

ACTIVE BEAM LEARNING FOR FULL-DUPLEX WIRELESS SYSTEMS

Jeong Min Kong and Ian P. Roberts

Wireless Lab, Department of Electrical and Computer Engineering, UCLA
{jeongminkong,ianroberts}@ucla.edu

ABSTRACT

In this paper, we present a novel active beam learning method for in-band full-duplex wireless systems, that aims to design transmit and receive beams which suppress self-interference and maximize the sum spectral efficiency. Rather than rely on explicit estimation of the downlink, uplink, and/or self-interference channels like in most existing work, our method instead actively probes all three channels through measurements of SNR and INR over a fixed number of time slots. Then, once this probing concludes, all collected probing measurements are used to design transmit and receive beams which serve downlink and uplink in a full-duplex fashion. We realize this active beam learning scheme through a network of LSTMs and DNNs, which learns to design each probing beam pair and subsequently extract and record valuable information from each probing measurement such that near-optimal serving beams can be designed following the probing stage. Simulation indicates that our method reliably suppresses self-interference while delivering near-maximal SNR on the downlink and uplink with merely 3–10 probing time slots, while exhibiting robustness to measurement noise and the structure of the self-interference channel.

1. INTRODUCTION

Upgrading base stations with the ability to transmit downlink and receive uplink at the same time and over the same bandwidth—i.e., in-band full-duplex capability—can yield higher data rates, lower latency, broader coverage, and enhanced sensing capabilities, paving the way toward 6G networks [1]. The key to realizing full-duplex 6G networks lies in effectively mitigating the *self-interference* inflicted by a given base station’s transmitter onto its own receiver, which would otherwise prohibitively degrade the uplink signal quality. Analog and digital self-interference cancellation schemes have proved capable of realizing full-duplex in conventional low-frequency base stations [2, 3], but these schemes are less suitable in current 5G and emerging 6G radios, as they scale unfavorably to many antennas, wide bandwidths, and high carrier frequencies [1, 4]. Motivated by this, recent works [5–9] have harnessed beamforming to cancel the self-interference coupled between the transmit

and receive antenna arrays of a massive multiple input multiple output (MIMO) or millimeter wave (mmWave) wireless system. These schemes have proven capable of reducing self-interference to below the noise floor and are thus a promising route to unlocking full-duplex wireless systems but exhibit noteworthy practical shortcomings.

Most notably, almost all existing full-duplex beam designs rely on explicit estimation of the users’ downlink and uplink channels as well as the MIMO self-interference channel. In massive MIMO and mmWave systems, the number of antennas can be on the order of dozens or even hundreds, and estimating these high-dimensional channels would thus be resource-expensive and impractical. This is evident even in today’s half-duplex massive MIMO and mmWave systems, which circumvent explicit estimation of downlink and uplink channels via codebook-based beam alignment procedures [10]. To our knowledge, the full-duplex beam designs in [8, 9] are the only ones which do not rely on explicit channel estimation but rather power measurements across the downlink, uplink, and/or self-interference channels. While experimentally validated, these methods in [8, 9] require around 50–500 measurements per downlink-uplink user pair, which may make them unsuitable in 6G networks, given their resource constraints.

In this paper, we propose a novel active beam learning method to design near-optimal transmit and receive beams of a full-duplex base station, without explicit channel estimation and with only a few measurements. Our approach involves the base station first jointly probing the downlink, uplink, and self-interference channels over a *fixed* number of time slots, from which it extracts implicit information about all three channels to design the final serving beams. Structured as a sequential decision-making process, the base station *actively* designs the probing beams in each time slot based on its prior probing measurements, allowing it to tailor each probing sequence to the particular user pair and self-interference channel realization. Notably, this method bypasses any explicit channel estimation and instead relies solely on power measurements to design effective probing and serving beams. Simulation results demonstrate that with only 3–10 probing time slots, our scheme can deliver near-maximal spectral efficiency on the downlink and uplink.

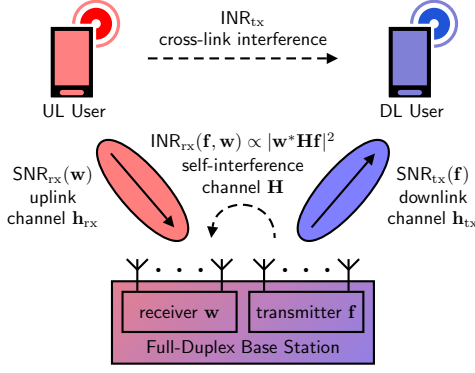


Fig. 1. A full-duplex base station transmits downlink to one user while receiving uplink from another user at the same time and same frequency.

2. SYSTEM MODEL

In this work, we consider a base station serving a single-antenna downlink user and a single-antenna uplink user simultaneously over the same frequency, as depicted in Fig. 1. The base station utilizes two separate antenna arrays, one for transmission and the other for reception; the transmit array has N_t antennas and analog beamforming weights $\mathbf{f} \in \mathbb{C}^{N_t \times 1}$, and the receive array has N_r antennas and analog beamforming weights $\mathbf{w} \in \mathbb{C}^{N_r \times 1}$. The signal-to-noise ratios (SNRs) of the downlink and uplink are

$$\text{SNR}_{\text{tx}}(\mathbf{f}) = \frac{P_{\text{tx}}^{\text{BS}} \cdot |\mathbf{h}_{\text{tx}}^* \mathbf{f}|^2}{P_{\text{n}}^{\text{UE}}} \quad (1)$$

$$\text{SNR}_{\text{rx}}(\mathbf{w}) = \frac{P_{\text{tx}}^{\text{UE}} \cdot |\mathbf{w}^* \mathbf{h}_{\text{rx}}|^2}{P_{\text{n}}^{\text{BS}}}, \quad (2)$$

where $P_{\text{tx}}^{\text{BS}}$ is the transmit power of the base station, P_{n}^{UE} is the noise power at the downlink user, $\mathbf{h}_{\text{tx}} \in \mathbb{C}^{N_t \times 1}$ is the downlink channel, $P_{\text{tx}}^{\text{UE}}$ is the transmit power of the uplink user, P_{n}^{BS} is the noise power at the base station, and $\mathbf{h}_{\text{rx}} \in \mathbb{C}^{N_r \times 1}$ is the uplink channel.

Since the uplink and downlink are operating at the same frequency band, the base station's transmit array inflicts so-called self-interference upon its own receive array across the MIMO channel $\mathbf{H} \in \mathbb{C}^{N_r \times N_t}$. While still the topic of active research, a plausible model for the self-interference channel \mathbf{H} , backed by measurements [9], is a Rician fading model

$$\mathbf{H} = \sqrt{\frac{\kappa}{\kappa + 1}} \bar{\mathbf{H}} + \sqrt{\frac{1}{\kappa + 1}} \tilde{\mathbf{H}}. \quad (3)$$

Here, κ is the Rician factor, $\bar{\mathbf{H}}$ is the static part of the self-interference channel caused by near-field coupling between the antenna arrays, and $\tilde{\mathbf{H}}$ is a time-varying component that stems from unpredictable environmental factors such as reflections. The strength of self-interference coupled by a

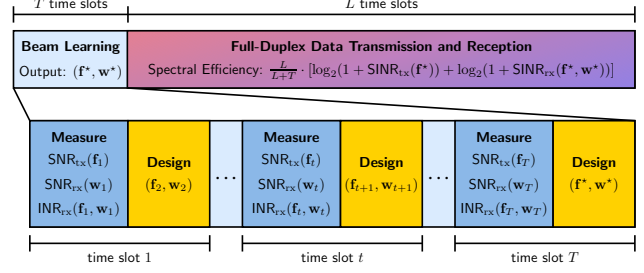


Fig. 2. Our envisioned active beam learning solution.

given transmit beam \mathbf{f} and receive beam \mathbf{w} can be captured by the interference-to-noise ratio (INR) given by

$$\text{INR}_{\text{rx}}(\mathbf{f}, \mathbf{w}) = \frac{P_{\text{tx}}^{\text{BS}} \cdot |\mathbf{w}^* \mathbf{H} \mathbf{f}|^2}{P_{\text{n}}^{\text{BS}}}. \quad (4)$$

Note that there is also interference induced by the uplink user on the downlink user, as shown in Fig. 1; the INR of this cross-link interference will be denoted as INR_{tx} . While self-interference depends on the transmit and receive beams at the base station, the cross-link interference depends only on the users and the channel between them.

Downlink and uplink signal-to-interference-plus-noise ratios (SINRs), which account for cross-link interference and self-interference, respectively, can be expressed as

$$\text{SINR}_{\text{tx}}(\mathbf{f}) = \frac{\text{SNR}_{\text{tx}}(\mathbf{f})}{1 + \text{INR}_{\text{tx}}} \quad (5)$$

$$\text{SINR}_{\text{rx}}(\mathbf{f}, \mathbf{w}) = \frac{\text{SNR}_{\text{rx}}(\mathbf{w})}{1 + \text{INR}_{\text{rx}}(\mathbf{f}, \mathbf{w})}. \quad (6)$$

The achievable downlink, uplink, and sum spectral efficiencies, with transmit beam \mathbf{f} and receive beam \mathbf{w} , are then

$$R_{\text{tx}}(\mathbf{f}) = \log_2(1 + \text{SINR}_{\text{tx}}(\mathbf{f})) \quad (7)$$

$$R_{\text{rx}}(\mathbf{f}, \mathbf{w}) = \log_2(1 + \text{SINR}_{\text{rx}}(\mathbf{f}, \mathbf{w})) \quad (8)$$

$$R(\mathbf{f}, \mathbf{w}) = R_{\text{tx}}(\mathbf{f}) + R_{\text{rx}}(\mathbf{f}, \mathbf{w}). \quad (9)$$

3. ACTIVE BEAM LEARNING FOR FULL-DUPLEX

In this section, we aim to design \mathbf{f} and \mathbf{w} in order to maximize the sum spectral efficiency $R(\mathbf{f}, \mathbf{w})$, without assuming knowledge of the downlink, uplink, or self-interference channels *a priori*. As mentioned in the introduction, most prior approaches to this problem assume these channels have been reliably estimated beforehand and thus ignores the resources consumed by such estimation. In contrast, we consider a resource-constrained setting, wherein the base station is allotted a *fixed* number of time slots to take measurements of downlink, uplink, and self-interference before designing the beams used for full-duplex data transmission and reception. Constraining our design to a fixed overhead

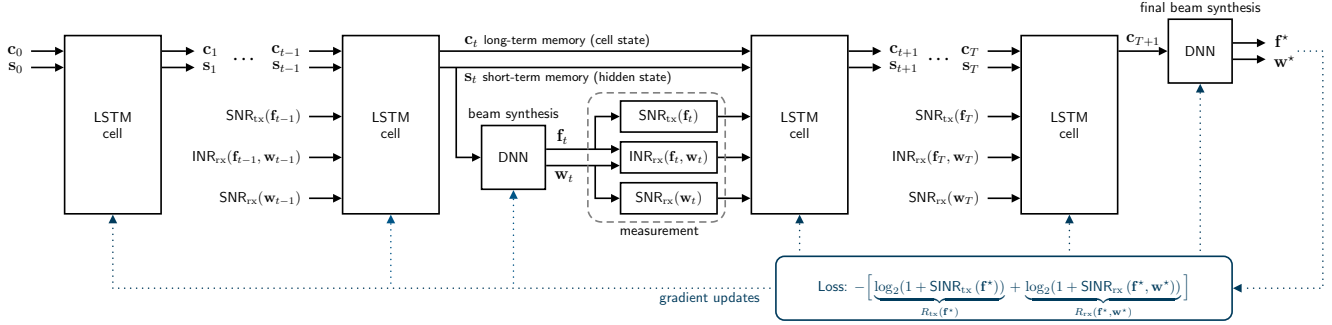


Fig. 3. An LSTM-based implementation of our proposed active beam learning solution for full-duplex wireless systems.

will thus make it more readily adopted in real-world wireless networks.

Our envisioned approach to this problem is shown in Fig. 2, where the base station is allotted T time slots before designing the beams it uses to serve downlink and uplink for L time slots. In the first time slot, the proposed *beam learning* process begins with the base station measuring $\text{SNR}_{\text{tx}}(\mathbf{f}_1)$, $\text{SNR}_{\text{rx}}(\mathbf{w}_1)$, and $\text{INR}_{\text{rx}}(\mathbf{f}_1, \mathbf{w}_1)$ using some initial *probing* beam pair $(\mathbf{f}_1, \mathbf{w}_1)$, followed by designing the next probing beam pair $(\mathbf{f}_2, \mathbf{w}_2)$ based on these initial measurements. In the next time slot, the base station uses these new probing beams to measure the SNRs and the INR, and then similarly generates the next probing beams $(\mathbf{f}_3, \mathbf{w}_3)$ based on *all* prior measurements. This process is repeated for all T probing time slots, after which the base station uses information gathered from the entire probing stage to design the *serving* transmit and receive beams $(\mathbf{f}^*, \mathbf{w}^*)$, which ideally maximizes the sum spectral efficiency. Designing the sequence of probing beams $\{(\mathbf{f}_t, \mathbf{w}_t)\}_{t=1}^T$ in the aforementioned *active* manner [11] will prove capable of designing near-optimal $(\mathbf{f}^*, \mathbf{w}^*)$, even when the probing overhead T is relatively small.

3.1. Problem Formulation

Inspired by the work of [11], our envisioned active beam learning problem, as just described and illustrated in Fig. 2, can be formulated as

$$\max_{\{\mathcal{F}_t(\cdot)\}_{t=0}^{T-1}, \mathcal{G}(\cdot)} R_{\text{tx}}(\mathbf{f}^*) + R_{\text{rx}}(\mathbf{f}^*, \mathbf{w}^*) \quad (10a)$$

$$\text{s.t. } (\mathbf{f}_{t+1}, \mathbf{w}_{t+1}) = \mathcal{F}_t(\mathbf{f}_{1:t}, \mathbf{w}_{1:t}, \mathbf{y}_{1:t}) \quad \forall t \quad (10b)$$

$$(\mathbf{f}^*, \mathbf{w}^*) = \mathcal{G}(\mathbf{f}_{1:T}, \mathbf{w}_{1:T}, \mathbf{y}_{1:T}) \quad (10c)$$

$$\|\mathbf{f}_t\|_2^2 = \|\mathbf{f}^*\|_2^2 = N_t \quad \forall t = 1, \dots, T \quad (10d)$$

$$\|\mathbf{w}_t\|_2^2 = \|\mathbf{w}^*\|_2^2 = N_r \quad \forall t = 1, \dots, T, \quad (10e)$$

where $\mathbf{y}_t = [\text{SNR}_{\text{tx}}(\mathbf{f}_t), \text{SNR}_{\text{rx}}(\mathbf{w}_t), \text{INR}_{\text{rx}}(\mathbf{f}_t, \mathbf{w}_t)]$ are the measurements taken with probing beams \mathbf{f}_t and \mathbf{w}_t during time slot t . Our goal in this problem is to find the series of

functions $\{\mathcal{F}_t(\cdot)\}_{t=0}^{T-1}$ and the function $\mathcal{G}(\cdot)$ that produces a final serving beam pair $(\mathbf{f}^*, \mathbf{w}^*)$ which maximizes the sum spectral efficiency. Here, $\mathcal{F}_t(\cdot)$ is a function that outputs the probing beams for time slot $t + 1$ based on all measurements collected through time t , whereas $\mathcal{G}(\cdot)$ is a function that designs the final serving beams based on all T measurements from the probing phase. Note that the same functions $\{\mathcal{F}_t(\cdot)\}_{t=0}^{T-1}$ and $\mathcal{G}(\cdot)$ are to be used across user pairs and self-interference channel realizations.

By developing an *active* probing strategy in this fashion, our aim is to collect a unique sequence of measurements tailored specifically to the user pair being served. This will allow the probing sequence to implicitly account for the structure of the self-interference channel in relation to those users' downlink and uplink channels. In turn, each probing sequence will be more targeted to the particular user pair and self-interference channel realization (than a non-active approach) and will thus give our active beam learning solution the potential to excel with only a few probing measurements. It is also worth noting that our proposed design is based solely on noncoherent SNR and INR measurements, lending itself well to certain practical deployments.

3.2. LSTM-Based Active Beam Learning

To implement our active beam learning solution and (approximately) solve problem (10), we adopt the long short-term memory (LSTM) architecture illustrated in Fig. 3, taking inspiration from [11]. LSTMs are especially well-suited for our active beam learning task, as they can effectively learn to extract valuable information from each probing measurement and retain a record of this extracted information, allowing us to leverage all prior probing measurements when designing subsequent probing beams or the final serving beams. In the remainder of this section, we will explain our LSTM-based approach and describe how the architecture shown in Fig. 3 is trained to solve problem (10).

As shown in Fig. 3, each LSTM cell indexed at time slot t takes as input the following three sets of parameters: (i) the downlink SNR, uplink SNR, and INR measured with prob-

ing beams \mathbf{f}_t and \mathbf{w}_t ; (ii) the LSTM hidden state at time t , denoted by \mathbf{s}_t ; and (iii) the LSTM cell state at time t , denoted by \mathbf{c}_t . The LSTM cell then outputs the updated hidden and cell states, denoted as \mathbf{s}_{t+1} and \mathbf{c}_{t+1} , respectively, and feeds the new hidden state into a deep neural network (DNN) to synthesize probing beams ($\mathbf{f}_{t+1}, \mathbf{w}_{t+1}$) for use in the next time slot, $t + 1$. As the hidden state acts as a memory that carries relevant information from both recent and past measurements, the described procedure synthesizes probing beams ($\mathbf{f}_{t+1}, \mathbf{w}_{t+1}$) based on all prior measurements. The aforementioned process is repeated sequentially across $T - 1$ time slots. Following the final measurements at time slot T , the last cell state \mathbf{c}_{T+1} , which functions as a long-term memory holding valuable information gathered throughout the entire probing phase, is input to another DNN to synthesize the serving beams \mathbf{f}^* and \mathbf{w}^* . Assuming we have models or datasets of the downlink, uplink, and self-interference channels, the model parameters can be optimized to maximize the sum spectral efficiency through backpropagation. Then, after training, the model can be deployed to actively generate probing beams for a given user pair and self-interference channel realization and, based on these probing measurements, can output a final beam pair ($\mathbf{f}^*, \mathbf{w}^*$) for serving downlink and uplink in a full-duplex fashion, realizing the solution envisioned in Fig. 2.

4. EVALUATION

To evaluate our method, we consider a 28 GHz full-duplex base station equipped with two 8-element half-wavelength linear arrays spaced apart horizontally 10λ , which it uses to serve single-antenna, line-of-sight users located uniformly between -60° and 60° in azimuth. We model the static component of the self-interference channel $\bar{\mathbf{H}}$ with the near-field spherical-wave model [12] and the time-varying component $\tilde{\mathbf{H}}$ as a Rayleigh fading channel drawn independently from user pair to user pair. To account for measurement errors that inevitably arise in practice, we assume (during training and evaluation) that the probing measurements (in dB) are corrupted by zero-mean Gaussian noise with variance σ^2 .¹ To satisfy the power constraints of (10d) and (10e), we perform appropriate normalizations to \mathbf{f}_t and \mathbf{w}_t at each time slot. Additionally, we normalize \mathbf{h}_{tx} , \mathbf{h}_{rx} , and \mathbf{H} such that the maximum achievable downlink and uplink SNRs are 10 dB and the maximum possible INR is 40 dB. Lastly, for simplicity, we assume that there is no cross-link interference present between the downlink and uplink users.

We use TensorFlow to implement our LSTM-based active beam learning solution. The dimensions of the LSTM hidden and cell states, \mathbf{s}_t and \mathbf{c}_t , are set to 512. The probing

¹Specifically, we assume all SNR and INR measurements are of the form $[\text{SNR}_{\text{tx}}(\mathbf{f}_t)]_{\text{dB}} = [\text{SNR}_{\text{tx}}(\mathbf{f}_t)]_{\text{dB}} + \mathcal{N}(0, \sigma^2)$, where $\mathcal{N}(0, \sigma^2)$ is Gaussian with mean zero and variance σ^2 and $[x]_{\text{dB}} = 10 \log_{10}(x)$.

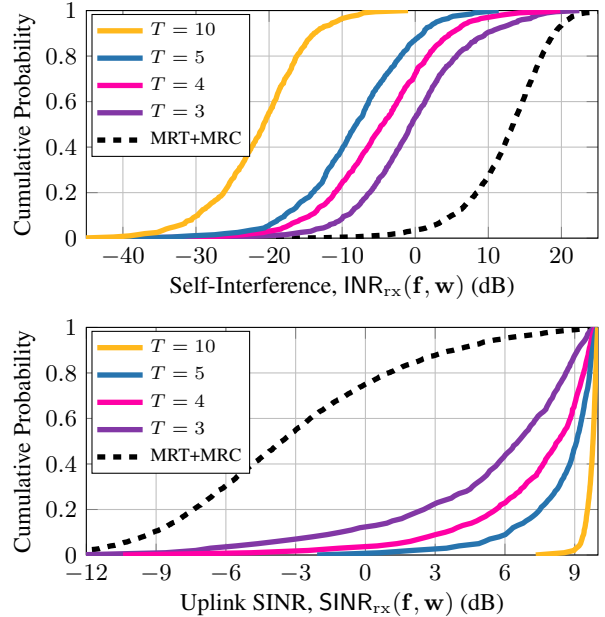


Fig. 4. Empirical CDFs of INR_{rx} and SINR_{rx} over 1000 random user pairs, when $\sigma^2 = 0.4$ and $\kappa = 7$ dB.

beam synthesis and final beam synthesis DNNs comprise of three hidden layers, each containing 1024 neurons with ReLU activation. Note that the output dimension of each of these DNNs is $2 \times (N_t + N_r)$ in order to account for both the real and imaginary components of the complex beamforming weights. During training, we utilize Adam optimizer with a learning rate of 10^{-4} and a mini-batch of size 128. As described, we set the loss function to $-R(\mathbf{f}^*, \mathbf{w}^*)$ to optimize the model to maximize the sum spectral efficiency.

We compare our approach against two baselines: (i) maximum ratio transmission (MRT) plus maximum ratio combining (MRC), which steers beams directly toward each user to maximize their SNRs, but ignores self-interference, and (ii) the full-duplex capacity, achieved by MRT+MRC in the absence of self-interference. Fig. 4 presents the empirical cumulative distribution functions (CDFs) of INR_{rx} and SINR_{rx} over 1000 random user pairs, when the measurement noise variance is $\sigma^2 = 0.4$ and the self-interference channel Rician factor is $\kappa = 7$ dB. The first key observation to make is that self-interference is often 10–20 dB above the noise floor with MRT+MRC, making it unsuitable for full-duplex operation. In contrast, active beam learning is able to design serving beams ($\mathbf{f}^*, \mathbf{w}^*$) that are far more robust to self-interference with only 3–10 time slots for probing. With $T = 10$, remarkably all 1000 randomly generated user pairs enjoy self-interference below the noise floor. This translates to higher uplink SINRs, which closely hugs the upper bound of 10 dB as the allotted probing time is increased to $T = 10$. Although MRT+MRC maximizes SNR,




Fig. 5. Sum spectral efficiency as a function of κ (averaged over 1000 random user pairs), when $\sigma^2 = 0.4$.

its high self-interference leads to poor uplink SINR.

In Fig. 5, we further investigated the impact of the structure of the self-interference channel \mathbf{H} on performance. Fixing $\sigma^2 = 0.4$, we plot the sum spectral efficiency as a function of the self-interference channel Rician factor κ , averaged over 1000 random user pairs. Recalling (3), increasing κ results in a more deterministic, spherical-wave self-interference channel, while a lower value leads to a more random, Gaussian channel. From the plot, it is evident that our method achieves higher sum spectral efficiency as κ increases; we attribute this to two main reasons. First, from the dashed MRT+MRC curve, we can see that a spherical-wave channel is structurally more favorable on average than a Gaussian one, also observed in [5]. Second, the performance of active beam learning increases with κ due to the simple fact that the channel becomes more deterministic, making it easier to learn probing strategies and final synthesis of $(\mathbf{f}^*, \mathbf{w}^*)$. With that being said, even when κ is low and \mathbf{H} is heavily Gaussian, active beam learning can still provide performance approaching the full-duplex capacity with sufficient probing time T . Altogether, these numerical results suggest that active beam learning is capable of intelligently probing to strategically design serving beams $(\mathbf{f}^*, \mathbf{w}^*)$ which suppress self-interference without compromising downlink and uplink SNR.

5. CONCLUSION AND FUTURE WORK

This work introduced an active beam learning scheme that aims to maximize the sum spectral efficiency of a full-duplex system. This is realized using a network of LSTMs and DNNs which learns to actively probe the downlink, uplink, and self-interference channels, rather than explicitly estimate them. We demonstrated that our approach is able to

suppress self-interference to below noise and attain sum spectral efficiencies that approach the full-duplex capacity when the probing time is sufficiently large. Our results also illustrate its robustness to measurement noise and to the structure of self-interference. Valuable future work may adapt our approach to scale more favorably to many users or to exploit temporal or spatial correlations in the system.

6. REFERENCES

- [1] B. Smida, A. Sabharwal, G. Fodor, G. C. Alexandropoulos, H. A. Suraweera, and C.-B. Chae, “Full-duplex wireless for 6G: Progress brings new opportunities and challenges,” *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2729–2750, Sep. 2023.
- [2] V. Singh, S. Mondal, A. Gadre, M. Srivastava, J. Paramesh, and S. Kumar, “Millimeter-wave full duplex radios,” in *Proc. ACM MobiCom*, Apr. 2020.
- [3] B. Smida, R. Wichman, K. E. Kolodziej, H. A. Suraweera, T. Riihonen, and A. Sabharwal, “In-band full-duplex: The physical layer,” *Proc. IEEE*, vol. 112, no. 5, pp. 433–462, May 2024.
- [4] I. P. Roberts, J. G. Andrews, H. B. Jain, and S. Vishwanath, “Millimeter-wave full duplex radios: New challenges and techniques,” *IEEE Wireless Commun.*, vol. 28, no. 1, pp. 36–43, Feb. 2021.
- [5] I. P. Roberts, S. Vishwanath, and J. G. Andrews, “LONESTAR: Analog beamforming codebooks for full-duplex millimeter wave systems,” *IEEE Trans. Wireless Commun.*, vol. 22, no. 9, pp. 5754–5769, Sep. 2023.
- [6] R. López-Valcarce and M. Martínez-Cotelo, “Full-duplex mmWave MIMO with finite-resolution phase shifters,” *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 8979–8994, Nov. 2022.
- [7] K. Satyanarayana, M. El-Hajjar, P. Kuo, A. Mourad, and L. Hanzo, “Hybrid beamforming design for full-duplex millimeter wave communication,” *IEEE Trans. Veh. Technol.*, pp. 1394–1404, Feb. 2019.
- [8] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, “STEER: Beam selection for full-duplex millimeter wave communication systems,” *IEEE Trans. Commun.*, vol. 70, no. 10, pp. 6902–6917, Oct. 2022.
- [9] I. P. Roberts, Y. Zhang, T. Osman, and A. Alkhateeb, “Real-world evaluation of full-duplex millimeter wave communication systems,” *IEEE Trans. Wireless Commun.*, vol. 23, no. 9, pp. 10 803–10 819, Sep. 2024.
- [10] Y. Heng *et al.*, “Six key challenges for beam management in 5.5G and 6G systems,” *IEEE Commun. Mag.*, vol. 59, no. 7, pp. 74–79, Jul. 2021.
- [11] F. Sohrabi, T. Jiang, W. Cui, and W. Yu, “Active sensing for communications by learning,” *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1780–1794, Jun. 2022.
- [12] J.-S. Jiang and M. A. Ingram, “Spherical-wave model for short-range MIMO,” *IEEE Trans. Commun.*, vol. 53, no. 9, pp. 1534–1541, Sep. 2005.