

A Resolution-Adaptive 8 mm² 9.98 Gb/s 39.7 pJ/b 32-Antenna All-Digital Spatial Equalizer for mmWave Massive MU-MIMO in 65nm CMOS

Oscar Castañeda¹, Zachariah Boynton², Seyed Hadi Mirfarshbafan¹,
Shimin Huang², Jamie C. Ye², Alyosha Molnar², and Christoph Studer¹

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

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

Abstract—All-digital millimeter-wave (mmWave) massive multi-user multiple-input multiple-output (MU-MIMO) receivers enable extreme data rates but require high power consumption. In order to reduce power consumption, this paper presents the first resolution-adaptive all-digital receiver ASIC that is able to adjust the resolution of the data-converters and baseband-processing engine to the instantaneous communication scenario. The scalable 32-antenna, 65 nm CMOS receiver occupies a total area of 8 mm² and integrates analog-to-digital converters (ADCs) with programmable gain and resolution, beamspace channel estimation, and a resolution-adaptive processing-in-memory spatial equalizer. With 6-bit ADC samples and a 4-bit spatial equalizer, our ASIC achieves a throughput of 9.98 Gb/s while being at least 2× more energy-efficient than state-of-the-art designs.

I. INTRODUCTION

Future wireless communication systems are expected to operate at millimeter-wave (mmWave) frequencies to exploit large amounts of unoccupied contiguous bandwidth [1]. Multiple-input multiple-output (MIMO) can be utilized to combat the strong path loss at mmWave frequencies while enabling improvements in spectral efficiency by supporting multi-user (MU) communication [2]. However, the concoction of wide bandwidths and large antenna arrays at the basestation (BS) results in a number of serious implementation challenges.

Power consumption is a major concern in the design of such mmWave massive MU-MIMO systems. While hybrid analog-digital BS architectures have been proposed to address this issue [3], they are limited in their spatial-multiplexing capabilities resulting in a spectral efficiency loss. In contrast, all-digital BS architectures, with one radio-frequency (RF) chain and dedicated analog-to-digital converters (ADCs) for each antenna, are able to realize the full potential of massive MU-MIMO while achieving comparable RF power consumption to hybrid architectures [4], [5]. However, in order to achieve comparable power consumption to hybrid solutions, all-digital BS architectures must rely on low-resolution ADCs [6] and low-resolution digital baseband processing [7]. As it has been shown in [8], [9], the required resolution of ADCs

and baseband processing depends on the system conditions, such as the number of user equipments (UEs), the modulation scheme, and the channel propagation scenario. This implies that energy-efficient all-digital BS architectures should dynamically adapt the ADC and baseband-processing resolutions to the instantaneous communication scenario, as opposed to hard-wiring those parameters for the worst case during design time.

Contributions: We propose the first *resolution-adaptive* ASIC for all-digital mmWave massive MU-MIMO BSs. In order to maximize energy efficiency, our design is able to dynamically adapt the ADC and baseband-processing resolutions to the instantaneous communication scenario. The 8 mm² 65 nm CMOS ASIC supports 32 BS antennas and up to 16 UEs. The ASIC tightly integrates an array of time-interleaved (TI) successive-approximation register (SAR) ADCs, a beamspace channel estimation (BEACHES) engine, and a spatial equalizer. To operate at low resolution while achieving high spectral efficiency, the spatial equalizer builds upon finite-alphabet equalization [7], an emerging baseband-processing paradigm that is implemented using an energy-efficient standard-cell-based processing-in-memory (PIM) architecture. We provide measurement results to demonstrate that our resolution-adaptive ASIC outperforms existing equalizers in terms of energy and area efficiency while achieving highest-in-class throughput.

II. MMWAVE MASSIVE MU-MIMO

A. System Model

We consider the mmWave massive MU-MIMO uplink, as illustrated in Figure 1, where U single-antenna UEs transmit data to a B -antenna BS. The baseband channel is modeled using the frequency-flat input-output relation $\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{n}$, where $\mathbf{y} \in \mathbb{C}^B$ contains the signals received at the BS antennas, $\mathbf{H} \in \mathbb{C}^{B \times U}$ is the MIMO channel matrix, $\mathbf{s} \in \mathcal{S}^U$ contains the UE-transmit symbols taken from the constellation \mathcal{S} (e.g., 16-QAM), and $\mathbf{n} \in \mathbb{C}^B$ is complex Gaussian noise with variance N_0 per entry. The UE-transmit symbols s_u , $u = 1, \dots, U$, are zero mean with variance $E_s \sigma_u^2$, and we assume ± 3 dB UE-side power control so that $\max_u \{\sigma_u^2 \|\mathbf{h}_u\|_2^2\} / \min_u \{\sigma_u^2 \|\mathbf{h}_u\|_2^2\} = 4$ with \mathbf{h}_u being the u th column of \mathbf{H} [7]. The received vector \mathbf{y} is quantized by the ADCs resulting in \mathbf{z} , which is then used to estimate \mathbf{H} and generate estimates $\hat{\mathbf{s}}$ of the transmit vector \mathbf{s} .

This work 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 and ECCS-1824379. Contact author: O. Castañeda (e-mail: caoscar@ethz.ch)

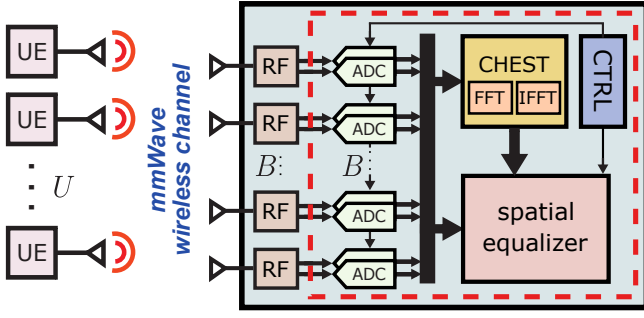


Fig. 1. Overview of the considered all-digital mmWave massive MU-MIMO system in which U UEs communicate with a B -antenna BS. Our ASIC includes the parts within the dashed red rectangle: an ADC array, a channel estimation (CHEST) engine, and a spatial equalizer. The resolutions of the ADCs and equalizer can be adapted to the communication scenario.

B. Beamspace Channel Estimation

The channel estimate $\hat{\mathbf{H}}$ is generated during a training phase. Each UE transmits a pilot symbol $\phi \in \mathcal{S}$ known to the BS, while the other UEs remain silent, resulting in $\mathbf{y} = \phi \mathbf{h} + \mathbf{n}$, where \mathbf{h} is the channel vector of the pilot-transmitting UE. The least squares (LS) estimate of the channel vector is then computed as $\hat{\mathbf{h}} = \mathbf{z}/\phi$. The LS estimate can then be transformed into the so-called *beamspace domain* by taking the discrete Fourier transform (DFT) of $\hat{\mathbf{h}}$. Since beamspace channel vectors are typically sparse in mmWave systems [1], we can use the BEACHES (short for beamspace channel estimation) algorithm [10] to denoise the channel vectors, resulting in significant performance improvements.

C. Finite-Alphabet Equalization

At the extremely high bandwidths offered by mmWave frequencies, even linear spatial equalizers of the form $\hat{\mathbf{s}} = \mathbf{W}^H \mathbf{z}$ result in power-hungry hardware [7]. To minimize power consumption without sacrificing spectral efficiency, finite-alphabet equalization [7] proposes to allow per-UE scaling factors (contained in a vector μ), so that the entries of the equalization matrix \mathbf{W}^H can be represented with low-resolution numbers in \mathbf{X}^H (e.g., 5 bits or lower). With an equalization matrix of the form $\mathbf{V}^H = \text{diag}(\mu) \mathbf{X}^H$, the main complexity of equalization now lies in computing $\mathbf{X}^H \mathbf{z}$, which can be implemented with low power digital circuitry [7], [11].

III. VLSI ARCHITECTURE

Figure 1 shows the top-level architecture of our $B = 32$ antenna all-digital receiver ASIC. The design contains an analog front-end with $2B = 64$ ADCs to quantize the I/Q receive vector \mathbf{y} into \mathbf{z} . The vector \mathbf{z} is used to estimate the symbols s_u transmitted by the $U \leq 16$ UEs with a specialized architecture of the parallel processor in associative content-addressable-memory (PPAC) [12]. The ADCs are programmable to support different resolutions in \mathbf{z} , while PPAC supports different resolutions in \mathbf{X}^H and in \mathbf{z} . The quantized vector \mathbf{z} can also be used to extract channel estimates using the BEACHES engine.

As detailed in Section III-B, the PPAC spatial equalizer operates bit-serially, i.e., it only processes one bit from all \mathbf{z} -entries per clock cycle. Since PPAC only operates one bit per

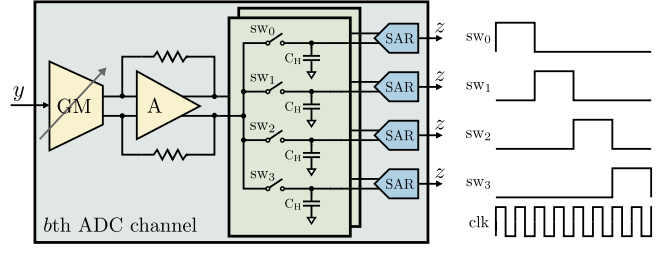


Fig. 2. Block diagram of an ADC channel and clocks for 6-bit samples.

\mathbf{z} -entry at a time, it can be tightly coupled with a SAR ADC that converts one bit per \mathbf{z} -entry at a time. As a single instance of PPAC is not sufficient to support the high bandwidths offered by mmWave systems, especially when implemented in a 65 nm technology node, we deployed four PPAC instances. The ADCs are TI by the same factor of four so that each PPAC instance is tightly coupled with its own ADC array. This co-design of ADC array and spatial equalizer enables one to easily scale our system to larger bandwidths with (i) more silicon area and/or (ii) a more advanced technology node, which would also significantly increase the per-instance throughput. Moreover, both the ADC array and PPAC can be scaled to more BS antennas or UEs. However, our current ASIC is limited to $B = 32$ and $U \leq 16$ due to silicon area and pin-I/O constraints.

A. ADC Array

The ADC array comprises $2B = 64$ channels (32 I/Q channels). As shown in Figure 2, each ADC channel consists of a two-stage programmable gain amplifier (PGA) and four TI-SAR ADCs. The two-stage PGA provides programmable gain, which can be set to amplifications from $1\times$ to $32\times$ in powers of two. The first stage of the PGA is a cascode amplifier with switchable transconductance cells to provide programmable gain; the second stage is a fixed-gain inverter-based amplifier. Sampling switches, controlled by non-overlapping sampling clocks, are used to TI the four fully-differential SAR ADCs.

The ADC array quantizes \mathbf{y} into \mathbf{z} using mid-rise quantization with either 3 bits or 6 bits of resolution: For the {3,6}-bit case, {1,2} clock cycles are used for sampling, requiring a total of {4,8} clock cycles per sample. During the sampling clock cycles, the associated PPAC instance is idle. Reducing the ADC resolution either enables higher sampling rates (and, hence, higher throughput) or lowers power consumption.

B. PPAC-Based Spatial Equalizer

Figure 3 illustrates our spatial equalizer, which is a version of the PPAC architecture [7] specialized for spatial equalization. PPAC is an all-digital, standard-cell-based PIM architecture for matrix-vector-product-like tasks, in which each memory bit-cell contains both storage (a latch) and logic (an XNOR). PPAC uses a so-called PPAC row to compute an inner-product between two vectors with one-bit entries. To support multi-bit operation, spatial replication is used for one of the vectors (a row of \mathbf{X}^H) and bit-serial operation for the other (\mathbf{z}). Our PPAC implementation supports 1 bit to 4 bits of resolution in \mathbf{X}^H ; the inputs to each PPAC row can be muted to save power

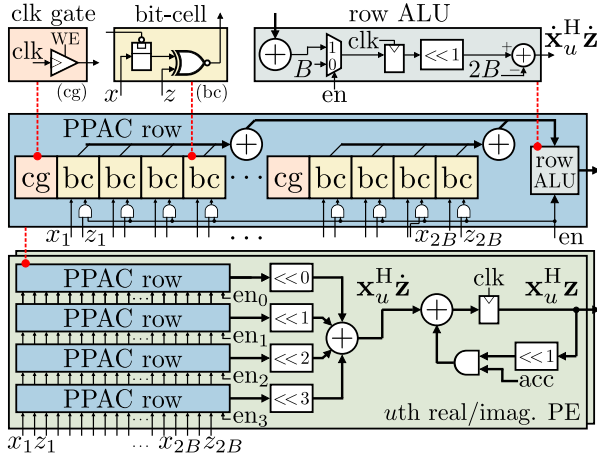


Fig. 3. Block diagram of the specialized PPAC for mmWave spatial equalization with mid-rise inputs. The architecture uses one PPAC row to compute an inner-product between two 1-bit vectors, $\hat{\mathbf{x}}_u$ and $\hat{\mathbf{z}}$. Multi-bit vectors \mathbf{x}_u and \mathbf{z} are processed using spatial replication and bit-serial operation, respectively.

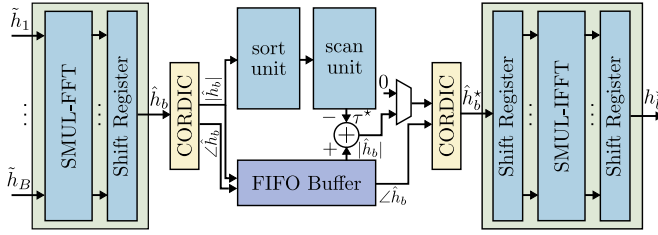


Fig. 4. Block diagram of the BEACHES engine. A streaming multiplierless (SMUL) FFT and IFFT are used for the beamspace transform. The sort and scan units determine the MSE-optimal denoising threshold τ^* , which is applied on the magnitudes (extracted with a CORDIC) of the beamspace channel vector.

when operating with fewer than 4 bits per \mathbf{X}^H -entry. It also supports up to 8 bits of resolution in \mathbf{z} : To operate a q -bit \mathbf{z} , a PPAC instance needs to run for q clock cycles.

The results in [11] have shown that the PIM-based PPAC can achieve lower power consumption than a conventional, non-PIM architecture when performing equalization with low-resolution matrices. In contrast to [11], our PPAC implementation operates with mid-rise quantizers, is programmable in the number of equalization matrix \mathbf{X}^H bits, and is tightly coupled to the ADC array. In addition, this is the first silicon realization of a PPAC-based hardware accelerator.

C. BEACHES Engine

BEACHES [10] is a mmWave channel denoising algorithm that performs soft-thresholding on the magnitudes of the estimated beamspace channel vector. The algorithm relies on Stein's unbiased risk estimator to apply a near mean-square error (MSE)-optimal denoising threshold. BEACHES uses a fast Fourier transform (FFT) to transform the channel estimate $\hat{\mathbf{h}}$ to beamspace $\hat{\mathbf{h}}_b$. The denoised vector $\hat{\mathbf{h}}^*$ is then transformed back to the antenna-domain with an inverse FFT (IFFT). Our BEACHES implementation is adopted from the architecture in [10] and illustrated in Figure 4. Unlike [10], we use a streaming multiplierless (SMUL) FFT architecture from [13] for both the FFT and IFFT. Our design is capable of denoising a new channel vector $\hat{\mathbf{h}}$ every $B = 32$ clock cycles.

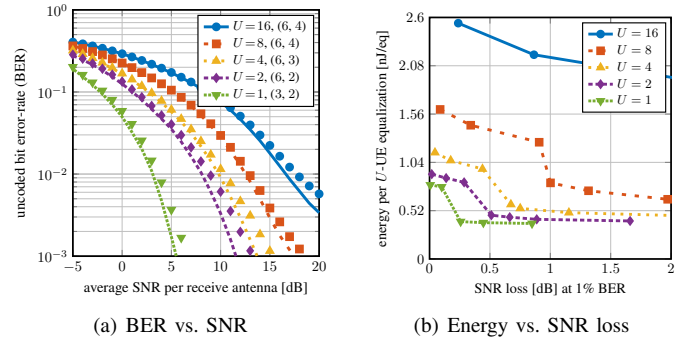


Fig. 5. Performance of the implemented 32-antenna ASIC over a LoS mmMAGIC UMi mmWave channel in terms of (a) bit-error rate (BER) with 16-QAM and (b) ADC array and spatial equalizer energy with QPSK. In (a), lines correspond to the performance of infinite-resolution ADCs and L-MMSE equalizer, markers to the ADC and equalizer resolution given in parentheses.

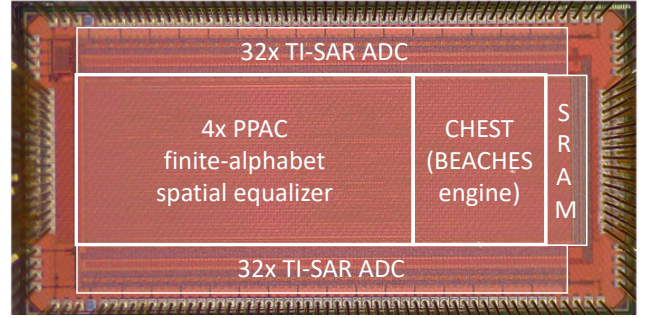


Fig. 6. Micrograph of the 8 mm² all-digital receiver ASIC in 65 nm CMOS.

IV. ASIC IMPLEMENTATION RESULTS

A. Bit Error-Rate (BER) Performance

Figure 5(a) shows the uncoded BER of our $B = 32$ receiver ASIC for different UE numbers U with line-of-sight (LoS) mmMAGIC UMi channels generated at a carrier frequency of 60 GHz using QuaDRiGA [14]. Our ASIC's performance (markers) achieves that of an infinite-resolution linear minimum MSE (L-MMSE) equalizer (lines) within 1 dB SNR loss measured at 1% BER. For common massive MU-MIMO load factors $\beta = U/B \leq 1/8$ ($U \leq 4$), where linear equalization achieves near-optimal BER performance, our ASIC can operate at reduced ADC and equalization resolutions.

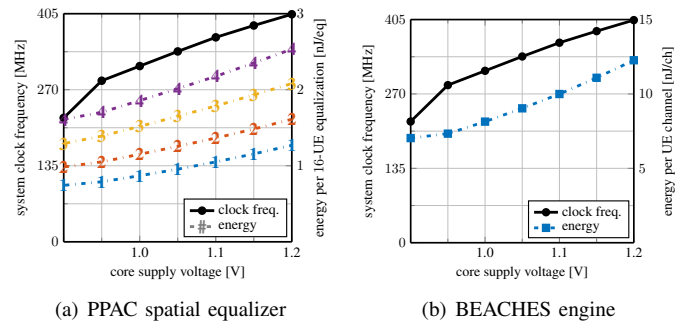


Fig. 7. System clock frequency and energy versus the supply voltage for the (a) PPAC spatial equalizer and (b) BEACHES engine. The ADC resolution is kept constant at 6 bits and the PPAC resolution is indicated by the markers.

The ADC and equalizer resolutions affect both the ASIC's error-rate performance and energy efficiency, as illustrated in Figure 5(b) for QPSK transmission and LoS channels: Markers represent the per-equalization energy and performance (in terms of SNR loss at 1% BER) for Pareto-optimal configurations of ADC and equalizer resolution; e.g., the left-most marker of each curve corresponds to the maximum-resolution configuration of 6 ADC bits and 4 equalizer bits. Figure 5(b) shows that for an SNR loss of 0.5 dB, lower load factors require lower resolution, which yields lower per-equalization energy, saving up to $2\times$ energy compared to the maximum-resolution configuration. In situations that allow for higher SNR losses, the resolutions and, hence, the energy, can be reduced even further.

B. ASIC Measurements and Comparison

Figure 6 shows the 8 mm^2 ASIC fabricated in TSMC 65 nm GP. The ADC array occupies 1.79 mm^2 , the PPAC-based spatial equalizer 2.41 mm^2 , and the BEACHES engine 0.94 mm^2 . The ASIC also includes an SRAM to perform high-speed tests. At nominal 1.0 V supply and a temperature of 300 K, the ASIC achieves a maximum clock frequency of 312 MHz with the critical path in the spatial equalizer. This clock frequency corresponds to 156 MS/s and 312 MS/s with 6-bit and 3-bit samples, respectively, as well as to a throughput of 9.98 Gb/s for 16 UEs transmitting 16-QAM with 6-bit ADC samples. Under these conditions, reducing the equalizer resolution by 1 bit saves an average of 51 mW in the spatial equalizer, which corresponds to 18% of the equalizer's power when operating at its maximum resolution of 4 bits. The ADC array consumes 106 mW with the PGAs configured to unit gain. Furthermore, the BEACHES engine is able to process 9.75 M channel vectors per second at 79 mW. Figure 7 shows the equalizer and BEACHES energy when performing voltage-frequency scaling.

Table I compares our ASIC to existing massive MU-MIMO equalizers. Our implementation is the only one including ADCs and supporting programmable resolution in both the ADC array and equalizer. After technology normalization, our design achieves $6\times$ higher throughput, $2\times$ higher area efficiency, and $2\times$ lower energy per bit when operating at maximum resolution and taking into account the ADC array area and power consumption. Thanks to the resolution-adaptivity of our ASIC, we can further reduce its energy or increase its throughput, up to doubling the throughput and tripling the energy efficiency when operating at the lowest-resolution configuration of 3 ADC bits and 1 equalization bit.

V. CONCLUSIONS

We have shown the first resolution-adaptive all-digital spatial equalizer ASIC for mmWave massive MU-MIMO. Our design includes an ADC array to support 32 I/Q baseband channels, a processing-in-memory-based spatial equalizer, and a channel estimation engine. Our ASIC measurements demonstrate that resolution-adaptivity enables highest-in-class throughput while being more than $2\times$ energy and area efficient compared to state-of-the-art equalizer designs, even with the ADC array power consumption and area included.

TABLE I
COMPARISON WITH STATE-OF-THE-ART MASSIVE MIMO EQUALIZERS

	This work			Tang [15]	Tang [16]	Jeon [17]	Wen [18]	Liu [19]
Max. BS antennas B	32	32	32	128	128	256	256	128
Max. UEs	16	16	16	32	16	32	32	8
Max. mod. [QAM]	16	16	16	256	256	256	256	64
ADC & CHEST integrated	yes	yes	yes	no	no	no	no	no
Resolution-adaptive	yes	yes	yes	no	no	no	no	no
ADC, EQ resolution	6.4	6.4	3.1	N/A	N/A	N/A	N/A	N/A
Includes ADC area & power	yes	no	yes	no	no	no	no	no
Technology [nm]	65	65	65	40	28	28	40	65 ^a
Core supply [V]	1.0	1.0	1.0	0.9	1.0	0.9	1.1	1.2
Core area [mm ²]	4.20	2.41	4.20	0.58	2.0	0.37	0.73	1.62
Max. frequency [MHz]	312	312	312	425	569	400	290	500
Power [mW]	396	290	246	221	127	151	87	120
Throughput [Gb/s]	9.98	9.98	20.0	2.76	1.80	0.354	1.96	1.5
Throughput ^b [Gb/s]	9.98	9.98	20.0	1.70	0.775	0.153	1.21	1.5
Area eff. ^b [Gb/s/mm ²]	2.38	4.15	4.76	1.11	0.072	0.076	0.626	0.926
Energy ^b [pJ/b]	39.7	29.0	12.3	260	380	2835	96.9	80.0

^a65nm LP, ^btechnology normalized to 65nm at nominal core supply

REFERENCES

- [1] T. S. Rappaport *et al.*, *Millimeter Wave Wireless Communications*. Prentice Hall, 2015.
- [2] E. G. Larsson *et al.*, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [3] A. Alkhateeb *et al.*, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 831–846, Oct. 2014.
- [4] K. Roth *et al.*, "Achievable rate and energy efficiency of hybrid and digital beamforming receivers with low resolution ADC," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 9, pp. 2056–2068, Sep. 2017.
- [5] P. Skrimponis *et al.*, "Power consumption analysis for mobile mmWave and sub-THz receivers," in *IEEE 6G Wireless Summit*, Mar. 2020.
- [6] S. Jacobsson *et al.*, "Throughput analysis of massive MIMO uplink with low-resolution ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 4038–4051, Jun. 2017.
- [7] O. Castañeda *et al.*, "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.
- [8] H. Yan *et al.*, "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.
- [9] M. Abdelghany *et al.*, "Towards all-digital mmWave massive MIMO: Designing around nonlinearities," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Feb. 2018, pp. 1552–1557.
- [10] S. H. Mirfarshbafan *et al.*, "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.
- [11] O. Castañeda *et al.*, "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.
- [12] —, "PPAC: A versatile in-memory accelerator for matrix-vector-product-like operations," in *IEEE Int. Conf. ASAP*, Jul. 2019, pp. 149–156.
- [13] S. H. Mirfarshbafan *et al.*, "SMUL-FFT: A streaming multiplierless fast Fourier transform," *IEEE Trans. Circuits Syst. II*, Mar. 2021.
- [14] S. Jaeckel *et al.*, "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.
- [15] W. Tang *et al.*, "A 0.58 mm^2 2.76 Gb/s 79.8pJ/b 256-QAM massive MIMO message-passing detector," in *IEEE Symp. VLSI C.*, Jun. 2016.
- [16] W. Tang *et al.*, "A 1.8Gb/s 70.6pJ/b 128×16 link-adaptive near-optimal massive MIMO detector in 28nm UTBB-FDSOI," in *IEEE ISSCC*, Feb. 2018, pp. 224–226.
- [17] C. Jeon *et al.*, "A 354 Mb/s 0.37 mm² 151 mW 32-user 256-QAM near-MAP soft-input soft-output massive MU-MIMO data detector in 28nm CMOS," *IEEE SSC. Lett.*, vol. 2, no. 9, pp. 127–130, Sep. 2019.
- [18] C.-C. Wen *et al.*, "A 1.96Gb/s massive MU-MIMO detector for next-generation cellular systems," in *IEEE Symp. VLSI C.*, Jun. 2020.
- [19] L. Liu *et al.*, "Energy- and area-efficient recursive-conjugate-gradient-based MMSE detector for massive MIMO systems," *IEEE Trans. Signal Process.*, vol. 68, pp. 573–588, Jan. 2020.