

Resolution-Adaptive All-Digital Spatial Equalization for mmWave Massive MU-MIMO

Oscar Castañeda¹, Seyed Hadi Mirfarshbafan¹, Shahaboddin Ghajari²,
Alyosha Molnar², Sven Jacobsson³, Giuseppe Durisi⁴, and Christoph Studer¹

¹*Department of Information Technology and Electrical Engineering, ETH Zürich, Zürich, Switzerland*

²*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA*

³*Ericsson Research, Gothenburg, Sweden*

⁴*Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden*

Abstract—All-digital basestation (BS) architectures for millimeter-wave (mmWave) massive multi-user multiple-input multiple-output (MU-MIMO), which equip each radio-frequency chain with dedicated data converters, have advantages in spectral efficiency, flexibility, and baseband-processing simplicity over hybrid analog-digital solutions. For all-digital architectures to be competitive with hybrid solutions in terms of power consumption, novel signal-processing methods and baseband architectures are necessary. In this paper, we demonstrate that adapting the resolution of the analog-to-digital converters (ADCs) and spatial equalizer of an all-digital system to the communication scenario (e.g., the number of users, modulation scheme, and propagation conditions) enables orders-of-magnitude power savings for realistic mmWave channels. For example, for a 256-BS-antenna 16-user system supporting 1 GHz bandwidth, a traditional baseline architecture designed for a 64-user worst-case scenario would consume 23 W in 28 nm CMOS for the ADC array and the spatial equalizer, whereas a resolution-adaptive architecture is able to reduce the power consumption by $6.7\times$.

I. INTRODUCTION

Millimeter-wave (mmWave) communication [1], [2] offers vast amounts of unused spectrum and is considered a key technology component of fifth-generation (5G) and beyond 5G wireless systems. A major challenge of mmWave communication is the high path loss [3], which can be mitigated with massive multiple-input multiple-output (MIMO) [4]. Massive multi-user (MU) MIMO is able to compensate for the high path loss via fine-grained beamforming while supporting concurrent communication with multiple user equipments (UEs) in the same frequency band. A practical realization of mmWave massive MU-MIMO basestations (BSs), however, faces serious implementation challenges that are caused by the large number of BS antennas and the large communication bandwidth.

Two prominent solutions that lower the cost and power consumption of mmWave massive MU-MIMO systems are (i) hybrid digital-analog architectures [5], [6] in which the number of radio-frequency (RF) chains is less than the number of antennas and (ii) all-digital architectures that rely on nonlinear RF chains and low-resolution data converters [7]–[9]. In this paper, we focus on all-digital architectures that are able to exploit the full potential of massive MU-MIMO and offer

higher spectral efficiency with comparable power consumption to hybrid architectures [10], [11]. However, in order to keep the system costs and power consumption of all-digital architectures within practical limits, it is indispensable to rely on low-resolution data converters [12] as well as on low-resolution digital baseband processing methods [13].

A. Contributions

In this paper, we show that for all-digital massive MU-MIMO systems with B BS antennas and U UEs, the system's load factor defined as $\beta \triangleq U/B$, as well as other specifics of the communication system (e.g., the modulation and channel conditions) do not only determine the spectral efficiency [14] and robustness against RF impairments [12], [15], [16], but also open up new means to optimize the power consumption of all-digital spatial equalization architectures. More specifically, if the number of active UEs is low, then it is possible to reduce the resolution of the analog-to-digital converters (ADCs) and the resolution of finite-alphabet equalizers, as well as the number of active BS antennas, without noticeably degrading the system's error-rate performance. We demonstrate that such *resolution-adaptive* all-digital massive MU-MIMO BS architectures are key to enabling up to two orders-of-magnitude power savings in the ADC array and the all-digital spatial equalizer.

Previous work [17]–[19] focused on adapting the resolution of each individual ADC to the channel state. In contrast, we (i) adapt the resolutions of both the ADC array *and* the spatial equalizer to the instantaneous communication scenario and (ii) study their individual impact on power consumption.

B. Notation

Boldface uppercase and lowercase letters represent matrices and column vectors, respectively. For a matrix \mathbf{A} , the conjugate transpose is \mathbf{A}^H , the k th column is \mathbf{a}_k , and the entry on the k th row and ℓ th column is $A_{k,\ell}$. The $N \times N$ identity matrix is \mathbf{I}_N . For a vector \mathbf{a} , the k th entry is a_k , the ℓ_2 -norm is $\|\mathbf{a}\|_2$, the real part is $\Re\{\mathbf{a}\}$ and the imaginary parts is $\Im\{\mathbf{a}\}$. The ℓ_∞ -norm and ℓ_∞ -norm are defined as $\|\mathbf{a}\|_\infty \triangleq \max_k |a_k|$ and $\|\mathbf{a}\|_\infty \triangleq \max\{\|\Re\{\mathbf{a}\}\|_\infty, \|\Im\{\mathbf{a}\}\|_\infty\}$. Expectation with respect to the random vector \mathbf{x} is denoted by $\mathbb{E}_{\mathbf{x}}[\cdot]$.

II. SYSTEM MODEL

We consider the uplink of a mmWave massive MU-MIMO system in which U single-antenna UEs transmit to a resolution-

The work of OC, SHM, AM, and CS was supported in part by ComSenTer, one of six centers in JUMP, a SRC program sponsored by DARPA. The work of CS was also supported by an ETH Research Grant and by the US NSF under grants CNS-1717559, CNS-1955997, and ECCS-1824379.

Contact author: O. Castañeda; e-mail: caoscar@ethz.ch

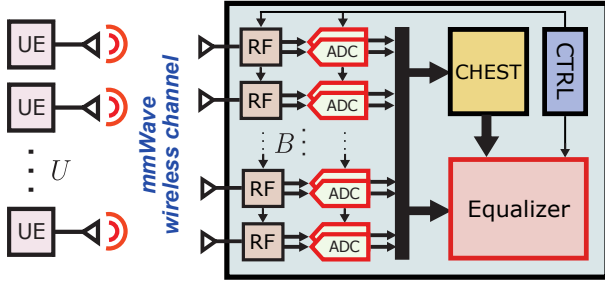


Fig. 1. System overview of a resolution-adaptive all-digital mmWave massive MU-MIMO BS equipped with B antennas, a pair of ADCs per RF chain, and a digital equalizer. A controller (CTRL) adapts the number of active BS antennas B' as well as the resolution of the ADCs and equalizer based on the communication scenario in order to minimize the system's power consumption.

adaptive all-digital BS with B antennas as illustrated in Figure 1. At the BS, we model the nonlinear distortion artifacts caused by finite-resolution ADCs and finite-alphabet spatial equalization. Our goal is to characterize the performance/power trade-offs of such a resolution-adaptive all-digital equalization architecture.

A. System Model

We model wireless transmission with the following frequency-flat input-output relation: $\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{n}$. Here, $\mathbf{y} \in \mathbb{C}^B$ is the (unquantized) receive vector at the BS, $\mathbf{H} \in \mathbb{C}^{B \times U}$ is the MIMO channel matrix, $\mathbf{s} \in \mathbb{S}^U$ is the transmit vector in which each entry corresponds to the per-UE transmit symbol taken from a constellation set \mathcal{S} (e.g., 16-QAM), and $\mathbf{n} \in \mathbb{C}^B$ models i.i.d. circularly-symmetric complex Gaussian noise with a per-entry variance of N_0 . In what follows, we assume that the UE transmit symbols s_u , $u = 1, \dots, U$, are independent and zero mean, each with a variance of $E_s \sigma_u^2$. We further assume ± 3 dB power control across UEs so that $\max_u \{\sigma_u^2 \|\mathbf{h}_u\|^2\} / \min_u \{\sigma_u^2 \|\mathbf{h}_u\|^2\} = 4$ [13]. For this model, the average receive signal-to-noise ratio (SNR) is

$$\text{SNR} \triangleq \frac{\mathbb{E}_{\mathbf{s}} [\|\mathbf{H}\mathbf{s}\|_2^2]}{\mathbb{E}_{\mathbf{n}} [\|\mathbf{n}\|_2^2]} = \frac{E_s \sum_{u=1}^U \sigma_u^2 \|\mathbf{h}_u\|_2^2}{N_0 B}, \quad (1)$$

where $\mathbf{h}_u \in \mathbb{C}^B$, $u = 1, \dots, U$, is the channel for the u th UE. We further consider that the BS can control the number of active antennas $B' \leq B$, so that only a subset $\hat{\mathbf{y}} \in \mathbb{C}^{B'}$ of the signals in \mathbf{y} are further processed. In our case, the BS will select the B' contiguous antennas at the center of the array.

B. Analog-to-Digital Conversion

In order to reduce the power consumption of all-digital BS architectures, practical systems will have to rely on low-resolution data converters [7]–[9]. To take into account the quantization artifacts caused by such low-resolution ADCs, we assume that the receive (active) signal $\hat{\mathbf{y}}$ passes through $2B'$ quantizers, B' for the in-phase and quadrature (IQ) components, resulting in the quantized receive vector \mathbf{z} , where $\mathbf{z} = Q_q(\Re\{\hat{\mathbf{y}}\}) + jQ_q(\Im\{\hat{\mathbf{y}}\})$. Here, the quantization function Q_q is applied entry-wise to its argument and models a q -bit uniform midrise quantizer with step size Δ as, e.g., in [8]

$$Q_q(y) = \begin{cases} \Delta \lfloor \frac{y}{\Delta} \rfloor + \frac{\Delta}{2}, & \text{if } |y| < \Delta 2^{q-1} \\ \frac{\Delta}{2} (2^q - 1) \frac{y}{|y|}, & \text{if } |y| > \Delta 2^{q-1}. \end{cases} \quad (2)$$

We use the step size Δ that minimizes the mean-squared error (MSE) between the quantizer's input y and output $Q_q(y)$ assuming that the ADCs' input is a circularly-symmetric complex Gaussian random variable [20] with variance equal to

$$\sigma_{\text{ADC}}^2 = N_0 + E_s \max_{b \in \{1, \dots, B'\}} \sum_{u=1}^U \sigma_u^2 |H_{b,u}|^2, \quad (3)$$

which is the maximum variance across all active receive antennas. In practice, the step size would be adjusted by an automatic gain control circuit.

C. Finite-Alphabet Equalization

Given the extremely large bandwidths available at mmWave frequencies, linear equalization is preferred to keep implementation complexity within reasonable bounds. A linear equalizer generates estimates $\hat{\mathbf{s}}$ of the transmit vector \mathbf{s} according to $\hat{\mathbf{s}} = \mathbf{W}^H \mathbf{z}$, where $\mathbf{W}^H \in \mathbb{C}^{U \times B'}$ is a suitably-designed equalization matrix. The equalization matrix \mathbf{W}^H is typically designed to minimize the post-equalization MSE defined as $MSE \triangleq \mathbb{E}_{\mathbf{s}, \mathbf{n}} [\|\hat{\mathbf{s}} - \mathbf{s}\|^2]$, which results in the widely-used linear minimum MSE (L-MMSE) equalization matrix

$$\mathbf{W}^H = \left(\tilde{\mathbf{H}}^H \tilde{\mathbf{H}} + \frac{N_0}{E_s} \mathbf{I}_U \right)^{-1} \tilde{\mathbf{H}}^H, \quad (4)$$

where $\tilde{\mathbf{H}} \in \mathbb{C}^{B' \times U}$ is the channel estimated during a training phase. While common L-MMSE equalizer implementations use high-resolution numbers (typically 10-bit to 12-bit) for the entries of \mathbf{W}^H (see, e.g., [21]), it was shown in [13] that the corresponding hardware can be power hungry in large-bandwidth applications, as it is the case for mmWave systems.

In order to reduce the power consumption of the equalization step $\hat{\mathbf{s}} = \mathbf{W}^H \mathbf{z}$, the work in [13] proposes *finite-alphabet equalization*, which uses carefully-designed low-resolution equalization matrices of the form $\mathbf{V}^H = \text{diag}(\boldsymbol{\mu}) \mathbf{X}^H$. Here, the entries of the matrix $\mathbf{X}^H \in \mathbb{C}^{U \times B'}$ are taken from a reduced-precision alphabet \mathcal{X} (e.g., represented using 1-bit to 4-bit numbers). The vector $\boldsymbol{\mu}$ contains post-equalization scaling factors that are represented with more bits (e.g., 10-bit). In order to obtain unbiased estimates of the transmitted symbols, the entries of the scaling vector $\boldsymbol{\mu}$ are set to [22]

$$\mu_u = (\mathbf{x}_u^H \tilde{\mathbf{h}}_u)^{-1}, \quad u = 1, \dots, U, \quad (5)$$

and the resulting estimates are computed as follows:

$$\hat{s}_u = \mathbf{v}_u^H \mathbf{z} = \mu_u \mathbf{x}_u^H \mathbf{z}, \quad u = 1, \dots, U. \quad (6)$$

Since the vectors \mathbf{x}_u^H , $u = 1, \dots, U$, contain low-resolution entries, computing the inner products $\mathbf{x}_u^H \mathbf{z}$ in (6) can be accomplished with low-resolution digital circuitry that occupies small area and consumes low power [13], [23]. Post-equalization scaling only needs to be calculated once per UE when computing the estimates \hat{s}_u , $u = 1, \dots, U$, and therefore does not dominate the equalizer's power consumption.

In [13], two methods are proposed to determine the entries of the low-resolution matrix \mathbf{X}^H , offering a trade-off between the complexity of calculating \mathbf{X}^H and the error-rate performance of the resulting finite-alphabet equalizer. In this work, we will use the finite-alphabet L-MMSE (FL-MMSE) method, which has lower complexity but exhibits inferior performance for 1-bit

and 2-bit resolutions. As detailed in [13], FL-MMSE computes the rows \mathbf{x}_u of \mathbf{X}^H by uniformly quantizing the rows \mathbf{w}_u of the L-MMSE equalization matrix \mathbf{W}^H in (4), according to $\mathbf{x}_u = Q_k(\Re\{\mathbf{w}_u\}) + jQ_k(\Im\{\mathbf{w}_u\})$, where $Q_k(\cdot)$ is given by (2) with $\Delta = \|\mathbf{w}_u\|_\infty 2^{1-k}$. Finally, the scaling factors μ_u , $u = 1, \dots, U$, are calculated from the vectors \mathbf{x}_u^H as in (5).

III. POWER MODELS

We now introduce our model for the total power consumption P of the ADC array and the spatial equalizer given by

$$P = P_{\text{ADC}} + P_{\text{EQ}}, \quad (7)$$

where P_{ADC} and P_{EQ} stand for the power consumption of the ADC array and the spatial equalizer, respectively. These two power models depend on the number of ADC bits q , the number of FL-MMSE bits k , the number of UEs U , the number of active BS antennas B' , and the ADC sampling rate f_s .

A. Power Model for the ADC Array

To model the power consumption P_{ADC} of the ADC array consisting of $2B'$ active converters that sample the I/Q baseband signals, we use Walden's figure of merit (FoM_W) [11] to obtain

$$P_{\text{ADC}} = \text{FoM}_W \times 2^q \times 2B' \times f_s. \quad (8)$$

In what follows, we will use a FoM_W of 70.8 fJ per conversion step, which was provided in [24] for a 28 nm CMOS successive-approximation register ADC with $f_s = 2$ GS/s.

B. Power Model for the Finite-Alphabet Equalizer

To model the power consumption P_{EQ} of a finite-alphabet equalizer, we consider the 28 nm CMOS design presented in [23], which supports a baseband sampling rate of $f_s = 2$ GS/s. This design consists of multiple time-interleaved instances of the Parallel Processor in Associative Content-addressable memory (PPAC) [25], a spatial architecture that uses $4kB'U$ processing-in-memory bit-cells to compute a complex-valued matrix-vector product between \mathbf{X}^H and \mathbf{z} in q clock cycles.

Since the equalization operation in (6) is independent for each UE, the power consumption is directly proportional to U . While the number of PPAC bit-cells is directly proportional to k and B' , varying k and B' also affects the power consumption of the logic in PPAC that processes the bit-cells' outputs. Modeling the effect of k and B' on this surrounding logic is non-trivial. However, since the bit-cells represent a significant component of the design, it is reasonable to use a linear model as in [7], [26] for the total power of the equalizer. As a matter of fact, the power consumption of the PPAC implementations in [23] scales linearly with k . We approximate the equalization power of PPAC as directly proportional to B' . Finally, the throughput of PPAC is inversely proportional to the ADC resolution q . Consequently, the number of time-interleaved PPAC instances required to reach the target $f_s = 2$ GS/s is directly proportional to q , just like P_{EQ} . Putting it all together, we arrive at

$$P_{\text{EQ}} = (2.44k - 0.48) \times \frac{q}{7} \times \frac{B'}{256} \times \frac{U}{16} \times \frac{f_s}{2 \text{ GS/s}} \quad [\text{W}]. \quad (9)$$

Here, the term $2.44k - 0.48$ is obtained from a linear fit for the PPAC power results provided in [23], which are for a system with $B = B' = 256$, $U = 16$, $q = 7$, and $k \in \{1, 2, 3\}$.

We note that the model in (9) only considers the power consumption for applying equalization as in $\hat{\mathbf{s}} = \mathbf{V}^H \mathbf{z}$. We do *not* consider the power consumption for preprocessing, i.e., computing the L-MMSE matrix in (4) and extracting the FL-MMSE matrix. While equalization must be carried out for every received vector at baseband sampling rate, preprocessing is required only once per channel coherence interval, which can last for more than 10^4 samples in mmWave systems [27].

IV. RESULTS

A. Simulation Setup and Performance Metrics

We consider a mmWave massive MU-MIMO uplink system in which a BS with a $B = 256$ antenna $\lambda/2$ -spacing uniform linear array (ULA) is receiving data from $U \in \{1, 2, 4, 8, 16, 32, 64\}$ UEs transmitting at $f_s = 2$ GS/s and a carrier frequency of 60 GHz. We perform Monte-Carlo simulations using line-of-sight (LoS) and non-LoS channel matrices generated with QuaDRiGA [28] using the mmMAGIC UMi scenario. We model situations in which the UEs are randomly placed at distances ranging from 10 m to 100 m in front of the BS, within a 120° sector, and a minimum angular separation of 1° . We furthermore assume ± 3 dB per-UE power control as described in Section II. We estimate the channel using a pilot-based least squares (LS) estimator followed by denoising with beamspace channel estimation (BEACHES) [29].

In order to evaluate the effect that the number of ADC bits q , finite-alphabet equalization bits k , and active BS antennas B' have on the system's error-rate performance and power consumption, we vary $q \in \{1, 2, \dots, 8, \infty\}$, $k \in \{1, 2, \dots, 6, \infty\}$, and $B' \in \{232, 233, \dots, 256\}$. A value of $q = \infty$ corresponds to bypassing the ADCs, i.e., $\mathbf{z} = \hat{\mathbf{y}}$; setting $k = \infty$ corresponds to unbiased L-MMSE equalization without quantization. For each configuration of q , k , B' , and U , we quantify the error-rate performance by finding the minimum SNR at which 1% uncoded bit-error rate (BER) is attained, and compute the SNR loss (in decibels) with respect to the ideal case in which $q = k = \infty$ and $B' = B = 256$. For each configuration of q , k , B' , and U , we also determine the resulting power consumption $P = P_{\text{ADC}} + P_{\text{EQ}}$ using the models in (8) and (9). For all of the considered configurations with a fixed number of UEs U , we find the Pareto-optimal envelope of the designs that achieve the lowest power consumption under the same SNR loss, or that achieve the lowest SNR loss at the same power consumption, i.e., the lower envelope of the (SNR loss, power P) data points generated from all considered (q, k, B') configurations. The results are shown in Figure 2 and discussed next.

B. SNR Operating Point vs. Power Trade-Offs

Figure 2(a) demonstrates that the proposed resolution-adaptive architecture enables one to trade-off performance (measured in terms of the SNR loss) versus the ADC and equalizer power P . For a fixed UE load U , we see that by allowing a larger SNR loss, the BS power can be reduced considerably. Indeed, by lowering the ADC and equalizer resolutions q and k , respectively, as well as the number of active BS antennas B' , power consumption is reduced, whilst requiring a higher SNR to achieve the same target BER of 1%. This implies that for situations in which a larger SNR




Fig. 2. ADC array and spatial equalizer power consumption $P = P_{\text{ADC}} + P_{\text{EQ}}$ for Pareto-optimal configurations in a $B = 256$ BS antenna system. For each UE load U , we vary the number of ADC bits q , equalization bits k , and active BS antennas B' . Our results for both LoS and non-LoS channels with QPSK and 16-QAM communication reveal that the power consumption can be reduced by two orders of magnitude depending on the communication scenario.




Fig. 3. ADC array and spatial equalizer power consumption $P = P_{\text{ADC}} + P_{\text{EQ}}$ for Pareto-optimal configurations as the number of UEs U varies. The baseline corresponds to a system designed for the worst case of $U = 64$ with 16-QAM that is only able to adapt the equalizer's active processing elements to the UE load U . Resolution-adaptive designs can achieve up to $22\times$ power savings, depending on the UE load U , modulation scheme, and SNR loss at 1% BER.

loss can be tolerated, one can drastically reduce the system's power consumption. Consider, for example, Figure 2(a) with $U = 8$ UEs. For an SNR loss of 0 dB, the power consumption would be $P = 26.6$ W whereas for an SNR loss of 1.5 dB, the power consumption is only $P = 1.5$ W. Hence, such resolution-adaptive architectures can reduce the power consumption by one order of magnitude for a given UE load.

When comparing Figure 2(a) with Figure 2(b), we furthermore observe that equalization of QPSK symbols instead of 16-QAM can further reduce the power consumption, which implies that simpler equalization problems require less work. Furthermore, when comparing Figure 2(b) with Figure 2(c), we can see that the same performance and power trade-off persists under LoS and non-LoS propagation conditions, respectively. This indicates that resolution-adaptive architectures provide similar gains for different mmWave propagation scenarios.

C. Communication Scenario vs. Power Trade-Offs

Figure 2 also illustrates how the total power consumption P spans a wide range as the load factor, modulation scheme, and propagation conditions vary. Hence, fully exploiting all of these dynamic aspects in the BS' power consumption requires a *resolution-adaptive* architecture (cf. Figure 1) that is able to dynamically adapt to these factors.

In order to further demonstrate the potential of such a resolution-adaptive BS architecture, we perform the following

experiment. First, we consider a baseline BS architecture with $B' = B = 256$ (active) antennas that can adapt *only* its equalizer to the UE load U as follows: Since linear equalization is an independent problem per UE, a linear equalizer architecture can easily be adapted so that the equalizer power P_{EQ} is directly proportional to U . However, the number of ADC bits q and equalization bits k for this baseline architecture is fixed so that the worst SNR loss experienced across all different considered scenarios (load factors, modulation schemes, and propagation conditions) is below 0.1 dB. In our simulations, this worst-case scenario is $U = 64$ with 16-QAM (under both LoS and non-LoS propagation conditions), and requires $q = 7$ and $k = 6$. The power consumption P of such a baseline architecture is shown with a solid line in Figure 3.

In contrast, the proposed resolution-adaptive architecture is not only able to optimize its equalization power P_{EQ} as a function of the number of UEs U , but it can also dynamically adapt the ADC resolution q , equalization resolution k , and the number of active antennas B' . The power consumption of such resolution-adaptive architecture is shown with a dashed line in Figure 3(a) when operating with 16-QAM under LoS propagation. Clearly, a resolution-adaptive architecture has to support the same worst-case scenario as the baseline and, hence, the power consumption is the same for $U = 64$ UEs. However, once we consider smaller load factors β , the resolution-adaptive architecture reduces the power consumption by $2.2\times$ for $U = 16$

and reaches maximum power savings of $14\times$ at $U = 1$. In Figure 3(b), the modulation scheme changes from 16-QAM to QPSK. The baseline architecture cannot exploit this change, whereas a resolution-adaptive design can readjust q , k , and B' to reduce power consumption by at least $1.8\times$ when $U = 64$. Furthermore, Figure 3(c) illustrates how further power savings can be achieved by allowing a larger SNR loss of 0.5 dB, significantly reducing power by $6.7\times$ when $U = 16$ and up to a factor of $22\times$ when $U = 1$. Finally, the shaded areas in Figure 3 show how much power is consumed by the ADC array and equalizer of the resolution-adaptive design. We observe that the equalizer dominates the total power consumption, except for very small load factors of $U \leq 2$.

V. CONCLUSIONS

We have proposed a resolution-adaptive ADC array and all-digital spatial equalization architecture for mmWave massive MU-MIMO, which is able to optimize the system's power consumption depending on the load factor, modulation scheme, and channel conditions. By combining low-resolution ADCs with finite-alphabet equalization, the power and performance (measured as the minimum SNR required to achieve a target BER) of all-digital architectures is tunable with respect to the ADC bits q , equalization bits k , and number of active BS antennas B' . We have shown that resolution adaptivity enables a wide trade-off region between power consumption and performance, with a power consumption ranging from 83 W for $U = 64$ UEs to only 0.15 W for $U = 1$ UE in a 256-antenna mmWave massive MU-MIMO BS supporting a bandwidth of 1 GHz. In summary, our results demonstrate that resolution adaptivity is a critical component of power-efficient all-digital mmWave massive MU-MIMO BS designs.

We see the following avenues for future work. Augmenting our analysis with the impact on power and performance of RF circuitry (including low-noise amplifier, mixer, etc.) would provide a more comprehensive assessment of the benefits of resolution-adaptive BS designs. A theoretical analysis of the trade-offs shown in this paper is part of ongoing work.

REFERENCES

- [1] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, May 2013.
- [2] T. S. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang, "Overview of millimeter wave communications for fifth-generation (5G) wireless networks—with a focus on propagation models," *IEEE Trans. Antennas Propag.*, vol. 65, no. 12, pp. 6213–6230, Dec. 2017.
- [3] T. S. Rappaport, R. W. Heath Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Prentice Hall, 2015.
- [4] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [5] F. Sohrabi and W. Yu, "Hybrid digital and analog beamforming design for large-scale antenna arrays," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 501–513, Apr. 2016.
- [6] R. W. Heath Jr., N. González Precic, S. Rangan, W. Roh, and A. Sayeed, "An overview of signal processing techniques for millimeter wave MIMO systems," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 436–453, Feb. 2016.
- [7] M. Abdelghany, A. A. Farid, U. Madhow, and M. J. W. Rodwell, "Towards all-digital mmWave massive MIMO: Designing around nonlinearities," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1552–1557.
- [8] S. Jacobsson, U. Gustavsson, G. Durisi, and C. Studer, "Massive MU-MIMO-OFDM uplink with hardware impairments: Modeling and analysis," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1829–1835.
- [9] L. Xu, X. Lu, S. Jin, F. Gao, and Y. Zhu, "On the uplink achievable rate of massive MIMO system with low-resolution ADC and RF impairments," *IEEE Commun. Lett.*, vol. 23, no. 3, pp. 502–505, Jan. 2019.
- [10] H. Yan, S. Ramesh, T. Gallagher, C. Ling, and D. Cabric, "Performance, power, and area design trade-offs in millimeter-wave transmitter beamforming architectures," *IEEE Circuits Syst. Mag.*, vol. 19, no. 2, pp. 33–58, May 2019.
- [11] P. Skrimponis, S. Dutta, M. Mezzavilla, S. Rangan, S. H. Mirfarshbafan, C. Studer, J. Buckwalter, and M. Rodwell, "Power consumption analysis for mobile mmWave and sub-THz receivers," in *IEEE 6G Wireless Summit*, Mar. 2020.
- [12] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, "Throughput analysis of massive MIMO uplink with low-resolution ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 4038–4051, Jun. 2017.
- [13] O. Castañeda, S. Jacobsson, G. Durisi, T. Goldstein, and C. Studer, "Finite-alphabet MMSE equalization for all-digital massive MU-MIMO mmWave communication," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 9, pp. 2128–2141, Sep. 2020.
- [14] C. Jeon, R. Ghods, A. Maleki, and C. Studer, "Optimality of large MIMO detection via approximate message passing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Oct. 2015, pp. 1227–1231.
- [15] U. Gustavsson, C. Sanchéz-Perez, T. Eriksson, F. Athley, G. Durisi, P. Landin, K. Hausmair, C. Fager, and L. Svensson, "On the impact of hardware impairments on massive MIMO," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2014, pp. 294–300.
- [16] E. Björnson, M. Matthaiou, and M. Debbah, "Massive MIMO with non-ideal arbitrary arrays: Hardware scaling laws and circuit-aware design," *IEEE Trans. Wireless Commun.*, vol. 14, no. 8, pp. 4353–4368, Apr. 2015.
- [17] N. Liang and W. Zhang, "Mixed-ADC massive MIMO," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 4, pp. 983–997, May 2016.
- [18] J. Choi, B. L. Evans, and A. Gatherer, "Resolution-adaptive hybrid MIMO architectures for millimeter wave communications," *IEEE Trans. Signal Process.*, vol. 65, no. 23, pp. 6201–6216, Dec. 2017.
- [19] D. Verenzuela, E. Björnson, and M. Matthaiou, "Optimal per-antenna ADC bit allocation in correlated and cell-free massive MIMO," *IEEE Trans. Commun.*, vol. 69, no. 7, pp. 4767–4780, Jul. 2021.
- [20] J. Max, "Quantizing for minimum distortion," *IRE Trans. Inf. Theory*, vol. 6, no. 1, pp. 7–12, Mar. 1960.
- [21] M. Wu, B. Yin, G. Wang, C. Dick, J. R. Cavallaro, and C. Studer, "Large-scale MIMO detection for 3GPP LTE: Algorithms and FPGA implementations," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 916–929, Oct. 2014.
- [22] O. Castañeda, S. Jacobsson, G. Durisi, T. Goldstein, and C. Studer, "Soft-output finite alphabet equalization for mmWave massive MIMO," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, May 2020.
- [23] O. Castañeda, S. Jacobsson, G. Durisi, T. Goldstein, and C. Studer, "High-bandwidth spatial equalization for mmWave massive MU-MIMO with processing-in-memory," *IEEE Trans. Circuits Syst. II*, vol. 67, no. 5, pp. 891–895, May 2020.
- [24] W. Jiang, Y. Zhu, C.-H. Chan, B. Murmann, and R. P. Martins, "A 7-bit 2 GS/s time-interleaved SAR ADC with timing skew calibration based on current integrating sampler," *IEEE Trans. Circuits Syst. I*, vol. 68, no. 2, pp. 557–568, Feb. 2021.
- [25] O. Castañeda, M. Bobbitt, A. Gallyas-Sanhueza, and C. Studer, "PPAC: A versatile in-memory accelerator for matrix-vector-product-like operations," in *IEEE Int. Conf. Appl.-specific Syst., Archit. and Processors (ASAP)*, Jul. 2019, pp. 149–156.
- [26] J. Xu, W. Xu, and F. Gong, "On performance of quantized transceiver in multiuser massive MIMO downlinks," *IEEE Wireless Commun. Lett.*, vol. 6, no. 5, pp. 562–565, Oct. 2017.
- [27] T. E. Bogale, L. B. Le, and X. Wang, "Hybrid analog-digital channel estimation and beamforming: Training-throughput tradeoff," *IEEE Trans. Commun.*, vol. 63, no. 12, pp. 5235–5249, Dec. 2015.
- [28] S. Jaeckel, L. Raschkowski, K. Börner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.
- [29] S. H. Mirfarshbafan, A. Gallyas-Sanhueza, R. Ghods, and C. Studer, "Beamspace channel estimation for massive MIMO mmWave systems: Algorithm and VLSI design," *IEEE Trans. Circuits Syst. I*, vol. 64, no. 12, pp. 5482–5495, Dec. 2020.