforms developed by research communities such as OpenMilli, MiRa and NI platform consumes 10-20 watts [1, 5, 47], far more than what a camera or an entire low-power WiFi module consumes. Commercial mmWave chipsets, such as Qualcomm QCA9500, consumes several watts (excluding phased arrays power consumption) [39]. The high power consumption of these radios is due to the high power consumption of RF components operating at mmWave frequencies. For example, a power amplifier and mixer operating at 24 GHz consumes about 2.5W and 1W respectively [7, 13, 14]. In addition, mmWave radios perform beam searches that make the hardware more complex than traditional radios. In particular, phased arrays, consisting of amplifiers and phase shifters, excessively increase the power consumption of these radios.

Expensive hardware: Another disadvantage of mmWave is the fact that present mmWave components are expensive. For example, mmWave components such as amplifiers, mixers and phase shifters each costs \$220, \$70, \$150, respectively [7, 10, 13, 14]. Therefore, a full mmWave radio cost hundreds of dollars, which is far more than what a typical WiFi module costs. On the other hand, today's low-power IoT sensors, such as cameras, cost only tens of dollars and consume less than a watt. Furnishing a camera with a radio which costs hundreds of dollars and consumes tens of watts does not seem feasible nor economically viable.

Beam searching: A major limitation of mmWave technology is that their signals decay very quickly with distance, requiring mmWave radios to focus their power into narrow beams to achieve long range communication. Due to this limitation, mmWave communication is only possible when the transmitter's and receiver's beams are aligned. Recent research projects propose different beam searching techniques for aligning mmWave beams [26, 38, 48]. However, these techniques are not suitable for low-cost, low-power IoT devices, due to their computation complexity, energy requirements, and hardware costs. Specifically, these techniques search for the direction of the best beam alignment, which makes both communication protocols and hardware very complex, unsuitable for low-cost, low-power IoT devices.

In this paper, we address these fundamental challenges, and as a result, we develop a mmWave network for low cost, low-power IoT devices. Our mmWave network works in both dynamic and stationary environments. Such a network can be used in different applications. For example, it can be used in smart homes to connect IoT sensors (cameras, TVs, etc.) to a home hub. It can also enable wireless connectivity to surveillance cameras in public areas such as malls, banks, libraries, and parks.¹ In fact, this wireless network

can even be used in autonomous cars to connect their high data rate cameras and sensors to their in-vehicle access points.²

To eliminate the need for beam searching, we design a modulation technique that exploits the high attenuation of mmWave signal to modulate the signal over the air. Traditional mmWave systems view the high attenuation as a harmful phenomenon that the radio has to compensate for by using highly directional antenna which requires beam searching. In contrast, we show that we can leverage the directionality property to create modulation over the air. This eliminates the need for beam searching as well as simplifying the hardware. In other words, instead of modulating the signal first and then transmitting it to the beam direction with the best channel quality, we intelligently transmit a sine wave to different beams, and since each beam experiences different attenuations, the signal is modulated over the air. We will show that this approach enables us to design a new architecture for mmWave radios, a far more efficient and cost-effective architecture for imminent IoT applications.

Contributions: This paper makes the following contributions:³

- We design and build the first low-cost, low-power mmWave hardware platform, which operates as a daughterboard for Raspberry Pi. We believe this can help advanced mmWave research in the networking community.⁴
- We design a new communication modulation scheme, which eliminates the need of beam searching in mmWave radios. This will make adaptation of mmWave communication easier and less costly, paving the way toward many new applications for mmWave technology.
- We demonstrate the capability of our design in enabling a mmWave network for IoT applications. Specifically, we show that, in a representative IoT setup, mmX provides wireless links with SNR of 10dB or more to all nodes even at 18 meters. The maximum data rate of mmX's node is 100 Mbps and it consumes 1.1 W. This results in an energy efficiency of 11nj/bit, which is even lower than existing WiFi modules [22]

2 BACKGROUND

The term Millimeter wave (mmWave) refers to very high frequency RF signals. mmWave technology promises to revolutionize wireless networks by offering multi-GHz of unlicensed bandwidth, which is far more than the bandwidth allocated to today's WiFi and cellular networks [35]. However, in contrast to traditional wireless systems, mmWave frequencies suffer from a large path loss, and therefore mmWave radios use directional antennas to focus the signal power in a narrow beam. Such directional antennas are implemented using phased arrays. A phased array is an array of antennas, each with a phase shifter.⁵ A Phase shifter control the phase of the signal on each antenna which enables creating and steering a beam electronically [27]. Since mmWave communication is only possible when the transmitter's and receiver's beams are aligned, existing mmWave

¹ HD video streaming requires 8-10 Mbps application bit rate.

² Autonomous cars will be equipped with at least 8 cameras for a 360-degree surrounding coverage [16].

³This work does not raise any ethical issues.

⁴Our current mmX's node costs \$110. However, this can be dramatically reduced with mass production.

⁵Since the wavelength of mmWave signal is very small, mmWave antennas are small and many of them can be packed into a small area.

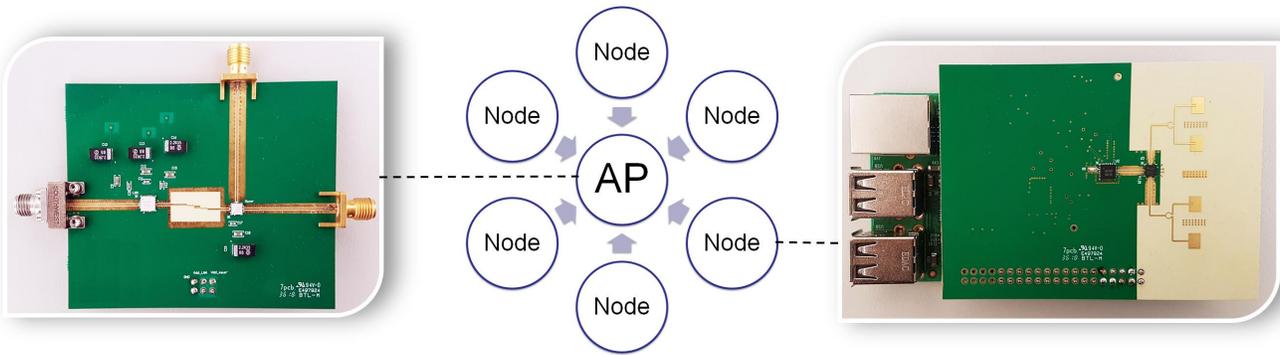


Figure 1: mmX platform. Multiple mmX's nodes transmit their data to a single AP. The figure also shows our custom designed mmX's IoT nodes and AP.

systems use phased arrays and different searching techniques to find the direction for the best beam alignment. Note, as shown in Figure 2, when the line-of-sight path is not blocked, the node needs to direct its beam toward the AP, and when the line-of-sight path is blocked, the node needs to direct its beam toward a reflector in the environment to use an indirect path. Past measurement studies show that in mmWave communication, typically there are a few paths [42] between two nodes. Therefore, existing radios need to search for the directions of the available paths and pick the best one. This requires the radios to have phased arrays and beam searching algorithms, which makes designing mmWave networks more challenging than designing traditional wireless networks equipped with omni-directional antennas.

3 RELATED WORK

Related work can be classified into three areas:

mmWave Communication: Recently, there has been significant interest in performing research on mmWave communications. However, much of the past work focuses on applications that require very high-data rate link, while having substantial energy and computing power. For example, the systems presented in [17, 23] utilize mmWave technology in data centers to enable high throughput links between server racks. There is also work in using mmWave for 5G applications [20, 36, 37]. Finally, some other systems use mmWave to enable high data rate for VR application to stream high-data rate videos from PC to VR headset [3, 4, 18]. In contrast, mmX focuses on designing a mmWave network for low-power, low-cost IoT devices. Unfortunately, such IoT devices cannot use existing mmWave radios and protocols, due to their complexity, cost and power consumption.

mmWave Radio Platforms: In the past few years, a variety of mmWave radio platforms have been proposed by academia and industry. However, these platforms are very costly and power hungry which makes them unsuitable for IoT applications. For example, National Instrument mmWave system costs \$134K and consumes more than 20 watts [1, 2]. Similarly, OpenMilli and Mira costs a few thousand dollars and consumes several watts [5, 47]. The main reason for the high cost and high power of these platforms is that

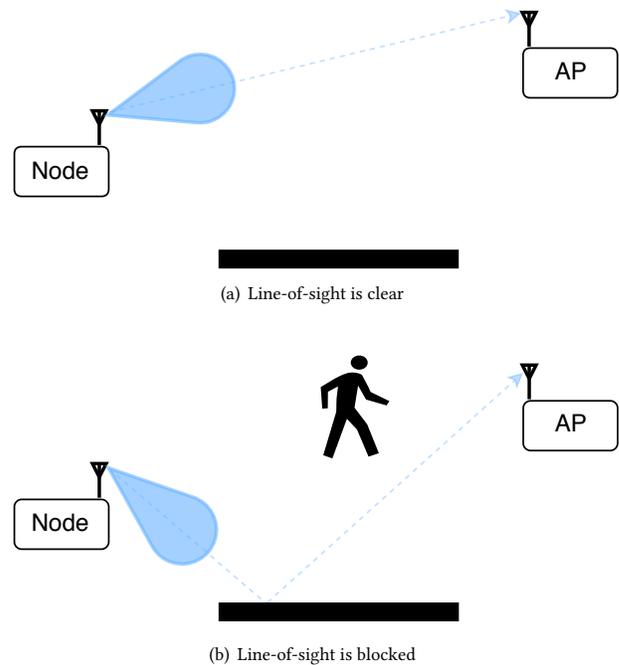


Figure 2: mmWave Communications. Existing mmWave devices need to search for the best path to AP, and align their beams toward it to establish a communication link.

they target applications which require very high-data rate links (Gbps). Hence, they are not suitable for low-power, low-cost IoT devices. In contrast, mmX's radio platform targets low-power, low cost applications, and since it has a very simple architecture, it provides lower cost and lower power consumption than existing mmWave platforms. In Section 10, we compare mmX with existing mmWave platforms in terms of cost, power consumption, bitrate, range, carrier frequency, and energy efficiency.

mmWave Beam Alignment: There is a large literature on mmWave beam search and alignment [26, 44, 46]. Much of the

work proposed beam searching algorithms that are not fast enough to enable mobile applications since they exhaustively search for the best beam alignment [21, 41]. Some previous work leverages sparse recovery algorithms (such as compressive sensing and Sparse Fourier transform) to speed up the search for the best beam alignment [6, 19, 24, 40]. However, existing beam searching techniques require expensive and lossy phased arrays with multiple phase shifters and amplifiers to steer the beam. Unfortunately, phased arrays are expensive to build and have high power consumption, which makes them unsuitable for IoT applications. Further, existing beam searching techniques are either complex or require feedback from the access point, and as a result, significantly increase the power consumption and computation requirement of the node. In contrast, mmX introduces a new modulation technique that eliminates the need for beam searching, enabling IoT device to communicate to the AP without searching for the best beam alignment.

4 SYSTEM OVERVIEW

mmX is a mmWave wireless network for low-cost, low-power IoT devices. It enables IoT devices to communicate using mmWave spectrum, without placing any strains on today's WiFi and LTE spectrum. Figure 1 shows mmX's setup, where multiple IoT nodes transmit data to a single receiver, which we refer to as the Access Point (AP). mmX operates in two phases: initialization and transmission. In the initialization phase, the AP allocates channels to IoT nodes. The bandwidth of an allocated channel depends on the data rate requirement of the IoT node. For example, if a device needs to stream an HD video, a few MHz of bandwidth must be allocated to it. In the transmission phase, the nodes send their data to the AP. Due to the directional property of mmWave communication, mmX performs a spatial reuse of the spectrum to make the spectrum usage more efficient. This allows multiple nodes to communicate to an AP, simultaneously, without creating any interference for other devices.

Over the next few sections, we will discuss the components that contribute to the design of mmX. We start by explaining the two key challenges in using mmWave technology for low-power, low-cost IoT devices, and how we overcome them. Then, we explain how mmX supports multiple nodes and enables them to communicate to a single access point, simultaneously.

5 COST AND ENERGY CHALLENGE

To overcome the high cost and energy inefficiency of existing mmWave radios, we introduce a new mmWave radio architecture. Specifically, our design utilizes joint ASK-FSK modulation to minimize the number of costly and power hungry mmWave components while providing robust communication links. ASK-FSK modulation combines two simple modulations: Amplitude-shift keying (ASK) and Frequency-shift keying (FSK). ASK modulation represents digital data as variations in the amplitude of a carrier sine wave. Specifically, it transmits a sine wave with high amplitude to present symbol '1' and transmits a sine wave with low amplitude to present symbol '0'. FSK modulation represents digital data as variations in the carrier frequency.

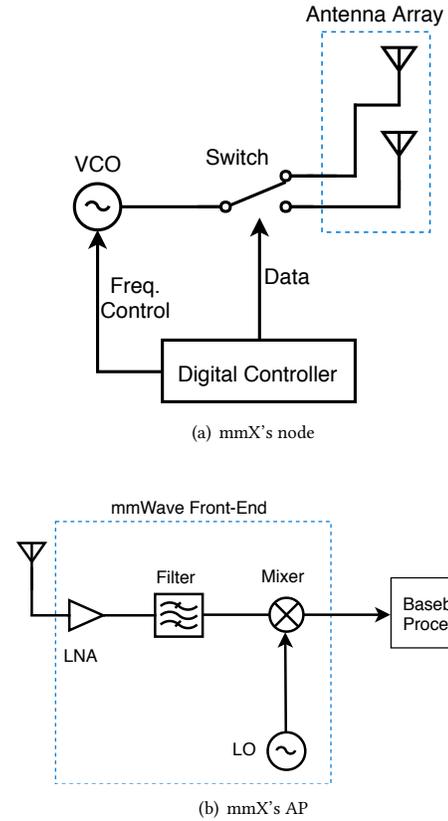


Figure 3: The block diagram of the mmX's node and mmX's AP. The architecture of mmX's node is very simple, making it ideal for low-power, low-cost IoT devices

Due to their simplicity, these modulation schemes offer a simple, low-power architecture. Traditional wireless systems avoid using ASK or FSK modulation because of their low spectral efficiency. This is mainly due to the fact that traditional wireless systems use omni-directional antennas and any inefficiency of a node in spectrum usage will impact other users too. As a result, these systems avoid using these simple modulations. However, since mmWave systems use directional antennas, these systems can perform spatial reuse of the spectrum. This will allow mmWave systems to use the spectrum much more efficiently and hence the low spectral efficiency of FSK or ASK modulation will not impact other users. This is the reason why these modulations are also used in optical and laser communication which are directional [45]. In the following subsections, we explain how these modulations enable simple architecture for mmX's IoT node and access points.

5.1 mmX's IoT Node

The block diagram of the mmX's node is shown in Figure 3 (a), which includes a mmWave section and a digital controller. The digital controller (i.e., a micro-controller) is used to control parameters of the mmWave section and to generate the data stream. The mmWave section includes only two active mmWave components:

a Voltage Controlled Oscillator (VCO) and a Single Pole Double Throw (SPDT) switch. The VCO generates a carrier signal (i.e., a sine wave) at the desired carrier frequency, and feeds it to the SPDT. The SPDT switches the signal between the two antenna arrays. Both the SPDT and the VCO are controlled by the digital controller. The device can choose different channels by changing the center frequency of the VCO. It can also change the data rate by changing the switching speed of the SPDT. Finally, the output of the SPDT is fed to an antenna array, which can create beams to different directions. Instead of using a phased array to create and steer a beam, we design an antenna array that can create beams toward different directions without using any costly phased shifters. In Section 6.2, we will explain our antenna design in more detail. The architecture of the mmX's node is very simple and low-power. Hence, it can be incorporated into low-power, low-cost IoT nodes.

5.2 mmX's Access Point

Figure 3 (b) shows the block diagram of the mmX's AP, which includes two main blocks: a mmWave down-converter and a baseband processor. The down-converter first amplifies the received mmWave signal with a Low Noise Amplifier (LNA). The LNA is placed at the first stage to reduce the total noise figure of the receiver. To reduce the possible interference from the out of band sources, the output of the LNA is fed to a filter. To avoid using costly filters, mmX exploits a microstrip coupled line filter, which is designed on the PCB board without any additional components. After amplification and filtering, the mixer multiplies the signal with a sine wave, generated by a local oscillator (LO). The LO signal is generated using a Phased lock loop (PLL). An PLL operating at mmWave frequency is costly and power hungry. Subsequently, we use a sub-harmonic mixer, which itself doubles the LO frequency. As a result, mmX uses a low-cost PLL operating at a much lower frequency. Finally, the output of the mixer is fed to a baseband processor to digitize and decode the down-converted signal. The AP architecture is simple and low-cost. This architecture can be used for a myriad of applications, including smart home hubs, autonomous cars and much more.

6 BEAM SEARCHING CHALLENGE

As explained in Section 2, mmWave radios use directional antennas to focus their power, and they need to search for the best direction of the beam. In today's radios, the beam steering is implemented using phased arrays. Phased arrays are an array of antennas, each with a phase shifter that controls the phase of the signal on the antenna. By modifying the phase of the signal, mmWave radios can create and steer a beam electronically. This will allow them to search for the best direction to direct their beam toward.

Unfortunately, phased arrays are costly and power hungry which makes them unsuitable for low-power IoT devices. A phased array with even a small number of antennas (8 elements⁶) consumes more than a watt and costs a few hundred dollars [5, 8, 11]. In addition, existing beam searching protocols are too complex for low power IoT applications. First, when these protocols are searching, they need multiple feedbacks from the AP, which significantly increases the power consumption of the nodes, and second, regular mobility

and environmental changes means that the beam must perform a continuous search, which is time-consuming and increases the power consumption of an IoT device. Therefore, we need to design a technique that enables IoT devices to communicate to an AP, without using any phased array and beam searching mechanism. One naïve approach is to use an antenna array with a fixed beam, and then ask the user to point the device towards the access point. Unfortunately, in this scenario, when the line-of-sight path gets blocked, the signal will be completely lost and the device will not be able to communicate with the access point. Another approach is to have an antenna array that can create multiple fix beams toward different directions, and pick the one which provides the highest SNR at the AP. However, this approach requires the access point to provide the IoT device feedback on which beam to pick. Moreover, due to mobility and environmental change, the AP needs to provide continuous feedback, which significantly increases the power consumption of the node and its complexity. Ideally, we want to employ a technique that enables an IoT device to communicate to an AP without requiring costly phased arrays and avoids the need for beam searching techniques. For the remainder of this section, we introduce our technique, named Over The Air Modulation (OTAM), which effectively addresses this problem.

6.1 Over The Air Modulation (OTAM)

As described in Section 5, mmX's nodes use Amplitude-shift Keying (ASK) modulation to communicate to the AP. Due to its simplicity, this modulation scheme offers a simple, low-power architecture suitable for IoT applications. However, to be able to communicate, the node needs to search for the best beam direction. To avoid the complexity of phased arrays and beam searching algorithms, while providing a robust wireless communication link, we propose a new architecture and modulation technique called Over The Air Modulation (OTAM). OTAM exploits the high attenuation property and directionality requirement of mmWave communication to create ASK modulation over the air, and as a result, does not need to search for the best beam direction.

Figure 4 illustrate how OTAM works. Our OTAM technique integrates the beam selection into data modulation. Specifically, instead of first creating an ASK signal and then choosing the best beam direction to transmit, OTAM sends a sine wave (carrier signal) to different beams depending on the value of data. For example, when the data bit is "1", Beam 1 (blue) is selected and when the data bit is "0", Beam 0 (red) is selected. Depending on the value of the data, the carrier signal is either transmitted using Beam 1 or Beam 0. Due to the directionality and channel property of mmWave communication, each transmitted signal will experience a different path loss, thus the AP will receive a sine wave in which its amplitude is modulated by the channel.

To understand why this approach does not require beam searching and works even with mobility and time-varying environments, let's consider two different scenarios, as shown in Figure 4: 1) when the line-of-sight (LoS) is clear, and 2) when the LoS is blocked.

Figure 4(a) shows the first scenario when there is no blockage in the LoS path. In this example, the signal sent through Beam 1 experiences much lower attenuation than the signal sent through Beam 0. This is because Beam 1 uses the direct path while Beam 0 relies

⁶Each element of phased array requires one LNA/PA and one Phase shifter

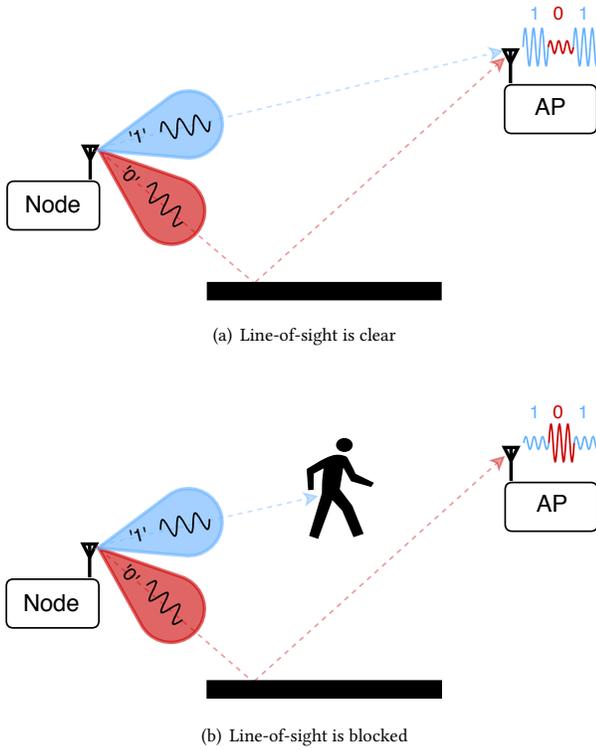


Figure 4: Illustrative example of Over The Air Modulation (OTAM) technique. mmX’s nodes exploit blockage limitation and high attenuation property of mmWave to create ASK modulations over the air, eliminating the need for beam searching.

on the reflection from the environment. Therefore, by switching between these two beams, the node can modulate the amplitude of the carrier frequency and create ASK signal at the receiver. For example, if the node want to transmit bit stream: 101, it sends its carrier signal to Beam 1, then it switches to beam 0 and finally it switches back to Beam 1. The receiver receives a carrier signal where its amplitude is modulated by the path loss. Because the loss of the two paths are sufficiently different, the receiver can easily decode the bits by monitoring the signal amplitude.

Now let’s consider the second scenario. As shown in Figure 4(b), due to mobility or environmental change, if the LoS path gets blocked by an object, the signal from Beam 1 will be attenuated much more than the signal from Beam 0 which relies on the reflection from the environment. Therefore, although the node transmits a pure sine wave (i.e., a carrier signal), the AP receives an ASK modulated signal since the two paths experience different losses. Note that in this scenario, as shown in Figure 4 (b), all bits are inverted. Therefore, in order to decode the bits, a few training bits are used at the beginning of each packet. Specifically, similar to most wireless communication systems, each mmX’s packet has known preamble bits. These bits are used to distinguish the signal of Beam 0 from Beam 1. Finally, it is worth mentioning that the OTAM technique works since mmWave signal attenuation in LoS path, NLoS

path and blockage are significantly different. In fact, past work has shown that NLoS paths typically experiences 10-20 dB higher attenuation than LOS path, and a blocked path typically experiences 10-15 dB higher attenuation than NLoS path [4]. Therefore, when both LoS and NLoS paths are available, the SNR will be 10-20 dB. When the LoS path is blocked and only the NLoS path is available, SNR reduces to 10-15 dB. Finally, when the LoS path is available and the NLoS path is blocked, the SNR can be up to 35 dB.

6.2 Orthogonal Beam Patterns

So far, we have described how mmX’s nodes communicate to mmX’s AP by transmitting a simple sine wave to two different beams which experience different path loss. However, there is a chance that the two paths experience similar loss, and as a result, signal levels will be the same and the AP will not be able to decode the signal using ASK demodulation.

Figure 5(a) shows an example of such a scenario. Here, the AP is in the middle of two beams and therefore the two NLoS paths experience similar attenuation. To prevent such scenarios from happening, we need to carefully design mmX’s beam patterns. Essentially, we need to design two radiation beams that are orthogonal to each other while they cover a large area. Orthogonality means each beam has nulls at the main direction of the other beam. Figure 5(b) shows our proposed beam patterns. The direction of the first beam (Beam 1) is on the broadside direction and perpendicular to the transmitter board, and the second beam (Beam 0) is divided into two directions. Further, Beam 0 has a null on the broadside direction, as such, Beam 1 and Beam 0 are orthogonal to each other. To implement these two beams mmX uses two different antenna arrays. Each antenna array includes two patch antennas. The array with the broadside beam (Beam 1) excites the patches with the same phase, while the array with null on the broadside (Beam 0) excites the two patches with 180° phase difference. The 180° phase difference creates a null in the broadside and produces two peaks at about $\pm 30^\circ$. In addition, the distance between antenna elements corresponding to Beam 1 is properly designed to create a null at $\pm 30^\circ$, so that the two beams are orthogonal to each other. It is worth mentioning that using the orthogonal beam pattern not only reduces the probability of getting similar losses for the two beams but also increases the coverage angle. Therefore, using orthogonal beam increases the robustness of our system and also allows us to cover wider angles.

6.3 Joint ASK-FSK Modulation

To this point, we have explained how we can improve the performance of mmX in decoding the signal by designing orthogonal beam patterns. However, our empirical results show that there is still a small chance ($<10\%$) that the received power from Beam 1 and Beam 0 experiences the same loss. In this case, the receiver will not be able to differentiate the difference between the two levels and cannot decode the bits. To solve this issue, we propose to combine ASK and FSK modulations, where the signal is decoded using both amplitude and frequency differences. Specifically, the frequency of the tone transmitted by Beam 1 will be slightly different from the frequency of the tone transmitted by Beam 0. The slight change in the frequency of the carrier signal can be simply implemented

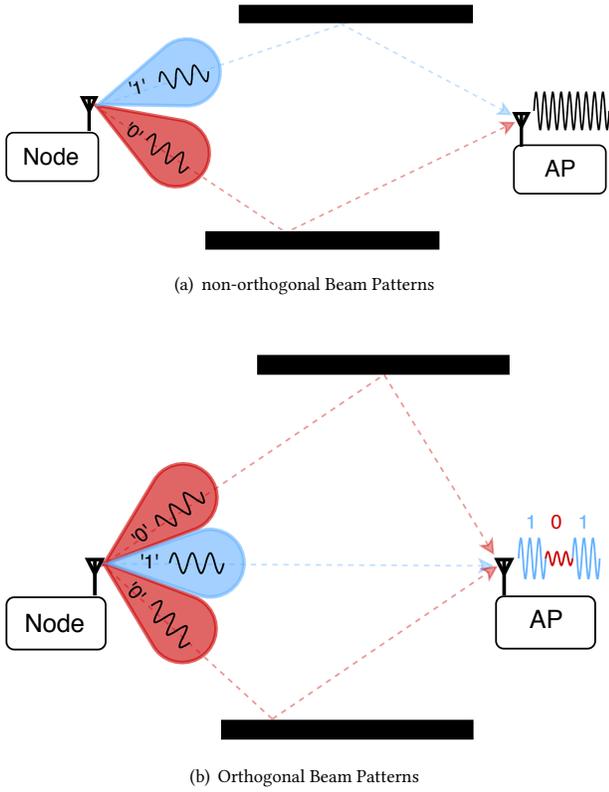


Figure 5: Non-orthogonal versus orthogonal beam patterns: In the orthogonal beam pattern, we split Beam 0 to two parts. This reduces the chance of experiencing the same path loss for '0' and '1' signal.

by changing the control voltage of the VCO. Note that FSK or ASK alone is not sufficient to decode the signal in all scenarios. Specifically, when the signals of Beam 0 and Beam 1 experience different path losses, the signal for one beam may be completely lost, and hence FSK demodulation does not work. In these cases, the signal is decoded using ASK demodulation, as shown in Figure 9(a). On the other hand, in rare cases, the signals of Beam 0 and Beam 1 might experience similar path losses. Therefore, the signals of Beam 0 and Beam 1 have similar amplitude, and ASK demodulation does not work. In these cases, the signal is decoded using FSK demodulation, as shown figure 9(b)) Therefore, utilizing joint ASK-FSK modulations is essential in order to decode the signal in all scenarios.

7 SUPPORTING MULTIPLE NODES

We explained how a single mmX node communicates to an AP. In this section, we explain how mmX enables multiple nodes to communicate to an AP, simultaneously. mmX uses spacial-division and frequency-division multiplexing to enable simultaneous communication to all nodes. In the following section, we discuss them in greater details.

(a) Frequency Division Multiplexing (FDM) mmX divides the available spectrum between nodes depending on their data rate demand. For example, the available unlicensed spectrum at 24 GHz and 60 GHz are 250 MHz and 7GHz wide, respectively. These bands are wide enough to support many nodes while providing each with 10-100s of MHz channel bandwidth. The channels are specified by the AP to each node in the initialization stage. The initialization takes place only once using a WiFi or Bluetooth module.

(b) Spatial Division Multiplexing (SDM) In scenarios where the total demanded bandwidth by the nodes is more than the available spectrum, mmX uses SDM to support all nodes, simultaneously. The directionality property of mmWave communication allows mmX to perform a spatial reuse of the spectrum, making the spectrum usage much more efficient. Specifically, since the number of paths between a node and AP is sparse and the signal is directional, most nodes do not create interference to each other over the air. However, if no multiplexing technique is used, the received signals are combined at the AP's antenna and interfere with each other. As a result, a spatial multiplexing technique is required in order to separate the signal. There are two different techniques for doing this: Hybrid MIMO Array and Time Modulated Array (TMA).

Hybrid MIMO Array: In this approach, the AP uses multiple mmWave chains connected to one or multiple arrays which create independent beams toward different directions [29]. This allows the AP to reuse the spectrum by performing Multiple-Input and Multiple-Output (MIMO), and hence enabling multiple nodes to communicate to an AP using the same frequency channel. However, since this architecture requires multiple mmWave chains, it is power hungry and costly for IoT applications.

Time Modulated Array (TMA): Instead of using multiple mmWave chains to separate the signals, another approach is to use TMA [34]. In this approach, an array of antennas, with each element connected to a switch, is used. The outputs of these switches are combined and fed to a single mmWave chain. By using a proper switching sequence, the signals on the same frequency channel, but arriving from different directions, can be shifted to different frequency channels. In other words, TMA hashes the signals arriving from different directions into different frequency bands, as shown in Figure 6. To understand how TMA works, let's consider an antenna array with N elements. The output of TMA for a signal arriving from direction θ can be written as:

$$y(\theta, t) = r(\theta, t) \cdot \sum_{n=0}^{N-1} w_n(t) \cdot e^{j \frac{\omega_0}{c} n d \cdot \sin \theta}, \quad (1)$$

where $r(\theta, t)$ is the arriving signal, ω_0 is the carrier frequency, c is the speed of light, d is the spacing between the elements of the array and $w_n(t)$, is a periodic signal that controls the switches and can be presented as follows:

$$w_n(t) = \begin{cases} 1 & 0 \leq t_n^{on} < t < t_n^{off} \leq T_p \\ 0 & otherwise \end{cases}. \quad (2)$$

Since $w_n(t)$ is a periodic signal, it can be represented by its Fourier series as follows:

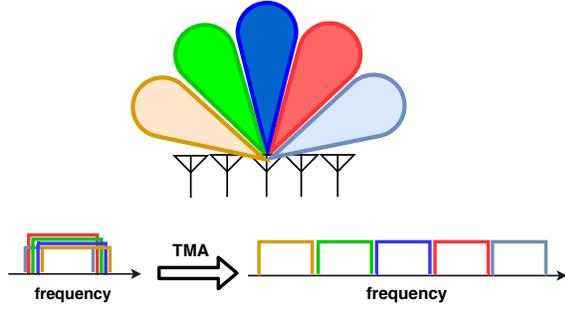


Figure 6: Time Modulated Array (TMA). TMA enables the AP to separate the signals arriving from different directions and map them to different channels.

$$w_n(t) = \sum_{m=-\infty}^{\infty} a_{mn} \cdot e^{j\omega_p t}, \text{ where} \quad (3)$$

$$a_{mn} = \frac{1}{T_p} \int_0^{T_p} w_n(t) \cdot e^{jm\omega_p t} dt.$$

By substituting (3) into (1), we obtain:

$$y(\theta, t) = r(\theta, t) \cdot \sum_{m=-\infty}^{\infty} e^{j(\omega_0 + m\omega_p)t} \sum_{n=0}^{N-1} a_{mn} \cdot e^{j\frac{\omega_0}{c} n d \sin \theta}. \quad (4)$$

This expression indicates that the received signal is copied at the carrier frequency (i.e., at ω_0) and at the harmonics of the switches' control signal (i.e., at $m\omega_p$). However, only one copy has significant amplitude and the rest are negligible (20-30dB weaker [25]). In fact, the center frequency of the strongest copy depends on the direction from which the signal arrives (θ). Therefore, by using TMA, the signals on the same frequency channel, but arriving from different directions, will be shifted to different frequency channels. This enables the AP to simultaneously communicate to multiple nodes using the same frequency channel.

8 IMPLEMENTATION

mmX's implementation has two primary components: IoT nodes and an access point. In the following, we explain the implementation of each component in more details.

8.1 mmX's IoT node

The block diagram of the mmX's nodes is shown in Figure 3, which includes a mmWave section and a control unit. A Raspberry Pi 3 is used as a control board. We implemented the mmWave section on a printed circuit board (PCB) using off-the-shelf components, as shown in Figure 1. The designed mmWave board can be used as a daughter board for the Raspberry Pi. The data is transferred from the Raspberry Pi to the mmWave board through the SPI communication port. This will enable robust and real-time data stream from the Raspberry Pi to the mmWave section. The mmWave board has two main components: VCO and an RF switch. For the VCO, we use HMC 533 from Analog Devices [9] which has a wide tunable frequency range, covering the entire 24 GHz ISM band. The

maximum output power of this component is 12 dBm, which eliminates the need for a power amplifier. For the switch, we have used ADRF 5020 SPDT from Analog Devices [12]. The switch has low insertion loss (<2dB) and high isolation (65 dB) between output ports which provides a great performance for our application.

As mentioned in Section 5, the two outputs of the SPDT switch are connected to two antenna arrays that have orthogonal radiation beams, which means each array has a null at the direction of the main beam of the other. The radiated power by the antenna is 10 dBm which complies with FCC regulations. Each antenna array includes two patch antennas. The first array excites the patches with the same phase, creating the broadside beam (Beam 1). The second array excites the two patches with 180° phase difference. The 180° phase difference creates a two-arm beam (Beam 0) which has a null in the broadside and produces two peaks at about $\pm 30^\circ$. The measured radiation patterns of the designed array is shown in Figure 8, where Beam 1 is directed toward the broadside, orthogonal to Beam 0. The Beam 0 is pointing toward ± 30 and has a null on the broadside direction.

8.2 mmX's Access Point

The block diagram of the mmX's AP is shown in Figure 3, which includes a mmWave down-converter board and a baseband processor. For the baseband processor, we used N210 USRP from TI with CBX daughter-board, which covers DC to 6 GHz RF carrier. The mmWave down-converter first amplifies the received 24 GHz signal with an LNA. We used HMC 751 from Analog Devices, which provides about 25 dB gain with only 2 dB noise figure at 24 GHz. The LNA is placed at the first stage to reduce the total noise figure of the receiver. To reduce the possible interference from the out of band sources, we designed a coupled line microstrip filter. The center frequency of the filter is at 24 GHz and the insertion loss at the passband is 5 dB. For the LO generator, we used the evaluation kit for ADF5356, generating a 10GHz signal which will be doubled by the sub-harmonic mixer. We use HMC264LC3B as a sub-harmonic mixer, which down convert the 24 GHz received signal to 4 GHz. For the AP's antennas, we designed and fabricated dipole antennas working at 24 GHz, with 5 dB gain and 3 dB beam width of 62 degree.

9 EXPERIMENTAL RESULTS

We evaluated the performance of mmX in both line-of-sight and non-line-of sight scenarios. We ran experiments in a lab area with standard furniture such as desks, chairs, computers and closets.

9.1 Microbenchmarks

Transmitter Performance: As described in Section 8, the node's radio has only two components: a VCO which generated the carrier signal, and a switch. Figure 7 shows the frequency of the VCO versus its control voltage. The VCO covers 23.95 GHz to 24.25 GHz by tuning the control voltage from 3.5 V to 4.9 V. The provided frequency range covers the entire 24 GHz ISM band, therefore mmX's node can tune its frequency to any channel assigned to it by the AP. The figure also shows that the frequency can be slightly altered by changing the control voltage. This allows mmX

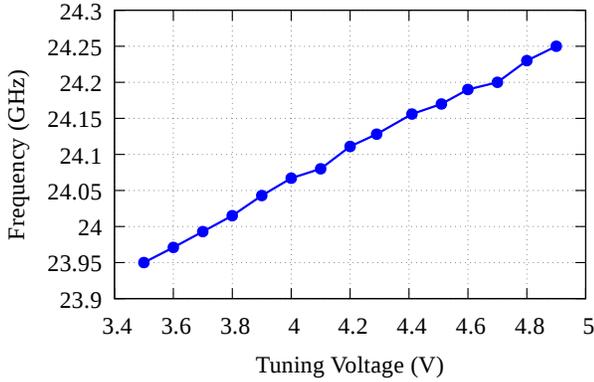


Figure 7: VCO's carrier frequency versus its control voltage. mmX's hardware platform operates over a wide range of frequency, covering an entire 24GHz ISM band.

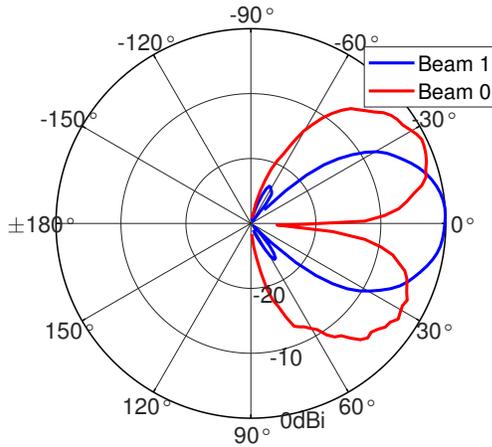


Figure 8: Measured beam patterns of mmX's node. Beam 0 and Beam 1 are orthogonal to each other (i.e., Beam 0 has a null at the peak of Beam 1, and Beam 1 has nulls at the peaks of Beam 0).

to slightly vary the frequency in order to perform the joint ASK-FSK modulation.

The maximum operating frequency of the RF switch is 100 MHz, which limits the data rate of mmX's nodes to 100 Mbps. This is much higher than what most high-data rate applications require. For example, HD video cameras require only 8-10 Mbps. Finally, mmX's node consumes 1.1 W which results in an energy efficiency of 11nj/bit at 100 Mbps. We believe that mmX's data rate and energy efficiency can be further improved by using a faster RF switch or designing an application-specific integrated circuit (ASIC).

Node's Antenna Performance: mmX's nodes use two orthogonal beams. Figure 8 shows the measured azimuth radiation patterns of the antenna arrays for the node. The antennas are designed and fabricated on RO4835 substrate and their radiation patterns are measured by a near field antenna measurement facility

in an anechoic chamber. As can be seen in Figure 8, Beam 1 has a peak at the broadside ($\theta = 0$), and Beam 0 has peaks at $\pm 30^\circ$. Moreover, Beam 0 has a very low magnitude on the main lobe of the Beam 1 and vice versa. Therefore, the two radiation beams are orthogonal to each other and create the minimum overlap. The elevation radiation pattern of each beam is similar to a single patch antenna (i.e., a wide beam with 3 dB bandwidth of 65°). This allows the node to work at different height with respect to the AP. The azimuth 3 dB beamwidth of each beam is 40° . Our results show that the node's field of view is 120° in front side of the node, and the maximum range is 18 m. Note that, depending on the use case, one can design narrower beams to improve the range at the cost of narrower field of view. Furthermore, one can easily extend the node's field of view to the back side of the node by incorporating additional patch antennas.

Joint ASK-FSK modulation: Figure 9(a) shows an example of a measured signal at the AP. As shown in the figure, the signal can be decoded using ASK demodulation. However, as explained in Section 6, there is a possibility that the paths for Beam 0 and Beam 1 experiences similar attenuation, and hence, the amplitude of the carrier signal is the same for bit 0 and bit 1. In this case, the SNR of the ASK signal will be very low for the AP to decode it. Figure 9(b) shows an example of a measured signal in such a scenario. Although our results show that the possibility of this happening is minor ($< 10\%$), mmX addresses this problem by joint ASK-FSK modulation such that the AP can always decode the signal. As illustrated in Figure 9(b), the frequency has slightly altered between the bits, and hence the signal can easily be decoded using FSK demodulation in this case.

9.2 mmX's SNR Performance

First, we evaluate the performance of mmX and the OTAM scheme in enabling robust mmWave links between nodes and AP. We conduct experiments in a $6m \times 4m$ room. We place mmX's AP on one side of the room and we place a mmX's node at random locations and heights. For each location, the orientation of the mmX's node (respect to the AP) is randomly picked between -60 and 60 degrees. We also asked people to walk around. In order to block the signal, one person was blocking the line-of-sight path between the node and the AP for the entire duration of the experiment. We then measured the SNR at the AP for two different scenarios: (1) without OTAM, in which the mmX's node utilizes only Beam 1 and transmit an ASK-FSK signal.; (2) with OTAM, in which the mmX's node utilizes both beams to create an ASK-FSK signal over the air, as explained in Section 6. Note that in the first scenario, the modulation is done at the node while in the second scenario, the modulation is done over the air.

Figure 10 plots the results of this experiment. The figure shows the SNR (at the AP) for different node locations in two different scenarios. Figure 10 (a) indicates that when the node sends the modulated signal through Beam 1, there are many locations with SNRs below 5 dB. On the other hand, Figure 10 (b) shows that for the same locations, when the node uses the OTAM scheme to modulates the signal over the air, the SNR is significantly improved. Specifically, with OTAM, mmX achieves SNRs of more than 11 dB

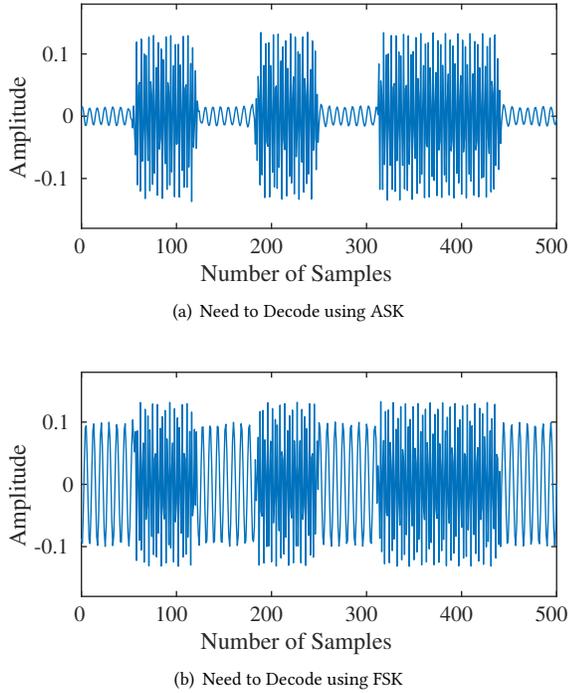


Figure 9: An example of the measured signal received at the AP. (a) Beam 0 and Beam 1 paths experience different losses, the AP can decode the signal using ASK demodulation. (b) Beam 0 and Beam 1 experiences similar losses. Combining FSK with ASK helps in decoding the signal.

in almost all locations, enabling a very low BER. These results show that the OTAM scheme significantly improves the link SNR and enables mmX’s nodes to effectively communicate with the AP without needing to search for the best beam.

9.3 mmX’s BER Performance

Next, we evaluate the performance of mmX in terms of bit-error-rate (BER). As in the previous experiment, we measure the SNR from 30 different locations, heights and orientations of nodes in the same testbed. Then, we compute the BER by substituting the SNR measurements into standard BER tables based on the ASK modulation [43].

Figure 11 shows the CDF of the BER for two scenarios: 1) without OTAM and 2) with OTAM as described in Section 9.2. The figure shows that without OTAM, the median and 90th percentile BER are 10⁻⁵ and 0.3, respectively. The figure also shows that OTAM significantly improves the BER of mmX network. Specifically, with OTAM, the median and 90th percentile BER are 10⁻¹² and 10⁻³, respectively. This physical BER is acceptable for most wireless applications and it can be reduced even further by using an error correction coding scheme.

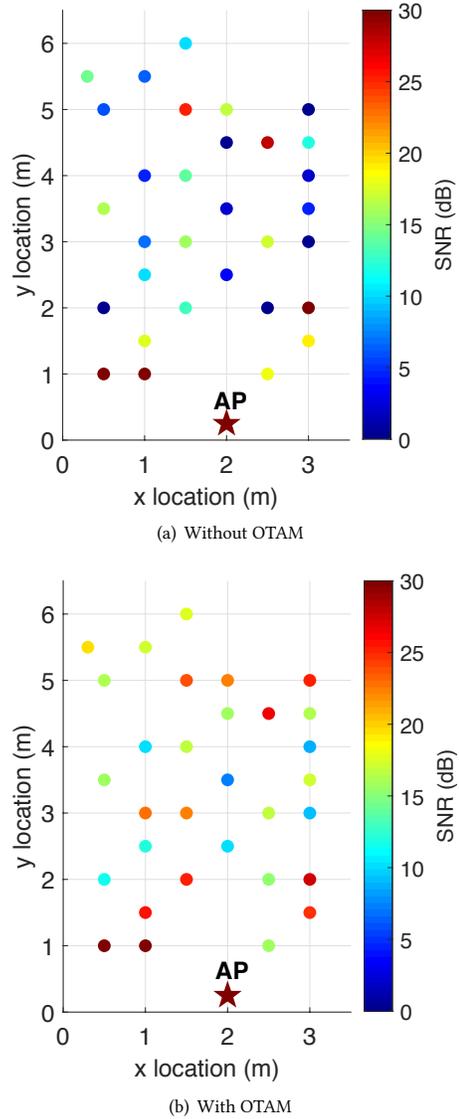


Figure 10: SNR of mmX’s nodes at the AP. SNR of mmX’s nodes at the mmX’s AP for two different scenarios. (1) without OTAM, where the node chooses Beam 1 and transmits ASK signal, and (2) with OTAM, where the node creates ASK signal over the air. The figure shows that without OTAM, for many cases the SNR is below 5 dB, resulting in a higher bit error rate. In contrast, with OTAM, mmX achieves SNR of more than 10 dB in all locations

9.4 mmX’s Range Performance

We now explore the impact of distance between the mmX’s node and AP on the SNR. We measure the SNR of the received signal at the AP while we change the distance between the AP and the node. For each location, we run experiments for two different scenarios: 1) the node is facing toward the AP, where the center beam has a line-of-sight toward the AP, and 2) the node is not facing toward the AP. Figure 12 shows the results for this experiment. As anticipated,

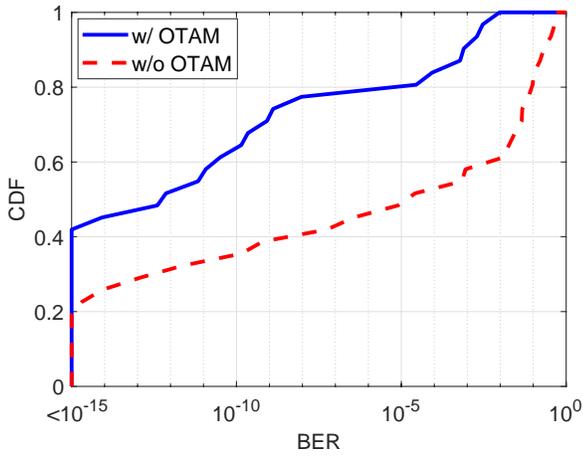


Figure 11: mmX's BER Performance. OTAM scheme significantly improves the BER.

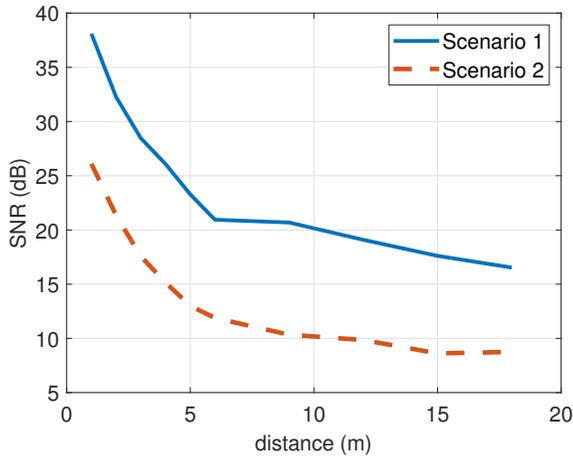


Figure 12: mmX's coverage. SNR at the AP versus distance between AP and the node for two different scenarios: 1) node is facing toward the AP, and 2) node is not facing toward the AP.

increasing the distance reduced the SNR. However, even at 18 meters, mmX provides SNRs of more than 15 dB, which is sufficient to achieve BER of lower than 10^{-8} . The figure also shows that the SNR slightly degrades when the node does not face toward the AP. This is expected since the beams are orthogonal, and in this case, only one arm of the side beam is directed to the AP. However, even at 18 meters, mmX still achieves SNRs as high as 9 dB. These results show that mmX provides a robust wireless link even when the AP and nodes are 18 m far from each other, which is enough to connect IoT sensors to an access point in a smart home, autonomous cars and many other IoT applications.

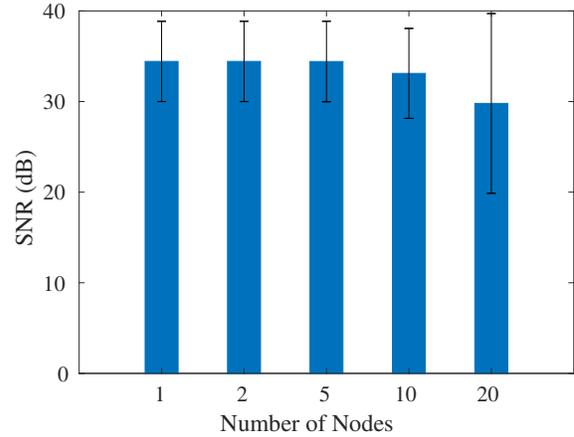


Figure 13: mmX's multi-node performance. Number of nodes transmitting simultaneously versus their SNR at the AP. mmX enables robust communication links even when 20 nodes simultaneously transmit to a single AP.

9.5 mmX's Network Performance

So far, we have evaluated the performance of mmX when a single node communicates to the AP at any point in time. We now evaluate the performance of mmX as a network, when multiple nodes communicate with the AP, simultaneously. We place the AP in one side of the room, and we place the nodes in random locations and orientations. We measure the SNR of their signal (at the AP) while multiple nodes communicate with the AP, simultaneously. We run 100 experiments. Due to limitations of USRPs, we cannot capture the entire bandwidth occupied by all nodes, and hence we do not implement Spatial Division Multiplexing (SDM) in hardware. However, we collect measurements from sub-bands (i.e., 25MHz occupied by each node), and we combine them in post-processing in order to simulate the effect of Frequency Division Multiplexing (FDM) and Spatial Division Multiplexing (SDM). Figure 13 shows the result of this experiment. As the number of nodes which simultaneously transmit increases, their SNR slightly decrease. This is expected since they create some interference for each other. However, even when 20 sensors transmit simultaneously, their average SNR is higher than 29 dB. Such results show that mmX enables a robust mmWave network for multiple nodes, even when they communicate simultaneously.

10 DISCUSSION

In this section, we compare mmX with existing wireless systems such as WiFi, Bluetooth, and other mmWave platforms. Specifically, we compare these systems in terms of cost, power consumption, throughput, range, carrier frequency, and energy efficiency. Table 1 shows the results of this comparison.

Comparison with other mmWave platforms: Past mmWave platforms such as MiRa and OpenMili cost a few thousand dollars and consume a few tens of watts. On the other hand, mmX costs only \$110 and consumes 1.1 watt. However, past mmWave platforms

	mmX	MiRa[5]	OpenMili/Pasternack [32, 47]	WiFi (802.11n) [15, 22]	Bluetooth
Carrier Frequency	mmWave(24 GHz)	mmWave(24 GHz)	mmWave(60 GHz)	2.4 GHz	2.4 GHz
Cost	\$110	\$7,000	\$8,000	\$10	\$10
Power Consumption	1.1 W	11.6 W	5 W (w/o phased array)	2.1 W	0.029 W
Transmission Power	10 dBm	10 dBm	12 dBm	30 dBm	5 dBm
Bandwidth	250 MHz	250 MHz	1 GHz	70 MHz	1 MHz
PHY-layer Bitrate	100 Mbps (at 18m)	1 Gbps (at 18m)	1.3 Gbps	120 Mbps (at 18m)	1 Mbps
Energy efficiency (nJ/bit)	11	11.6	3.8	17.5	29
Range	18 m	100 m	11 m	50 m	10 m

Table 1: Comparison of mmX with existing mmWave platforms and other wireless systems

provide Gbps throughput while mmX provides 100 Mbps. Therefore, existing mmWave platforms are suitable for applications which require multiple Gbps throughput, while having substantial energy and computing power. In contrast, mmX targets applications that require less than 100 Mbps throughput while they are low power and low cost.

Comparison with WiFi and Bluetooth: The main advantage of mmX compared to WiFi (ex. 802.11n) is that it utilizes mmWave spectrum (24 GHz) rather than WiFi spectrum (2.4 GHz). Utilizing mmWave spectrum will remove a huge strain from today's WiFi spectrum. Furthermore, as shown in Table 1, the power consumption and bitrate of mmX is in the same range as WiFi. Note that the reported WiFi performance is for an ideal scenario. In fact, most of today's WiFi networks have much lower performance since their spectrum is overloaded. Therefore, mmX provides similar performance to ideal WiFi network while utilizing mmWave spectrum. Although, current mmX prototype costs more than existing WiFi modules, the cost can be significantly reduced in mass production.

In comparison with Bluetooth, mmX provides much higher bitrate. Specifically, Bluetooth provides only 1Mbps which is not sufficient for many IoT applications. On the other hand, mmX provides up to 100 Mbps.

11 CONCLUSION

This paper introduces mmX, a low-power, low-cost mmWave network for IoT devices. In particular, mmX overcomes fundamental challenges that prevent existing mmWave systems from being used in low-power, low-cost IoT devices. mmX introduces the first mmWave low-power hardware platform which operates as a daughterboard for RaspberryPi. We believe that this platform helps advance mmWave research in the IoT domain. In addition, mmX introduces OTAM, a novel technique to modulate the signal over the air. OTAM eliminates the need for costly phased array and beam searching techniques, making adaptation of mmWave communication easier and less costly. Finally, mmX can be used in many applications (e.g., smart home, autonomous cars, etc.) to connect sensors to an access point without placing any strain on today's WiFi spectrum.

ACKNOWLEDGMENT

We thank Shashank Goel for his help in programming of Raspberry Pi boards. We also thank our shepherd, Fadel Adib, and the anonymous SIGCOMM reviewers for their feedback and in-sights. We thank NSERC for their support.

REFERENCES

- [1] 2015. 71-76 GHz Millimeter-wave Transceiver System. National Instruments.
- [2] 2018. Introduction to the NI mmWave Transceiver System Hardware. National Instruments. <http://www.ni.com/white-paper/53095/en/>
- [3] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2016. Cutting the cord in virtual reality. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. ACM, 162–168.
- [4] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2017. Enabling High-Quality Untethered Virtual Reality. In *NSDI*. 531–544.
- [5] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. 2016. Poster: A millimeter wave software defined radio platform with phased arrays. In *MobiCom*. ACM, 419–420.
- [6] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. 2016. Millimeter wave communications: From point-to-point links to agile network connections. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. ACM, 169–175.
- [7] Analog Devices [n. d.]. *GaAs pHEMT MMIC Power Amplifier, DC - 28 GHz*. Analog Devices. v04.0218.
- [8] Analog Devices [n. d.]. *HMC342, GaAs MMIC Low Noise Amplifier, 13-25 GHz*. Analog Devices. v01.0907.
- [9] Analog Devices [n. d.]. *HMC533, MMIC VCO w/ Divide-by-16, 23.8-24.8 GHz*. Analog Devices. v00.0405.
- [10] Analog Devices [n. d.]. *HMC644A, GaAs MMIC 5-Bit digital phase shifter, 15 - 18.5 GHz*. Analog Devices. v00.0516.
- [11] Analog Devices [n. d.]. *HMC933, 470° Analog Phase shifter, 18 - 24 GHz*. Analog Devices. v02.0211.
- [12] Analog Devices 2017. *ADRF5020, 100 MHz to 30 GHz, Silicon SPDT Switch*. Analog Devices. Rev. A.
- [13] Analog Devices 2017. *HMC8191, 6 GHz to 26.5 GHz, Wideband I/Q Mixer*. Analog Devices. Rev. B.
- [14] Analog Devices 2018. *HMC815B, 21 GHz to 27 GHz, GaAs, MMIC, I/Q Upconverter*. Analog Devices. Rev. 0.
- [15] Cisco. [n. d.]. *Wireless Mesh Constraints*. ([n. d.]). https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/7-3/design/guide/Mesh/Mesh_chapter_011.pdf
- [16] Tesla Company. [n. d.]. *Full Self-Driving Hardware on All Cars*. ([n. d.]). <https://www.tesla.com/autopilot>
- [17] Yong Cui, Shihan Xiao, Xin Wang, Zhenjie Yang, Shenghui Yan, Chao Zhu, Xiang-Yang Li, and Ning Ge. 2018. Diamond: Nesting the data center network with wireless rings in 3-d space. *IEEE/ACM Trans. Networking* 26, 1 (2018), 145–160.
- [18] Mohammed Elbamby, Cristina Perfecto, Mehdi Bennis, and Klaus Doppler. 2018. Toward Low-Latency and Ultra-Reliable Virtual Reality. *IEEE Network* 32, 2 (2018), 78–84.
- [19] Mohammed E Eltayeb, Ahmed Alkhateeb, Robert W Heath, and Tareq Y Al-Naffouri. 2015. Opportunistic beam training with hybrid analog/digital codebooks for mmWave systems. In *GlobalSIP*. IEEE, 315–319.
- [20] Zhen Gao, Linglong Dai, De Mi, Zhaocheng Wang, Muhammad Ali Imran, and Muhammad Zeeshan Shakir. 2015. MmWave massive-MIMO-based wireless

- backhaul for the 5G ultra-dense network. *IEEE Wireless Communications* 22, 5 (2015), 13–21.
- [21] Muhammad Kumail Haider and Edward W Knightly. 2016. Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: protocol design and system implementation. In *MobiHoc*. ACM, 61–70.
- [22] Daniel Halperin, Ben Greenstein, Anmol Sheth, and David Wetherall. 2010. Demystifying 802.11N Power Consumption. In *HotPower*.
- [23] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. 2011. Augmenting data center networks with multi-gigabit wireless links. In *SIGCOMM*, Vol. 41. ACM, 38–49.
- [24] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast millimeter wave beam alignment. In *SIGCOMM*. ACM, 432–445.
- [25] Chong He, Xianling Liang, Bin Zhou, Junping Geng, and Ronghong Jin. 2015. Space-division multiple access based on time-modulated array. *IEEE Ant. Wirel. Prop. Lett.* 14 (2015), 610–613.
- [26] Bin Li, Zheng Zhou, Weixia Zou, Xuebin Sun, and Guanglong Du. 2013. On the efficient beam-forming training for 60GHz wireless personal area networks. *IEEE Trans. Wirel. Comm.* 12, 2 (2013), 504–515.
- [27] Robert J Mailloux. 2017. *Phased array antenna handbook*. Artech house.
- [28] Michael Miller. 2015. *The internet of things: How smart TVs, smart cars, smart homes, and smart cities are changing the world*. Pearson Education.
- [29] Andreas F Molisch, Vishnu V Ratnam, Shengqian Han, Zheda Li, Sinh Le Hong Nguyen, Linsheng Li, and Katsuyuki Haneda. 2017. Hybrid beamforming for massive MIMO: A survey. *IEEE Communications Magazine* 55, 9 (2017), 134–141.
- [30] Evan Nisselson, Abigail Hunter-Syed, and Sadhana Shah. 2017. 45 billion cameras by 2022 fuel business opportunities. In *tech. rep., LDV Capital, New York, NY*.
- [31] Yong Niu, Yong Li, Depeng Jin, Li Su, and Athanasios V Vasilakos. 2015. A survey of millimeter wave communications (mmWave) for 5G: opportunities and challenges. *Wireless Networks* 21, 8 (2015), 2657–2676.
- [32] Pasternack. [n. d.]. 60 GHz Development System, Transmit (Tx), Low Phase Noise Development. ([n. d.]). <https://www.pasternack.com/60-ghz-development-system-low-phase-noise-pem009-kit-tx-p.aspx>
- [33] Zhouyue Pi and Farooq Khan. 2011. An introduction to millimeter-wave mobile broadband systems. *IEEE communications magazine* 49, 6 (2011).
- [34] Lorenzo Poli, Paolo Rocca, Giacomo Oliveri, and Andrea Massa. 2011. Harmonic beamforming in time-modulated linear arrays. *IEEE Tran. Ant. Prop.* 59, 7 (2011), 2538–2545.
- [35] Sundeep Rangan, Theodore S Rappaport, and Elza Erkip. 2014. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proc. IEEE* 102, 3 (2014), 366–385.
- [36] Theodore S Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N Wong, Jocelyn K Schulz, Mathew Samimi, and Felix Gutierrez Jr. 2013. Millimeter wave mobile communications for 5G cellular: It will work! *IEEE access* 1, 1 (2013), 335–349.
- [37] Wonil Roh, Ji-Yun Seol, Jeongho Park, Byunghwan Lee, Jaekon Lee, Yungsoo Kim, Jaeweon Cho, Kyungwhoon Cheun, and Farshid Aryanfar. 2014. Millimeter-wave beamforming as an enabling technology for 5G cellular communications: theoretical feasibility and prototype results. *IEEE Communications Magazine* 52, 2 (February 2014), 106–113.
- [38] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Roshan Shyamsunder, Owen Torres, Daniel Uvaydov, Josep Miquel Jornet, Edward Knightly, Dimitrios Koutsonikolas, Dimitris Pados, and Zhi Sun. 2019. X60: A programmable testbed for wideband 60 ghz wlans with phased arrays. *Computer Communications* 133 (2019), 77–88.
- [39] Swetank Kumar Saha, Tariq Siddiqui, Dimitrios Koutsonikolas, Adrian Loch, Joerg Windmer, and Ramalingam Sridhar. 2017. A detailed look into power consumption of commodity 60 GHz devices. In *WoWMoM*. 12–15.
- [40] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2018. Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. ACM, 257–271.
- [41] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz indoor networking through flexible beams: A link-level profiling. In *SIGMETRICS*, Vol. 43. ACM, 71–84.
- [42] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *NSDI*. 193–206.
- [43] Qinghui Tang, Sandeep KS Gupta, and Loren Schwiebert. 2005. BER performance analysis of an on-off keying based minimum energy coding for energy constrained wireless sensor applications. In *IEEE Internat. Conf. Com.*, Vol. 4. IEEE, 2734–2738.
- [44] Y Ming Tsang, Ada SY Poon, and Sateesh Addepalli. 2011. Coding the beams: Improving beamforming training in mmwave communication system. In *IEEE GLOBECOM*. IEEE, 1–6.
- [45] CH Yeh, Yun-Fu Liu, Chi-Wai Chow, Y Liu, PY Huang, and Hon Ki Tsang. 2012. Investigation of 4-ASK modulation with digital filtering to increase 20 times of direct modulation speed of white-light LED visible light communication system. *Optics Express* 20, 15 (2012), 16218–16223.
- [46] Wenfang Yuan, Simon MD Armour, and Angela Doufexi. 2015. An efficient and low-complexity beam training technique for mmWave communication. In *PIMRC*. IEEE, 303–308.
- [47] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. 2016. OpenMili: a 60 GHz software radio platform with a reconfigurable phased-array antenna. In *MobiCom*. ACM, 162–175.
- [48] Anfu Zhou, Leilei Wu, Shaoqing Xu, Huadong Ma, Teng Wei, and Xinyu Zhang. 2018. Following the shadow: Agile 3-D beam-steering for 60 GHz wireless networks. In *IEEE INFOCOM*. IEEE, 2375–2383.
- [49] Yibo Zhu, Xia Zhou, Zengbin Zhang, Lin Zhou, Amin Vahdat, Ben Y. Zhao, and Haitao Zheng. 2014. Cutting the Cord: A Robust Wireless Facilities Network for Data Centers. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking (MobiCom '14)*. ACM, 581–592.