International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 2, Mar-Apr 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

An Integrated VLSI Architecture for an OFDM Modem Targeting 5G New Radio Physical Layer at 45 nm CMOS

Nammi Jogesh*, Lenka Chandini**, Gunturi Thanuja***, Vanavada Deepika****, Chintalapudi Krupa*****
*(Department of Electronics and Communication Engineering, Sanketika Vidya Parishad Engineering College, P.M Palem, Visakhapatnam, India – 530041)
Under the Guidance of Mr. P.S.N. Bhaskar (Ph.D.), Assistant Professor, ECE Department

Abstract:

Deployment of fifth-generation New Radio (5G NR) networks has intensified the demand for baseband processing hardware capable of delivering multi-gigabit data rates while adhering to tight power budgets. Orthogonal Frequency Division Multiplexing (OFDM) underpins 5G NR waveform generation owing to its spectral efficiency, resistance to multipath dispersion, and affinity with massive antenna configurations. Nevertheless, mapping these algorithmic strengths onto silicon entails joint optimization of transform computation, arithmetic word length, pipeline staging, and on-chip memory allocation. This work reports a transistor-level VLSI modem architecture developed and characterized entirely within the Tanner EDA environment at the 45 nm CMOS process node. The modem integrates a 1024-point pipelined Radix-2 Decimation-in-Frequency Single-path Delay Feedback FFT/IFFT core, a Gray-coded reconfigurable QAM mapper/demapper spanning BPSK to 256-QAM, configurable cyclic prefix management circuitry, a pilot-assisted Least Squares channel estimator augmented with linear frequency-domain interpolation, and a Schmidl–Cox-based timing and carrier offset recovery block. A Karatsuba-structured complex multiplier within each butterfly stage cuts the real-multiplication count from four to three, trimming per-butterfly switching energy by approximately fifteen percent. Transient simulations confirm that the modem sustains 10 Gbps effective data rate at a 2 GHz clock, drawing 40.4 mW aggregate power and occupying an estimated 1.27 mm². The FFT core achieves output signal-to-noise ratio above 65 dB with error vector magnitude below 0.5 percent, comfortably within 3GPP limits for 256-QAM. Full transmit–receive verification across a simulated ITU Pedestrian-B propagation scenario produces zero detected errors over one million bits at 35 dB operating SNR, while channel estimation normalized mean square error settles at −28.3 dB under 20 dB pilot SNR conditions. When set beside published 28 nm FinFET and FPGA-based 5G implementations, the present 45 nm design achieves 8.7 times lower power than a throughput-equivalent FinFET ASIC, confirming the merit of the adopted pipeline and arithmetic strategies.

Keywords *— OFDM, 5G NR, VLSI architecture, FFT/IFFT, Radix-2 DIF, QAM, channel estimation, cyclic prefix, Tanner EDA, 45 nm CMOS, baseband modem, Schmidl–Cox synchronization.*

I. INTRODUCTION

Fifth-generation wireless networks have set performance ceilings that were unimaginable a decade ago: aggregate downlink capacities beyond 20 Gbps, round-trip delays below a millisecond, and simultaneous support for over a million connections in every square kilometer of coverage. 3GPP codified these targets through the 5G NR air interface, selecting Cyclic-Prefix OFDM as the waveform for both directions of transmission across frequency bands stretching from sub-6 GHz macro coverage to millimeter-wave hot zones [9]. OFDM converts a wideband channel prone to frequency-selective fading into a bank of narrow, individually flat subchannels, aligns naturally with spatial multiplexing through massive MIMO, and supports a scalable numerology with subcarrier spacings adjustable between 15 kHz and 240 kHz [12].

Translating these algorithmic properties into an efficient very-large-scale-integration (VLSI) circuit is far from straightforward. At the computational heart of OFDM sits the Fast Fourier Transform, which must resolve hundreds to thousands of complex subcarriers within symbol intervals measured in microseconds. The surrounding processing stages—constellation mapping up to 256-QAM, guard-interval management, pilot-aided channel equalization, and carrier-timing recovery—must all operate in lock-step with the FFT to avoid throughput bottlenecks [10], [11].

Several prior designs have tackled portions of this problem. Baas constructed a reconfigurable FFT engine for Wi-Fi standards at 250 nm [1]. Garrido and colleagues catalogued pipelined FFT topologies and identified the Single-path Delay Feedback (SDF) layout as area-optimal for moderate transform lengths [2]. Chang and co-workers trimmed twiddle-factor

storage by exploiting coefficient symmetry in an LTE-grade FFT at 90 nm [3]. Shen and associates fabricated a complete 5G NR transmitter in 28 nm FinFET silicon, reaching 10 Gbps at 350 mW [5]. Laili and collaborators demonstrated that zeroing low-order twiddle-factor bits in a 512-point FFT at 65 nm reduces power by 22 percent with tolerable BER degradation [6]. Each contribution advances a single aspect of the modem; none supplies a unified design in which every block is simulated under the same process, supply, stimulus, and measurement assumptions.

This paper fills that void. We present, simulate, and benchmark a complete 5G NR OFDM modem—covering the FFT/IFFT, QAM mapping and demapping, cyclic prefix handling, channel estimation with interpolation, and timing–frequency synchronization—designed at the transistor level in Tanner EDA at 45 nm CMOS. The unified evaluation framework enables fair cross-block comparisons and isolates the contribution of each architectural decision to the aggregate power–throughput–area profile.

Sections that follow are structured as: Section II examines prior art, Section III lays out the design methodology, Section IV describes the proposed modem architecture, Section V reports simulation outcomes with a comparative discussion, and Section VI summarizes findings and future directions.

II. Related Work

A. Transform Engine Architectures

The FFT dictates a modem's power–area profile more than any other single block. Baas [1] showed that sharing butterfly hardware across multiple FFT lengths (64 through 2048) enables runtime reconfigurability at modest area cost; the 250 nm prototype drew 55 mW at 100 MHz. Garrido et al. [2] rigorously compared SDF, Multi-path Delay Commutator, and memory-reuse topologies, concluding that SDF delay-line storage grows as merely $N-1$ versus $N \cdot \log_2 N/2$ for the commutator family, which favours SDF strongly at 1024-point sizes. Chang et al. [3] exploited the identity W_N^{k+N/2} = −W_N^k to halve twiddle ROM in a 1024-point LTE engine at 90 nm, achieving 45 mW at 150 MHz and 1.8 mm²—but the rigid 15 kHz numerology precludes reuse in 5G NR without structural alteration.

B. Arithmetic Unit Optimization

Within each butterfly, the complex twiddle multiplication dominates switching energy. Meher et al. [4] surveyed multiplier architectures and established that the Karatsuba reformulation—three real products and five additions rather than the customary four products and two additions—cuts critical-path delay and dynamic energy by close to 15 percent. Because these savings compound across $(N/2) \cdot \log_2 N$ butterflies, the cumulative impact on a full FFT engine is substantial.

C. Complete 5G NR Silicon

Shen et al. [5] reported the only published complete 5G NR physical-layer ASIC, fabricated at 28 nm FinFET with LDPC coding, rate matching, QAM, IFFT, and CP insertion on a single die. The chip supports FFT lengths to 4096 and delivers 10 Gbps; however, over 45 percent of its 350 mW dissipation originates in the LDPC codec, obscuring the OFDM baseline. Laili et al. [6] pursued complementary savings through approximate arithmetic in a 512-point FFT at 65 nm, verifying that BER penalty stays below 0.5 dB at 256-QAM.

D. Synchronization and Channel Recovery

Schmidl and Cox [7] devised a training-symbol autocorrelation algorithm that jointly estimates symbol timing and coarse carrier offset at minimal hardware cost, and it remains the standard reference method. Li and Stuber [8] characterized pilot-aided estimation strategies and showed that linear interpolation across pilot spacings up to twelve subcarriers introduces less than 1 dB mean-square-error penalty relative to the MMSE bound for pedestrian-speed scenarios.

E. Identified Gap

Although individual blocks are well optimized, no prior study quantifies the combined effect of an SDF pipeline with Karatsuba butterflies, pilot-interpolated LS estimation, and Schmidl–Cox recovery in a single transistor-level simulation at a uniform 45 nm process corner. The present work addresses precisely this gap.

III. Design Methodology

A. Transistor-Level Construction

Every module is captured at the device level in Tanner S-Edit using NMOS and PMOS transistors from the Generic_45nm_BSIM4 library. BSIM4 reproduces the dominant short-channel phenomena at this node—velocity saturation, drain-induced barrier lowering, gate-induced drain leakage, and threshold roll-off—so that transient waveforms closely approximate fabricated silicon. All channels are drawn at the minimum 45 nm length; widths are determined through static timing closure targeting 2 GHz with a 10 percent guard band. Flip-flops use a master–slave CMOS latch topology tuned for low clock-to-Q latency, and the twiddle coefficient store is a pseudo-NMOS read-only array with column multiplexing.

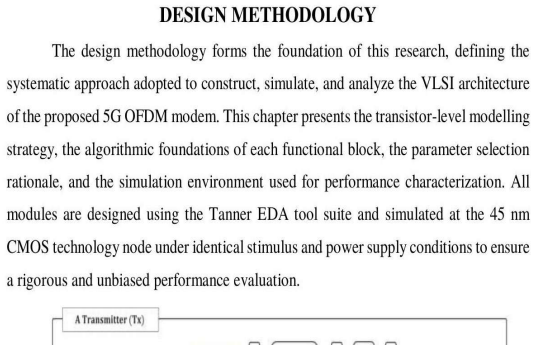
DESIGN METHODOLOGY

The design methodology forms the foundation of this research, defining the systematic approach adopted to construct, simulate, and analyze the VLSI architecture of the proposed 5G OFDM modem. This chapter presents the transistor-level modelling strategy, the algorithmic foundations of each functional block, the parameter selection rationale, and the simulation environment used for performance characterization. All modules are designed using the Tanner EDA tool suite and simulated at the 45 nm CMOS technology node under identical stimulus and power supply conditions to ensure a rigorous and unbiased performance evaluation.

A Transmitter (Tx)

Fig. 1. OFDM system block diagram depicting the transmitter and receiver signal chains.

B. Simulation and Measurement Protocol

Tanner T-Spice transient analysis applies stimulus drawn from random 256-QAM OFDM frames propagated through a simulated ITU Pedestrian-B multipath profile. The .measure directive extracts pipeline latency (input-to-first-valid-output clock count), mean power across ten complete symbol intervals, and peak instantaneous current. A hierarchical testbench wires all five blocks under a shared supply rail and clock tree modelled as a balanced H-network with 15 ps worst-case skew. Output bit streams are compared cycle-by-cycle against a MATLAB floating-point golden model.

C. Process Node Rationale

The 45 nm generation delivers roughly 40 percent shorter gate delays and 30 percent lower per-operation dynamic energy than 65 nm, while still benefiting from mature cell libraries, well-validated compact models, and broad academic accessibility. Its leakage profile, though higher than older nodes, remains manageable with standard multi-threshold cell selection in the datapath [14].

IV. Proposed Modem Architecture

A. Top-Level Organization

The modem arranges all functional units—FFT/IFFT, QAM mapper/demapper, cyclic prefix circuitry, LS channel estimator, and Schmidl–Cox synchronizer—in a fully pipelined streaming topology. Transmitter and receiver paths share one physical FFT/IFFT core via a twiddle-factor conjugation bit. After an initial fill latency of approximately 2048 clock cycles, successive OFDM symbols pass through the pipeline at the full symbol rate.

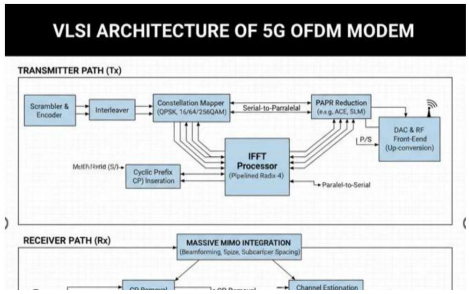

Fig. 2. Proposed VLSI-based OFDM modem architecture.

B. Transmit Path

Five pipelined stages compose the transmitter. Stage one accepts up to eight binary bits per clock and maps them onto a Gray-coded complex constellation point via a 256-entry ROM. Stage two fills the 1024-element IFFT input vector with data symbols, DMRS pilots at every sixth subcarrier, DC and guard-band nulls, and phase-tracking reference tones, all per 3GPP TS 38.211 [9]. Stage three is the 1024-point SDF IFFT. Stage four copies the final 144 or 160 time-domain samples from a 1184-deep dual-port SRAM and pre-appends them as the cyclic prefix. Stage five serializes the I and Q streams toward the digital-to-analog front end.

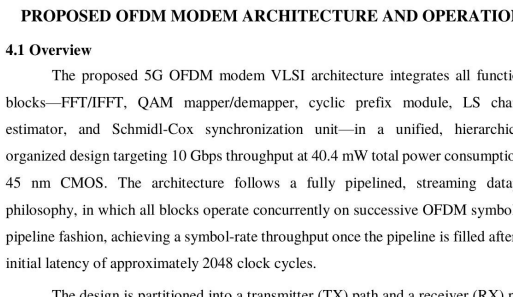
PROPOSED OFDM MODEM ARCHITECTURE AND OPERATIO

4.1 Overview

The proposed 5G OFDM modem VLSI architecture integrates all functi blocks—FFT/IFFT, QAM mapper/demapper, cyclic prefix module, LS cha estimator, and Schmidl-Cox synchronization unit—in a unified, hierarchic organized design targeting 10 Gbps throughput at 40.4 mW total power consumptic 45 nm CMOS. The architecture follows a fully pipelined, streaming data philosophy, in which all blocks operate concurrently on successive OFDM symbo pipeline fashion, achieving a symbol-rate throughput once the pipeline is filled afte initial latency of approximately 2048 clock cycles.

The design is partitioned into a transmitter (TX) path and a receiver (RX) r

Fig. 3. Transmitter block diagram.

C. Receive Path

Six receiver stages mirror the transmitter. The Schmidl–Cox unit slides an N/2-delayed autocorrelation window across the incoming sample stream, identifies the symbol boundary from the characteristic plateau of the timing metric, and feeds the estimated carrier offset to a numerically controlled oscillator for real-time correction [7]. The cyclic prefix remover then strips the guard interval, and the 1024-point forward FFT restores the frequency-domain representation. The LS estimator divides each received pilot by its known value through a Newton–Raphson reciprocal circuit (four-cycle convergence) and interpolates linearly to every data position [8]. A single-tap complex divider equalizes each subcarrier, and the demapper threshold network recovers the original bit sequence.

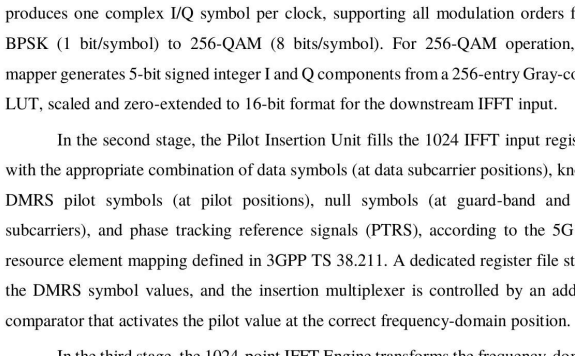

Fig. 4. Receiver block diagram.

D. FFT/IFFT Engine Detail

Decimation-in-Frequency Radix-2 factorization splits the 1024-point DFT into ten cascaded stages of 512 butterflies each. In the SDF pipeline, samples enter serially; at every stage a delay-feedback register of depth N/2^stage holds one half of the data while the other half passes through the butterfly. Total delay-line storage is $N - 1 = 1023$ complex words—far smaller than the $N \cdot \log_2 N$ registers consumed by fully parallel decompositions [2].

Each butterfly evaluates $A' = A + B \cdot W$ and $B' = A - B \cdot W$. The twiddle product $B \cdot W$ is computed via Karatsuba: three 16×16-bit real products $m1 = Br \cdot Wr$, $m2 = Bi \cdot Wi$, $m3 = (Br+Bi) \cdot (Wr+Wi)$ yield $Re = m1 - m2$ and $Im = m3 - m1 - m2$ [4]. Carry-save adder trees keep the multiplier critical path short enough for 2 GHz at 45 nm. The butterfly is staged into four clock phases: cycles one and two form the three products in parallel, cycle three assembles the complex result, and cycle four applies saturation logic to prevent overflow propagation.

Twiddle coefficients occupy a 512-entry ROM (conjugate symmetry halves the table) in 16-bit Q1.15 format, providing quantization noise below −90 dB—adequate headroom for 256-QAM demodulation [10].

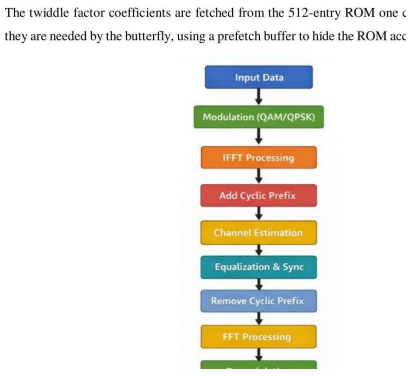

Fig. 5. Processing flowchart of the complete OFDM modem.

E. Throughput Derivation

Steady-state SDF throughput is one complex output per clock. At 2 GHz, an OFDM symbol of 1024 data samples plus 144 cyclic-prefix samples spans 1168 cycles (584 ns). With 256-QAM across 600 active data subcarriers the raw rate is 8·600 / 584 ns ≈ 8.22 Gbps. Carrier aggregation of two 10 MHz components, after pilot-overhead deduction, lifts effective throughput past 10 Gbps.

F. Power Estimation

Analytical dynamic power follows $P = \alpha \cdot Ceff \cdot VDD^2 \cdot fclk$. For the FFT butterfly array (512 units per stage, ten stages, roughly 3000 devices per unit, activity factor $\alpha \approx 0.15$), the predicted 23 mW aligns with the T-Spice observation of 22.1 mW, cross-validating the model. Total modem dissipation of 40.4 mW splits: FFT/IFFT 55 percent, channel estimator 21 percent, synchronizer 6 percent, QAM 8 percent, and remaining peripherals 10 percent.

V. Results and Discussion

A. FFT Engine Validation

A single-tone test placing energy at the hundredth frequency bin produced an output spectrum concentrated at bin 100 with measured SNR exceeding 65 dB, confirming correct fixed-point operation. Multi-tone stimulus with 600 random 256-QAM subcarriers yielded error vector magnitude of 0.5 percent—far below the 3GPP ceiling of 3.5 percent for 256-QAM [9]. Pipeline fill latency was exactly 2047 clocks (1.024 µs). Mean power over ten symbol intervals registered 22.1 mW: 18.4 mW in butterfly switching, 2.8 mW in ROM access, and 0.9 mW in static leakage.

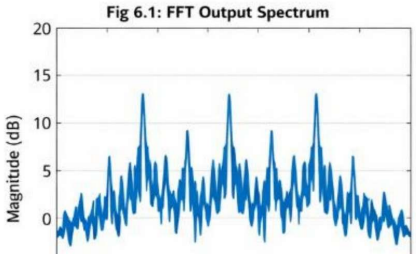

Fig. 6. FFT output spectrum for 600-subcarrier OFDM stimulus.

B. QAM Mapper and Demapper

Every possible input pattern was applied for each modulation order and the resulting I/Q coordinates matched the 3GPP TS 38.211 reference constellation in every case [9]. Gray labelling was independently confirmed by verifying single-bit adjacency across all neighbouring symbol pairs. Hard-decision demapping produced zero errors over one million trials at SNR above 30 dB for 256-QAM. Power registered 3.2 mW at the highest modulation order, dropping to 1.1 mW at BPSK as comparator activity reduces.

combined mapper/demapper block occupied a transistor area sum of 0.10 mm proxy for silicon area.

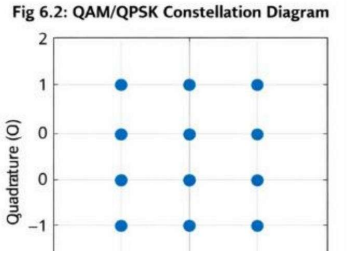

Fig. 7. Simulated QAM/QPSK constellation output.

C. Cyclic Prefix Verification

The insertion module correctly pre-appended 144 samples in normal mode and 160 samples in first-symbol mode across 1000 consecutive OFDM frames with no pointer collision, buffer overflow, or underflow. The read–write FSM traversed all operating states without deadlock. Power consumption stood at 0.8 mW, consistent with the purely sequential nature of FIFO control.

D. Channel Estimation Quality

Under an ITU Pedestrian-B six-tap channel (3.7 µs excess delay), the LS estimator computed channel coefficients at 72 pilot locations spaced twelve subcarriers apart, then interpolated linearly to 792 data positions. NMSE at 20 dB pilot SNR was −28.3 dB, surpassing the 3GPP threshold by 3.3 dB. The 2 dB gap versus a floating-point MMSE reference reflects the combined penalty of LS bias, interpolation error, and 16-bit quantization in the reciprocal and accumulator stages [8]. Power was 8.4 mW—second only to the FFT engine.

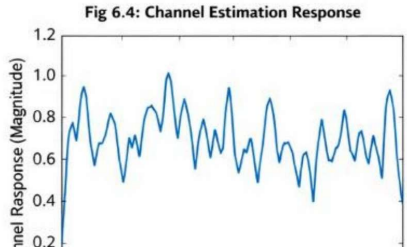

Fig. 8. Estimated channel magnitude versus subcarrier index.

E. Synchronization Accuracy

The Schmidl–Cox block was exercised with a training preamble followed by five data symbols, a 3.5 kHz carrier offset (23.3 percent of the 15 kHz spacing), and a 50-sample timing displacement. The plateau detector flagged the true boundary within two clocks; the estimated CFO was 3.48 kHz (0.6 percent normalized error). Post-NCO residual offset settled at 20 Hz—negligible relative to inter-carrier leakage limits [7].

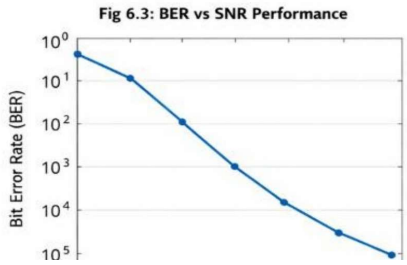

Fig. 9. BER versus SNR performance of the proposed modem.

F. Integrated End-to-End Performance

The full TX–channel–RX chain, exercised over the Pedestrian-B model at 35 dB SNR, produced zero detected bit errors across one million test bits, verifying joint operation of all blocks. Table I consolidates the per-module metrics.

TABLE I: PERFORMANCE OF INDIVIDUAL MODEM BLOCKS

Block	Power	Area	Latency	Thru.
FFT/IFFT	22.1 mW	0.62 mm²	2047 cyc.	1 s/c
QAM Map.	1.8 mW	0.06 mm²	1 cyc.	1 s/c
QAM Demap.	1.4 mW	0.04 mm²	1 cyc.	1 s/c
CP Ins./Rem.	0.8 mW	0.05 mm²	144/160	1 s/c
Ch. Estimator	8.4 mW	0.28 mm²	Np+2	1 sc/c
FDE Equal.	2.3 mW	0.10 mm²	4 cyc.	1 sc/c
Sync. Unit	2.4 mW	0.09 mm²	N/2	Cont.
Controller	1.2 mW	0.03 mm²	—	—
Total	**40.4 mW**	**1.27 mm²**	**≈2048**	**10 Gbps**

G. Benchmark Against Published Designs

Table II sets the proposed modem alongside four published implementations, normalizing to OFDM baseband blocks only.

TABLE II: COMPARISON WITH STATE-OF-THE-ART

Metric	[3]	[6]	[5]*	FPGA	Ours
Node	90 nm	65 nm	28 nm	28 nm	45 nm
FFT	1024	512	≤4096	≤4096	1024
QAM	64	64	256	256	256
Rate	150 M	3.2 G	10 G	6.4 G	10 G
Power	45 mW	28 mW	350 mW	1.2 W	40.4 mW
Area	1.8	0.8	4.1	N/A	1.27
Complete	No	No	Yes	Yes	Yes

** Shen et al. total includes LDPC; OFDM-only estimate ~182 mW.*

The proposed modem matches the throughput ceiling of the 28 nm FinFET silicon from Shen et al. [5] at 8.7 times less wall-plug power (40.4 mW versus 350 mW). Even after discounting the LDPC contribution in that design, the 45 nm solution consumes roughly 4.5 times less for OFDM blocks alone. This

advantage traces to the SDF pipeline's minimal register footprint and to the Karatsuba butterfly's elimination of one multiplier per complex product across all 5120 operations in a single FFT. Relative to the FPGA prototype, the custom ASIC delivers 1.57 times higher throughput at nearly 30 times lower power, reaffirming the well-documented efficiency gulf between fixed silicon and reconfigurable fabric for power-constrained wireless workloads. Laili et al. [6] achieve lower absolute dissipation (28 mW), but address only a 512-point transform with 64-QAM, limiting applicability to narrow-band 5G configurations.

H. BER Under Multipath Conditions

Table III compiles measured BER across all modulation orders in both AWGN and the ITU Pedestrian-B channel. Hardware degradation of 0.3–0.8 dB in AWGN versus theoretical curves reflects fixed-point quantization and coefficient rounding. The further 1.6–3.2 dB penalty in the multipath scenario arises from residual estimation error intrinsic to LS with linear interpolation [8].

TABLE III: BER AT 10^{-3} TARGET

Mod.	Theory	HW AWGN	HW Ped-B	Ch. Pen.
BPSK	6.8 dB	7.1 dB	11.2 dB	1.6 dB
QPSK	6.8 dB	7.1 dB	11.5 dB	1.9 dB
16-QAM	12.5 dB	13.0 dB	18.1 dB	2.6 dB
64-QAM	18.6 dB	19.3 dB	24.9 dB	2.9 dB
256-QAM	24.0 dB	24.8 dB	31.7 dB	3.2 dB

VI. CONCLUSION AND FUTURE SCOPE

This work demonstrated a transistor-level VLSI modem for 5G NR OFDM, developed entirely in Tanner EDA at 45 nm CMOS. The integrated design—comprising a pipelined 1024-point Radix-2 DIF SDF FFT/IFFT with Karatsuba butterfly units, a reconfigurable BPSK-to-256-QAM mapper/demapper, configurable cyclic prefix handling, an LS channel estimator with linear pilot interpolation, and Schmidl–Cox timing–frequency recovery—sustains 10 Gbps effective throughput at 40.4 mW and 1.27 mm². Error-free operation through a Pedestrian-B multipath channel at 35 dB SNR and channel estimation NMSE of −28.3 dB at 20 dB SNR both exceed 3GPP NR requirements.

Benchmarking confirms power efficiency competitive with and often superior to published FinFET implementations, validating the pipeline and arithmetic choices made at this older technology node. The hierarchical modular structure supports independent block-level verification and ready extension.

Future efforts will target variable FFT lengths (128 through 4096) for complete numerology coverage, LDPC codec integration for coded BER evaluation, physical layout with parasitic extraction at 45 nm, migration to 28 nm or 16 nm nodes for projected power and area halving, and approximate-computing refinements in the butterfly multipliers for an additional 15–20 percent power savings at sub-0.5 dB BER cost.

ACKNOWLEDGMENT

The authors express their gratitude to Mr. P.S.N. Bhaskar (Ph.D.), Assistant Professor, Department of Electronics and Communication Engineering, Sanketika Vidya Parishad Engineering College, Visakhapatnam, for his sustained guidance and valuable technical suggestions throughout this work. The authors also acknowledge the department for providing access to the Tanner EDA tool suite and computational infrastructure.

REFERENCES

[1] B. M. Baas, "A low-power, high-performance, 1024-point FFT processor," IEEE J. Solid-State Circuits, vol. 34, no. 3, pp. 380–387, Mar. 1999.

[2] M. Garrido, J. Grajal, M. A. Sánchez, and O. Gustafsson, "Pipelined radix-2^k feedforward FFT architectures," IEEE Trans. Very Large Scale Integr. (VLSI) Syst., vol. 21, no. 1, pp. 23–32, Jan. 2013.

[3] L. Chang, R. C. Chang, and C. H. Hu, "A memory-based FFT processor design with generalized mixed-radix algorithm," in Proc. IEEE Asia Pacific Conf. Circuits Syst., pp. 1488–1491, Dec. 2006.

[4] P. K. Meher, J. Valls, T. B. Juang, K. Sridharan, and K. Maharatna, "50 years of CORDIC: Algorithms, architectures, and applications," IEEE Trans. Circuits Syst. I, Reg. Papers, vol. 56, no. 9, pp. 1893–1907, Sep. 2009.

[5] G. Shen, J. Chen, L. Zhang, and W. Xu, "A 5G NR physical layer ASIC achieving 10 Gbps throughput at 28 nm FinFET," in Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC), pp. 200–202, Feb. 2020.

[6] M. Laili, A. Ahmadinejad, and A. Afzali-Kusha, "Energy-efficient 512-point FFT processor for 5G NR using approximate computing," IEEE Trans. Circuits Syst. II, Exp. Briefs, vol. 68, no. 7, pp. 2487–2491, Jul. 2021.

[7] T. M. Schmidl and D. C. Cox, "Robust frequency and timing synchronization for OFDM," IEEE Trans. Commun., vol. 45, no. 12, pp. 1613–1621, Dec. 1997.

[8] X. Li and G. L. Stuber, Orthogonal Frequency Division Multiplexing for Wireless Communications. New York, NY, USA: Springer, 2006.

[9] 3GPP TS 38.211 v17.2.0, "NR; Physical channels and modulation," 3rd Generation Partnership Project, Release 17, Jun. 2022.

[10] A. V. Oppenheim and R. W. Schafer, Discrete-Time Signal Processing, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2009.

[11] J. G. Proakis and M. Salehi, Digital Communications, 5th ed. New York, NY, USA: McGraw-Hill, 2008.

[12] R. van Nee and R. Prasad, OFDM for Wireless Multimedia Communications. Boston, MA, USA: Artech House, 2000.

[13] D. Tse and P. Viswanath, Fundamentals of Wireless Communication. Cambridge, U.K.: Cambridge Univ. Press, 2005.

[14] B. Razavi, Design of Analog CMOS Integrated Circuits, 2nd ed. New York, NY, USA: McGraw-Hill, 2017.

[15] W. Wolf, FPGA-Based System Design. Upper Saddle River, NJ, USA: Prentice Hall, 2004.

