

# DecodeX: Exploring and Benchmarking of LDPC Decoding Across CPU, GPU, and ASIC Platforms

Zhenzhou Qi<sup>1</sup>, Yuncheng Yao<sup>2</sup>, Yiming Li<sup>1</sup>, Chung-Hsuan Tung<sup>1</sup>, Junyao Zheng<sup>1</sup>  
Danyang Zhuo<sup>2</sup>, and Tingjun Chen<sup>1</sup>

Departments of Electrical and Computer Engineering<sup>1</sup> and Computer Science<sup>2</sup>, Duke University

## Abstract

Emerging virtualized radio access networks (vRANs) demand flexible and efficient baseband processing across heterogeneous compute substrates, where low-density parity-check (LDPC) decoding in the Physical layer remains one of the most performance- and energy-critical workloads. In this paper, we present DecodeX, a unified benchmarking framework for evaluating LDPC decoding implementations across different hardware platforms. DecodeX integrates a comprehensive suite of LDPC decoder implementations, including kernels, APIs, and test vectors for CPUs (FlexRAN), GPUs (Aerial and Sionna-RK), and ASICs (ACC100), and can be readily extended to additional architectures and configurations. Using DecodeX, we systematically characterize how different platforms orchestrate computation—from threading and memory management to data movement and accelerator offload—and quantify the resulting decoding latency under varying Physical layer parameters. Our observations reveal distinct trade-offs in parallel efficiency and offload overhead, showing that accelerator gains strongly depend on these insights, we discuss how cross-platform benchmarking can inform adaptive scheduling and co-design for future heterogeneous vRANs, enabling scalable and energy-efficient baseband processing for NextG wireless systems.

## CCS Concepts

• **Networks** → **Network performance analysis; Wireless access points, base stations and infrastructure.**

## Keywords

Virtualized radio access network (vRAN); baseband processing; heterogeneous computing

## ACM Reference Format:

Zhenzhou Qi<sup>1</sup>, Yuncheng Yao<sup>2</sup>, Yiming Li<sup>1</sup>, Chung-Hsuan Tung<sup>1</sup>, Junyao Zheng<sup>1</sup>, Danyang Zhuo<sup>2</sup>, and Tingjun Chen<sup>1</sup>. 2026. **DecodeX**: Exploring and Benchmarking of LDPC Decoding Across CPU, GPU, and ASIC Platforms. In *The 27th International Workshop on Mobile Computing Systems and Applications (HotMobile '26)*, February 25–26, 2026, Atlanta, GA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3789514.3792034>

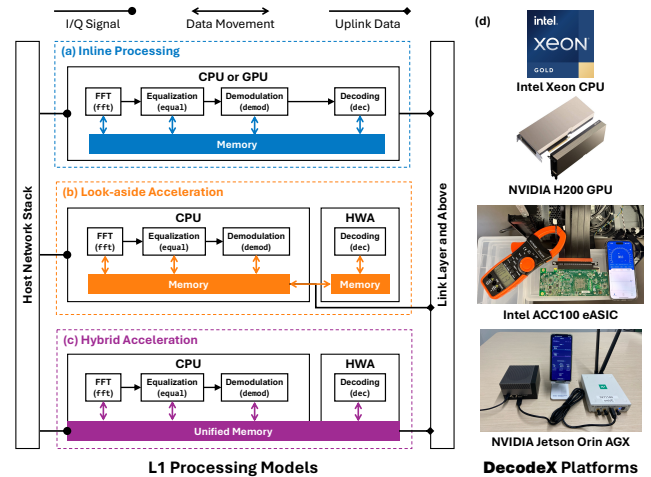


This work is licensed under a Creative Commons Attribution 4.0 International License. *HotMobile '26, Atlanta, GA, USA*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2471-8/2026/02

<https://doi.org/10.1145/3789514.3792034>



**Figure 1:** (a)–(c) Overview of L1 processing models with different hardware accelerators (HWAs): Inline Processing, Lookaside Acceleration, and Hybrid Acceleration. (d) Representative DecodeX implementations of these models: DecodeCPU-FlexRAN (Intel Xeon CPU), DecodeGPU-Aerial (NVIDIA H200 GPU), DecodeASIC-ACC100 (Intel eASIC), and DecodeGPU-SionnaRK (NVIDIA Jetson Orin AGX).

## 1 Introduction

Virtualized radio access networks (vRANs) are redefining how next-generation cellular systems deploy baseband processing [9, 15, 21–23, 25]. By decoupling Layer-1 (L1) functions from dedicated hardware and running them as software on general-purpose or accelerated platforms, vRANs enable flexible orchestration and cloud-native elasticity. However, this shift also exposes severe computational and latency challenges, particularly for the Physical (PHY) layer forward error correction (FEC) blocks such as LDPC decoding, which dominate the processing time in the 5G stack.

Recent advances in heterogeneous computing have introduced multiple candidate substrates for vRAN baseband processing acceleration: specialized software libraries on multicore CPUs (e.g., Intel FlexRAN [12]), GPU-based solutions from both open-source and vendor-optimized frameworks (e.g., OAI CUDA LDPC [20], NVIDIA’s Aerial [6] and Sionna-RK [19]), and ASICs (e.g., Intel ACC100 [26]). Despite their growing availability, comprehensive benchmarking and system-level understanding of how these platforms execute LDPC decoding, interact with memory hierarchies, and trade off latency versus flexibility remains largely unexplored.

In this paper, we present DecodeX, a cross-platform exploration of LDPC decoding acceleration across CPU, GPU, and ASIC platforms. As a comprehensive benchmarking framework, DecodeX

allows for unified evaluation across different kernel implementations under a consistent interface, each named `Decode{HW}-{Impl}`. We examine representative implementations of DecodeX, including DecodeCPU-FlexRAN, DecodeGPU-Aerial, DecodeGPU-SionnaRK, and DecodeASIC-ACC100, which collectively capture the diversity of software and hardware acceleration models in modern vRAN deployments. DecodeX is readily extensible to additional configurations such as DecodeCPU-OAI and DecodeGPU-BBDev, making it a versatile tool for future benchmarking and cross-platform analysis.

We perform a detailed analysis of how each system organizes its decoding pipeline: from threading and single instruction, multiple data (SIMD) parallelization on CPUs, offload management and queue-based scheduling in DPDK BBDev, to CUDA kernel launches and memory transfers on GPUs. Extensive profiling results quantify LDPC decoding latency as a function of modulation and coding scheme (MCS), physical resource block (PRB), and signal-to-noise ratio (SNR) across different DecodeX implementations. Our measurements reveal that accelerator performance is influenced not only by raw compute throughput but also by data-movement efficiency and task granularity. To improve GPU utilization and latency efficiency, we extended NVIDIA’s PyAerial framework by introducing a parallel LDPC processing interface, enabling simultaneous decoding of multiple transport blocks across CUDA streams.

Overall, DecodeX provides a comparative, system-level characterization of LDPC decoding across heterogeneous vRAN systems. Beyond benchmarking, it offers actionable insights for future vRAN baseband and compute co-design, where dynamic resource allocation across CPUs, GPUs, and ASICs will be essential to balancing latency, energy efficiency, and flexibility. *To the best of our knowledge, this work presents the first comprehensive study of LDPC decoding performance across different platforms, supported by a complete suite of unit tests and benchmarking scripts open-sourced at [1].*

## 2 Related Work and Background

The evolution of vRANs has driven extensive research on leveraging heterogeneous computing platforms to accelerate 5G L1 workloads. Prior works such as CloudRIC [25], Hydra [10], and Savannah [23] demonstrate that general-purpose CPUs can execute full 5G baseband pipelines efficiently, while accelerator-based designs explore offloading compute-intensive kernels, such as forward error correction (FEC), to specialized hardware. However, these solutions often face challenges in balancing latency, data-movement overhead, and resource sharing across heterogeneous devices, motivating the need for a systematic cross-platform analysis.

**5G L1 and LDPC decoding.** In 5G NR, each slot consists of 14 OFDM symbols [2], and depending on the MCS and bandwidth, multiple transport blocks (TBs) may be transmitted per frame. A TB corresponds to a data unit scheduled by the MAC layer and may span one or more slots. For data channels, 5G NR adopts LDPC codes based on two standardized base graphs (BG1 and BG2) [3], each optimized for different code rates and lifting sizes. These quasi-cyclic codes compactly represent the parity-check matrix through circulant shifts, avoiding storage of large dense matrices. At the receiver, rate de-matching reconstructs code blocks (CBs) [16] from TBs, followed by iterative belief-propagation decoding across thousands of variable and check nodes. The LDPC decoding configuration

depends on several key parameters. The MCS defines the modulation order ( $Q_m$ ) and determines the code rate after rate matching (via puncturing or repetition). The *lifting size* ( $Z_c$ ) expands the base graph into the full parity-check matrix. A single TB may be segmented into multiple CBs, with the number of decoding iterations adapting to channel quality and SNR. LDPC decoding typically dominates the L1 processing latency as it involves fine-grained data dependencies, irregular memory access, and iterative message-passing updates.

**Heterogeneous acceleration.** To meet real-time constraints, diverse acceleration strategies have been explored: CPU implementations rely on SIMD vectorization [12], eASICs like Intel’s ACC100 exploit deeply pipelined and deterministic architectures, and GPUs leverage massive parallelism for high iteration counts. While GPUs provide exceptional throughput [24, 27, 29], discrete devices often follow a *lookaside offload model*, requiring frequent host-device transfers that increase energy and latency overhead. Integrated platforms such as NVIDIA Jetson mitigate these penalties through unified memory, but their limited compute capacity restricts scalability for datacenter or multi-cell vRAN deployments.

DecodeX differs from prior efforts by providing a unified, cross-platform benchmarking and profiling framework for LDPC decoding across CPU, GPU, and ASIC architectures. It systematically dissects each platform’s execution pipeline, from threading and queue scheduling to memory movement and kernel launches, quantifying decoding latency across varying MCS, SNR, and PRB allocations.

## 3 Design of DecodeX

This section describes the design of DecodeX, focusing on how LDPC decoding is implemented across heterogeneous compute platforms. We detail how each platform organizes computation, manages memory, and interfaces with accelerators, highlighting differences in threading, offload mechanisms, and execution models.

### 3.1 DecodeCPU-FlexRAN

The DecodeCPU-FlexRAN platform represents a CPU-based inline LDPC decoding environment built upon Intel’s FlexRAN framework illustrated in Fig. 1(a). To compile and execute the FlexRAN FEC Software Development Kit (SDK) [12], several Intel toolchains are required, including the `oneAPI Base Toolkit` and the `HPC Toolkit`. Depending on the operating system kernel version and system environment, either the Intel `icc` (classic) or `icx` (LLVM-based) compiler must be used to ensure compatibility with the vectorized kernels in the FEC SDK. These compilers enable efficient vectorization through Intel’s Advanced Vector Extensions (AVX), which form the foundation of the high-performance LDPC decoder in FlexRAN. The FlexRAN LDPC decoder provides multiple architecture-specific implementations optimized for AVX2/AVX512.

Since the CPUs in our testbed support AVX512F, we focus on AVX512, which leverages wide-lane SIMD parallelism to accelerate the layered belief-propagation decoding steps. Each decoding thread operates on 512-bit vector registers, processing 32 LLR values concurrently per iteration. Parallelism at the task level is straightforward for CPU-based LDPC processing. The smallest schedulable unit is a TB, and to fully utilize available CPU resources, multiple TBs can be decoded in parallel across independent cores. In

practice, decoding  $N_{tb}$  transport blocks concurrently maps to  $N_{tb}$  physical CPU cores (or hyperthreads), enabling near-linear scaling until memory bandwidth saturation. This design serves as a baseline reference for comparison with hardware-accelerated decoding platforms, described next.

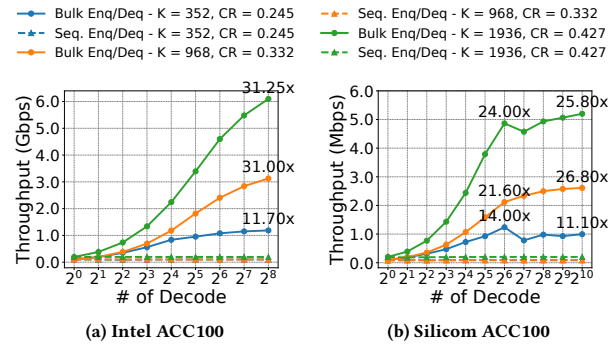
### 3.2 DecodeGPU-Aerial

While DecodeCPU-FlexRAN demonstrates the flexibility of CPU-based baseband processing, its performance is fundamentally limited. To improve the processing latency, recent research has explored offloading computationally intensive 5G NR L1 tasks, such as LDPC decoding, to GPUs. Earlier implementations, such as the CUDA-based LDPC decoder in OAI [20] follow a *lookaside* model, where decoded data must traverse the host-device boundary for each transport block. Although this approach improves raw decoding throughput, frequent PCIe transfers and high GPU power draw introduce considerable latency and energy overheads. To address these challenges, NVIDIA’s Aerial platform re-architects the entire L1 processing pipeline as a fully GPU-resident data path. This *inline* acceleration enables Aerial to execute all L1 functions—from FFT and equalization to LDPC decoding—directly within GPU, minimizing data movement and realizing a practical pathway toward GPU-based vRAN deployments.

NVIDIA Aerial is a cloud-native, GPU-accelerated vRAN software stack designed to execute the full 5G NR baseband pipeline, encompassing L1 and portions of MAC layer (L2) functions directly on the GPU. Its core components include: (i) **cuPHY**, which implements 5G NR L1 primitives such as FFT/IFFT, channel estimation, MIMO equalization, and LDPC encoding/decoding; (ii) **cuMAC**, which provides scheduling and control integration between MAC and L1 layers; and (iii) **pyAerial**, a set of Python APIs for rapid prototyping and testing of GPU-based signal-processing pipelines. Together, these modules provide a unified runtime, where baseband workloads are expressed as CUDA kernels operating over shared GPU memory. Aerial adopts an *inline acceleration* model in which baseband data is transferred directly from the NIC or radio front-end into GPU memory (HWA memory) using GPUNetIO via GPUDirect RDMA [17]. Once in device memory, the entire L1 chain—from OFDM demodulation to LDPC decoding and rate matching—is executed within the GPU context without host intervention. This eliminates intermediate PCIe copies and enables microsecond-level kernel scheduling through CUDA streams and graphs. Fig. 1(a) illustrates this architecture, where the entire data path remains fully in the GPU, and only decoded transport blocks are returned to the host for higher-layer processing.

### 3.3 DecodeASIC-ACC100

DecodeASIC-ACC100 integrates a hardware LDPC decoder using Intel’s ACC100 eASIC, which is dedicated to FEC acceleration in vRAN deployments and supports multiple functions, including LDPC encoding and decoding for 5G NR. This design follows the *lookaside acceleration* model, as illustrated in Fig. 1(b), where data is explicitly transferred between the host CPU and the accelerator through system memory. Access to the accelerator is managed through the Data Plane Development Kit (DPDK) [28], an open-source framework for user-space packet and data-plane processing. As a baseband device,



**Figure 2: Throughput comparison between sequential and bulk LDPC enqueue and dequeue pipelines on the Intel and Silicom ACC100 platforms.**

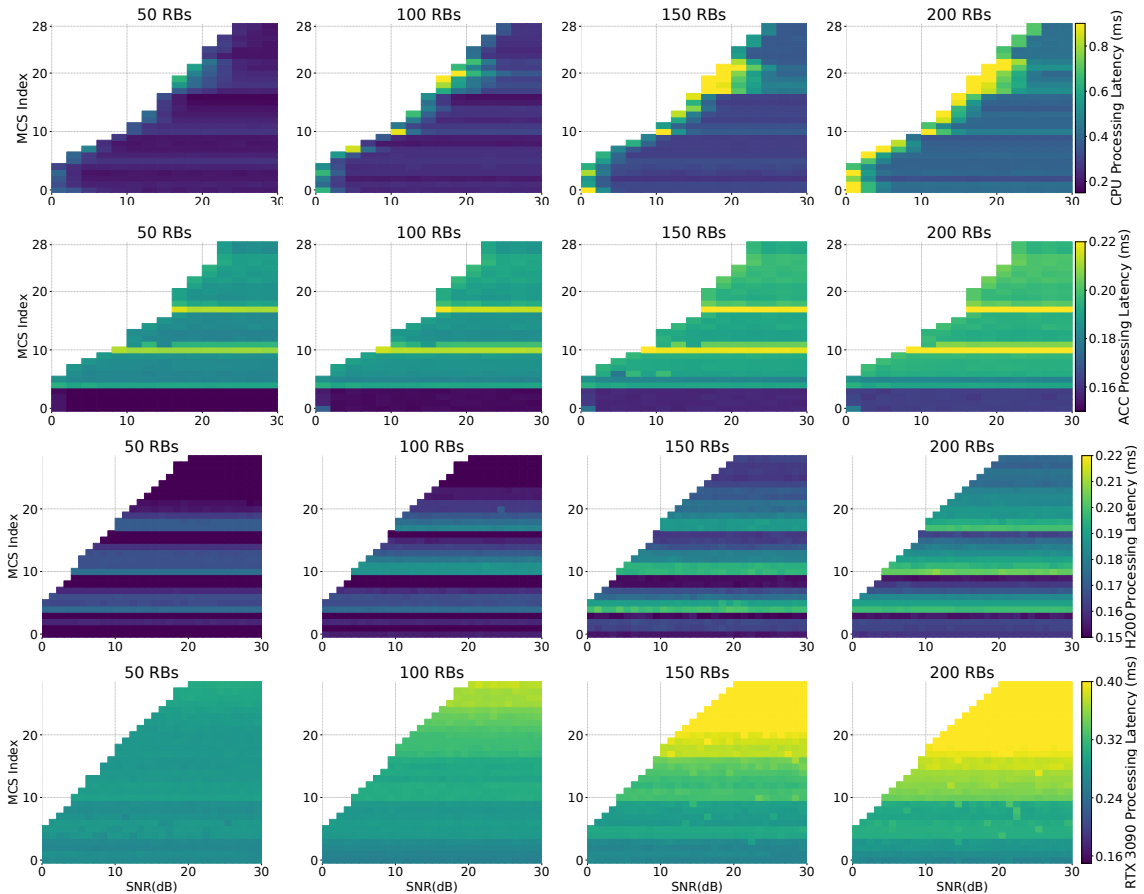
the ACC100 [26] exposes a polling-mode driver (PMD) interface through DPDK, enabling low-latency interaction between user-space applications and the hardware accelerator. In this section, we briefly describe how DecodeASIC-ACC100 interacts with the ACC100 through DPDK’s `rte_bbdev` APIs, configure LDPC decoding parameters, and performance profiling.

**Driver initialization and configuration.** The ACC100 driver functions form a subset of DPDK’s `rte_bbdev` API (“baseband device” abstraction). The high-level workflow for initializing and executing LDPC decoding via DPDK includes device setup, operation configuration, and runtime execution. During device setup, the application detects the accelerator using `rte_bbdev_info_get()` and initializes communication queues via `rte_bbdev_setup_queues()` and `rte_bbdev_queue_configure()`. The device is then activated by `rte_bbdev_start()`. Prior to execution, input and output buffers are allocated, and operation descriptors are prepared, where each operation structure specifies the buffer pointers and decoding parameters (see §2). Once operations are configured, they are enqueued to the device through `rte_bbdev_enqueue_ops()`, and completion is monitored via polling using `rte_bbdev_dequeue_ops()`. A successful dequeue indicates the completion of decoding for one TB.

**Runtime integration.** At runtime, each OFDM symbol following demod (i.e., the log-likelihood ratio output, or LLR outputs) is reorganized into an `rte_bbdev_op_data` structure following the BBDev specification, and transferred to the ACC100 device via `rte_memcpy()`. To maximize parallelism, DecodeASIC-ACC100 dispatches each demodulated symbol immediately after demod using `rte_bbdev_enqueue_ops()`, and once all enqueue operations are issued and completed, bulk dequeue requests are triggered using `rte_bbdev_dequeue_ops()` to collect decoded results for the entire frame. This *parallel enqueue-dequeue strategy* enables continuous data flow, allowing subsequent software tasks (e.g., fft for the next OFDM symbol) to execute concurrently with ongoing hardware decoding. As shown in Fig. 2, fully exploiting the ACC100’s intrinsic parallelism through bulk enqueue/dequeue improves decoding throughput by up to 30× compared to sequential processing, and the same trend exists across both Intel and Silicom ACC100 variants.

### 3.4 DecodeGPU-SionnaRK

The NVIDIA Jetson Orin AGX [18] platform represents a class of edge AI systems-on-chip (SoCs) that integrate CPU, GPU, and



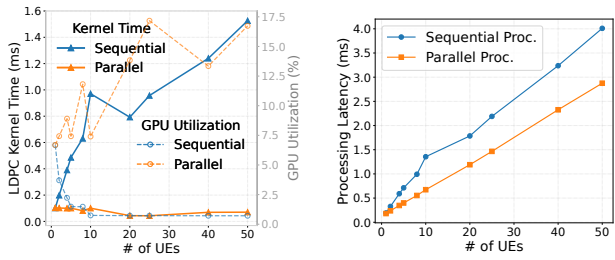
**Figure 3: LDPC decoding latency across heterogeneous platforms under varying L1 configurations.** Each heatmap shows per-TB decoding latency (ms) as a function of MCS index and SNR level, for four PRB allocations: (*1st row*) CPU (FlexRAN, AVX512); (*2nd row*) ACC100 lookaside acceleration; (*3rd row*) H200 decoding via NVIDIA’s Aerial; and (*4th row*) RTX 3090 decoding via NVIDIA’s Aerial.

memory subsystems within a unified architecture. Unlike discrete accelerator platforms that rely on PCIe-based data transfers, Jetson adopts a *unified memory* design in which the CPU and GPU share a single physical DRAM space managed by a coherent memory controller. This design eliminates explicit host-device data copies and allows intermediate baseband data to be accessed directly by both CPU and GPU threads. From a memory management perspective, it behaves as a hybrid between inline and lookaside architectures, as illustrated in the bottom panel of Fig. 1. While data still moves between CPU and GPU contexts, the overhead is significantly reduced compared to traditional split-memory systems since both processors operate on the same physical memory region. This property enables efficient acceleration for baseband processing tasks. DecodeGPU-SionnaRK executes LDPC decoding through the Sionna-RK [19] framework, which provides a TensorFlow- and CUDA-based implementation of the 5G NR LDPC chain optimized for embedded GPU platforms. By leveraging CUDA streams and unified memory, LDPC decoding can be performed *hybrid* with other baseband kernels—such as FFT, channel equalization, and demodulation—without intermediate buffer copies. As depicted in Fig. 1(c), demodulated LLRs are written directly into unified memory

and immediately consumed by the LDPC decoder kernel, forming a continuous GPU-resident data path.

## 4 Evaluation

**Implementation.** We run the FlexRAN FEC SDK on a dual-socket Intel Xeon Gold 6348 server with 112 hardware threads (2.6 GHz, AVX512-enabled), equipped with 84 MB L3 cache. The system runs Ubuntu 22.04 (kernel 6.8.0) with Intel oneAPI Base and HPC toolkits and the `icx` compiler for AVX512-optimized LDPC kernels. This platform serves as the CPU baseline. For full inline acceleration, we also perform LDPC unit tests using NVIDIA Aerial (v25.2-cuBB) on a dual-socket AMD EPYC 9355 server (128 hardware threads, AVX512-enabled) equipped with four NVIDIA H200 NVL GPUs running Ubuntu 24.04 (kernel 6.8.0). To broaden our evaluation, we extend the analysis to additional discrete GPUs, including the NVIDIA RTX 3090, enabling cross-generation comparison of decoding latency and kernel efficiency. The same server hosts two LDPC ASIC accelerators: one Intel ACC100 and one Silicom-branded equivalent device, both accessed through the DPDK BBDev drivers. We use an NVIDIA Jetson Orin AGX edge platform running Ubuntu 22.04 (kernel 5.15) with CUDA-accelerated LDPC decoding through the




**Figure 4: PyAerial LDPC decoding performance under sequential and parallel processing. (a) GPU kernel execution time and utilization, excluding host-device data movement. (b) Overall LDPC decoding latency, including CPU orchestration and data movement.**

Sionna-RK (v1.2.1) toolkit. This platform executes all LDPC decoding kernels directly in unified system memory.

**LDPC decoding latency.** Fig. 3 summarizes the LDPC decoding latency across the CPU-, GPU-, and ACC100-based platforms for varying PRB sizes, MCS levels, and SNR values. We exclude Jetson Orin AGX results from this comparison since its workstation-class CPU and GPU resources are not directly comparable to server-grade acceleration environments. We observe that both ACC100 and GPU acceleration substantially reduce decoding time compared to CPU-only execution. On the CPU baseline, latency varies significantly with SNR, as poor channel conditions increase the number of required decoding iterations. In contrast, ACC100 and GPU implementations exhibit minimal SNR sensitivity.

Distinct scaling behaviors emerge across platforms. CPU decoding latency increases with PRB size, while ACC100 latency remains nearly constant from 50 to 200 PRBs, demonstrating that the accelerator’s deeply pipelined architecture maintains throughput even under higher data volumes. H200 GPU decoding via NVIDIA Aerial follows a similar trend, although PCIe orchestration overhead becomes more pronounced at smaller workloads. Across both accelerators, specific MCSs (notably MCS-10 and MCS-17) incur higher latency, as lower code rates introduce greater redundancy and decoding complexity. Overall, the latency profiles of ACC100 and H200 GPU align closely under sequential processing. With other GPU models, we observe that the RTX 3090 exhibits latency behavior intermediate between the CPU and H200 platforms; its performance remains MCS-dependent and consistently higher than that of the H200, underscoring the value of the DecodeX benchmark in identifying software-hardware mismatches and guiding future vRAN deployments across heterogeneous GPU architectures.

**Exploring GPU parallelism.** To exploit the parallelism on GPU, Fig. 4 compares DecodeGPU-Aerial performance under sequential and parallel LDPC decoding. In sequential processing, the data stream for each UE is decoded independently, one TB at a time; while parallel processing decodes  $n$  TBs simultaneously across multiple CUDA grids and blocks. For fairness, the total data volume remains constant across both settings; for instance, with 10 UEs and 200 PRBs assigned, sequential processing assigns 20 PRBs per UE and decodes them sequentially, whereas parallel processing decodes all PRBs assigned to the 10 UEs concurrently. When considering only the GPU kernel execution time in Fig. 4(a), parallel decoding achieves up to 20 $\times$  smaller latency than sequential decoding but with increased GPU utilization. Sequential processing maintains



**Figure 5: LDPC decoding latency on DecodeGPU-SionnaRK across different iteration counts, with varying number of information bits ( $K$ ) and code rates (CR).**

a nearly constant utilization regardless of UE count, while parallel execution scales utilization linearly with the number of active streams. This observation highlights that host-device data transfer and orchestration costs dominate the overall latency, reinforcing the advantage of *inline processing* designs such as Aerial, where data remains resident on the GPU to minimize transfer overhead.

## 5 Discussions

**DecodeGPU-SionnaRK on Jetson Orin AGX.** We also present profiling results for DecodeGPU-SionnaRK as shown in Fig. 5. Since the current implementation does not yet support the full MCS set with variable code rates, we evaluate three representative cases:  $K = 1,936$  and  $K = 4,224$  with a code rate of 0.33, and  $K = 8,440$  with a code rate 0.88. Here,  $K$  denotes the number of information bits per TB. The overall results resemble a similar trend observed in the CPU-based decoding results (see Fig. 3): larger transport block sizes lead to higher processing latency, while varying the amount of redundancy (code rate) has a relatively smaller impact; and the processing latency also increases with the iteration count, which corresponds to lower-SNR decoding conditions. Overall, Jetson Orin AGX exhibits similar scaling characteristics to CPU decoding, as both the embedded GPU and CPU in this research-oriented platform have limited computational throughput compared to discrete accelerators such as the ACC100 or datacenter-grade GPUs.

**Optimizing GPU-based L1 processing.** Our preliminary investigation of DecodeGPU reveals that further optimization can be achieved by refining how GPU kernels and CUDA streams are organized across the L1 processing pipeline [5, 14, 31]. Future optimization may focus on fusing adjacent kernels or reusing persistent kernels to reduce kernel bring-up latency. Another direction is to improve stream-level parallelism by dynamically allocating multiple CUDA streams per group of subcarriers, thereby increasing concurrency across OFDM symbols and fully exploiting the subcarrier-level parallelism inherent in the 5G NR waveform. Additionally, exploring more fine-grained task scheduling within the GPU’s streaming multiprocessors could help maintain high utilization under variable TB sizes and traffic conditions.

**Energy-aware compute resource allocation.** Allocation of heterogeneous compute resources must account not only for RAN throughput and latency, but also for the energy consumed by computation and data movement, where the latter often dominates under low compute utilization. Our recent work *Nexus* [22] explored dynamic and energy-aware baseband scheduling across CPU and ACC100 accelerators, demonstrating that joint optimization of

computation and data-transfer pipelines can significantly improve energy proportionality in heterogeneous vRAN systems. Complementary to this direction, prior work [4] investigates adaptive resource provisioning in large-scale radio deployments for balancing latency and energy efficiency. Unlike CPUs or dedicated accelerators, GPUs exhibit highly dynamic power behavior that depends on kernel concurrency, memory access intensity, and streaming multiprocessor utilization. Future work may explore approaches that map kernel configurations (e.g., MCS, iteration count) to power consumption [13] to enable adaptive and energy-aware scheduling.

**Coexistence with ML workloads.** A frequently raised concern pertains to the high capital cost and energy efficiency of GPUs when dedicated solely to baseband processing. However, such concerns can be mitigated if the same hardware is shared dynamically with other compute-intensive tasks when the traffic load is low. Recent work [8] also advocates for a converged AI-cellular architecture, motivating the co-location of communication and computation workloads to support industrial-scale AI applications. For instance, during non-URLLC periods or under light traffic conditions, GPU resources used for L1 tasks can be time- or spatially-partitioned to support large ML model training and/or inference workloads. This flexible multi-tenancy model highlights the potential of GPU-based architectures to unify communication, sensing [7, 11, 30], and computation workloads in the vRAN context.

## 6 Conclusions

We presented DecodeX, a unified benchmarking framework that systematically evaluates LDPC decoding performance across heterogeneous computing platforms, including CPUs, GPUs, and ASICs. Through careful implementation and detailed profiling, we highlighted how different system architectures organize their L1 processing pipelines, and how factors such as kernel launch overhead, data-movement orchestration shape overall decoding latency and determinism. As an open and extensible benchmark, DecodeX provides a foundation for future research on cross-platform baseband processing, facilitating energy-aware compute resource management for heterogeneous vRAN systems.

## Acknowledgments

This work was supported in part by NSF grants CNS-2211944, ECCS-2434131, CNS-2443137, CNS-2450567, and OAC-2503010, and NVIDIA Academic Grants. This work was also supported in part by the Center for Ubiquitous Connectivity (CUbIC), sponsored by Semiconductor Research Corporation (SRC) and Defense Advanced Research Projects Agency (DARPA) under the JUMP 2.0 program.

## References

- [1] 2025. DecodeX: Exploring and benchmarking of LDPC decoding across CPU, GPU, and ASIC platforms. <https://github.com/functions-lab/DecodeX>.
- [2] 3GPP. 2019. 5G; NR; *Physical Layer Procedures for Control*. Technical Specification (TS) 38.213. 3rd Generation Partnership Project (3GPP). [https://www.etsi.org/deliver/etsi\\_ts/138200\\_138299/138213/15.06.00\\_60/ts\\_138213v150600p.pdf](https://www.etsi.org/deliver/etsi_ts/138200_138299/138213/15.06.00_60/ts_138213v150600p.pdf) Version 15.6.0.
- [3] 3GPP. 2020. 5G; NR; *Multiplexing and channel coding*. Technical Specification (TS) 38.212. 3rd Generation Partnership Project (3GPP). [https://www.etsi.org/deliver/etsi\\_ts/138200\\_138299/138212/16.02.00\\_60/ts\\_138212v160200p.pdf](https://www.etsi.org/deliver/etsi_ts/138200_138299/138212/16.02.00_60/ts_138212v160200p.pdf) Version 16.2.0.
- [4] Parthiban Annamalai, Minsung Kim, Anurag Khandelwal, and Lin Zhong. 2025. Exploring intelligent dynamic resource provisioning for elastic massive MIMO vRAN. In *Proc. ACM HotMobile'25*.
- [5] Paramvir Bahl, Matthew Balkwill, Xenofon Foukas, Anuj Kalia, Daehyeok Kim, Manikanta Kotaru, Zhihua Lai, Sanjeev Mehrotra, Bozidar Radunovic, Stefan Saroiu, et al. 2023. Accelerating Open RAN research through an enterprise-scale 5G testbed. In *Proc. ACM MobiCom'23*.
- [6] NVIDIA Corporation. 2025. *NVIDIA Aerial™ CUDA-Accelerated RAN*. <https://github.com/NVIDIA/aerial-cuda-accelerated-ran>
- [7] Xinran Fang, Wei Feng, Yunfei Chen, Ning Ge, and Yan Zhang. 2022. Joint communication and sensing toward 6G: Models and potential of using MIMO. *IEEE Internet of Things Journal* 10, 5 (2022), 4093–4116.
- [8] Xenofon Foukas and Bozidar Radunovic. 2025. The future of the industrial AI edge is cellular. In *Proc. ACM HotMobile'25*.
- [9] Zhihui Gao, Zhenzhou Qi, and Tingjun Chen. 2024. Mambas: Maneuvering analog multi-user beamforming using an array of subarrays in mmWave networks. In *Proc. ACM MobiCom'24*.
- [10] Junzhi Gong, Anuj Kalia, and Minlan Yu. 2023. Scalable distributed massive MIMO baseband processing. In *Proc. USENIX NSDI'23*.
- [11] David Hunt, Kristen Angell, Zhenzhou Qi, Tingjun Chen, and Miroslav Pajic. 2024. MadRadar: A black-box physical layer attack framework on mmWave automotive FMCW radars. In *Proc. ISOC NDSS'24*.
- [12] Intel. 2024. FlexRAN-FEC-SDK-Modules. <https://github.com/intel/FlexRAN-FEC-SDK-Modules>.
- [13] Lopamudra Kundu, Xingqin Lin, and Rajesh Gadiyar. 2025. Toward energy efficient RAN: From industry standards to trending practice. *IEEE Wireless Communications* 32, 1 (2025), 36–43.
- [14] Xingqin Lin, Lopamudra Kundu, Sebastian Cammerer, Yan Huang, Chris Dick, Charles Santhosam, Rajesh Gadiyar, Reinhard Wiesmayr, and Christoph Studer. 2025. AI-native 6G: Empowering intelligent RAN with accelerated compute. *IEEE Wireless Communications* 32, 6 (2025), 11–14.
- [15] Xingqin Lin, Lopamudra Kundu, Chris Dick, and Soma Velayutham. 2023. Embracing AI in 5G-advanced toward 6G: A joint 3GPP and O-RAN perspective. *IEEE Communications Standards Magazine* 7, 4 (2023), 76–83.
- [16] Xingqin Lin and Namyoon Lee. 2021. *5G and Beyond: Fundamentals and Standards* (1st ed.). Springer, Cham, Switzerland.
- [17] NVIDIA. 2025. DOCA Documentation. <https://docs.nvidia.com/doca/sdk/doca-gpunetio/index.html>.
- [18] NVIDIA. 2025. NVIDIA Jetson Orin. <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/>.
- [19] NVIDIA. 2025. Sionna Research Kit (RK). <https://nvlabs.github.io/sionna/rk/quickstart.html>.
- [20] OpenAirInterface. 2025. CUDA LDPC in OAI. [https://gitlab.eurecom.fr/oai/openairinterface5g/-/tree/develop/openair1/PHY/CODING/nrLDPC\\_decoder\\_LYC?ref\\_type=heads](https://gitlab.eurecom.fr/oai/openairinterface5g/-/tree/develop/openair1/PHY/CODING/nrLDPC_decoder_LYC?ref_type=heads).
- [21] Zhenzhou Qi, Zhihui Gao, Chung-Hsuan Tung, and Tingjun Chen. 2023. Programmable millimeter-wave MIMO radios with real-time baseband processing. In *ACM WINTeCH'23*.
- [22] Zhenzhou Qi, Chung-Hsuan Tung, Zhihui Gao, and Tingjun Chen. 2025. NEXUS: Efficient and scalable multi-cell mmWave baseband processing with heterogeneous compute. *arXiv preprint arXiv:2509.04625* (2025).
- [23] Zhenzhou Qi, Chung-Hsuan Tung, Anuj Kalia, and Tingjun Chen. 2024. Savannah: Efficient mmWave baseband processing with minimal and heterogeneous resources. In *Proc. ACM MobiCom'24*.
- [24] Leonardo Lo Schiavo, Jose A Ayala-Romero, Andres Garcia-Saavedra, Marco Fiore, and Xavier Costa-Perez. 2024. Yinyangran: Resource multiplexing in GPU-accelerated virtualized RANs. In *Proc. IEEE INFOCOM'24*.
- [25] Leonardo Lo Schiavo, Gines Garcia-Aviles, Andres Garcia-Saavedra, Marco Gramaglia, Marco Fiore, Albert Banchs, and Xavier Costa-Perez. 2024. CloudRIC: Open radio access network (O-RAN) virtualization with shared heterogeneous computing. In *Proc. ACM MobiCom'24*.
- [26] Silicom Ltd. 2023. Lisbon P2 ACC100 FEC accelerator extended temp. <https://www.silicom-usa.com/wp-content/uploads/2023/09/Lisbon-P2-ACC100-FEC-Accelerator-Extended-temp-1v1.pdf>.
- [27] Chance Tarver, Matthew Tonnemacher, Hao Chen, Jianzhong Charlie Zhang, and Joseph R Cavallaro. 2020. GPU-based LDPC decoding for vRAN systems in 5G and beyond. In *Proc. IEEE ISCAS'20*.
- [28] The Linux Foundation. 2023. Data plane development kit (DPDK). <https://doc.dpdk.org/guides-21.11/tools/testbbdev.html>.
- [29] Davide Villa, Imran Khan, Florian Kaltenberger, Nicholas Hedberg, Ruben Soares da Silva, Stefano Maxenti, Leonardo Bonati, Anupa Kelkar, Chris Dick, Eduardo Baena, et al. 2025. X5G: An open, programmable, multi-vendor, end-to-end, private 5G O-RAN testbed with NVIDIA ARC and OpenAirInterface. *IEEE Transactions on Mobile Computing* (2025).
- [30] Zehao Wang, Yue-Kai Huang, Ezra Ip, Zhenzhou Qi, Gil Zussman, Dan Kilper, Koji Asahi, Hideo Kageshima, Yoshiaki Aono, and Tingjun Chen. 2023. Field trial of coexistence and simultaneous switching of real-time fiber sensing and coherent 400 GBe in a dense urban environment. *JIT* 42, 4 (2023), 1304–1311.
- [31] Xincheng Xie, Wentao Hou, Zerui Guo, and Ming Liu. 2025. Building massive MIMO baseband processing on a single-node supercomputer. In *Proc. USENIX NSDI'25*.