

Compiler-Assisted Speculative Sampling for Accelerated LLM Inference on Heterogeneous Edge Devices

Alejandro Ruiz y Mesa^{*†}, Guilherme Korol[‡], Moritz Riesterer[‡], João Paulo C. de Lima^{†§} Jeronimo Castrillon^{†§}

[†] Dresden University of Technology, Dresden, Germany

[‡]NXP Semiconductors, Munich, Germany

[§]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden, Germany

[†]alejandro.ruiz_y_mesa@mailbox.tu-dresden.de, {joao.lima, jeronimo.castrillon}@tu-dresden.de,

[‡]{guilherme.korol, moritz.riesterer}@nxp.com

Abstract—LLM deployment on resource-constrained edge devices faces severe latency constraints, particularly in real-time applications where delayed responses can compromise safety or usability. Among many approaches to mitigate the inefficiencies of sequential token-by-token generation, Speculative Decoding (SD) has emerged as a promising technique. However, SD at the edge is hindered by two major challenges: (1) integrating SD into a compiler-based workflow without sacrificing performance or programmability, and (2) exploiting the heterogeneous compute resources of modern SoCs through carefully designed partitioning strategies. This work addresses these challenges by using an analytical cost model that explores heterogeneous hardware configurations and guides coarse-grained partitioning of LLM subgraphs, particularly with edge-typical short input sequence lengths. The cost model predicts when speculative sampling and heterogeneous execution are jointly beneficial and is validated on an edge device featuring a hexacore Cortex-A CPU and a Mali GPU, revealing up to $1.68\times$ speedup for translation tasks, closely matching analytic expectations.

Index Terms—Edge computing, Large language model, Speculative sampling, Heterogeneous execution, IREE, SoC

I. INTRODUCTION

Enabling Large Language Models (LLMs) on edge devices is essential for applications requiring privacy, robustness, and low latency, such as offline health assistants or real-time translation [2]. However, deploying LLMs on edge platforms faces severe challenges under the scarce compute, power, and memory available. Modern System-on-Chip (SoC) platforms integrate multiple heterogeneous Processing Units (PUs), such as multi-core CPUs, mobile GPUs, and accelerators, creating opportunities for workload acceleration. Still, transformer-based LLMs present a particularly challenging workload due to their autoregressive decoding, which generates one token at each forward pass and imposes hard sequential dependencies. Addressing these challenges, therefore, requires a combination of algorithmic-, system-, and hardware-level optimization techniques.

Speculative Decoding (SD), in particular, has demonstrated remarkable improvements in the generative pipeline on high-

end GPUs, achieving speedups up to $6.5\times$ compared to traditional autoregressive decoding [3], [4]. SD converts savings from traditional algorithmic optimizations (e.g., quantization, pruning, and distillation) into end-to-end latency reductions by relying on an inexpensive draft model to reduce the number of costly target-model decoding steps, with performance determined by the draft cost, verification cost, and acceptance rate. This approach naturally lends itself to migrating LLM inference from cloud to edge platforms.

A key question is how to integrate SD into the software stack of an edge system in a way that preserves productivity for developers while enabling meaningful hardware-aware optimization. Without such careful integration, SD is not inherently beneficial: naive adoption or poor hardware mapping can negate its advantages and even increase end-to-end latency. In this work, we build upon IREE [6], an open-source ML compiler to provide a unified representation of SD that supports explicit device affinities, coarse-grained heterogeneous mapping, and Ahead-of-Time (AOT) compilation across multiple backends. On top of that, we leverage the analytical cost model from [3] to guide LLM partitioning and mapping across heterogeneous PUs. Concretely, our contributions are the following:

- We apply an analytical cost model to (i) determine when to use speculative sampling and to (ii) map it onto heterogeneous PUs for efficient LLM inference acceleration;
- We present compiler-level abstractions of speculative sampling that enable device placement decisions for heterogeneous edge execution, comparing monolithic and modular compilation strategies;
- We demonstrate that our optimized heterogeneous mappings can outperform homogeneous CPU execution by up to $1.68\times$ by reducing speculation overhead;
- We validate the proposed cost model using a hexacore Cortex-A55 CPU and a Mali-G310 GPU integrated on an NXP i.MX95 SoC, for translation tasks, with a 4% deviation from analytical expectations.

^{*}Work realized during internship at NXP Semiconductors.

II. BACKGROUND AND RELATED WORK

This section builds on insights into computational bottlenecks inherent in LLM inference, SD as an algorithmic optimization technique, and the role of ML compilers in enabling efficient deployment across diverse hardware platforms. Toward the end, we review related work on accelerating neural network inference and highlight how our approach differs from existing techniques.

A. LLM Inference Bottlenecks

Inference in transformer-based LLMs proceeds in two phases, each presenting distinct computational challenges [7]. The prefill phase processes the entire input sequence in parallel, populating the Key-Value (KV) cache across all transformer layers. In contrast, the decode phase generates tokens sequentially, with each token depending on the previous model state. As such, opportunities for parallelism are limited. Furthermore, the hidden dimension of the LLM, denoted as d , plays a pivotal role in shaping the computational characteristics of the model [1]. The sequence length (S_L) is often categorized relative to d : short sequences ($S_L \ll d$) are dominated by the linear layers, whereas the attention layers dominate in long sequences ($S_L \gg d$). This distinction is crucial for understanding performance bottlenecks.

B. Speculative Decoding

SD accelerates autoregressive generation by reducing the number of full forward passes needed from the LLM (referred to as the target model). The process is conceptually divided into two phases. In the *speculation phase*, the system generates a sequence of candidate tokens using a computationally inexpensive mechanism. This may be a smaller transformer, a distilled or pruned version of the original network, or a specialized predictive module, and is designed to approximate the next-token distribution of the target model at a significantly lower computational cost. In the subsequent *verification phase*, the larger target model evaluates the drafted tokens in a single forward pass, in a parallel way similar to the prefill phase of the LLM inference. Speculated tokens are accepted or rejected according to a selection criterion. Fig. 1 illustrates standard incremental decoding (top) and two speculative strategies: sequence-based (middle) and tree-based (bottom). Orange segments indicate speculation by a lightweight model, while green segments denote verification by the target model.

Speculative sampling [3] is a form of SD that avoids the training of a drafter mechanism, since it employs a smaller transformer model as the speculation strategy. The structural similarity between drafter and target yields sufficiently correlated logits to produce meaningful acceptance rates, enabling substantial speedups on server hardware. On edge devices, however, acceptance behavior is more complex. Quantization alters the relative probabilities of the vocabulary distribution, resulting in acceptance rates lower than those observed with full-precision models [8].

In speculative sampling, the achieved speedup depends not only on the alignment between the drafter and target

model distributions, but also on the workload, the task, and the hardware-software environment in which the accelerated LLM is executed. A formalism modeling the hardware impact on the speculative sampling speedup is developed in [3]. It presents the speedup S as a function of the acceptance rate's expected value (α , the mean proportion of accepted tokens), the speculated draft length (γ , the sequence length of the speculated tokens), and the cost coefficient $c = t_{\text{draft}}/t_{\text{target}}$, a hardware- and software-dependent parameter representing the ratio between the latency of a single drafter forward pass (t_{draft}) and that of the target model (t_{target}). The speedup S is given by (1).

$$S(\alpha, \gamma, c) = \frac{1 - \alpha^{\gamma+1}}{(1 - \alpha)(\gamma c + 1)} \quad (1)$$

Moreover, the condition $c < \alpha$ must hold to achieve any speedup at all. This guarantees that there is at least one value of γ that yields a speedup greater than 1. Furthermore, the optimal draft length γ^* that maximizes the speedup depends on both the hardware configuration and the quality of the drafting mechanism [3].

C. Related Work

Traditional algorithmic optimizations (e.g., quantization, pruning, and knowledge distillation) reduce compute and memory requirements, effectively producing a cheaper draft model and/or improving draft-target agreement. Recent systems (e.g., Sequoia [16], DuoDecoding [17], and Dove-tail [18]) further show that the optimal draft/verifier configuration depends on hardware characteristics and device heterogeneity, motivating joint algorithm-system co-design. However, these approaches primarily target high-end GPU platforms rather than edge-grade SoCs. Complementary work on mapping neural workloads onto heterogeneous compute resources, including polyhedral optimization and design-space exploration, appears in frameworks like Tiramisu [19], ZigZag [20], and FlexInfer [21]. Related efforts on the edge explore algorithm-hardware co-design on reconfigurable plat-

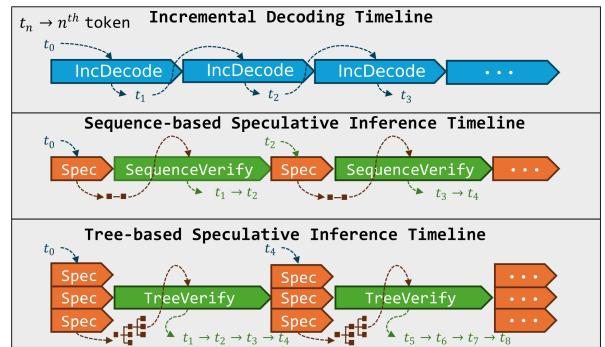


Fig. 1: Three generative pipelines: standard sampling (top, in blue), and two variants of SD: sequential drafting (center) and tree-based drafting (bottom). Adapted from [15].

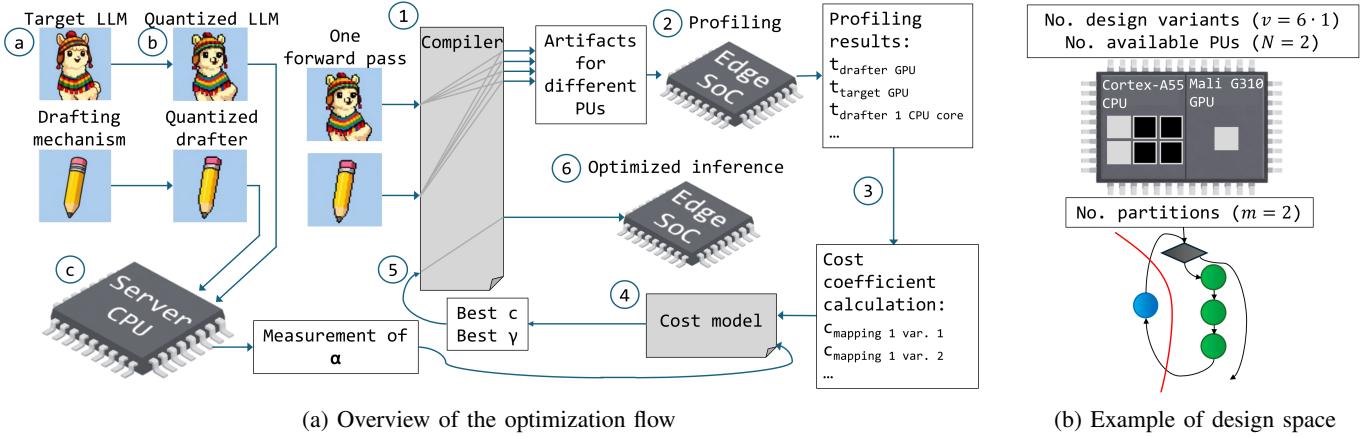


Fig. 2: Overview of heterogeneous mapping workflow for speculative sampling on edge devices.

forms, combining algorithmic optimizations, such as pruning and early-exit [10], [11].

Neural network multi-tenancy represents another relevant research direction. Frameworks like MAGMA [22] and Adyna [23] address how concurrent models contend for limited compute and memory resources. Such systems reveal structural similarities to speculative sampling pipelines, where drafter and target models form a multi-tenant workload sharing limited edge resources. However, these works focus on the orchestration of independent models, rather than coupled models within a pipeline. Furthermore, they evaluate against accelerator models or simulators rather than on silicon, as done in this work.

Our work differs since (1) we target an edge SoC rather than high-end GPU platforms, addressing the unique constraints of resource-limited devices; (2) we provide the first exploration of coarse-grained heterogeneous partitioning strategies for speculative sampling, leveraging high-level compiler frontend abstractions to express device placement decisions; (3) we apply an analytical cost model that jointly optimizes algorithmic acceleration (speculative sampling) and hardware mapping decisions for heterogeneous edge execution; and (4) we validate our approach on real silicon rather than relying on simulators, providing empirical evidence of the method accuracy.

III. HETEROGENEOUS MAPPING FRAMEWORK FOR SPECULATIVE SAMPLING

This section presents a systematic approach to evaluate mappings of LLMs with speculative sampling on edge devices. It begins by providing an overview of the complete compilation and optimization workflow in Sec. III-A, followed by the formulation of the heterogeneous mapping problem as a Design Space Exploration (DSE) in the Sec. III-B and III-C. Sec. III-D details how abstractions can be aligned with hardware-aware execution to preserve productivity while enabling good performance.

A. Overview of the Compilation and Optimization Workflow

Our approach combines offline profiling, analytical modeling, and compiler-assisted code generation (Fig. 2a). We evaluate speculative sampling [3] for its training-free nature, though the workflow is generalizable to other sequential SD techniques. Our workflow consists of the following steps:

a) *Offline quantization* (Sec. III-C): Model quantization is performed offline. In Fig. 2a, from ① to ⑤ the target model and the drafter model are quantized with different schemes that match the available arithmetic on the target SoC. In ⑥, the acceptance rate α is measured for the different combinations of quantized target-drafter mechanisms.

b) *Profiling* (Sec. III-B and III-C): We compile forward passes of both models targeting all PUs ①, profile them on hardware to measure t_{draft} and t_{target} ②, and calculate the cost coefficient $c = t_{\text{draft}}/t_{\text{target}}$ for each design variant ③ (Fig. 2b).

c) *Exploration* (Sec. III-B): With α and c values in hand, the analytical cost model ④ is evaluated to determine the optimal draft length γ and device mapping for each hardware configuration (see ⑤).

d) *Compilation and Inference* (Sec. III-D): The complete SD pipeline is compiled and executed using IREE in ⑥.

B. Design Space Encoding, Exploration, and Evaluation

We formulate the heterogeneous execution of an LLM with speculative sampling as a DSE problem over static spatial mappings of the computational graph: each subgraph is assigned ahead of time to a PU and no dynamic remapping is considered. Additionally, the search concentrates on coarse-grained partitions that separate the drafter (speculation mechanism model) from the target model (LLM to be accelerated), a modeling choice motivated by the fact that the cost coefficient c (the cost of the speculation phase) must be lowered by, for example, decreasing the latency of the drafter model.

We express the size of the design space as $v \cdot N^m$, where v is the number of design variants, N the number of distinct PUs available for assignment, and m is the number of subgraph partitions. In this formulation, a *design variant* corresponds to a unique combination of cores, shaders, or PEs available

across all PUs, which we count by placing its unordered configurations in bijective correspondence with the Cartesian product

$$v = \prod_i n_i,$$

where n_i denotes the number of available cores, shaders, or PEs in PU_i . For example, Fig. 2b (top) depicts a CPU with six physical cores and a GPU with a single shader that yields $v = 6 \times 1 = 6$ distinct hardware configurations for mapping subgraphs. In practice, only a subset of these resources may be accessible at runtime, for example two active CPU cores as shown in gray. Moreover, for the configuration illustrated in Fig. 2b, assuming two PUs available for mapping ($N = 2$) (top) and two graph partitions ($m = 2$, one for the drafter and one for the target, bottom), the size of the design space becomes $v = 6 \cdot 2^2 = 24$ possible mappings.

This space grows quickly if any of its factors increase: finer-grained partitioning (larger m) increases the exponent, adding more types of PUs (larger N) raises the base, and incorporating more cores, shaders, and PEs (larger v) multiplies the size of the space. In practice, therefore, exhaustive exploration becomes infeasible on richer platforms; for this reason, and to keep the study focused, we deliberately keep v , N , and m small and rely on the analytical cost model presented in (1) to guide the search.

To evaluate each mapping, we find the γ value that yields the highest acceleration S , for a given cost coefficient c , that encodes the speculation computational cost of a given mapping, and a given acceptance rate α , that encodes the drafter model speculation quality. Both α and c are measured empirically as described next.

C. Measuring Acceptance Rates and Cost Coefficients

We determine empirically the acceptance rate α using a 16-core server-grade CPU. Although α is a model-dependent value, reflecting how closely the drafter model approximates the larger target model, it remains hardware-independent¹. Concretely, to estimate α , we employed the Spec-Bench, a benchmark designed to evaluate SD performance across multiple tasks [24]. The dataset comprises 480 samples distributed among thirteen tasks.

In edge deployments, such task variability is lower because models are tailored to a single domain and typically benefit from knowledge distillation or task-specific fine-tuning. For this reason, we focus on the translation task. Additionally, due to the nature of translation, the length of the generated tokens tends to closely match the input sequence length, which is typically short (a few tens of tokens representing a sentence).

As quantization is a key enabler of ML systems on edge devices, we evaluated the acceptance rate α of multiple model configurations. These included:

- fully quantized target-drafter pairs,

¹Slight deviations may occur when different devices handle rounding and precision differently, potentially leading to numerical discrepancies or precision divergence arising from quantization.

- semi-quantized combinations, and
- unquantized FP16 counterparts as reference models.

All quantization schemes are static and implemented in w8a8 precision using the Intel Neural Compressor [25].

We compute c by compiling a single forward pass of both the drafter and target models using IREE 3.6.0, evaluating all combinations of the LLVM and SPIR-V backends independently, and benchmarking their respective runtimes. The ratio between these measured latencies directly yields c , thereby quantifying the relative computational cost of the speculation phase with respect to the target model for each homogeneous and heterogeneous configuration.

Given the measured acceptance rate α and c values, the cost model given by (1) is used to estimate the expected speedup and the optimal speculative length γ for each device pairing (homogeneous and heterogeneous). We then calculate the best-performing mapping and then validate the predicted speedup experimentally on hardware, comparing the resulting acceleration against a CPU-only, non-speculative baseline.

D. Matching Abstractions of Speculative Sampling and Heterogeneous Execution

Selecting the abstraction level for spatial partitioning onto heterogeneous PUs presents a tension: speculative sampling is most naturally expressed in high-level ML frameworks, whereas heterogeneous partitioning requires low-level control. We address this by rising hardware partitioning into the compiler frontend and integrating it directly within the speculative sampling pipeline. This unified, high-level formulation preserves the programmability of PyTorch-level code while enabling targeted acceleration (e.g., mapping the speculative phase onto an appropriate PU).

IREE, backed by Multi-Level Intermediate Representation (MLIR), enables mixing high-level dialects (e.g., Torch IR) with lower-level dialects (e.g., Flow), allowing spatial partitioning to be expressed at the PyTorch level using custom operators resolved to physical devices during lowering. This enables capturing the entire pipeline within a single optimized module (Fig. 3), reducing glue code while supporting end-to-end optimization. However, monolithic AOT compilation increases complexity, particularly for procedural logic with data-dependent control flow.

Alternatively, Fig. 4 shows a modular design where target and drafter are compiled independently while procedural logic executes in the serving mechanism. This simplifies heterogeneous execution but introduces stricter boundaries between modules, reflected by additional API calls (thick black arrows), which add runtime overhead. We evaluate both alternatives in the following section.

IV. EXPERIMENTAL SETUP AND EVALUATION

Tab. I summarizes the experimental configuration adopted in this work. It outlines the workload characteristics, task definition, optimization objective, SD strategy, and the heterogeneous hardware resources considered during evaluation. In our setup, greedy sampling is used across all experiments and

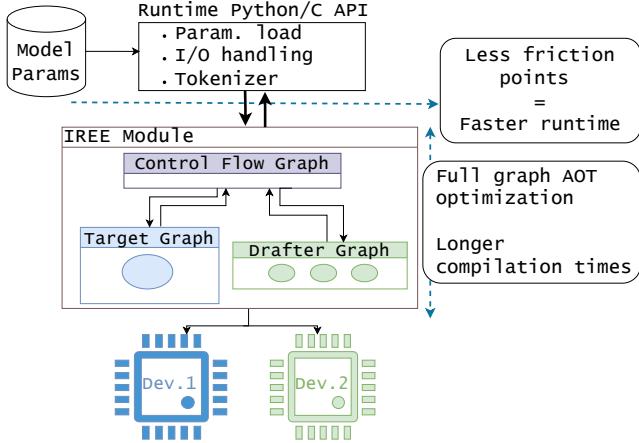


Fig. 3: Monolithic approach: single IREE module with target, drafter, and control flow subgraphs with heterogeneous device affinities.

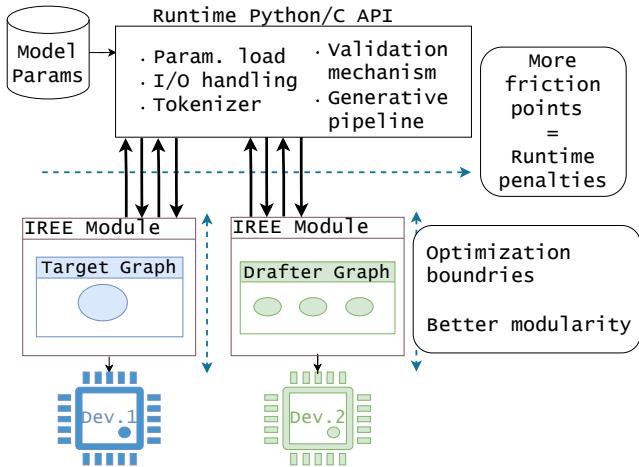


Fig. 4: Modular approach: separate IREE modules for model arithmetic with control flow in the serving platform.

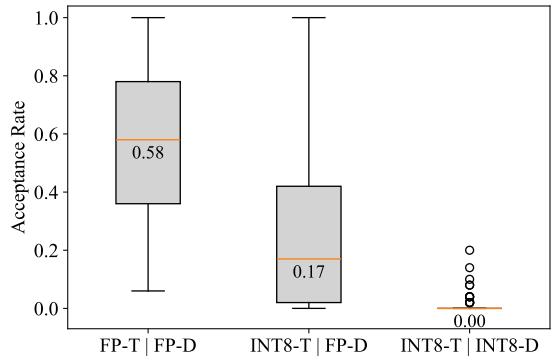
TABLE I: Experimental settings

Setting	Description
Workload	Short sequence lengths ($S_L \ll d$)
Task	Translation (from Spec-Bench benchmark [24])
Optimization objective	Minimize latency
SD technique	Speculative sampling with target: Llama 3.2 3B and drafter: Llama 3.2 1B.
Heterogeneous edge hardware	Mali-G310 GPU and Hexacore Cortex-A55 CPU on NXP i.MX95 SoC

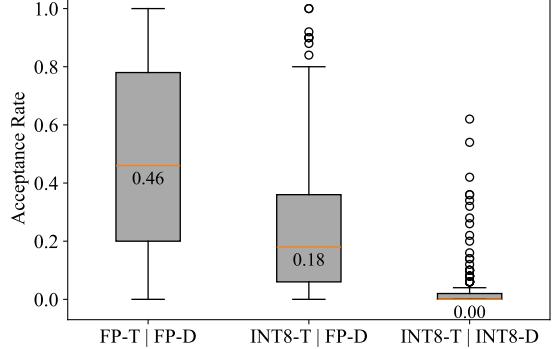
no KV cache is enabled. We employ the Llama 3.2 family, specifically Llama 3.2 1B as drafter and Llama 3.2 3B as target, following prior work demonstrating that alignment of the training data distributions between drafter and target is beneficial for draft acceptance [3].

A. Empirical Estimation of Acceptance Rate

The acceptance rate α serves as a required input parameter for the analytical cost model given by (1). While improving α



(a) Translation acceptance rates comparison



(b) Entire dataset acceptance rates comparison

Fig. 5: Acceptance rate α distribution for different quantization schemes: FP (FP32), T (target), D (drafter).

through better quantization schemes [8] or training quantized drafters directly [4] falls outside the scope of this work, we must characterize its behavior to validate our heterogeneous mapping methodology on real hardware. Evaluating these techniques on edge hardware represents a direction for future work; our contributions lies in how to exploit α effectively.

Fig. 5 shows the impact of quantization on α for the translation task (Fig. 5a) and the full Spec-Bench dataset (Fig. 5b). The vertical axes represent the acceptance rate α , while each box represents the distribution of α for three pairs of models: from left to right: an unquantized FP16 drafter and target, a partially quantized configuration, where just the target model is quantized², and a fully quantized INT8 setup. As the level of quantization increases, the boxes shift downward, showing a consistent reduction in the median value for α . In particular, the fully quantized pairing median collapses toward $\alpha \approx 0$, indicating that almost no drafted tokens are accepted. In contrast, the unquantized configuration achieves a median value of $\alpha = 0.58$.

These results confirm prior work [8]: as quantization increases, it introduces a distributional mismatch between drafter and target models, and the median α decreases substantially

²Note that the fully unquantized configuration, as well as the partially quantized variant in which only the drafter is quantized, were excluded from the experiments due to memory limitations during initialization on our hardware-software setup.

for speculative sampling on edge devices. This pattern is consistent across the entire Spec-Bench dataset, as shown in Fig. 5b, where the median falls along with increased quantization. The subsequent analysis concentrates on the translation task using the semiquantized configuration, for two reasons: (1) it provides a realistic edge deployment scenario balancing memory constraints with model accuracy, and (2) its broad distribution of acceptance rates, with α values spanning from 0% to 100%, enables an exploration of the full range of speculative sampling behavior.

B. Computation of the Cost Coefficient

We individually measured the latency of one forward pass for both the target and drafter models as a function of the input sequence length and then calculated the cost coefficient c , summarized in Fig. 6. The horizontal axes represent the input sequence length, and each curve depicts one design variant. Fig. 6a shows c values for homogeneous mappings on CPU, whereas Fig. 6b shows heterogeneous mappings where the drafter is executed on the GPU³. The regions with $c > 1$ are shaded in red to indicate infeasible configurations in which one forward pass of the drafter is slower than the target model. These infeasible cases arise mainly when the system uses three to six CPU cores in the heterogeneous mapping (i.e., drafter model mapped on the GPU).

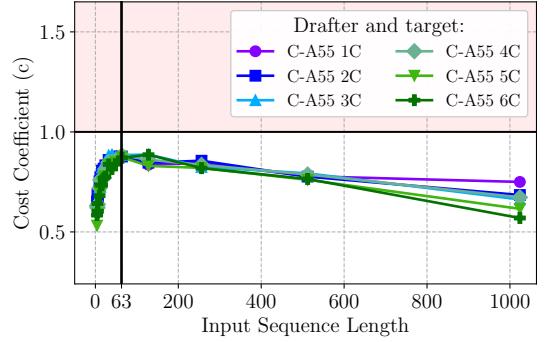
A notable case appears in the design variant with only a single CPU core available (purple curves in both plots). When performing a heterogeneous mapping, the cost coefficient c decreases substantially, from approximately 0.80 (purple curve at the top for an input sequence length of 63)⁴ to 0.41 at the same sequence length, shown at the bottom. This improvement arises because the GPU executes the drafter model roughly three times faster than a single CPU core, significantly reducing the computational cost of the speculative phase. As a result, this configuration emerges as the most favorable candidate for speculative sampling on the evaluated SoC using the Spec-Bench dataset.

C. Estimation of the Speedup with the Cost Model

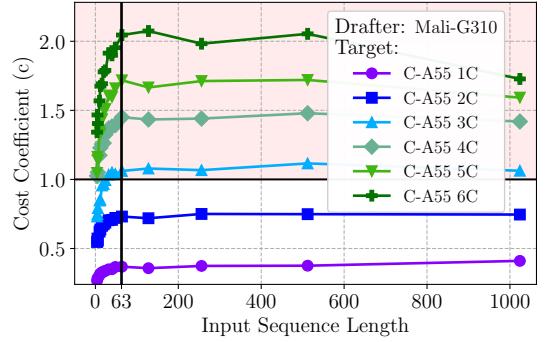
The expected speedup is estimated using Eq. (1) for two acceptance rates: a high value of $\alpha = 0.90$ corresponding to the 90th percentile and a low value of $\alpha = 0.17$ corresponding to the median. The results are reported in Tab. II and Tab. III, respectively. For $\alpha = 0.90$ (Tab. II), speculative sampling combined with heterogeneous execution yields substantial improvements, particularly in design variant 1, which corresponds to a configuration with one CPU core available for mapping and the GPU (recall that a design variant represents a unique combination of cores, shaders, or PEs across all the PUs). This variant reaches a speedup of 1.68 \times using $\gamma = 5$. The correct selection of the draft sequence length (γ) is

³We do not map the quantized target model onto the GPU because it does not support `INT8` datatype; any `INT8` tensor is promoted to `FP32`, adding overhead and diminishing the benefits of quantization.

⁴This input sequence length corresponds to the average length of the translation task in the Spec-Bench dataset.



(a) Measured cost coefficient (homogeneous execution)



(b) Measured cost coefficient (heterogeneous execution)

Fig. 6: Cost coefficients c for homogeneous (a) and heterogeneous (b) mappings as a function of input sequence length. The black vertical line indicates $S_L = 63$, the average input sequence length for the translation task in the Spec-Bench C-A55 nC: Cortex-A55 CPU, n cores.

crucial: each design variant achieves its best performance with a different γ value, ranging from 0 (no speculative sampling) to 5.

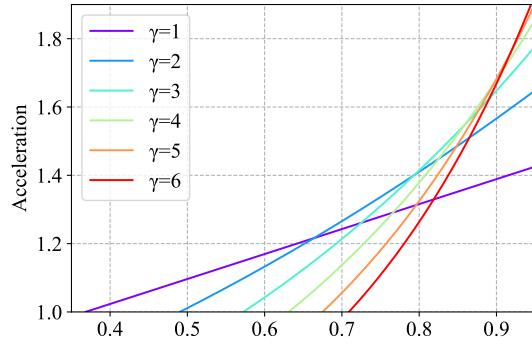
TABLE II: Estimated speedup for $\alpha = 0.90$ and $S_L = 63$

Design Variant	Speculative Sampling	Heterogeneous Execution	Speedup [\times]
1	Yes ($\gamma = 5$)	Yes	1.68
2	Yes ($\gamma = 2$)	Yes	1.10
3	No	NA	1
4	No	NA	1
5	Yes ($\gamma = 1$)	No	1.02
6	No	NA	1

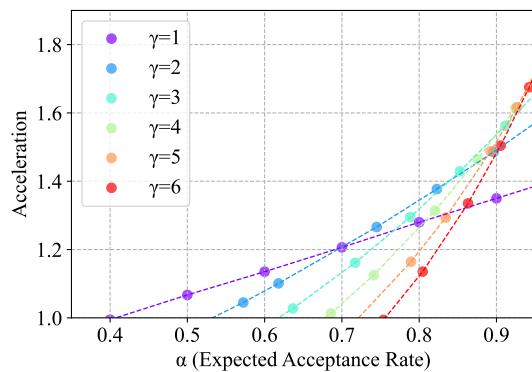
TABLE III: Estimated speedup for $\alpha = 0.17$ and $S_L = 63$

Design Variant	Speculative Sampling	Heterogeneous Execution	Speedup [\times]
1–6	No	NA	1

In contrast, design variants with a higher number of available cores (i.e., 3, 4, and 6 cores upwards) are expected to



(a) Estimated acceleration



(b) Measured acceleration

Fig. 7: Measured acceleration introduced by speculative sampling with different drafted token lengths γ executed heterogeneously: Target model mapped on one CPU core; drafter model, on GPU. The input sequence length is set to 63 and we report the average after ten runs.

have a performance decline when both speculative sampling and heterogeneous execution are applied. Therefore, these methods should be avoided in such configurations. For the design variant holding five CPU cores, there is a modest speedup, but heterogeneous mapping should not be applied. If the performance gain is too small, there is a risk that, in a real system, the improvement becomes negligible, especially when accounting for deployment overheads. As a result, we discourage the use of heterogeneous mapping in this scenario. However, for $\alpha = 0.17$ (Tab. III), even the shortest draft length $\gamma = 1$ becomes detrimental and neither speculative sampling nor heterogeneous execution yields a speedup in any design variant.

D. Validation and Discussion

The underlying acceleration trend and its empirical validation are jointly illustrated in Fig. 7, where the horizontal axes represent the acceptance rate α and the vertical axes report the predicted or achieved acceleration S , in Figs. 7a and 7b, respectively. In both cases, we present the heterogeneous mapping in which the quantized target model is executed on a single CPU core, while the unquantized drafter model

is assigned to the GPU. The results for an input sequence length of 63 are shown. The different curves correspond to different draft lengths γ , with longer drafts (orange and red curves) offering greater acceleration at sufficiently high α , while becoming detrimental when α falls. This explains why the semiquantized configuration with an acceptance rate of $\alpha = 0.17$ in Tab. III yields no measurable speedup: the acceptance rate is simply too low for speculative sampling to be beneficial. Fig. 7b presents the measured acceleration, showing that the expected acceleration (Fig. 7a) is achieved with an α approximately 4% higher than estimated. Despite this shift, the empirical curve remains consistent with the analytic prediction.

The evaluated heterogeneous configuration (drafter on GPU, target on single CPU core) arises from hardware-software constraints (e.g., IREE lacking INT8 support for Mali backends, quantization effects on α leading to semiquantized setups). Alternative strategies (full-GPU execution, dynamic scheduling) either exceed the memory budget of the platform, damage drafter accuracy, or require unavailable runtime capabilities. The analytical cost model (Eq. 1) abstracts hardware characteristics via c to determine when heterogeneous execution is beneficial. The methodology generalizes to other platforms through profiling c and measuring task-specific α .

Runtime constraints in IREE 3.6.0 prevented deploying a monolithic graph with heterogeneous device affinities. Consequently, we compiled separate graphs (Fig. 4) orchestrated via IREE’s runtime API. This introduces interface overhead that may explain the measured 4% deviation, but provided necessary flexibility for heterogeneous execution.

V. CONCLUSION AND OUTLOOK

This work showed how compiler-assisted heterogeneous mapping can accelerate LLM inference with speculative sampling on resource-constrained edge devices. An analytical cost model provided a basis for deciding when speculative sampling and heterogeneous execution are worthwhile under fixed hardware budgets, and model validation showed a 4% deviation from predictions. Exposing low-level abstractions into the compilation flow facilitated combining speculative sampling with heterogeneous execution, improving productivity with minimal performance penalties. Under favorable conditions (with a predicted $\alpha = 0.90$ and measured $\alpha = 0.94$), Llama 3.2 3B achieved up to a $1.68\times$ speedup.

Future work should: (1) integrate improved quantization techniques to achieve higher α values; (2) validate the cost model with additional edge SoCs and other SD techniques, as discussed in Sec. III-A; (3) extend the model to incorporate finer-grained partitioning would allow mapping computation onto NPUs and other accelerators with restricted operation support; and (4) improve the runtime memory allocator for heterogeneous execution of monolithic graphs with mixed device affinities to allow shared GPU-CPU scheduling.

ACKNOWLEDGMENT

This result is part of the IPCEI ME/CT and is funded by the European Union Next Generation EU, the German Federal Ministry for Economic Affairs and Energy, the Bavarian Ministry of Economic Affairs, Regional Development and Energy, the Free State of Saxony with the help of tax revenue based on the budget approved by the Saxon State parliament and the Free and Hanseatic City of Hamburg. The authors acknowledge the financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the program Center of Excellence for AI-research “Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig”, project identification number: ScaDS.AI.

REFERENCES

- [1] S. Kim, C. Hooper, T. Wattanawong, M. Kang, R. Yan *et al.*, “Full Stack Optimization of Transformer Inference,” in *Proc. Architecture and System Support for Transformer Models (ASSYST @ ISCA)*, 2023.
- [2] Y. Zheng, Y. Chen, B. Qian, X. Shi, Y. Shu, and J. Chen, “A Review on Edge Large Language Models: Design, Execution, and Applications,” *ACM Computing Surveys*, Feb. 2025, Art. no. 3719664, doi: 10.1145/3719664.
- [3] Y. Leviathan, M. Kalman, and Y. Matias, “Fast Inference from Transformers via Speculative Decoding,” in *Proc. International Conference on Machine Learning (ICML)*, 2023, pp. 19274–19286.
- [4] Y. Li, F. Wei, C. Zhang, and H. Zhang, “EAGLE: Speculative Sampling Requires Rethinking Feature Uncertainty,” in *Proc. International Conference on Machine Learning (ICML)*, 2024.
- [5] L. Deng, G. Li, S. Han, L. Shi, and Y. Xie *et al.*, “Model Compression and Hardware Acceleration for Neural Networks: A Comprehensive Survey,” *Proceedings of the IEEE*, vol. 108, no. 4, pp. 485–532, 2020.
- [6] Linux Foundation, “IREE,” <https://iree.dev/>, 2025. Accessed: Sept. 11, 2025.
- [7] Y. Hu, Z. Liu, Z. Dong, T. Peng, B. McDanel, and S. Q. Zhang, “Speculative Decoding and Beyond: An In-Depth Survey of Techniques,” *arXiv:2502.19732*, Mar. 2025, unpublished.
- [8] Y. Zhang, W. Zhao, X. Han, T. Zhao, W. Xu, H. Cao, and C. Zhu, “Speculative Decoding Meets Quantization: Compatibility Evaluation and Hierarchical Framework Design,” *arXiv:2505.22179*, May 2025, unpublished.
- [9] M. Li, Y. Liu, X. Liu, Q. Sun, X. You, H. Yang, Z. Luan, L. Gan, G. Yang, and D. Qian, “The Deep Learning Compiler: A Comprehensive Survey,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 3, pp. 708–727, Mar. 2021.
- [10] G. Korol, M. G. Jordan, M. B. Rutzig, J. Castrillon, and A. C. S. Beck, “Pruning and Early-Exit Co-Optimization for CNN Acceleration on FPGAs,” in *Proc. Design, Automation and Test in Europe Conference (DATE)*, Antwerp, Belgium, Apr. 2023, pp. 1–6, doi: 10.23919/DATE56975.2023.10137244.
- [11] G. Korol, M. G. Jordan, M. B. Rutzig, J. Castrillon, and A. C. S. Beck, “Design Space Exploration for CNN Offloading to FPGAs at the Edge,” in *Proc. IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, Foz do Iguaçu, Brazil, Jun. 2023, doi: 10.1109/ISVLSI59464.2023.10238644.
- [12] T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, *et al.*, “TVM: An Automated End-to-End Optimizing Compiler for Deep Learning,” in *Proc. 13th USENIX Conf. on Operating Systems Design and Implementation (OSDI’18)*, Carlsbad, CA, USA, 2018, pp. 579–594.
- [13] N. Rotem, J. Fix, S. Abdulrasool, G. Catron, and S. Deng *et al.*, “Glow: Graph Lowering Compiler Techniques for Neural Networks,” *arXiv:1805.00907*, 2019, unpublished.
- [14] C. Lattner, M. Amini, U. Bondhugula, A. Cohen, A. Davis, *et al.*, “MLIR: Scaling Compiler Infrastructure for Domain Specific Computation,” in *Proc. IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2021, pp. 2–14.
- [15] X. Miao, G. Oliaro, Z. Zhang, X. Cheng, Z. Wang, *et al.*, “SpecInfer: Accelerating Large Language Model Serving with Tree-based Speculative Inference and Verification,” in *Proc. 29th ACM Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Vol. 3, Apr. 2024, pp. 932–949.
- [16] Z. Chen, A. May, R. Svirchevski, Y. Huang, M. Ryabinin, *et al.*, “SEQUOIA: Scalable and Robust Speculative Decoding,” in *Proc. 38th Int. Conf. on Neural Information Processing Systems (NeurIPS)*, Vancouver, BC, Canada, 2024, Art. no. 4116, 33 pp.
- [17] K. Lv, H. Guo, Q. Guo, and X. Qiu, “DuoDecoding: Hardware-aware Heterogeneous Speculative Decoding with Dynamic Multi-Sequence Drafting,” *arXiv:2503.00784*, Mar. 2025, unpublished.
- [18] L. Zhang, Z. Zhang, Xubaizhou, R. Li, Z. Tian, *et al.*, “Dovetail: A CPU/GPU Heterogeneous Speculative Decoding for LLM Inference,” in *Proc. EMNLP*, Suzhou, China, Nov. 2025, pp. 17393–17406.
- [19] R. Baghdadi, J. Ray, M. B. Romdhane, E. D. Sozzo, A. Akkas, *et al.*, “Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code,” in *Proc. CGO*, 2019, pp. 193–205.
- [20] L. Mei, P. Houshmand, V. Jain, S. Giraldo, M. Verhelst, *et al.*, “ZigZag: Enlarging Joint Architecture-Mapping Design Space Exploration for DNN Accelerators,” *IEEE Transactions*
- [21] S. Na, G. Jeong, B. H. Ahn, A. Jezghani, and J. Young *et al.*, “FlexInfer: Flexible LLM Inference with CPU Computations,” in *Proc. 8th Conf. on Machine Learning and Systems (MLSys)*, 2025.
- [22] S. C. Kao and T. Krishna, “MAGMA: An Optimization Framework for Mapping Multiple DNNs on Multiple Accelerator Cores,” in *Proc. IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, Apr. 2022, pp. 814–830.
- [23] Z. Li, B. Yang, J. Li, T. Chen, X. Li, and M. Gao, “Adyna: Accelerating Dynamic Neural Networks with Adaptive Scheduling,” in *Proc. IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Mar. 2025, pp. 549–562.
- [24] H. Xia, Z. Yang, Q. Dong, P. Wang, and Y. Li *et al.*, “Unlocking Efficiency in Large Language Model Inference: A Comprehensive Survey of Speculative Decoding,” in *Findings of the Association for Computational Linguistics: ACL 2024*, Aug. 2024, pp. 7655–7671.
- [25] Intel Corporation, “intel/neural-compressor,” Oct. 2025. [Online]. Available: <https://github.com/intel/neural-compressor>