



Power Efficient Multi-Carrier Baseband Processing for 5G and 6G Wireless

Conference Paper**Author(s):**

Skrimponis, Panagiotis; Mirfarshbafan, Seyedhadi ; Studer, Christoph ; Rangan, Sundeep

Publication date:

2020-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000529081>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/IEEECONF51394.2020.9443544>

Power Efficient Multi-Carrier Baseband Processing for 5G and 6G Wireless

Panagiotis Skrimponis[†] Seyed Hadi Mirfarshbafan^b Christoph Studer^b Sundeep Rangan[†]

[†]NYU Tandon School of Engineering, Brooklyn, NY, USA

^bECE, Cornell Tech, New York, NY, USA

Abstract—Power consumption is one of the significant challenges in millimeter wave (mmWave) systems due to the need to support wide bandwidths and large numbers of antennas. This paper explores energy efficient implementations of the baseband trans-receiver components for a multi-carrier 3GPP New Radio (NR) system. The analysis considers key components including channel selection filters, digital beamforming and FFT engines for the OFDM processing. A methodology is presented for optimizing bit widths in various components, which is critical in low power designs. Fully digital and analog beamforming architectures are also compared. Preliminary power estimates are provided using a TSMC 28 nm process for a 400 MHz system at 28 GHz similar to 5G systems today and a hypothetical 1.6 GHz system at 140 GHz for potential 6G deployment.

I. INTRODUCTION

The vast amount of available bandwidth in millimeter wave (mmWave) bands offer the potential for low-latency communication and massive data rates [1]–[3]. With the deployment of the fifth generation (5G) communication system in the lower mmWave bands, the frequencies above 100 GHz have attracted significant interest for potential 6G use cases [4]–[7].

A key technical challenge for commercial devices in these bands is power consumption, which is particularly critical for the handheld and mobile devices. Due to the high isotropic path loss, mmWave communication systems rely on narrow, electrically steerable beams. To enable the beam steering, front-ends must support large numbers of antennas at very wide bandwidths [1], [8]. These front-end require large number of radio frequency (RF) chains as well as baseband processing at high sample rates, both of which increase the power significantly. In addition, RF circuits above 100 GHz are still in their infancy [9]–[12], and energy efficiency remains an even greater challenge than front-end devices below 100 GHz.

While there have been a large number of published results on the power consumption of individual components of mmWave front-ends, estimating and optimizing the overall system power is more challenging as there are trade-off and interactions between components. Analytic models for the

energy and performance trade-offs in mmWave MIMO front-ends can be found in [13], [14]. A key focus is on ADC resolution, as this has been suspected of being the largest power driver, particularly when supporting large numbers of streams at high data rate. The works [15], [16] attempted to obtain power estimates for 28 GHz and 37 GHz for analog, hybrid and fully digital front-ends using state-of-the-art published numbers for different components in 5G use cases. A similar methodology was followed in [17] that extended the analysis to 140 GHz using recent devices [9]–[12].

A limitation of works such as [15]–[17] is that the components were not optimized for the overall end-to-end power-performance trade-off. To overcome this issue, a recent work [18] developed a general methodology for performing this end-to-end power optimization given options for each component. The methodology was applied to a 140 GHz fully-digital receiver, with various device options were developed in each of the the critical RFFE components (e.g., LNA, Mixer, LO). With detailed circuit and system-level simulations the authors achieved a remarkable 70 – 80% reduction in power consumption compared to the baseline 140 GHz design in [17].

The broad scope of this paper is to extend the analysis in [17], [18] by providing initial power consumption estimates for the most important elements in the receiver baseband processing. We focus on a multi-carrier 3GPP New Radio (NR) like waveform [19]. The NR waveform is extremely flexible, and is well-designed for mobile devices, low latency and support for multiple users. However, baseband power consumption is a potential issue, particularly due to the OFDM processing and need to support channelization. Since the power consumption of any digital component depends mostly on the operating frequency and the bit resolution our analysis considers bit-width optimizations for the most critical baseband components.

Another focus of this paper is to compare the beamforming architectures for the receiver. Although most mobile devices have used phased arrays for RF beamforming in 5G, fully digital architectures can offer the fast beam search in mobile environments [20], which is valuable for fast recovery from blocking [21]–[23], and aggressive use of idle and DRX modes [24], [25]. Fully digital architectures observe all directions simultaneously [8], [26], [27] to reduce the initial access time by an order of magnitude [28], [29]. However, power consumption is a potential issue for fully digital architectures

The work of all authors was supported in part by the ComSenTer, a JUMP program sponsored by the Semiconductor Research Corporation. Additionally, the work of P. Skrimponis and S. Rangan was supported in part by NSF grants 1925079, 1302336, 1564142, and 1547332, and the industrial affiliates of NYU Wireless. The work of S. H. Mirfarshbafan and C. Studer was supported in part by Xilinx, Inc., and NSF grants ECCS-1408006, CCF-1535897, CCF-1652065, and CNS-1717559.

due to the need to support multiple ADCs and mixers. A key goal of the paper is to understand the tradeoffs between analog and fully digital front-ends.

II. SYSTEM MODEL

A. System Assumptions

The parameters for the analysis at 28 and 140 GHz systems are shown in Table I. For the 28 GHz systems, we assume a UE array with $N_{\text{rx}} = 8$ antennas, which is consistent with capacity simulations in [30] as well as recent UE designs in [23]. Even though arrays may have multiple arrays for 360 degree coverage, we assume that only one is on at a time [23]. For the 140 GHz system, we assume $N_{\text{rx}} = 16$ antennas. Due to the smaller wavelength, a 16 element array (e.g., 4×4 UPA) at 140 GHz can occupy a smaller total area compared to an 8 element array (e.g., 2×4 UPA) at 28 GHz.

We consider a 3GPP NR type system [19] with carrier aggregation [31]. NR standard is very flexible and provides a very good baseline for 6G evaluation. Carrier aggregation enables partition of the processing for each one of the component carriers (CCs) independently based on the capabilities of the UE. For the 28 GHz system, we let $N_{\text{CC}} = 4$ denote the number of CCs, with configuration as described in [32] and parameters that are commonly used in 5G systems today. Each CC uses OFDM processing with an FFT of size $N_{\text{FFT}} = 1024$, and sub-carrier spacing of 120 KHz. The CCs are spaced at 100 MHz with an occupied bandwidth of $B_{\text{CC}} = 95.04$ MHz, based on 66 resource blocks [32]. For the 140 GHz system, we assume $N_{\text{CC}} = 8$ CCs spaced at 200 MHz with an occupied bandwidth of $B_{\text{CC}} = 190.08$ MHz each. Assuming the same number of resource blocks while increasing the sub-carrier spacing to the maximum allowed in 3GPP NR, 240 KHz. The total occupied signal bandwidth is 1.6 GHz, four times the bandwidth of the 28 GHz system. For both systems, the sample rate is derived from the FFT size, number of component carriers, and sub-carrier spacing.

For the data plane, we consider that a mobile device needs to support $N_{\text{str}} = 2$ digital streams in each CC. Larger numbers of streams are likely not needed since the channels will lack spatial diversity. Although, polarization diversity is still possible.

B. Receiver Architectures

In this work, we analyze the power for both fully-digital and analog receivers as shown in Fig. 1. In the **fully-digital receiver (top panel)**, each RX antenna has an independent low noise amplifier (LNA), mixer and a pair of analog to digital converters (ADC). The signal from the ADCs is going through a CC selection digital circuit, which consists of a numerically controlled oscillator (NCO), a low pass filter and a $1/N_{\text{CC}}$ downsampler. The filters are running at a sample rate of $f_{\text{CC}} = f_s/N_{\text{CC}}$. The output of the CC selection circuit from all N_{rx} antenna for each CC, is fed into an N_{rx} -input N_{str} -output linear digital beamformer. Each of the N_{str} streams is then fed to an FFT-engine for the OFDM processing. As a result,

TABLE I
SYSTEM PARAMETERS USED IN THE ANALYSIS.

Parameter	Value		Remarks
	28 GHz	140 GHz	
Carrier frequency, f_c	28 GHz	140 GHz	
Number RX antennas, N_{rx}	8	16	Assume a single array, typically UPA.
Subcarrier spacing (SCS), [kHz]	120	240	
Number component carriers, N_{CC}	4	8	
Bandwidth per CC, B_{CC} [MHz]	95.04	190.08	Based on 66 occupied RBs per CC [32]
FFT size per CC, N_{FFT}	1024	1024	
Total bandwidth, $B_{\text{CC}}N_{\text{CC}}$ [GHz]	0.380	1.52	
Sample rate, f_s [GHz]	0.491	1.966	Based on FFT size, SCS and N_{CC}
OFDM symbol duration, T_{sym} [μ s]	8.92	4.46	Derived from SCS
Number digital streams per CC, N_{str}	2	2	More streams not needed due to lack of spatial diversity

the digital beamformer eliminates the need to run one FFT on each antenna – instead it is run on each stream.

For the **analog receiver architecture (lower panel)**, each RX antenna is first passed through an LNA and then one RF phase shifter is used for each stream to perform the analog beamforming. This produces a total of N_{rx} RF signals for each stream. The streams are then combined and a mixer is used to downconvert to baseband where there is one ADC pair for each stream. In each stream, there are N_{CC} CC selection filters to extract the component carriers. There is one OFDM FFT for each stream in each CC.

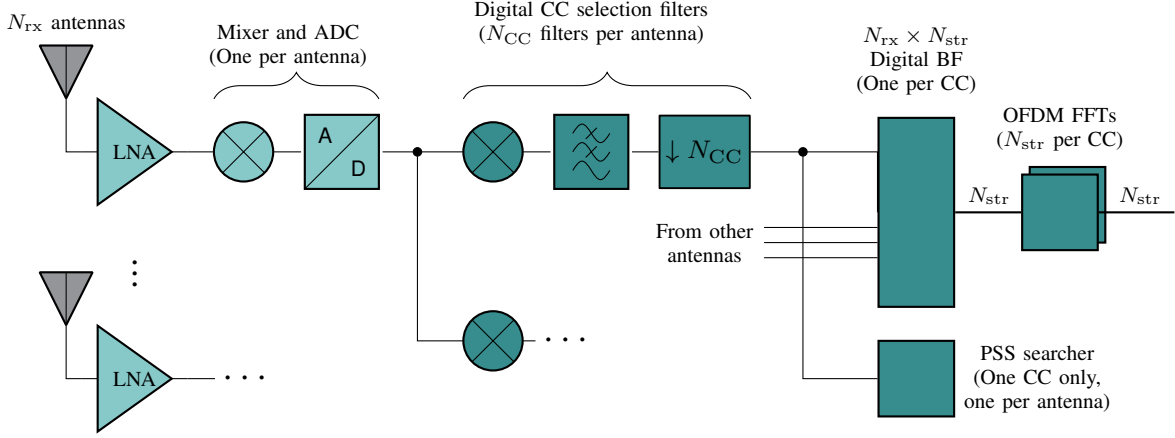
For the synchronization in 3GPP NR, the receiver uses the narrow-band primary synchronization signal (PSS) and secondary synchronization signal (SSS), which fit in a single CC [29]. The PSS searcher runs at low SNR and can therefore benefit from very low bit widths. In fact, simulations in [28] indicate that two bits are sufficient. In both architectures the PSS searcher runs for only one CC at a time. While in the analog receiver architecture the PSS searcher runs for each digital stream, in the fully-digital the PSS searcher runs for each antenna. This enables the fully-digital systems to look in multiple directions.

III. COMPONENT AND POWER CONSUMPTION MODELS

A. Component Carrier Downsampling Filter

Carrier aggregation with multiple component carriers (CCs) requires special signal processing techniques to extract the samples for the CC. While some systems prefer to use frequency-domain channelizers [33], in mmWave devices this can be impractical since the power consumption of the FFT can be very high. In this work, we assume a standard time-domain CC channel selection circuit: The samples from the ADC are first passed through a numerically controlled oscillator (NCO) that shift the CC to be centered around DC. Then, a decimator comprised of a digital finite impulse response (FIR)

(a) Fully digital RX



(b) Analog RX

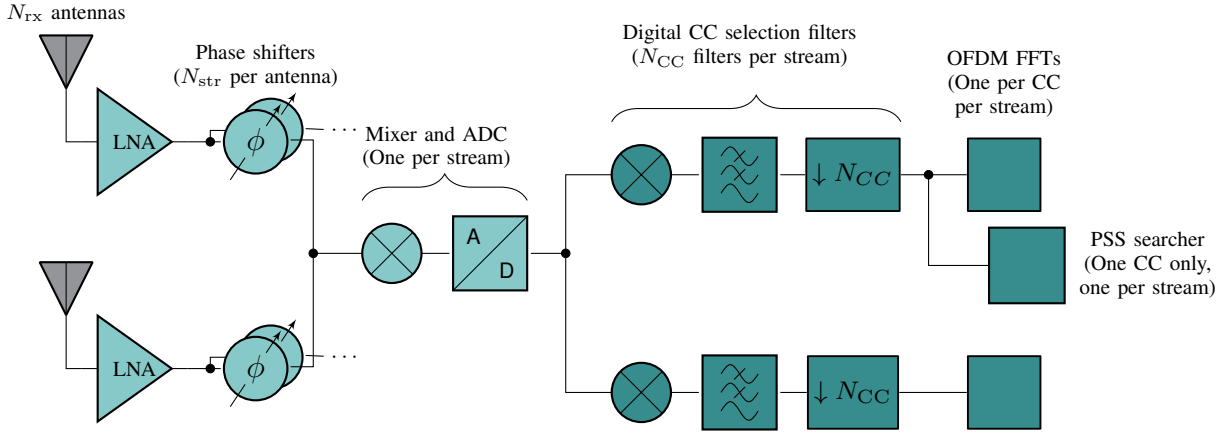


Fig. 1. High level architecture for a fully digital receiver (top panel) and analog receiver (bottom panel). Each architecture supports N_{rx} antennas, N_{CC} component carriers and N_{str} digital streams per component carrier. The light green boxes represent the analog components, and the dark green boxes represent the digital components. In the RF front-end, some components such as filters are not shown.

low pass filter (LPF) and a $1/N_{CC}$ downsampler, remove out-band noise from adjacent CCs and reduce the sample rate to $f_{cc} = f_s/N_{CC}$ per CC.

To design an efficient FIR filter we take into consideration that the power consumption of any digital component depend strongly on the bit-widths in the processing. Using the MATLAB DSP toolbox we design a FIR filter with 6-bit fixed-point real coefficients, maximum ripple in the pass-band up to $F_p = B_{CC}/2$ and maximum stop-band rejection for frequencies beyond $F_{st} = f_s/N_{CC}/2$. The frequency response of the optimized filter for the 28 GHz system is shown in Fig. 2, observing a stop-band rejection of $A_{st} = 29.19$ dB.

This stop-band rejection is sufficient for most cellular applications [34]. Since, with the development of the low-density parity check (LDPC) codes in 3GPP NR, most modulation and coding schemes up to 64-QAM can be decoded reliably with 25 dB [34]. Suppose the power spectral density (PSD) in the

desired CC is P . If the PSD in the adjacent carriers is also P , then the total aliased PSD from the $N_{CC} - 1$ adjacent carriers will be at most $(N_{CC} - 1)P/A_{st}$. Hence, the maximum SNR per antenna will be,

$$\frac{P}{P(N_{CC} - 1)/A_{st}} = \frac{A_{st}}{N_{CC} - 1}.$$

After beamforming, the receiver can achieve an SNR of $N_{rx}A_{st}/(N_{CC} - 1)$. For the 28 GHz system, $A_{st} = 29.19$ dB, $N_{CC} = 4$, $N_{rx} = 8$, we can obtain more than 33 dB maximum SNR post beamforming. Similar acceptable performance was found at 140 GHz with the same bit-widths and filter length.

To estimate the power of the filter, we observe that there are $2L$ multiply-and-accumulate (MAC) operations in each filter output (the factor of 2 due to I and Q). Each filter outputs at a rate of f_s/N_{CC} . In the digital design, there are $N_{rx}N_{CC}$ filters. Thus, the power consumption for the CC filters in the




Fig. 2. Frequency response of a possible CC downsampling filter for the 28 GHz system. The filter is implemented with 6-bit coefficients and 65 taps. The passband is $B_{CC}/2$ shown in the black dashed line and the stop-band begins at $f_s/N_{CC}/2$ shown in the the red dashed line.

fully digital case is

$$P_{LPF} = 2 E_{MAC} L N_{rx} f_s, \quad (1)$$

where E_{MAC} is the energy per real MAC. Similarly, for the analog beamforming case, we need one filter per digital stream and CC, and we obtain the power,

$$P_{LPF} = 2 E_{MAC} L N_{str} f_s. \quad (2)$$

In order to estimate the power consumption of the FIR filter, we use the Spiral code generation tool [35] to generate hardware description language (HDL) for a multiplier-less FIR filter. Then, by reducing the bit-widths of the accumulators and buffers we optimize the generated HDL code. The model of developed filter is then used in the next section for bit-accurate system-level simulations. Initial power estimations with the TSMC 28 nm process show that the fully-digital design needs $E_{MAC} \approx 100$ fJ per MAC at 28 GHz, and $E_{MAC} \approx 50$ fJ per MAC at 140 GHz, and the analog design needs $E_{MAC} \approx 178$ fJ per MAC at 28 GHz, and $E_{MAC} \approx 88$ fJ per MAC at 140 GHz. The analog design has a larger number of input bits from the ADC, and the 140 GHz systems requires less power since the SNR per antenna is lower. The total power consumption for the CC selection circuit is shown in Table II.

B. Digital beamformer

The fully-digital system requires the use of a digital beamformer to multiply the inputs from N_{rx} antennas into N_{str} digital streams. In each sample, it must perform $N_{rx}N_{str}$ complex MAC operations. Given that the sample rate for each CC is f_s , the total power consumption for all CCs is,

$$P_{BF} = E_{MAC} N_{rx} N_{str} f_s, \quad (3)$$

where E_{MAC} is the energy per complex MAC operation. Initial power estimations with the TSMC 28 nm process show that

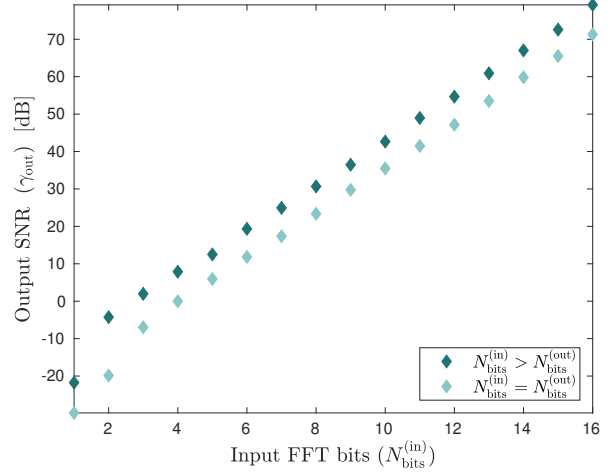


Fig. 3. Maximum output SNR (γ_{out}) for a fixed-point FFT implementation, over the input number of bits. Comparing two possible configurations of bit growth within the architecture.

$E_{MAC} \approx 1$ pJ per complex MAC was possible with 6-bit real inputs from the CC filters. The resulting power is shown in Table II.

C. PSS

The PSS detection is essentially a matched filter for the PSS sequences from the base station. Similarly to results in [28] we found that the PSS searcher can be implemented with very few bits. The PSS search algorithm exploits the cross-correlation properties of the m-sequences that are currently used in the 3GPP NR.

To reduce the power consumption of this module we quantize the input signals from the CC filters to 2-bit. Then, the Spiral code generation tool [35] is used to generate a multiplier-less correlation circuit. The original generated HDL code does not allow for variable length across the coefficients and accumulators. As a result, this will not produce a very efficient implementation in terms of power consumption. Using bit-width optimization for the coefficients, accumulators, and memories we develop a correlation circuit with 3-bit signed coefficients, and 5–8 bits for the accumulators.

To estimate the power of the PSS searcher, we observe that there are $L = 128$ complex MAC operations in each time-domain correlation. Each PSS searcher operates at a rate of f_s/N_{CC} frequency. In the digital design, there are N_{rx} PSS searchers. Assuming that the PSS searcher is running for only one CC at a time, the power consumption for the fully digital case is,

$$P_{PSS} = E_{MAC} L N_{rx} \frac{f_s}{N_{CC}}, \quad (4)$$

where E_{MAC} is the energy per complex MAC. Similarly, for the analog beamforming case, we need one PSS search per digital stream, obtaining the power,

$$P_{PSS} = E_{MAC} L N_{str} \frac{f_s}{N_{CC}}, \quad (5)$$

TABLE II
POWER CONSUMPTION ESTIMATES (IN mW) FOR THE FULLY DIGITAL AND ANALOG RECEIVER ARCHITECTURES WITH PARAMETERS IN TABLE I.

Component	Fully Digital			Analog		Remarks
	28 GHz	140 GHz		28 GHz	140 GHz	
		Design ⁽¹⁾	Design ⁽²⁾			
LNA	9.05	254.4	76.8	180	1628.8	N_{rx} LNAs required for both designs, but the power consumption is increased in the analog case since each LNA drives N_{str} phase shifters at $IL = 10$ dB.
Mixer	-	80	80	-	-	For the 28 GHz system we consider the mixer as passive devices. While for the 140 GHz fully-digital systems we use an active mixer design described in [18].
LO	80	63.57	76.8	20	196	The 28 GHz systems use $P_{LO} = 10$ dBm, while Design ⁽¹⁾ uses $P_{LO} = -11$ dBm, Design ⁽²⁾ uses $P_{LO} = -6$ dBm, and the 140 GHz analog system uses the $P_{LO} = 19.9$ dBm from [10].
ADC	8.18	65.43	130.86	16.4	65.4	N_{rx} 4-bit ADC pairs for fully-digital 28 GHz and Design ⁽¹⁾ , 5-bit for Design ⁽²⁾ , and one 8-bit ADC pair for the analog.
CC LPF	51.12	205	205	22.72	45.44	$N_{rx}N_{CC}$ filters for digital and $N_{str}N_{CC}$ filters for analog.
Digital BF	6.4	51.2	51.2	0	0	N_{CC} units each performing $N_{rx} \times N_{str}$ beamforming. Not used in analog case.
PSS Search	24.8	99	99	6.19	12.4	$N_{rx}N_{CC}$ time-domain correlations for digital and $N_{str}N_{CC}$ for analog.
OFDM FFT	114.68	458.75	458.75	114.68	458.75	Both analog and fully digital designs require $N_{str}N_{CC}$ FFTs in each OFDM symbol period.
Total	295	1278	1179	360	2407	

Initial power estimations with the TSMC 28 nm process show that the designs require $E_{MAC} \approx 197$ fJ per complex MAC operation. The total power consumption of the PSS searcher for each system is shown in Table II.

D. FFT

With the development of OFDM, the discrete Fourier transform (DFT) became essential part of any modern communication system. An efficient implementation of the DFT is the fast Fourier transform (FFT) algorithm. Since this is an essential part of the NR standard, the power and performance is critical for the receiver. A common technique to reduce the resolution is to use fixed-point arithmetic. Based on the number of FFT points $N_{FFT} = 1024$, we develop an implementation based on the Radix-4 butterfly architecture. To explore the the performance and power consumption trade-offs of FFT implementations we define a set of design options: (i) Growth Profile: this parameter controls the bit growth in each stage; (ii) Scaling Schedule: this parameter controls the scaling factor in each stage; (iii) Fractional Length: the number of binary point bits out of the total number of bits.

To measure the power consumption, we use the Spiral code generation tool [35], to generate a basic HDL fixed-point FFT implementation. In this implementation enables us to change the HDL accordingly and measure the effect of the design options. To measure the performance, we developed a fixed-point FFT MATLAB model that accurately depicts the hardware implementation. In Fig. 3 we show that effect of the bit-growth in the system performance. For a given input FFT bit resolution $N_{bits}^{(in)}$ we find the best configuration for the fractional length, and the optimal input distribution that gives the best performance. To find the optimal distribution we

generate i.i.d. random Gaussian samples with zero mean and variance σ^2 . We take both the fixed-point and floating point FFT using the MATLAB built-in function, and we measure the maximum performance by taking the linear estimate. By altering the variance we can find the optimal input distribution to the fixed-point FFT. For each fractional-bit length and total-bit length combination.

Based on our analysis, both analog and fully-digital systems require $N_{str}N_{CC}$ FFTs in each OFDM symbol period. The power consumption of the FFT can be calculated as follows,

$$P_{FFT} = E_{FFT} N_{str} f_s, \quad (6)$$

where E_{FFT} is the total energy per transform. Initial power consumption estimates indicate that for a 12-bit FFT in TSMC 28 nm, the $E_{FFT} \approx 128$ nJ per transform. This number is based on a Radix-4 implementation, with a growth profile equal to 1 (the data remains the same across the stages), and scaling schedule 2 (the data are being scale by 2 after each stage). The power consumption of the FFT can be quite high, compared to the other digital components. This FFT implementation contains a lot of memories and permutation buffers that dominate the total power consumption. Hence, reducing the number of input bits might have a bigger effect in performance than the power consumption. In the future other implementation of FFT will be explored.

IV. SIMULATION AND POWER CONSUMPTION ESTIMATES

To understand better the trade-off between power consumption and low-precision baseband processing we developed a simulation framework. The code for the simulations is provided as an open-source MATLAB package [36]. Currently, the repository includes models for low-resolution baseband

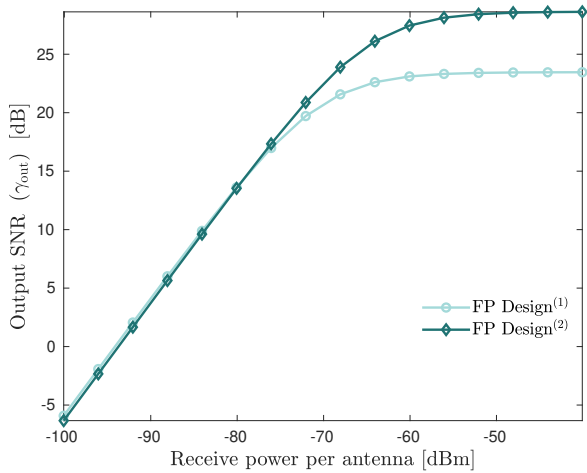


Fig. 4. Evaluation of Design⁽¹⁾ and Design⁽²⁾ in [18] using fixed-point 3GPP NR-like multi-carrier processing, 16 antenna elements at 140 GHz.

processing as described in Section III, and models for RFFE components at 140 GHz as described in [18]. Future releases will include models for RF components designed in other mmWave and THz frequencies and antenna element models. Following the documentation provided in the repository, a user can replicate the simulation results.

A. 3GPP NR-like Multi-Carrier Link-Layer Simulation

Similar to [17], [18] we consider a downlink system with a single NR base station (gNB), and a single UE device. For the UE we use the parameters shown in Table I. The gNB is equipped with an antenna array with 64 elements (e.g., 8×8 UPA), ideal RFFE, and baseband processing. The gNB can either use the entire wide-band bandwidth by performing carrier aggregation, or transmit a single component carrier. We consider a single path channel between the gNB and the UE with random gain and phase. For the 140 GHz simulation the RFFE of the UE contains the optimized devices described in [18]. More specifically uses the parameters of Design⁽¹⁾ and Design⁽²⁾.

For each slot and for each CC, the gNB generates a physical downlink shared channel (PDSCH) that contains the information and some control signals. In addition, a random angle-of-departure (AoD) and angle-of-arrival (AoA) in both azimuth and elevation angles is generated for each slot. The receiver uses the demodulation reference signals (DM-RS) and phase tracking reference signals (PT-RS) signals in PDSCH for practical channel estimation. The PT-RS introduced by the 3GPP NR aims to compensate the common phase error (CPE). The CPE is commonly generated by the phase noise of the LO, which can severely deteriorate the SNR. To perform coherent estimation of the CPE, the receiver is using the algorithm provided in [37]. The receiver uses the PSS and SSS for the directional search.

The UE process the data for each CC independently with non-linear RFFE and low-precision baseband processing.

Then, we measure the output SNR, on the received sub-carriers. In Fig. 4, we report the average output SNR for all CCs.

V. CONCLUSIONS

Power consumption is one of the most significant technical challenges in realizing mmWave and THz communication systems. In this paper, we develop a general framework for optimizing power consumption of a multi-antenna front-end including the power consumption in RF components, ADC and baseband processing. For the baseband processing, we assumed a multi-carrier NR waveform, a natural baseline for 6G systems. Our main result shows that after optimizations across the components, the FFTs for the OFDM processing represent a significant fraction of the overall power. For example, in a 140 GHz design with 2 GHz sampling, the OFDM FFTs consumed approximately 40% of the overall power – more than the ADCs, LO and LNAs combined. These results suggest that either new, more efficient implementations of the FFTs or alternate non-OFDM waveforms may be a valuable focus of future research.

REFERENCES

- [1] S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter-wave cellular wireless networks: Potentials and challenges,” *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, 2014.
- [2] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejewski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, and C. Gutterman, “Challenge: Cosmos: A city-scale programmable testbed for experimentation with advanced wireless,” in *Proc. 26th Annual International Conference on Mobile Computing and Networking*. ACM, 2020.
- [3] T. S. Rappaport, R. W. Heath Jr, R. C. Daniels, and J. N. Murdock, *Millimeter wave wireless communications*. Pearson Education, 2014.
- [4] T. S. Rappaport, Y. Xing, O. Kanhere, S. Ju, A. Madanayake, S. Mandal, A. Alkhatieb, and G. C. Trichopoulos, “Wireless communications and applications above 100 GHz: Opportunities and challenges for 6G and beyond,” *IEEE Access*, vol. 7, pp. 78 729–78 757, 2019.
- [5] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, “Toward 6G networks: Use cases and technologies,” *IEEE Communications Magazine*, vol. 58, no. 3, pp. 55–61, 2020.
- [6] H. Elayan, O. Amin, R. M. Shubair, and M.-S. Alouini, “Terahertz communication: The opportunities of wireless technology beyond 5G,” in *2018 International Conference on Advanced Communication Technologies and Networking (CommNet)*. IEEE, 2018, pp. 1–5.
- [7] I. F. Akyildiz, J. M. Jornet, and C. Han, “Terahertz band: Next frontier for wireless communications,” *Physical Communication*, vol. 12, pp. 16–32, 2014.
- [8] R. W. Heath, N. Gonzalez-Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, “An overview of signal processing techniques for millimeter wave MIMO systems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 3, pp. 436–453, 2016.
- [9] R. Wang, Y. Sun, M. Kaynak, S. Beer, J. Borngräber, and J. C. Scheytt, “A micromachined double-dipole antenna for 122–140 GHz applications based on a SiGe BiCMOS technology,” in *2012 IEEE/MTT-S International Microwave Symposium Digest*. IEEE, 2012, pp. 1–3.
- [10] A. Simsek, S.-K. Kim, and M. Rodwell, “A 140 GHz MIMO transceiver in 45 nm SOI CMOS,” 10 2018, pp. 231–234.
- [11] M. Urteaga, D. Scott, S. Krishnan, Y. Wei, M. Dahlstrom, Z. Griffith, N. Parthasarathy, and M. J. Rodwell, “G-band (140-220-GHz) InP-based HBT amplifiers,” *IEEE Journal of Solid-State Circuits*, vol. 38, no. 9, pp. 1451–1456, 2003.
- [12] K. K. Tokgoz, S. Maki, J. Pang, N. Nagashima, I. Abdo, S. Kawai, T. Fujimura, Y. Kawano, T. Suzuki, T. Iwai *et al.*, “A 120Gb/s 16QAM CMOS millimeter-wave wireless transceiver,” in *Proc. IEEE International Solid-State Circuits Conference-ISSCC*. IEEE, 2018, pp. 168–170.

- [13] S. He, C. Qi, Y. Wu, and Y. Huang, "Energy-efficient transceiver design for hybrid sub-array architecture MIMO systems," *IEEE Access*, vol. 4, pp. 9895–9905, 2016.
- [14] X. Gao, L. Dai, S. Han, I. Chih-Lin, and R. W. Heath, "Energy-efficient hybrid analog and digital precoding for mmWave MIMO systems with large antenna arrays," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 4, pp. 998–1009, 2016.
- [15] S. Dutta, C. N. Barati, D. Ramirez, A. Dhananjay, J. F. Buckwalter, and S. Rangan, "A case for digital beamforming at mmwave," *IEEE Transactions on Wireless Communications*, pp. 1–1, 2019.
- [16] 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 and Systems Magazine*, vol. 19, no. 2, pp. 33–58, 2019.
- [17] 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 *2020 2nd 6G Wireless Summit (6G SUMMIT)*, 2020, pp. 1–5.
- [18] P. Skrimponis, N. Hosseinzadeh, A. Khalili, E. Erkip, M. J. W. Rodwell, J. Buckwalter, and S. Rangan, "Towards Energy Efficient Mobile Wireless Receivers Above 100 GHz," *IEEE Access*, 2020.
- [19] 3GPP, "TS 38.300, NR and NG-RAN Overall Description; Stage 2," 2017.
- [20] C. Herranz, M. Zhang, M. Mezzavilla, D. Martin-Sacristán, S. Rangan, and J. F. Monserrat, "A 3GPP NR compliant beam management framework to simulate end-to-end mmwave networks," in *Proc. ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2018, pp. 119–125.
- [21] G. R. MacCartney, T. S. Rappaport, and S. Rangan, "Rapid fading due to human blockage in pedestrian crowds at 5G millimeter-wave frequencies," in *Proc. IEEE GLOBECOM*, 2017, pp. 1–7.
- [22] C. Slezak, V. Semkin, S. Andreev, Y. Koucheryavy, and S. Rangan, "Empirical effects of dynamic human-body blockage in 60 GHz communications," *IEEE Commun. Mag.*, vol. 56, no. 12, pp. 60–66, 2018.
- [23] V. Raghavan, V. Podshivalov, J. Hulten, M. A. Tassoudji, A. Sampath, O. H. Koymen, and J. Li, "Spatio-temporal impact of hand and body blockage for millimeter-wave user equipment design at 28 GHz," *IEEE Commun. Mag.*, vol. 56, no. 12, pp. 46–52, 2018.
- [24] S. H. A. Shah, S. Aditya, S. Dutta, C. Slezak, and S. Rangan, "Power Efficient Discontinuous Reception in THz and mmWave Wireless Systems," in *Proc. IEEE SPAWC*, 2019.
- [25] S. H. A. Shah, S. Aditya, and S. Rangan, "Power-efficient beam tracking during connected mode DRX in mmWave and sub-THz systems," Oct 2020.
- [26] W. B. Abbas, F. Gomez-Cuba, and M. Zorzi, "Millimeter wave receiver efficiency: A comprehensive comparison of beamforming schemes with low resolution ADCs," *IEEE Transactions on Wireless Communications*, vol. 16, no. 12, pp. 8131–8146, 2017.
- [27] J. Zhang, L. Dai, X. Li, Y. Liu, and L. Hanzo, "On low-resolution ADCs in practical 5G millimeter-wave massive MIMO systems," *IEEE Communications Magazine*, vol. 56, no. 7, pp. 205–211, 2018.
- [28] C. N. Barati, S. A. Hosseini, M. Mezzavilla, T. Korakis, S. S. Panwar, S. Rangan, and M. Zorzi, "Initial access in millimeter wave cellular systems," *IEEE Trans. Wireless Commun.*, vol. 15, no. 12, pp. 7926–7940, Dec. 2016.
- [29] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, "A tutorial on beam management for 3GPP NR at mmWave frequencies," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 173–196, 2018.
- [30] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE journal on selected areas in communications*, vol. 32, no. 6, pp. 1164–1179, 2014.
- [31] Z. Shen, A. Papasakellariou, J. Montojo, D. Gerstenberger, and F. Xu, "Overview of 3GPP LTE-advanced carrier aggregation for 4G wireless communications," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 122–130, 2012.
- [32] 3GPP, "TS 38.104, Base station (BS) radio transmission and reception," 2019.
- [33] G. Such, M. Schraml, and A. Knopp, "Frequency domain channelizer: A computationally efficient method for non-uniform channelization," *Proceedings of the GNU Radio Conference*, vol. 3, no. 1, 2018.
- [34] S. Lagen, K. Wanuga, H. Elkotby, S. Goyal, N. Patriciello, and L. Giupponi, "New radio physical layer abstraction for system-level simulations of 5G networks," *arXiv preprint arXiv:2001.10309*, 2020.
- [35] M. Puschel, J. M. F. Moura, J. R. Johnson, D. Padua, M. M. Veloso, B. W. Singer, Jianxin Xiong, F. Franchetti, A. Gacic, Y. Voronenko, K. Chen, R. W. Johnson, and N. Rizzolo, "Spiral: Code generation for DSP transforms," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 232–275, 2005.
- [36] P. Skrimponis and S. Rangan, "mmwComm: A millimeter wave communications package," <https://github.com/nyu-wireless/mmwComm>, 2020.
- [37] V. Syrjala, M. Valkama, N. N. Tchamov, and J. Rinne, "Phase noise modelling and mitigation techniques in OFDM communications systems," in *2009 Wireless Telecommunications Symposium*, 2009, pp. 1–7.