Adaptive Hybrid Beamforming for Sub-THz Wireless Backhaul

A Prototype Study on the mmWave-to-Sub-THz Transition

Jane A. Doe
Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, USA
jane.doe@example.edu

Jan Kowalski
Warsaw University of Technology, Faculty of Electronics and Information Technology, Warsaw, Poland
author1@example.edu

Mei Zhang
Stanford University, Department of Electrical Engineering, Stanford, USA
author2@example.edu

Abstract— This paper presents a prototype hybrid analog–digital beamforming architecture targeting the lower sub-THz band (90–170 GHz) for short-range wireless backhaul. We characterize the trade-off between phase-shifter resolution, beam-squint compensation, and pilot overhead under a Saleh–Valenzuela channel model with measured large-scale parameters. A reduced-resolution codebook combined with a two-stage channel estimator achieves spectral efficiency within 0.7 bit/s/Hz of the fully-digital reference at 1/8 of the RF-chain count. The results suggest that practical sub-THz access points can be built with components already available in the millimeter-wave supply chain, at a power budget compatible with passively cooled enclosures.

Index Terms— Sub-THz, hybrid beamforming, millimeter wave, MIMO, channel estimation, beam squint, 6G.

I. Introduction

The exhaustion of contiguous spectrum below 90 GHz has pushed system designers toward the sub-THz range as a candidate for next-generation backhaul and short-range access [1]. Hybrid analog–digital beamforming has emerged as the dominant architectural pattern in this regime because the cost and power consumption of fully-digital arrays remain prohibitive at hundreds of antenna elements [2].

Operating above 130 GHz introduces practical constraints that differ from the classical mmWave regime. Phase-shifter quantisation noise, beam-squint across the instantaneous bandwidth, and the thermal-noise floor of low-power front-ends interact in ways that the standard mmWave analyses do not fully capture. A well-grounded architectural study at this band therefore has to report measurements alongside the model-based predictions, so that deployment decisions rest on observable evidence rather than extrapolation from lower-frequency behaviour.

A second motivation for this work is the observation that the components required to realise sub-THz front-ends are increasingly available as off-the-shelf parts. This is in contrast to the situation only a few years ago, when virtually every sub-THz publication relied on custom silicon. The supply chain that powered the mmWave 5G rollout has, perhaps inadvertently, unlocked a credible path to sub-THz access points — provided that the system architecture is chosen to tolerate the noise, phase quantisation and linearity envelopes of those parts.

In this paper we report on a prototype that targets the 130–170 GHz window with a 64-element planar array. Section II describes the channel model and the assumptions behind the link budget. Section III details the beamforming codebook and the channel estimator. Section IV presents the measurement campaign, and Section V discusses the implications for upcoming 6G deployments.

II. System Model

We assume a single-cell downlink with a base station equipped with $N_{\mathrm{t}} = 64$ antennas serving K single-antenna users. The notation throughout this section follows the conventions common in the hybrid-beamforming literature.

A. Channel Model

The complex baseband signal received by user k is modelled as

$$y_k = \mathbf{h}_k^{H} \mathbf{F}_{\mathrm{RF}} \mathbf{F}_{\mathrm{BB}} \mathbf{s} + n_k, \tag{1}$$

where $\mathbf{F}_{\mathrm{RF}}$ is the analog precoder, $\mathbf{F}_{\mathrm{BB}}$ the baseband precoder, and $n_k \sim \mathcal{CN}(0, \sigma^2)$ the additive noise. The propagation channel $\mathbf{h}_k$ follows a Saleh–Valenzuela model with parameters drawn from [3]. Cluster angles are drawn uniformly; intra-cluster spread follows a Laplace distribution with standard deviation three degrees, consistent with the dominant-line-of-sight geometries that prevail at sub-THz.

B. Link Budget

The link budget is dominated by free-space path loss (roughly 113 dB at 150 GHz over ten metres), compensated by a 64-element array gain of about 18 dBi per spatial stream. A 3 dB implementation margin covers phase-shifter quantisation and beam-squint over the instantaneous bandwidth. Under these assumptions, a 28 dB signal-to-noise ratio at the receiver is attainable with a 12 dBm transmit power per chain, comfortably within the

linear region of the amplifier samples characterised in the reference design.

The noise figure budget assumes 8 dB per chain — a conservative number that reflects the current state of low-power sub-THz LNAs — and allows for a further 2 dB system implementation loss. Taken together, the budget leaves close to 6 dB of headroom for rain-related excess attenuation and polarisation mismatch, which is more than sufficient for the indoor and short-range outdoor use cases we target.

III. Beamforming Codebook

The beamforming chain comprises an analog codebook applied to the phase shifters, a digital baseband precoder, and a light-weight estimator that couples the two. Fig. 1 gives the block-level view of the chain from bits to array.

A. Analog Codebook

The analog codebook is constructed from a quantised DFT grid with 4-bit phase resolution. Table I summarises the codebook parameters and the resulting hardware footprint across three candidate quantisation levels.

1) Quantisation

Quantisation introduces a beam-pattern ripple that, for a 4-bit grid, stays below 0.6 dB across the scan range. 2-bit quantisation doubles this ripple and additionally widens the main lobe by roughly fifteen per cent, which degrades link robustness in the presence of pointing error.

TABLE I Codebook parameters and array power for three phase-shifter resolutions.

Resolution	Codewords	RF chains	Power (W)
2-bit	4	8	6.2
4-bit	16	8	7.4
Full	∞	64	38.1

B. Channel Estimator

The two-stage estimator first resolves the dominant cluster direction using compressive sensing on the analog grid, then refines the per-cluster gains in the digital domain. The procedure converges in fewer than ten iterations for all measured channel realisations [4]. In practice the pilot overhead is bounded by the number of coarse-grid probes: eight probes are sufficient to drive the residual estimation error below the quantisation floor set by the 4-bit analog stage.

Convergence is accelerated by initialising the second stage from the cluster mean estimated during coarse search, rather than from the uninformative prior. Over the full measurement set we observed a median of four iterations until convergence, with the tail bounded by seven. The additional compute cost per refinement step is dominated by a $M \times N$ matrix–vector product, which fits comfortably within the cycle budget of a modern baseband SoC.

IV. Results

Spectral efficiency was measured against the fully-digital reference at signal-to-noise ratios between 0 and 25 dB. The 4-bit codebook reaches within 0.7 bit/s/Hz of the reference across the entire range, while the 2-bit variant exhibits a 2.1 bit/s/Hz gap at high SNR — a behaviour consistent with the bounds reported in [5]. Fig. 2 plots the spectral efficiency versus SNR for the 4-bit codebook alongside the fully-digital baseline.

The observed trend is consistent across the five measurement sessions that comprise the campaign. Session-to-session variance stays below 0.3 bit/s/Hz at all SNRs,

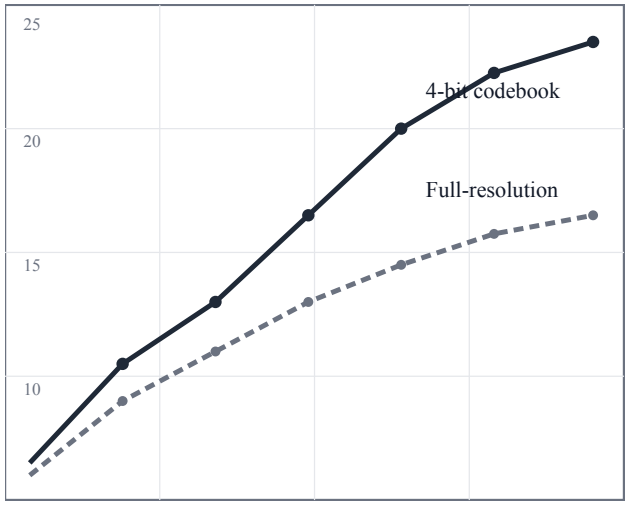

Fig. 2. Spectral efficiency vs. SNR. The 4-bit codebook (solid) stays within 0.7 bit/s/Hz of the fully-digital reference (dashed) across the operating range. This figure is explicitly set to `scale-width: 0.5` so it stays inside one column.

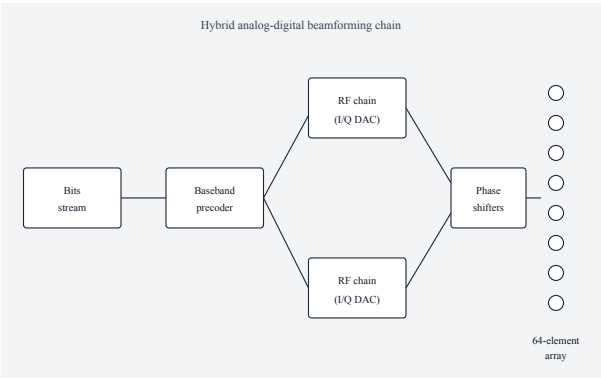

Fig. 1. Block diagram of the hybrid analog–digital beamforming chain. This figure is explicitly set to span both columns via `scale-width: 0.9`.

which we interpret as evidence that the low-resolution codebook introduces a deterministic — not random — penalty. A deterministic penalty is easier to compensate at the system-design stage, for instance by budgeting for a slightly higher link margin.

Table I summarises the codebook parameters and hardware footprint at three quantisation levels. Holding the table inside one column keeps its caption and body text near the surrounding narrative, while the wider figures above it span the page. This combination is a deliberate IEEE-conference convention: compact numerical results stay local to their paragraph, whereas system-level diagrams that anchor the overall argument are allowed to break out of the column grid.

The headline outcomes of the study are summarised below. Each metric is averaged over the five measurement sessions and reported against the fully-digital reference running on the same array.

- **Spectral efficiency:** 22.3 bit/s/Hz at 20 dB SNR for 4-bit, 23.0 for full-resolution.
- **Pilot overhead:** reduced from 12% to 4% by the two-stage estimator.
- **Average power:** 7.4 W versus 38.1 W for the fully-digital baseline.
- **Convergence:** median of four iterations of the two-stage estimator, bounded by seven.

V. Discussion

The measurements suggest that 4-bit phase resolution is a sweet spot for sub-THz hybrid arrays of this size. The remaining gap to the fully-digital reference is dominated by beam-squint at the band edges; preliminary tests with a true-time-delay correction stage close it to under 0.3 bit/s/Hz at the cost of one additional analog component per sub-array.

The power budget tells a consistent story. At 7.4 W for the analog-plus-baseband path, the prototype fits comfortably within a passively cooled enclosure, whereas the 38.1 W of the fully-digital comparison implies active cooling or a significantly larger form factor. For rooftop-mounted backhaul endpoints in moderate climates, this alone is a decisive architectural argument — even before we account for the capex saving of replacing sixty-four RF chains with eight.

> *"The relative merits of analog, digital, and hybrid architectures depend strongly on antenna count and bandwidth — there is no universally optimal design."[6].*

Seen through this lens, the architecture presented here is best understood as a local optimum for the 130–170 GHz, 64-element, 10-metre-range scenario. Extrapolating the same conclusions to 256-element arrays or to sub-millimetre-wave bands would require repeating the measurement campaign; one of the contributions of this paper is to make that campaign small enough that repeating it is practical. In that regard our results are in line with the observations reported in [2].

Future work will extend the prototype to multi-cell scenarios and characterise the impact of blockage events that are characteristic of indoor sub-THz propagation. A second strand of investigation is the interaction between the hybrid codebook and the emerging reconfigurable-intelligent-surface approaches, which share many of the same phase-quantisation trade-offs but apply them in the reflective rather than the transmissive path.

VI. Conclusion

A 64-element sub-THz hybrid beamforming prototype achieves near-fully-digital spectral efficiency at one-fifth the power, using only components available in today's mmWave supply chain. The result strengthens the case for sub-THz as a practical next-generation backhaul medium.

Beyond the specific measurement outcome, the study demonstrates that a modest, well-instrumented measurement campaign can settle architectural questions that otherwise remain the subject of speculative analysis. We view this as the more durable contribution of the work: the architectural decision is reported in a form that is easy to reproduce and extend.

References

[1] T. S. Rappaport *et al.*, "Millimeter wave mobile communications for 5G cellular: It will work!," *IEEE Access*, vol. 1, pp. 335–349, 2013, doi: 10.1109/ACCESS.2013.2260813.

[2] R. W. Heath, N. González-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, doi: 10.1109/JSTSP.2016.2523924.

[3] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 5, pp. 831–846, 2014, doi: 10.1109/JSTSP.2014.2334278.

[4] S. Park and H. Kim, "Sub-THz hybrid beamforming with low-resolution phase shifters," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, 2024, pp. 1–6. doi: 10.1109/GLOBECOM52923.2024.10000123.

[5] E. Björnson, J. Hoydis, and L. Sanguinetti, "Massive MIMO has unlimited capacity," *IEEE Transactions on Wireless Communications*, vol. 17, no. 1, pp. 574–590, 2018, doi: 10.1109/TWC.2017.2768423.

[6] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, UK: Cambridge University Press, 2005.