Bringing Millimeter Wave Technology to Any IoT Device

Mohammad Mazaheri
UCLA
USA

Rafael Ruiz
IMDEA Networks
Spain

Domenico Giustiniano
IMDEA Networks
Spain

Joerg Widmer
IMDEA Networks
Spain

Omid Abari
UCLA
USA

ABSTRACT

With the advancement of the Internet of Things (IoT), many devices will be connected to the Internet, enabling digital twin and smart home applications. However, currently, these IoT devices are operating at lower frequency bands of the wireless spectrum, typically ranging from a few hundred MHz (such as RFID and LoRa) to a few GHz (such as BLE and WiFi). As a result, the current IoT devices not only place a huge strain on these bands, but also cannot benefit from the large bandwidth available in the higher frequencies of the spectrum such as mmWave bands. In this paper, our goal is to bring mmWave technology to *existing* IoT devices so they can benefit from the advantages this technology offers, such as high network capacity, low interference, and Space Division Multiple Access. To this end, we design mmPlug, a novel plug-and-play module which is simple and energy-efficient. mmPlug can be easily connected to the antenna port of any IoT device, enabling it to operate in the mmWave band. mmPlug is compatible with different wireless technologies (such as WiFi, Lora, etc.) and does not require any modification to the circuit, firmware or communication protocols of the existing IoT devices. mmPlug achieves this by a novel design which can seamlessly be connected to the antenna port of the IoT device. We have implemented mmPlug on PCB and empirically evaluated its performance. Our results show that mmPlug enables existing IoT devices (such as WiFi and Lora) to operate at mmWave band while achieving accurate localization, uplink and downlink even when they are more than 30 m far from the access point.

ACM MobiCom '23, October 2–6, 2023, Madrid, Spain

ACM ISBN 978-1-4503-9990-6/23/10.
https://doi.org/10.1145/3570361.3613255

CCS CONCEPTS

• **Hardware** → **Wireless devices**; • **Networks** → **Mobile networks**.

KEYWORDS

Wireless; Millimeter wave (mmWave); Internet of Things (IoT); Low power;

ACM Reference Format:
Mohammad Mazaheri, Rafael Ruiz, Domenico Giustiniano, Joerg Widmer, and Omid Abari. 2023. Bringing Millimeter Wave Technology to Any IoT Device. In *The 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23), October 2–6, 2023, Madrid, Spain.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3570361.3613255

1 INTRODUCTION

Over the past several years, there has been significant interest in developing efficient millimeter-wave (mmWave) networks [21, 46, 47]. These networks have multiple advantages over traditional wireless networks which make them very attractive for emerging applications. First, they can provide very high network capacity by operating over a large bandwidth in the high-frequency spectrum bands (24 GHz and above) [17, 26, 40, 43]. Second, the use of highly directional mmWave antennas enables space division multiplexing and can provide connectivity to many nodes simultaneously over the same bandwidth, which results in much more efficient spectrum usage [6]. Finally, such directional antenna arrays together with the large bandwidth and short wavelength enable very accurate device localization [12, 32, 45].

In parallel, the number of Internet-of-Things (IoT) devices grows at fast pace. The current IoT networks at sub-6GHz band are suitable for a range of existing applications, but they have some limitations, specifically when the number of connected nodes increases. In fact, these networks may not always scale to emerging applications such as massive

IoT networks. In addition, the sub-6GHz band is heavily congested and the networks are facing spectrum crunch. mmWave technology can provide the large spectrum and directionality for frequency and spatial reuse needed to enable massive IoT. To resolve the issues with the lower frequency band we migrate the sub-6GHz to mmWave band, where spectrum is usually not congested and directional communication increases capacity through Space Division Multiple Access (SDMA). Therefore, it can enable very high density networks for the currently deployed IoT nodes that do not interfere with other deployed sub-6GHz networks. Despite the advantages of mmWave technology, most IoT devices do unfortunately not benefit from it due to the high complexity, cost and energy consumption of mmWave radios. Hence, there is a need for new designs which can bring mmWave technology to IoT devices.

To address this problem, researchers have designed mmWave modules and systems targeted for IoT devices [30, 31]. Although these systems enable low-power, low-cost mmWave communication, they have major limitations. First, they provide only an uplink, and do not support two-way communication. Second, they do not work as a plug-and-play module since they have their own protocols, and hence cannot easily be integrated and used in today's IoT devices. Ideally, a mmWave module for IoT devices should satisfy the following key requirements:

- **Plug-and-play Integration:** First, it must be a plug-and-play component which can be easily attached to existing wireless IoT nodes to allow the upgrade of deployed legacy networks and ensure rapid market adoption. Ideally, the integration should be as easy as disconnecting the IoT node's antenna, and connecting the mmWave module to the same port. This allows any existing IoT nodes to seamlessly operate at the mmWave band and benefit from its advantages of high network capacity, accurate localization and space division multiplexing.
- **Simple and Low-power:** The module needs to be simple and low-power, to meet the typical requirements of IoT devices. Unfortunately, the power consumption of most mmWave components is high. The problem is intrinsic as most of the components of a radio consume much more power when they run at such high frequencies. Furthermore, mmWave radios are costly compared to most IoT devices. Hence, the design of the mmWave module needs to be very simple and use only the minimum number of mmWave components to reduce cost and power consumption.
- **Compatibility:** Finally, the design must enable both uplink and downlink operation, since otherwise it will not be compatible with most IoT applications. It should also be seamlessly interoperable with different IoT technologies. In particular, it should be able to enable WiFi, LoRa and other IoT technologies to operate at mmWave bands. Being able to retrofit existing IoT devices allows to leverage chipset vendors, device integrators, end customers and operators of the lower frequency IoT ecosystem, to address mass market requirements.

With the above requirements satisfied, it is possible to bring mmWave technology and its advantages to potentially any IoT device as well as today's low-power low-frequency SDRs such as TinySDR [20]. To be able to fulfill these requirements, we design mmPlug, a mmWave module with a unique design, which allows IoT nodes to transmit and receive in the mmWave band using their original RF interface (WiFi, LoRa, etc.) *without any modifications* to their chipsets, firmware or protocol. For data transmission, the mmPlug module translates the original IoT signal to the mmWave band, and transmits it over the air. The mmWave signal is then received and decoded by a mmWave access point that supports the IoT communication protocol. For reception, the module receives the modulated mmWave signals from the mmWave access point, and translates them back to the IoT device's original low frequency band. mmPlug module does all this using only a *single mmWave amplifier* without any further active mmWave RF components, together with a novel passive PCB design. To achieve this, mmPlug introduces two key innovations:

1) Leveraging Mutual Coupling to Generate mmWave Signals. Mutual coupling occurs when an antenna gets too close to another antenna and impacts their performance. Typically, this behavior is considered harmful and is thus avoided (as much as possible) in RF design. In contrast, mmPlug leverages mutual coupling to enable an amplifier to generate and amplify a mmWave signal at a specific frequency. In particular, we design two patch antennas and connect them to the input and output ports of an amplifier. Then, by optimizing their mutual coupling, we create a loop between the output and input of the amplifier, causing it to resonate. To ensure that the amplifier resonates only at the target mmWave frequency, we carefully design the mutual coupling of the antennas such that they have the maximum coupling at the required frequency. Said differently, we optimize the mutual coupling between the antennas such that the system works as a bandpass loop which has the least attenuation at the required mmWave frequency. Hence, even a small signal such as noise will cause the amplifier to resonate and generates a tone at the desired mmWave frequency.

2) Leveraging Amplifier Saturation to Translate Frequencies. Once the amplifier starts resonating at the desired frequency, the resonating mmWave tone will continue growing until the amplifier reaches saturation. This is happening because of the loop caused by the mutual coupling, and the

amplification gain provided by the amplifier. Once the amplifier saturates, it enters a very non-linear region. This strong non-linearity will result in the intermodulation of the amplifier's input signal. Since in the downlink the amplifier input signal is the sum of the received modulated mmWave signal and the mmWave resonating signal, the intermodulation results in the multiplication of the two signals, which down-converts the received modulated mmWave signal to the lower frequency band. Finally, we design a passive PCB based diplexer to extract the low-frequency signal from the other signals in the loop, and feed it to the antenna port of the IoT device. A similar process is happening for the uplink mode, where the loop and amplifier translate the IoT signal to the mmWave band.

In this paper, we make the following contributions:

- We introduce the first low-power and plug-and-play mmWave module which enables any IoT device to operate at mmWave bands.
- We develop a technique which leverages (typically undesired) antenna coupling and amplifier saturation behavior to translate low frequency signals to mmWave and vice versa.
- We implement mmPlug using an off-the-shelf amplifier and standard PCB board. We experimentally evaluate the performance of our design and also demonstrate its practicality by integrating it with off-the-shelf WiFi and LoRa IoT devices. Our results show that even when IoT devices are 30 m away from a mmWave AP, mmPlug enables them to communicate robustly at mmWave frequency bands.
- We employ the high-speed Analog to Digital Converters (ADCs) within the mmWave AP for precise localization of mmPlug nodes. Our findings demonstrate the AP's capability to determine the Angle of Arrival for a node with a precision surpassing 1 degree. Furthermore, the distance error is under 40cm in an indoor environment.

Paper outline: We describe the overview of mmPlug in Section 2. We detail mmPlug's design in Section 3. The implementation is described in Section 4, followed by micro-benchmark experiments and the evaluation in Section 5. We discuss the operation of mmPlug in a network of IoT devices in Section 6 and the related work in Section 7, and conclude our work in Section 8.

2 SYSTEM OVERVIEW

mmPlug is a plug-and-play mmWave module for any IoT devices. It enables any existing and future IoT devices to communicate using mmWave spectrum. This enables IoT devices to benefit from the advantages of operating at mmWave frequencies, such as availability of high network capacity, accurate localization, and the capability of performing space division multiplexing by exploiting the sparse mmWave multipath and the directionality of mmWave antennas. Moreover, operating at mmWave band reduces the strain on the lower part of the spectrum as the number of IoT devices grows.

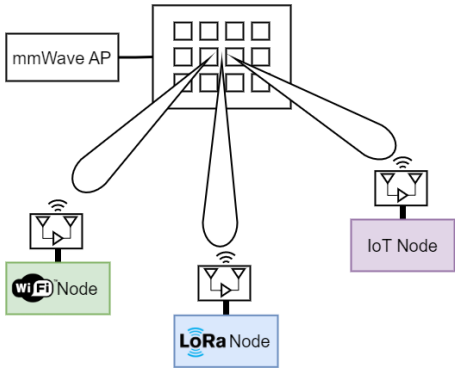

Figure 1: mmPlug Overview. mmPlug's module can easily be connected to the antenna port of off-the-shelf IoT devices (such as WiFi, LoRa, etc.), and enables them to operate at mmWave spectrum and communicate to a single mmWave AP.

Figure 1 shows mmPlug's setup, where multiple IoT nodes (each equipped with an mmPlug module) communicate to a single mmWave access point (AP) which supports different communication protocols (such as LoRa, WiFi, etc). mmPlug's module operates as a plug-and-play device and can be easily connected to IoT devices (instead of their antenna). All that is required is to remove the IoT device antenna, and connect an mmPlug's module to the same antenna port. mmPlug's module enables both uplink and downlink communication, and has on-board mmWave patch antennas. Moreover, due to the directional property of mmWave communication, the mmWave AP can perform space division multiplexing. This allows multiple nodes to communicate with an AP simultaneously, without creating any interference for other devices.

In the next sections, we will explain the details of the mmWave module to meet IoT requirements.

3 MMPLUG DESIGN

At a high-level, mmPlug module needs to translate the IoT signal (ranging from hundreds of MHz to a few GHz) to a suitable mmWave frequency and vice versa. However, for this module to be practical and ubiquitously deployable, it must be simple and energy-efficient, while allowing for plug and play operation. Hence, our design needs to use the minimum number of active components and chips. In particular, compared to components operating at lower frequencies,

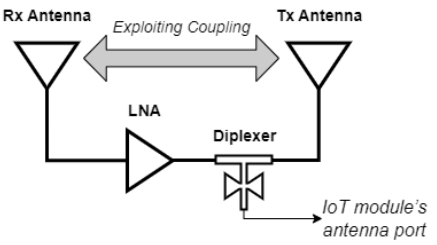

Figure 2: mmPlug's module design. mmPlug's module uses only one active component (i.e. LNA). It also includes two mmWave patch antennas and a diplexer which are passive and can be built using metal traces on the same PCB.

mmWave components cost significantly more, and consume much higher power. Unfortunately, this is a fundamental problem since the power consumption of RF circuits is proportional to their operating frequencies [37]. Hence, ideally, most part of the mmPlug's design should be passive and be implemented as metal traces on the Printed Circuit Board (PCB). To achieve this, we propose a novel design which uses a single active mmWave component. Figure 2 shows the block diagram of our design. It consists of a single LNA (Low Noise Amplifier). It also includes two patch antennas and a simple diplexer that allow to extract/inject the low frequency signal from/to the mmWave circuit. A diplexer is a passive component and can be implemented directly as metal traces on the PCB. In this design, the coupling between the antennas creates a feedback loop which causes the LNA to resonate. By engineering and reshaping the coupling effect, we make sure this resonance is at the required mmWave frequency.

In the uplink mode, the IoT signal is fed to the loop through the diplexer. Note that the diplexer works as a mmWave isolator to ensure that the mmWave signal does not leak to the IoT device, while the lower frequency signals pass through it. Due to the high loop gain, the LNA is already saturated and operates in a strong non-linear mode. Hence, it naturally mixes the IoT signal with the mmWave resonance signal. The result is a mmWave signal, modulated by the IoT signal. Finally, we design a patch antenna (TX antenna) to transmit the modulated mmWave signal. In the downlink mode, mmPlug operates similar to the uplink mode. In particular, its RX antenna receives a modulated mmWave signal from the AP. The non-linearity of the LNA mixes the received signal with the resonating mmWave signal, resulting in a modulated signal at IoT's low-frequency band. Finally, the diplexer at the output of the LNA passes the low-frequency signal to the IoT device. In the remainder of this section, we will discuss our design in more detail.

3.1 Leveraging Mutual Coupling of Antennas

The first step to translate IoT RF signals to mmWave and vice versa is to generate a mmWave carrier signal. To do so, our idea is to leverage mutual coupling between antennas to cause an LNA to resonate at a specific mmWave frequency. In particular, as shown in Figure 2, we connect the input and output of the LNA to two antennas which are placed very close to each other. These closely spaced antennas provide high mutual coupling. Generally, coupling is considered as an undesired phenomenon in any transmitter or receiver. However, here we take advantage of it to create a feedback loop between the output and input of the LNA, causing it to resonate. In particular, the mutual coupling between the antennas passes an attenuated version of the signal at the output of the LNA to its input. Hence, as long this attenuation is smaller than the gain of the LNA (i.e., the loop gain created by the LNA and mutual coupling of antennas is higher than unity), even a small signal such as thermal noise is enough to cause the LNA to resonate and thus create the resonating signal. Hence, we need to achieve sufficiently high mutual coupling between the antennas to allow the LNA to resonate. Conversely, if the mutual coupling becomes too strong, it significantly impacts the performance of the antennas as radiating and receiving elements. Therefore, the mutual coupling has to be meticulously designed to be high enough for a loop gain larger than one, while being low enough to not have a significant impact on the performance of the antennas. The amount of coupling can be adjusted through the spacing and relative location of two antennas. In particular, by increasing the distance between the antennas, their mutual coupling reduces, meaning that less signal power will be fed back to the input of the LNA from its output. So far, we explained how we can create a resonating signal using the antennas' mutual coupling and a single LNA. However, the question is how can we make sure that the frequency of the resonating signal is exactly at the required mmWave frequency and not any other frequency. To answer this question, we first provide some background on Scattering parameters (S-parameters) and then explain our solution in detail.

Background on Scattering Parameters: Scattering Parameters (also known as S-Parameters) are used to define the relationship between different ports of an RF component (or a system) in terms of amplitude and phase. In fact, S-Parameters can be used to describe the behavior of a simple system which has one port (such as antenna), or a more complex system which has many ports. In general, S_{nm} represents the power transferred from port m to port n of a system. For example, S_{21} represents the power transferred from port 1 to port 2, and S_{11} represent the power which is reflected by port 1 of the system. In particular, antennas

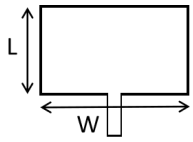

Figure 3: A typical patch antenna. it can be fabricated using metal traces on a PCB.

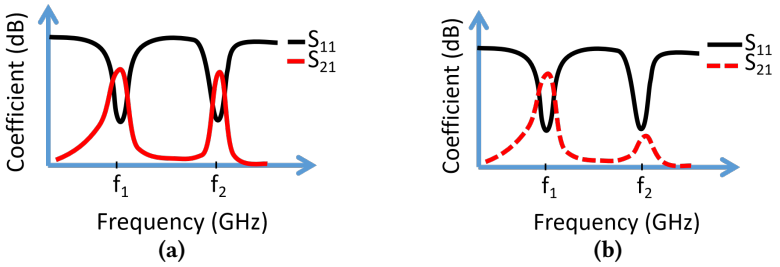

Figure 4: Return Loss Coefficient (S_{11}) and Mutual Coupling coefficient (S_{21}) of two antennas (a) without, and (b) with antenna location optimization. By optimizing the mutual coupling of the antennas, mmPlug ensures that the LNA and loop resonate only at f_1.

should have low S_{11} in their operating frequency, meaning that they radiate most of power without reflecting it back to the radio circuits.

We now explain our approach to configure the frequency of the resonating signal at exactly the required mmWave frequency and not other frequencies. The frequency of the resonating signal is determined by the frequency response of the two antennas. In particular, an antenna acts like a filter that passes the signal at a specific mmWave band and attenuates the other bands. Hence, the frequency of the resonating signal will be the one at which there is the highest mutual coupling between the antennas and also the antennas resonate. However, narrow-band antennas, such as microstrip patch antennas (Figure 3) typically resonate and radiate at two frequencies, f_1 and f_2. Figure 4(a) shows S_{11} of a typical microstrip patch antenna. As can be seen, the antenna have low S_{11} at both f_1 and f_2, meaning that they can resonate and operate at these frequencies. This is due to the fact that each microstrip antenna has a rectangular shape with width W and length L, as shown in Figure 3. Hence, at each antenna's resonant frequency, one edge resonates and the other one radiates into the space.[1] These resonant frequencies f_1 and f_2 are determined by L and $\frac{W}{2}$, respectively.[2]

To make sure that our LNA resonates at the required frequency, we design the antennas such that their main resonance frequency f_1 is at the required mmWave carrier frequency. At the same time, to ensure that the LNA does not resonate at the second (undesired) resonance frequency f_2, we need to locate the antennas in a way that their coupling is strong at f_1 and very weak at f_2. Figure 4(a) shows the mutual coupling (S_{21}) between two patch antennas which are placed next to each other. As it can be seen, the coupling is strong at both f_1 and f_2 which is not desired since it may cause the LNA to resonate at f_2. Hence, we need to find a way to place antennas next to each other such that their coupling (S_{21}) is strong at f_1 while very week at f_2 as shown in Figure 4 (b).

To explain how we can suppress the coupling at one resonant frequency without impacting the other one, we first need to understand how two antennas create mutual coupling. In order to have a strong coupling between two microstrip antennas, the peaks and nodes of the electric field on the first antenna edges should match the peaks and nodes of the electric field of the second antenna edges. In a microstrip antenna, a non-resonating edge will have a uniform electrical field distribution while the resonating edge has sinusoidal distribution. This means that at f_1, the edge W has an electrical field with uniform distribution and the edge L has an electrical field with sinusoidal distribution as shown in Figure 5(a). Similarly, at f_2, the edge L has an electrical field with uniform distribution, and the edge W has an electrical field with sinusoidal distribution as shown in Figure 5(b). Hence, to reduce the coupling at f_2 while keeping the coupling at f_1, all we need to do is to place the two antennas such that their W edges are facing each other with some displacement. Specifically, we slightly shift one of the antennas along the X-axis such that they partly face each other as shown in Figure 5(c) and (d). Consequently, at f_2, the field's peaks and nodes of one antenna does not match the peaks and nodes of the other antenna, which reduces the coupling at that frequency. At the same time, this shift does not impact the coupling of antennas at f_1 since at that frequency the electric field on the facing edge is linear and even with a shift the peaks and nodes of the electric fields of the first antenna and second antenna match as shown in Figure 5(c).

[1]Note that it is not practical to design a square patch antenna with $W = L$ for multiple reasons. Such an antenna would have a very low bandwidth which is not useful for the mmWave band. Moreover its efficiency and gain would be very low.

[2]Patch antennas also have a resonant frequency at f_0 which is smaller than f_1 and it is determined by W. However, the patch do not radiate at this frequency, and hence this frequency is not typically considered among the antenna's resonant frequencies[13].

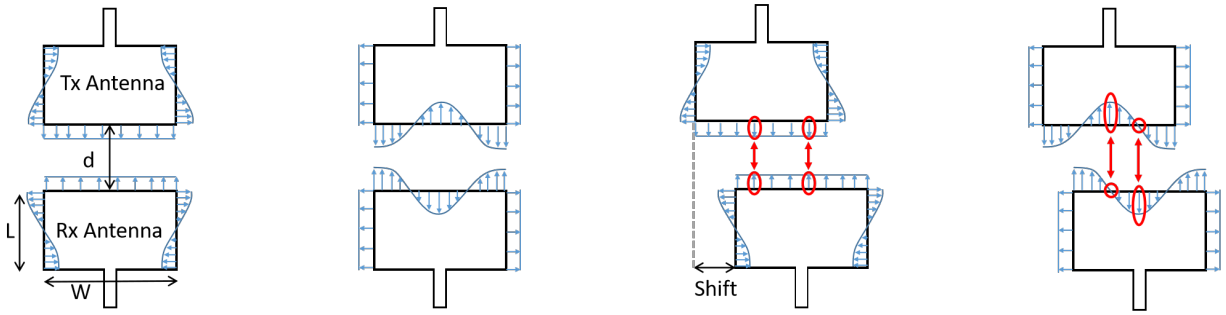

(a) at frequency f_1 **(b) at frequency f_2** **(c) at frequency f_1** **(d) at frequency f_2**

Figure 5: Patch Antenna's electric field at different operating frequencies and antennas arrangement. Patch antennas typically have two resonate frequency (f_1 and f_2) which are determined by by L and $\frac{W}{2}$, respectively. At f_1, the facing edges of the two antennas have uniform fields, and hence any shift in the antennas positions will not impact on the coupling at that frequency (see a and c). However, at f_2, the facing edges of the two antennas have sinusoidal field, and hence a shift will significantly reduce the coupling at that frequency since the peaks and nodes of the field do not match anymore (see b and d).

In summary, we connect the input and output of an LNA to two patch antennas. The mutual coupling between the antennas creates a feedback loop causing the LNA to resonate at a specific frequency. To guarantee that the resonance occurs at the required mmWave frequency, we adjust the spacing and orientation of the antennas in order to reshape their mutual coupling such that the coupling is strong at the required frequency and weak at the other frequency. We use this to generate a mmWave carrier signal. In the next section, we explain how we can use this carrier signal and the amplifier's saturation mode to translate IoT signal to the mmWave band and vice versa.

3.2 Leveraging Amplifier Saturation

mmPlug needs to upconvert IoT RF signals to mmWave for uplink, and downconvert mmWave signals to the original IoT frequency for downlink. To do so, the received mmWave signal must be multiplied by the carrier signal (i.e., the resonant signal created by the LNA and antenna loop) to generate the low-frequency IoT signal. It is worth clarifying that for up/down conversion we mix the IoT signal with a mmWave signal rather than creating the n^{th} harmonic of the IoT signal. To perform this multiplication, our approach is to rely on the intermodulation caused by the non-linear behavior of the LNA. LNAs are typically operated in a fairly linear region. However, in our design the loop gain (including both LNA gain and the antennas' mutual coupling) is larger than one, and thus the oscillating signal will continue growing in amplitude until the LNA enters in saturation mode. In saturation, the LNA acts as a strong non-linear component, which creates intermodulate the IoT signal with the resonating mmWave signal.

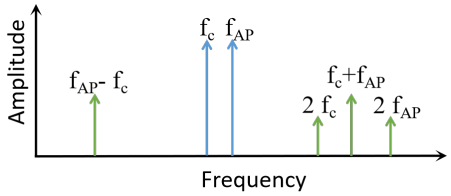

Figure 6: Leveraging Amplifier Saturation. The saturated LNA acts as a strong non-linear component which intermodulates the AP's mmWave signal (at f_{AP}) with the loop's resonating signal (at f_c). The result is a low-frequency $f_c - f_{AP}$ signal at the IoT band, and some other high frequency components which will be filtered out.

To analyze the output of the LNA and better understand the non-linear effect of the LNA, we use the Taylor expansion of a non-linear function. In particular, any non-linear function (Ψ) can be written as follow:

$$\Psi(x) = a_0 + a_1 \cdot x + a_2 \cdot x^2 + \cdots a_n \cdot x^n, \quad (1)$$

where a_n are the Taylor coefficients of function Ψ. Now let us consider two signals s_c (corresponding to the resonant frequency of the loop) and s_{AP} (corresponding to the received mmWave signal from the AP). When the sum of these signals is present at the input of the non-linear LNA, the resulting output can be presented as follows using Taylor expansion:

$$y = \Psi(s_c + s_{AP}) = a_0 + a_1 s_c + a_1 s_{AP} + a_2 (s_c + s_{AP})^2 + \cdots a_n (s_c + s_{AP})^n. \quad (2)$$

In the Taylor expansion, higher order coefficients (a_3 to a_n) are typically very small and hence can be ignored. As a result

the signal at the output of the LNA can be approximated as:

$$y = \Psi(s_c + s_{AP}) \approx a_0 + a_1 s_c + a_1 s_{AP} + a_2 s_c^2 + a_2 s_{AP}^2 + a_2 s_c s_{AP}. \quad (3)$$

Figure 6 shows the frequency components of this intermodulation. As it can be seen, the nonlinear behaviour of the LNA has multiplied the receiving mmWave signal (s_{AP}) with the oscillating mmWave signal (s_c) which results in a signal at the difference of their frequencies $f_{AP} - f_c$ and some other high-frequency components which are filtered out. Hence, by choosing the AP's frequency (f_AP) to be equal to $f_c + f_{IoT}$ and filtering the other high frequency components, the resulting signal will be exactly at the IoT's operating frequency (f_{IoT}). So far, we explained how mmPlug uses a saturated LNA to convert a received mmWave signal to an IoT signal (downlink). In a similar way, mmPlug can also upconverts the IoT signal to transmit a mmWave signal. By injecting a signal at f_{IoT} into the loop resonating at f_c, a signal at the AP's frequency of $f_c + f_{IoT}$ will be generated and transmitted via the antennas.

In summary, using a single wide-band LNA component and exploiting mutual coupling between patch antennas and LNA saturation behavior, we enable a simple and low-power design for mmPlug which can be easily connected to any IoT devices and enable them to communicate at a mmWave frequency band.

3.3 Frequency Stability

So far we have explained how mmPlug's module creates a tone at mmWave band and leverages the intermodulation to translate the IoT device's RF signal to the mmWave band and vise versa. However, the mmWave tone generated by the mmPlug module is not stable since the module does not use any reference clock or frequency stabilization technique. Hence, the frequency of the tone suffers from a large Carrier Frequency Offset (CFO) which varies by as much as 100 KHz even over a short time interval (such as the duration of a packet). This frequency variation requires correction, since otherwise it would cause a high packet error rate. Below, we explain how the mmWave AP can solve this issue in both uplink and downlink.

In the uplink, we leverage the fact that mmPlug's module transmits both the mmWave tone and the IoT signal ($s(t)$) at the mmWave band. Since both these two signals suffer from the same CFO variation, the AP can simply eliminate the CFO by mixing these two signals. In particular, we take the square of the received signals which results in:

$$\begin{aligned} y(t) &= \Big\{ cos\left[\omega(t)\cdot t\right] + s(t)\cdot cos\left[(\omega(t)+\omega_s)\cdot t\right] \Big\}^2 \\ &= cos^2\left[\omega(t)\cdot t\right] + s^2(t)\cdot cos^2\left[(\omega(t)+\omega_s)\cdot t\right] + \\ &\quad s(t)\cdot cos\left[(\omega(t)+\omega(t)+\omega_s)t\right] + \\ &\quad s(t)\cdot cos\left[\cancel{\omega(t)} - \cancel{\omega(t)} + \omega_s t\right] \\ &= cos^2\left[\omega(t)\cdot t\right] + s^2(t)\cdot cos^2\left[(\omega(t)+\omega_s)\cdot t\right] + \\ &\quad s(t)\cdot cos\left[(2\omega(t)+\omega_s)t\right] + s(t)\cdot cos(\omega_s t), \end{aligned} \quad (4)$$

where $\omega(t)$ is the varying mmWave frequency, ω_s is the IoT frequency, and $s(t)$ is the IoT's modulated signal. As indicated in this equation, the output of the squared signal has a low frequency component at ω_s which is the down converted IoT signal without any CFO. Hence, by using a low pass filter, we can extract the signal and resolve the CFO variation in the uplink communication.

For the downlink, we also rely on the AP to mitigate the CFO variation. Recall that the mmPlug module always transmits its carrier signal. The AP captures this signal and uses it to transmit the downlink signal with the same CFO variation as the mmPlug module. Hence, the received signal at the mmPlug module has the same CFO variation as the module's mmWave tone, and thus the intermodulated signal is CFO free at the IoT device's original band.

3.4 Diplexer Design

So far, we have explained how mmPlug leverages the mutual coupling of antennas to create a loop, and causes the LNA to generate a resonant frequency at the mmWave band. We also explained how mmPlug forces the LNA into saturation to convert modulated mmWave signals to the IoT frequency band and vice versa. However, our design needs to enable the IoT device to feed its low frequency signal to the loop or extract it from the loop, without impacting the loop itself. Otherwise, the IoT node's RF impedance would load mmPlug's loop impedance and reduce the loop gain, preventing the LNA and antennas from resonating. Said differently, we need to probe the loop to extract or feed the IoT signal without affecting the mmWave signal path.

To solve this problem, we design and integrate a compact, passive component called diplexer. Our diplexer has three ports as shown in Figure 7. Its first and second ports are simply connected together using a transmission line, and hence any signal (with any frequency) can pass through them. However, on the third port, we design and include a bow-tie structure, which is a combination of two open ended transmission lines. The length of the bow-tie and its location along the main transmission line is designed in a way that at mmWave frequency, the third port of the diplexer is isolated. Hence, mmWave signals can pass from port one to port two without entering into port three. However, at

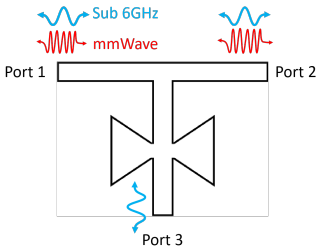

Figure 7: Diplexer Design. Both mmWave and low frequency signals can go from port 1 to port 2 and vice versa. However, only low frequency signals can be passed to/from port 3, whereas mmWave signals are isolated from port 3.

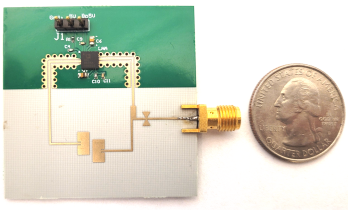

Figure 8: Our fabricated mmPlug's module. it includes an LNA, mmWave patch antennas, and a passive diplexer. The design is 5 x 4.5 cm, and has a SMA connector which can be easily connected to the antenna port of any IoT device.

lower frequencies the third port is not isolated anymore, and hence the low frequency signal can pass from the main transmission line (connecting port one and two) to the third port, and vice versa.

We integrate the designed Diplexer in our mmPlug module. In particular, we connect its port 1 to the LNA output, its port 2 to the Tx antenna, and its port 3 to the IoT module as shown in Figure 2. This design will isolate the IoT module from the mmWave loop and hence the IoT device will not interfere with the loop's resonant frequency, while enabling both uplink and downlink communication as described in Section 3.

4 IMPLEMENTATION

This section provides a detailed description of mmPlug's implementation and our testbed. The block diagram of mmPlug is shown in Figure 2, where an LNA is connected to two tightly coupled antennas and a diplexer is employed to connect the IoT device to the module. Note that the LNA is the only active component of our design. The rest of our design consists of passive components which have no DC power consumption and can be implemented using metal traces on the PCB. We optimize our design to operate in the 28 GHz

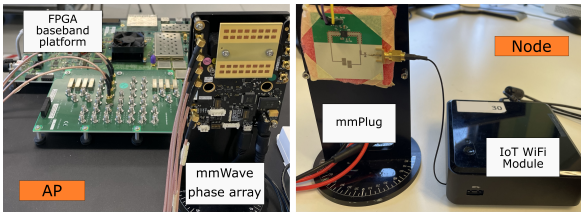

Figure 9: Experimental Tesbed. The end to end performance of mmPlug is evaluated through the programmable mmWave AP. On the node side mmPlug is connected to an IoT node for reception and transmission.

mmWave band. For the LNA, we use the ADL9005 [11] integrated circuit from Analog Devices which is a Wideband Low Noise Amplifier.[3]

We designed and simulated the antennas and the diplexer using the ANSYS HFSS software. We integrated the LNA with our patch antennas and the diplexer on a single small PCB. The final PCB design is fabricated on Rogers 4003 with 8mil thickness which has low loss for high frequency signals. Figure 8 shows the final version of our fabricated mmPlug. Its size is 5 x 4.5 cm which is small and simple enough to be connected to typical IoT devices (instead of their antennas) using their SMA connector. We did not specifically optimize the design for small size, and further miniaturization of the design is possible if needed by specific device form factors.

We perform microbenchmark and end-to-end evaluations to completely explore performance of mmPlug. For the microbenchmark, we investigate the received signal strength and frequency stability of mmPlug by using the Keysight mmWave platform [25] that has a signal generator to transmit a single tone and a Signal Analyzer to measure the received signal power accross spectrum. The end-to-end testbed measures the Packet Error Rate as well as SNR and RSSI. The testbed shown in Figure 9 consists of a Sivers EVK02001 28 GHz mmWave frontend [41] connected to an FPGA platform acting as a mmWave AP for the IoT nodes. This allows the AP to transmit and receive mmWave signals with any protocol and the required modulation. The Sivers EVK02001 has a maximum EIRP of 45 dBm, and includes a phased array antenna which can electronically steer its beam. Finally, for IoT nodes, we connect the mmPlug module to off the shelf IoT modules. Specifically, we use the M0 RFM96 [4] as a 433MHz LoRa module and the Atheros AR5B95 [1] as a 2.4GHz WiFi module.

[3]Note, ADL9005 is advertised as a 26.5 GHz LNA. However, it can operate even at 28 GHz with slightly lower gain as shown in its data sheet's measurement section.

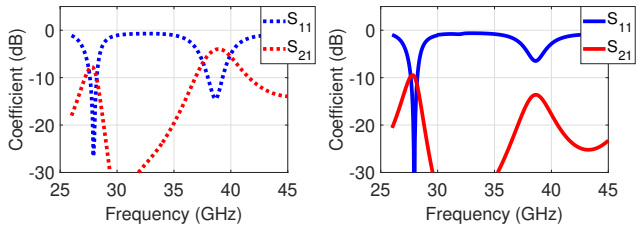

Figure 10: Our Patch Antenna's Performance (a) without and (b) with our optimized antenna placement. Placing the two antennas face to face (with some shift) successfully changes the mutual coupling such that the LNA resonates only at the desired frequency (i.e., 28 GHz).

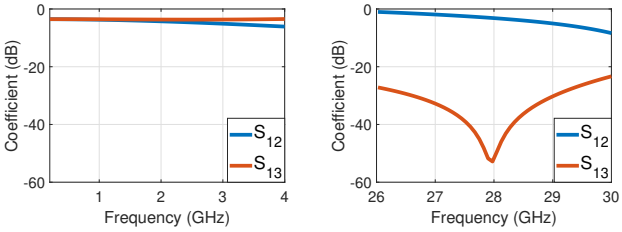

Figure 11: Our Diplexer performance over (a) low frequency bands and (b) mmWave bands. At low frequency, the design has high transmission coefficients, meaning that the signal can pass through the loop to the IoT device and vice versa. However, at mmWave frequency, it has a very low transmission coefficient between port 1 and 3, i.e., the IoT device is isolated from the resonating loop.

5 EXPERIMENTAL RESULTS

We now evaluate the performance of mmPlug and verify its capability to enable IoT communication at the mmWave frequency band. We conduct experiments in an indoor office with objects such as tables, chairs and shelves to provide a complex, realistic channel environment.

5.1 Micro-Benchmarks

5.1.1 Microstrip Antenna Performance: We first evaluate the performance of our microstrip antenna design using the HFSS Software. Figure 10(a) shows the matching (S_{11}) and coupling coefficient (S_{12}) of the two antennas without using our technique to modify their mutual coupling. This result shows that our antennas provide a matching coefficient (S_{11}) of less than -20 dB at our operating frequency of 28 GHz and our antennas thus resonate and operate very efficiently at that frequency. However, the figure also shows that our antennas have another resonant frequency at around 38 GHz. As discussed in Section 3.1, the coupling at the second resonance should be suppressed so that it becomes lower than the coupling at the desired frequency.

Figure 10(b) shows the matching (S_{11}) and coupling coefficient (S_{12}) of the two antennas using our technique to shift their relative placement with respect to each other to modify the mutual coupling. Comparing this figure with (a), we can see that by shifting one of the antennas, the coupling at the higher resonance frequency (38 GHz) is significantly reduced without impacting the coupling at our targeted operating frequency of 28 GHz. This result confirms that our approach effectively ensures that the LNA resonates only at the specific desired frequency.

5.1.2 Diplexer Performance: As discussed in Section 3.4, the diplexer has 3 ports. At mmWave, the signal should only pass between the LNA and antennas (port 1 and 2) and no signal should pass to port 3 where the IoT device is connected. The reason is to keep the mmWave signal circulating in the loop and avoid to change the loop impedance when an IoT device connects to it. However, the IoT signal whose frequency is much lower than mmWave should be able to enter or exit the loop.

We evaluate the performance of our diplexer using the HFSS Software as shown in Figure 11. Here, subfigure (a) shows its performance for the sub-4 GHz band, and subfigure (b) shows its performance for the mmWave band. In these plots, S_{12} is the transmission coefficient between port 1 and 2 (i.e., between LNA and TX antenna) and S_{13} is the transmission coefficient between port 1 and 3 (i.e. between LNA and the IoT connection port). The figure shows that S_{13} and S_{12} are high for the sub-4 GHz band, meaning that IoT signal can pass through the diplexer from the IoT device to the loop and vice versa as desired. However, at the 28 GHz mmWave band, S_{13} is very low (much lower than S_{12}). This means that mmWave signals can go between the LNA and the antenna in the mmPlug loop as desired, but they cannot pass the diplexer to enter port 3 which is connected to the IoT device. Hence, our diplexer effectively isolates the IoT device from the operation of the loop at mmWave frequency.

5.1.3 Loop Resonant Frequency: Next, we verify whether our idea of using the antenna's mutual coupling to cause the LNA to resonate works as expected. We power up a mmPlug module and measure the resonating signal using a N9020B MXA Signal Analyzer. Figure 12 shows the Signal Analyzer measurement which indicates a tone close to 28.17 GHz. This implies that mmPlug's module has successfully created a mmWave tone by using a single LNA and leveraging the antennas' mutual coupling. Note that the small deviation

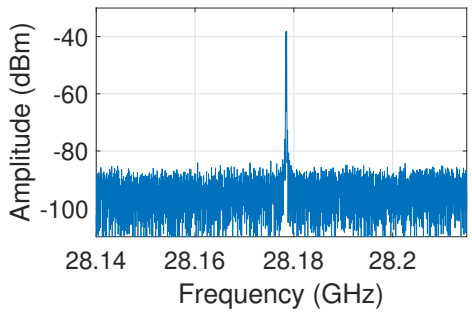

Figure 12: mmPlug's resonating signal. mmPlug successfully creates a mmWave tone around 28 GHz, using a single LNA and leveraging the mutual coupling of antennas.

from 28 GHz stems from the fabrication accuracy and tolerance. However, this does not create any problem since slight deviations in the resonant frequency can be compensated by adjusting the AP's carrier frequency, so that the frequency difference is exactly the frequency of the low frequency IoT band. Moreover, based on our measurements, by adjusting the biasing voltage of the LNA in a mmPlug's module, the resonate frequency can be slightly (up to 80 MHz) adjusted during the installation stage.

As we described in Section 3.3, the frequency of the generated mmWave tone is not very stable and varies over time. Hence, we proposed a technique to compensate for this carrier frequency variation. Here, we examine if our approach works. We connect a mmPlug's module to a signal generator. The signal generator generates a 40 MHz continuous wave (CW) and feed to to the mmPlug's module to tranmits it at mmWave band. Then we capture the downconverted signal received at the mmWave AP. Figure 13 shows the downconverted signal with and without our correction technique. As shown in the top figure, the signal suffers from a significant carrier frequency variations. However, our techniques is completely eliminating the problem as shown in the bottom figure.

5.1.4 Module Orientation. The mmPlug module includes onboard patch antennas which are used to transmit and receive mmWave signals. Here, we evaluate if the orientation of the module with respect to the AP impacts its performance. We place a mmPlug module at 5m with respect to the mmWave AP which is transmitting mmWave signals. We then measure the down-converted signal power at the output of the module for different orientations with respect to the AP. Figure 14 shows the results of this experiment. For simplicity, we normalize all received power values to the maximum received power. As expected, the maximum received power is when the module is facing toward the AP. However, even when the angle is 50 degrees, the signal power is reduced

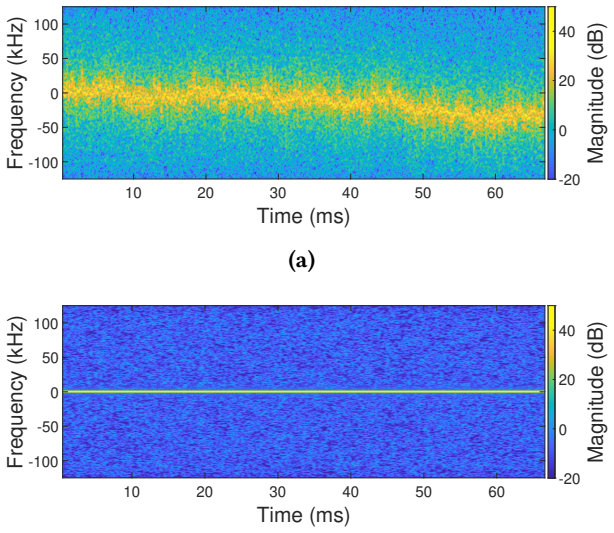

Figure 13: CFO Correction. A continues wave signal transmitted by mmPlug's module and received by our mmWave AP (a) without, and (b) with our CFO correction technique described in 3.3.

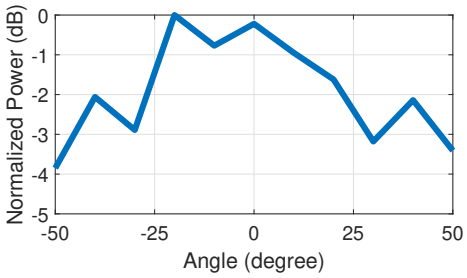

Figure 14: mmPlug's downconverted signal power versus its orientation respect to the AP. The orientation of the mmPlug module with respect to the AP can change by as much as 100 degrees (from −50° to 50°) without significant impact on the received power and communication performance.

only by 3-4 dB which is not significant. This means that the orientation of the mmPlug module relative to the AP can change by as much as 100 degrees without significant impact on the communication performance.

5.2 End-to-End Evaluation

We evaluate the end-to-end performance of mmPlug. To do so, we connect mmPlug's modules to off-the-shelf WiFi and LoRa modules and examine whether our mmWave AP can communicate to these modules at the mmWave band. Note

that we make no modifications to the protocol, firmware or chipset of the modules. We only disconnect their antennas and connect the mmPlug module to their antenna port.

First, we evaluate the performance of mmPlug in communicating to the WiFi device equipped with a mmPlug module. We place the WiFi node at different distances with respect to the mmWave AP and we measure the Packet Error Rate (PER) over 10,000 packets at the node side. We try different WiFi coding rate and modulation schemes. Figure15 shows the result for this experiment. The figure shows that even at the distance of 30 m, the PER is lower than 1% for even MCS3 with 16 QAM modulation. Note, a 1% PER is much lower than the PER of a typical WiFi networks. This result shows that mmPlug enables off-the-shelf WiFi modules to successfully communicate to a mmWave AP without making any changes to WiFi module's protocol, chipset or firmware.

Next we examine the performance of mmPlug in enabling a LoRa module to communicate to the mmWave AP. We connect a mmPlug module to the antenna port of the LoRa module and measure the PER at different distances. Based on our initial measurement, we found that even at 30 m, the PER was zero when calculated over 10,000 packets. Note, this is expected since LoRa is much more robust than WiFi and targets longer operating distances. Therefore, instead of measuring the PER, we measure the SNR and RSSI of the packets at the AP and the node respectively. Note that we cannot measure the SNR at the node since the LoRa module reports only RSSI. Figure 16 shows the result of this experiment. The SNR at the AP is higher than -8 dB even when the node is 30 m away from the AP. Note, this SNR is enough for LoRa to enable PER of less than 10^{-5} using LoRa spreading factor of 8 or higher [7]. In addition, the measured RSSI at the node is higher than -80 dBm even when the node is 30 m away from the AP. This RSSI is more than enough to establish LoRa communication with less than 10^{-5} PER. Note, we were not able to run experiments for longer distances due to space limitation. However, this result implies that even for longer distances, mmPlug can enable off-the-shelf LoRa modules to successfully communicate to a mmWave AP without making any changes to LoRa module's protocol, chipset or firmware.

5.3 Localization

One advantage of enabling IoT devices to operate at the mmWave band is the ability to accurately localize them. Here we evaluate the localization performance of the mmWave AP in detecting the angle of mmPlug-equipped IoT nodes in an indoor environment where possible multipath components also exist in the received signal space. To find the angle of the node, we implement the Angle of Arrival (AoA) detection algorithm (similar to [29, 36]) on our mmWave AP.

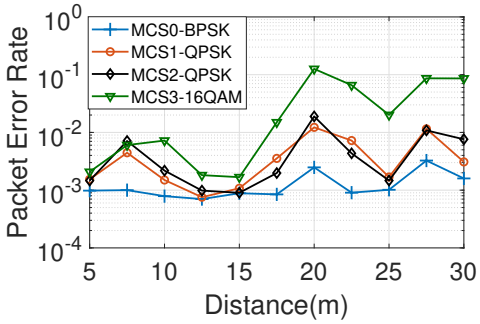

Figure 15: End-to-End Performance for the WiFi implementation. Packet Error Rate (PER) versus distance when a mmWave AP is communicating to an off-the-shelf WiFi device equipped with a mmPlug's module

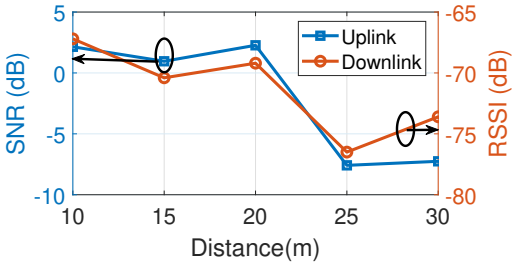

Figure 16: End-to-End Performance for the LoRa implementation. SNR and RSSI versus distance when a mmWave AP is communicating to an off-the-shelf LoRa device equipped with a mmPlug's module. Packet error rate (PER) is less than 10^{-5} for all distances.

Specifically, the AP uses a codebook with 21 beam patterns to estimate the AoA of the IoT node. We place the IoT node at different distances and angles respect to the AP while the AP estimates its AoA. Figure 17 (a) shows the CDF of error in AoA estimation for these measurements. This result shows that mmPlug enables the mmWave AP to estimate AoA of IoT nodes with accuracy of less than a degree. Note, this is possible since operating at mmWave enables the AP to create different beam patterns and use them to accurately estimate the AoA. The AP can also use the estimated AoA from two different locations to estimate the location of the IoT node using triangulation. Figure 17 (b) shows the accuracy in estimating the location when we performed our AoA measurements from two phased arrays spaced by 2 m, and then use it to estimate the location of the IoT node using triangulation. The AP can estimate the location of the IoT node with less than 40 cm error. This reported localization performance

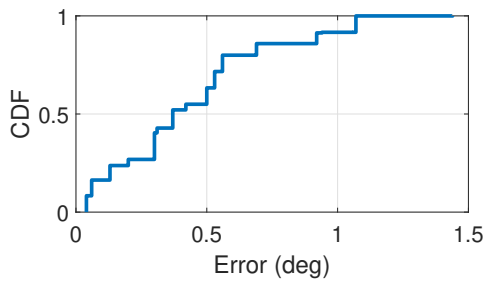

(a) AoA

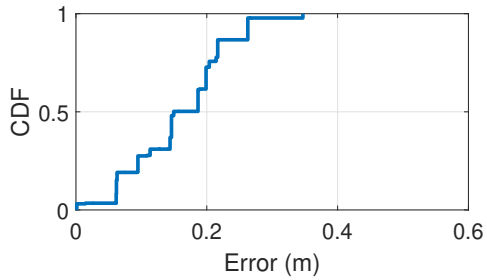

(b) Range

Figure 17: Localization performance.mmWave AP can accurately localize IoT devices which use mmPlug's module instead of their antennas.

is better than existing Lora localization and similar to the sophisticated WiFi localization techniques. Please note that the contribution of this paper is not advancing localization algorithms, however, these results show that enabling IoT devices to operate at mmWave band provides opportunities to accurately localize them. In fact, one can further improve our localization accuracy using state-of-the art localization techniques.

5.4 Energy Consumption

We now measure the energy consumption of the mmPlug module, and compare it with prior low power mmWave systems for IoT applications. The only active component used in mmPlug is an LNA which draws 59 mA.[4] Table 1 compares mmPlug with past low-power mmWave systems for IoT in terms of energy efficiency, and other factors such as data-rate, operating range and capabilities. For a fair comparison, the energy efficiency of all systems are reported for 100 Mbps data rate, and include only the energy consumption of their mmWave front-ends. mmPlug consumes 59 mA, with a voltage of 8.5V which results in 501.5mW power consumption. Considering 100Mbps data rate, the energy consumption of mmPlug is 5.015nj/b. The results shows that the energy efficiency of mmPlug is lower than mmX and slightly higher than mmTag. However, it worth mentioning that mmTag is a backscatter system and hence has significantly lower range and capabilities. In particular, mmPlug can enable an operating range of 30 m or more, depending on the SNR requirements of the IoT device. For a better understanding, a sub-6GHz IoT device like WiFi consumes 17.5nj/b which is higher than mmPlug. Furthermore, past systems such as mmX and mmTag provide only uplink communication, and require their own modulation schemes (such as FSK and ASK) and new MAC protocols. In contrast, mmPlug provides both uplink and downlink, and it operates as a plug-and-play module which can be seamlessly connected to existing and future commodity IoT devices.

[4]For comparison, a traditional approach to translate a low-frequency signal to mmWave band and vice versa is to use a VCO (such as HMC732 [10]), and up and down converters (such as ADMV1013 [9] and ADMV1014 [8]). However, such a design draws more than 1A in total which is significantly higher than the current consumption of a mmPlug's module.

6 DISCUSSION

In this work we explained how mmPlug enables IoT devices to operate in mmWave bands. Here we discuss how a network of IoT devices equipped with mmPlug modules will operate. Such an mmPlug IoT network not only enables communication based on the original PHY and MAC protocol of the IoT technology but also benefits from SDMA provided by mmWave technology. In particular, since our module works as a plug-and-play IoT add-on, devices do not even notice that they are operating in a mmWave band and can continue to take advantage of their own MAC protocol. At the same time, since they operate at the mmWave frequency over the air, the AP can simultaneously create beams toward multiple of them for SDMA. We will discuss both techniques in more detail bellow:

IoT Multiple Access Protocol: Today's IoT devices use different MAC protocols to enable multiple nodes to communicate to a single AP. One advantage of mmPlug is that it is technology agnostic, meaning that it can work with any IoT device regardless of its MAC protocol. Hence, even when mmPlug is used, the IoT devices can still use their own MAC protocol to support multiple nodes as shown in our evaluation.

SDMA: The SDMA is one of the advantages of using mmWave communication. It stems from directional communication at mmWave, where the AP can generate a beam and steer toward the proper direction to communicate with a device. This is possible since the mmWave signal has a very short wavelength and thus small antenna elements. Hence, by using a large array of antennas, the AP can create multiple narrow beams toward multiple IoT nodes. This enables the

System Name	Energy Efficiency (nJ/bit)	Operating Range (m)	Uplink	Downlink	Plug-and-Play Compatibility
mmPlug	5	> 30	Yes ✓	Yes ✓	Yes ✓
mmTag [31]	2.4	8	Yes ✓	No ✗	No ✗
mmX [30]	11	18	Yes ✓	No ✗	No ✗

Table 1: Performance comparison of mmPlug with past low-power mmWave communication systems for IoT.

AP to send different data to each beam, supporting all nodes simultaneously.

FDMA: The other advantage of mmWave is to provide a significantly large bandwidth that can be channelized and allocated to many users through FDMA. In this scheme, each user must be able to transmit and receive in the allocated frequency channel. The current design of mmPlug allows the node to change its working frequency for few MHz by changing the biasing voltage of LNA. However, it is not enough to cover the whole band width of mmWave band. As a future direction of this work, one can add full frequency tunability to mmPlug by adding varactor diodes to the microstrip antennas. By changing the biasing voltage of the varactors, the resonance frequency of mmPlug will change to the desired operating band at mmWave.

RTT: Finally, it is worth mentioning that due to compatibility of mmPlug with existing protocols such as WiFi and LoRa, one can also measure Round Trip Time (RTT) of packets to estimate the range of the IoT device from the access point.

7 RELATED WORK

Past work on mmWave mostly focuses on applications that require very high-data-rate links, while having substantial energy and computing power [22–24, 28]. For example, the systems presented in [14, 19] utilize mmWave technology in data centers to enable high throughput links between server racks. There is also work in using mmWave for 5G applications [16, 35, 38] and sensing applications [18, 33]. Finally, some other systems use mmWave to enable high data rates for VR applications to stream high quality video from a PC to VR headset [2, 15]. In contrast to past work, this paper focuses on designing a very low-power, simple module to bring mmWave technology to any IoT devices. The closest past work to this paper are the systems presented in [30, 31], which propose low-power mmWave communication systems. However, these systems only enable uplink communication and does not provide a downlink. Moreover, these systems have their own MAC protocols and modulation schemes, and hence they cannot be used as a plug-and-play module for existing IoT devices. There is also some work on mmWave backscatter tags for road sign assistance [42], and passive RFID tags at mmWave [5, 27, 34]. However, these system target accurate localization and sensing applications, and hence do not support uplink or downlink communications. Moreover, they cannot be connected to IoT devices as a plug-and-play module.

Finally, circuit non-linearity has been used in past work for other applications. For example, ReMix uses a circuit's non-linearity to mix two low-frequency signals, enabling in-body backscatter Communication and Localization [44]. Backdoor uses Alexa's circuit non-linearity to make its microphones hear inaudible sounds [39]. In contrast to past work, mmPlug forces an LNA to work in its saturation mode to enable strong non-linear behavior, translating mmWave signals to low-frequency IoT and vice versa. Moreover, mmPlug combines this behavior with mutual coupling of antennas. Some past works have proposed changing an antenna's matching design. For example, μMedIC has presented a re-configurable antenna design which can shift its resonant frequency, allowing efficient harvesting and communication [3]. However, this system focuses on the sub-GHz band, and does not use mutual coupling of antennas to generate resonating signals using an amplifier.

In contrast to all past work, mmPlug introduces the first plug-and-play module which brings mmWave technology to IoT devices. mmPlug is the first system which leverages an antenna's mutual coupling and an amplifier's saturation mode to translate mmWave signals to lower-frequency IoT signals using a single low-noise-amplifier.

8 CONCLUSION

In this paper, we presented mmPlug, a small, low-power, simple and plug-and-play device which brings mmWave technology to existing IoT nodes. mmPlug enables IoT devices to take advantage of the large bandwidth and short wavelength of mmWave signals and high directionality of mmWave antennas, which is essential to enable the vision of massive IoT deployments. The simple and low-power design is made possible through an innovative RF design with only a single active component, together with multiple passive components which can be fabricated using metal traces on the PCB. The mmPlug module can be simply connected to IoT devices instead of their antennas using an SMA connector. mmPlug achieves such design simplicity by leveraging the antennas' mutual coupling and amplifier saturation to translate mmWave signals to IoT RF signals and vice versa. Our evaluation and experimental results show that mmPlug enables robust mmWave communication links between a

mmWave AP and existing IoT devices (such as LoRa and WiFi), even when the AP and the IoT device are 30 m away.

REFERENCES

[1] [n. d.]. https://www.aliexpress.com/i/32998893674.html?gatewayAdapt=4itemAdapt.

[2] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2017. Enabling High-Quality Untethered Virtual Reality. In *NSDI*. 531–544.

[3] Mohamed R Abdelhamid, Ruicong Chen, Joonhyuk Cho, Anantha P Chandrakasan, and Fadel Adib. 2020. Self-reconfigurable micro-implants for cross-tissue wireless and batteryless connectivity. In *MobiCom'20: Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*.

[4] Adafruit. [n. d.]. Adafruit Feather M0 RFM96 LoRa Radio - 433MHz - RadioFruit. https://www.adafruit.com/product/3179.

[5] Ajibayo Adeyeye, Charles Lynch, Xuanke He, Sanghoon Lee, John D Cressler, and Manos M Tentzeris. 2021. Fully inkjet printed 60GHz backscatter 5G RFID modules for sensing and localization in Internet of Things (IoT) and digital twins applications. In *2021 IEEE 71st Electronic Components and Technology Conference (ECTC)*. IEEE, 1193–1198.

[6] Ansuman Adhikary, Ebrahim Al Safadi, Mathew K Samimi, Rui Wang, Giuseppe Caire, Theodore S Rappaport, and Andreas F Molisch. 2014. Joint spatial division and multiplexing for mm-wave channels. *IEEE Journal on Selected Areas in Communications* 32, 6 (2014), 1239–1255.

[7] Orion Afisiadis, Matthieu Cotting, Andreas Burg, and Alexios Balatsoukas-Stimming. 2019. On the error rate of the LoRa modulation with interference. *IEEE Transactions on Wireless Communications* 19, 2 (2019), 1292–1304.

[8] Analog Devices [n. d.]. *24 GHz to 44 GHz Wideband, Microwave Downconverter*. Analog Devices. Rev. A.

[9] Analog Devices [n. d.]. *24 GHz to 44 GHz, Wideband, Microwave Upconverter*. Analog Devices. Rev. B.

[10] Analog Devices [n. d.]. *WIDEBAND MMIC VCO WITH BUFFER AMPLIFIER 6 - 12 GHz*. Analog Devices. v02.0514.

[11] Analog Devices 2022. *Wideband, Low Noise Amplifier, Single Positive Supply*. Analog Devices. Rev. A.

[12] Roshan Ayyalasomayajula, Aditya Arun, Chenfeng Wu, Anees Shaikh, Shrivatsan Rajagopalan, Yige Hu, Shreya Ganesaraman, Christopher J Rossbach, Aravind Seetharaman, Emmett Witchel, et al. 2020. LocAP: Autonomous Millimeter Accurate Mapping of WiFi Infrastructure. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. 1115–1129.

[13] Constantine A Balanis. 2015. *Antenna theory: analysis and design*. John wiley & sons.

[14] 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.

[15] 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.

[16] 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.

[17] Yasaman Ghasempour, Claudio RCM Da Silva, Carlos Cordeiro, and Edward W Knightly. 2017. IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.

[18] Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11464–11473.

[19] 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.

[20] Mehrdad Hessar, Ali Najafi, Vikram Iyer, and Shyamnath Gollakota. 2020. TinySDR: Low-Power SDR Platform for Over-the-Air Programmable IoT Testbeds. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. 1031–1046.

[21] National Instruments. 2015. 71-76 GHz Millimeter-wave Transceiver System.

[22] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. 2021. Two beams are better than one: Towards reliable and high throughput mmWave links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 488–502.

[23] Ish Kumar Jain, Rohith Reddy Vennam, and Dinesh Bharadia. 2023. Towards Flexible Frequency-Dependent mmWave Multi-Beamforming. In *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications*. 132–132.

[24] Suraj Jog, Jiaming Wang, Junfeng Guan, Thomas Moon, Haitham Hassanieh, and Romit Roy Choudhury. 2019. Many-to-many beam alignment in millimeter wave networks. In *USENIX Conference on Networked Systems Design and Implementation, NSDI'19*.

[25] Keysight. [n. d.]. 5G R&D Test Bed. https://www.keysight.com/us/en/products/modular/reference-solutions/5g-waveform-generation-analysis-testbed-reference-solution.html.

[26] Jesus Omar Lacruz, Dolores Garcia, Pablo Jiménez Mateo, Joan Palacios, and Joerg Widmer. 2020. mm-FLEX: an open platform for millimeter-wave mobile full-bandwidth experimentation. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 1–13.

[27] Zhengxiong Li, Baicheng Chen, Zhuolin Yang, Huining Li, Chenhan Xu, Xingyu Chen, Kun Wang, and Wenyao Xu. 2019. Ferrotag: A paper-based mmwave-scannable tagging infrastructure. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 324–337.

[28] Sohrab Madani, Suraj Jog, Jesús Omar Lacruz, Joerg Widmer, Haitham Hassanieh, et al. 2021. Practical Null Steering in Millimeter Wave Networks.. In *NSDI*. 903–921.

[29] Dolores Garcia Marti, Jesus Omar Lacruz, Pablo Jimenez Mateo, Joan Palacios, Rafael Ruiz, and Joerg Widmer. 2021. Scalable Phase-Coherent Beam-Training for Dense Millimeter-wave Networks. *IEEE Transactions on Mobile Computing* (2021).

[30] Mohammad H Mazaheri, Soroush Ameli, Ali Abedi, and Omid Abari. 2019. A millimeter wave network for billions of things. In *Proceedings of the ACM Special Interest Group on Data Communication*. 174–186.

[31] Mohammad Hossein Mazaheri, Alex Chen, and Omid Abari. 2021. mmTag: a millimeter wave backscatter network. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 463–474.

[32] Joan Palacios, Paolo Casari, and Joerg Widmer. 2017. JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*. IEEE, 1–9.

[33] Akarsh Prabhakara, Diana Zhang, Chao Li, Sirajum Munir, Aswin C Sankaranarayanan, Anthony Rowe, and Swarun Kumar. 2022. Exploring mmWave Radar and Camera Fusion for High-Resolution and Long-Range Depth Imaging. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3995–4002.

[34] Pekka Pursula, Francesco Donzelli, and Heikki Seppa. 2011. Passive RFID at millimeter waves. *IEEE Transactions on Microwave Theory and Techniques* 59, 8 (2011), 2151–2157.

[35] 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.

[36] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent mmWave path tracking. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications*. 13–18.

[37] Behzad Razavi and Razavi Behzad. 2012. *RF microelectronics*. Vol. 2. Prentice hall New York.

[38] 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.

[39] Nirupam Roy, Haitham Hassanieh, and Romit Roy Choudhury. 2017. Backdoor: Making microphones hear inaudible sounds. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 2–14.

[40] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Owen Torres, Daniel Uvaydov, et al. 2017. X60: A programmable testbed for wideband 60 ghz wlans with phased arrays. In *Proceedings of the 11th Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization*. 75–82.

[41] SiversSemiconductors. [n. d.]. EVK02001. https://www.sivers-semiconductors.com/sivers-wireless/evaluation-kits/evaluation-kit-evk02001/.

[42] Elahe Soltanaghaei, Akarsh Prabhakara, Artur Balanuta, Matthew Anderson, Jan M Rabaey, Swarun Kumar, and Anthony Rowe. 2021. Millimetro: mmWave retro-reflective tags for accurate, long range localization. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 69–82.

[43] Anthony Ngozichukwuka Uwaechia and Nor Muzlifah Mahyuddin. 2020. A comprehensive survey on millimeter wave communications for fifth-generation wireless networks: Feasibility and challenges. *IEEE Access* 8 (2020), 62367–62414.

[44] Deepak Vasisht, Guo Zhang, Omid Abari, Hsiao-Ming Lu, Jacob Flanz, and Dina Katabi. 2018. In-body backscatter communication and localization. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. 132–146.

[45] Chenshu Wu, Feng Zhang, Beibei Wang, and KJ Ray Liu. 2020. mmTrack: Passive multi-person localization using commodity millimeter wave radio. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2400–2409.

[46] Ming Xiao, Shahid Mumtaz, Yongming Huang, Linglong Dai, Yonghui Li, Michail Matthaiou, George K Karagiannidis, Emil Björnson, Kai Yang, I Chih-Lin, et al. 2017. Millimeter wave communications for future mobile networks. *IEEE Journal on Selected Areas in Communications* 35, 9 (2017), 1909–1935.

[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 *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. 162–175.