

SLED: A Speculative LLM Decoding Framework for Efficient Edge Serving

Xiangchen Li
Virginia Tech
Blacksburg, Virginia, USA
lixiangchen@vt.edu

Dimitris Spatharakis
National Technical University of
Athens
Athens, Greece
dspatharakis@netmode.ntua.gr

Saeid Ghafouri
Queen's University Belfast
Belfast, Northern Ireland, UK
s.ghafouri@qub.ac.uk

Jiakun Fan
Virginia Tech
Blacksburg, Virginia, USA
jiakunfan@vt.edu

Hans Vandierendonck
Queen's University Belfast
Belfast, Northern Ireland, UK
h.vandierendonck@qub.ac.uk

Deepu John
University College Dublin
Dublin, Ireland
deepu.john@ucd.ie

Bo Ji
Virginia Tech
Blacksburg, Virginia, USA
boji@vt.edu

Dimitrios S. Nikolopoulos
Virginia Tech
Blacksburg, Virginia, USA
dsn@vt.edu

Abstract

The growing gap between the increasing complexity of large language models (LLMs) and the limited computational budgets of edge devices poses a key challenge for efficient on-device inference, despite gradual improvements in hardware capabilities. Existing strategies, such as aggressive quantization, pruning, or remote inference, trade accuracy for efficiency or lead to substantial cost burdens. This position paper introduces a new framework that leverages speculative decoding, previously viewed primarily as a decoding acceleration technique for autoregressive generation of LLMs, as a promising approach specifically adapted for edge computing by orchestrating computation across heterogeneous devices. We propose *SLED*, a framework that allows lightweight edge devices to draft multiple candidate tokens locally using diverse draft models, while a single, shared edge server verifies the tokens utilizing a more precise target model. To further increase the efficiency of verification, the edge server batches the diverse verification requests from devices. This approach supports heterogeneous devices and reduces server-side memory footprint by sharing a single upstream target model across devices. Our initial experiments with Jetson Orin Nano, Raspberry Pi 4B/5, and an edge server equipped with 4 Nvidia A100 GPUs indicate substantial benefits: $\times 2.2$ higher system throughput, $\times 2.8$ higher system capacity, and better cost efficiency, all without sacrificing model accuracy.

CCS Concepts

• **Computing methodologies** → **Distributed computing methodologies**; **Distributed artificial intelligence**; *Natural language*



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

SEC '25, Arlington, VA, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2238-7/25/12

<https://doi.org/10.1145/3769102.3770608>

processing; • **Computer systems organization** → *Distributed architectures*; • **Networks** → **Cloud computing**; • **General and reference** → **General conference proceedings**.

Keywords

Speculative Decoding, Large Language Models, Edge Computing, Distributed Inference, Token Verification, Resource-Aware Serving

ACM Reference Format:

Xiangchen Li, Dimitris Spatharakis, Saeid Ghafouri, Jiakun Fan, Hans Vandierendonck, Deepu John, Bo Ji, and Dimitrios S. Nikolopoulos. 2025. *SLED: A Speculative LLM Decoding Framework for Efficient Edge Serving*. In *The Tenth ACM/IEEE Symposium on Edge Computing (SEC '25)*, December 3–6, 2025, Arlington, VA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3769102.3770608>

1 Introduction

LLMs have revolutionized various domains, demonstrating remarkable capabilities in natural language understanding, generation, and complex reasoning [1]. Their widespread adoption has led to transformative applications in areas such as intelligent chatbots, content creation, code generation, and scientific discovery. However, the immense memory and compute footprint associated with state-of-the-art LLMs, often comprising billions or even trillions of parameters, pose significant challenges for deployment. These models typically demand powerful accelerators like GPUs and substantial memory, limiting their direct execution on resource-constrained devices. Deploying LLMs at the edge, closer to data sources and end-users, offers significant advantages including reduced latency, enhanced privacy, and lower bandwidth consumption[2]. Nevertheless, edge environments, characterized by limited memory, processing power, and energy budgets, present formidable obstacles to efficient LLM inference. Existing strategies to address these limitations include aggressive model compression techniques such as quantization [16, 32], pruning[19, 22], and knowledge distillation [10]. Other approaches involve distributed inference, where model

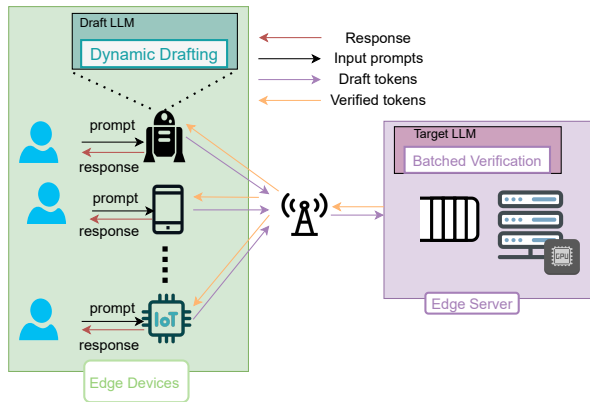


Figure 1: System overview of the proposed speculative LLM decoding framework for efficient edge serving.

layers are partitioned across multiple devices or between edge and cloud[29, 33], or full remote inference, where the entire computation is offloaded to a powerful central server[7, 31]. While these methods show some potential, they often come with trade-offs: compression can sacrifice model accuracy, distributed inference introduces synchronization overheads and is incompatible with heterogeneous edge devices, and remote inference negates the benefits of edge deployment, incurring non-negligible costs.

Speculative decoding [15] is a decoding acceleration technique that first generates multiple draft tokens using a relatively small draft model and then verifies them in a single forward pass using a larger, more accurate target model. By generating multiple tokens with the smaller model and validating them in a single pass on the larger model, speculative decoding significantly reduces the number of forward passes required on the large model, thereby accelerating the decoding process. This position paper introduces *SLED*, shown in Fig. 1, a novel approach that re-imagines speculative decoding as a paradigm specifically tailored for efficient LLM inference at the edge, by intelligently orchestrating computation across heterogeneous devices. Within the defined service area consisting of an edge server and multiple heterogeneous edge devices, and each edge device is equipped with its own lightweight LLMs scaled according to individual computational and memory resource capacities. These edge devices are responsible for serving diverse LLM-based applications such as intelligent personal assistants, text generation, and semantic analysis, among other tasks. Concurrently, a single, shared edge server, equipped with a more precise target model, efficiently batches and verifies these drafted tokens.

The advantages of the *SLED* are threefold:

- (1) Compared with inference solely on the edge device, the *SLED* improves the quality of response on the device by leveraging a larger target model on the server to verify draft tokens.
- (2) Compared with inference solely on the edge server, *SLED* reduces the monetary cost for edge users by limiting their use of server resources, requiring only token verification rather than full generation.

Table 1: Capacity of *SLED* and centralized serving system (Capacity = maximum number of supported devices at target response rate (devices))

System	RPi 4b (llama.cpp)	RPi 5 (llama.cpp)	Nvidia Jetson
<i>SLED</i>	18.30	5.24	19.53
Centralized serving	7.05	1.83	7.06
Capacity improvement	×2.60	×2.86	×2.77

- (3) It utilizes the computational resource of edge server to verified batched draft tokens from devices, rather than generate all tokens solely, enabling edge server support more edge devices simultaneously.

We compare the system capacity, the number of edge devices supported by the system, of *SLED* and a centralized serving system with the same response rate but different device types. From Tab. 1, we observe that compared with a centralized LLM serving system for the edge, the proposed *SLED* is capable of increasing the system capacity by 2.6 to 2.9 times.

Our key contributions are summarized as follows:

- We propose *SLED*, a novel speculative decoding framework specifically designed for heterogeneous edge computing environments by batching the verification requests from varying draft models on the server, enabling efficient LLM inference without accuracy degradation.
- In *SLED*, we propose and deploy the dynamic drafting scheme on edge devices. By dynamically requesting for verification according to the confidence score of the draft model, the edge devices can avoid unnecessary verification, hence reducing the communication rounds and improving the utilization of the server.
- We demonstrate through preliminary evaluation the substantial benefits of *SLED* in terms of ×2.2 more system throughput, ×2.8 more system capacity, and better cost-efficiency on diverse edge hardware.

The remainder of this paper is organized as follows: Section 2 reviews existing work in LLM inference for edge computing. Section 3 details the architectural design and key components of *SLED*. Section 4 presents our experimental setup and discusses the evaluation results. Finally, Section 5 concludes the paper and outlines future research directions.

2 Related Work

The efficient inference of LLMs on resource-constrained devices has been a focal point of research, broadly categorized into model compression techniques, distributed inference strategies, and remote offloading paradigms.

2.1 Model Compression and Lightweight Architectures

To enable LLMs to run on resource-constrained devices, significant efforts have been directed towards model compression. Quantization reduces the numerical precision of model parameters and activations to decrease memory footprint and accelerate computation

Table 2: Comparison of related work; Edge-Serving: Does the system support edge computing?; Heterogeneity: Is the heterogeneity of edge devices considered in the system design?; Lossless: Whether does the system deliver LLM service without any performance degradation? Scalable Model: Is the system capable of scaling model according to conditions without too much overhead?

System	Edge-Serving	Heterogeneity	Lossless	Scalable Model
EdgeShard[33]	✓	✓	✗	✗
Galaxy[29]	✓	✗	✓	✗
Orca[30]	✗	✗	✓	✓
vLLM[14]	✗	✗	✓	✓
FastServe[26]	✗	✗	✓	✓
AWQ[16]	✓	✓	✗	✗
MobileBERT[23]	✓	✓	✗	✗
<i>SLED</i>	✓	✓	✓	✓

[4, 16, 17]. Pruning identifies and removes redundant connections or neurons from the neural network without significant performance loss, resulting in sparser and smaller models [9]. Knowledge distillation involves training a smaller "student" model to mimic the behavior of a larger "teacher" model, thereby transferring knowledge and achieving comparable performance with a significantly smaller footprint [10]. Beyond these optimization techniques, research has also focused on designing inherently lightweight transformer architectures that are more efficient from the ground up, such as MobileBERT [23], Mamba [8] or other compact variants, often by optimizing attention mechanisms or reducing the number of layers and hidden dimensions. Recent efforts, such as EdgeLLM [27], have explored speculative decoding for on-device inference by introducing compute-efficient branch navigation and adaptive fallback strategies to reduce resource demands. However, co-locating both draft and target models on edge devices further strains limited resources, and the output quality remains constrained by the size of the deployable target model. Despite their advantages in reducing model size and computational demands, a common limitation of these model compression techniques is the inherent trade-off with model quality: aggressive compression often leads to a measurable decrease in accuracy compared to their full-sized counterparts.

2.2 Edge-Cloud/Server Offloading and Distributed Inference

Another line of research focuses on distributing LLM computation across multiple devices or partitioning tasks between edge and cloud/server infrastructure. Model partitioning schemes divide a large LLM into smaller sub-models, with different parts executed on different devices [28, 29, 33]. For example, EdgeShard[33] partitioned LLM into shards and deploy on distributed devices to benefit from the collaboration among edge devices and cloud server. This often involves pipeline parallelism or tensor parallelism techniques, where different stages or segments of the model's computation are assigned to different devices. While this allows larger models to run on resource-constrained setups, it introduces communication overheads and synchronization challenges, particularly for heterogeneous hardware and varying network conditions. Edge-cloud offloading dynamically decides which parts of the inference task

should be performed locally at the edge and which should be offloaded to a more powerful cloud server, often based on real-time resource availability, network bandwidth, and latency requirements [33]. These methods aim to balance the benefits of edge processing with the computational power of the cloud, but often require sophisticated orchestration and robust connectivity.

2.3 Pure Remote Inference

Pure remote inference, where the entire LLM resides on centralized data-center GPUs, represents a prevalent deployment paradigm due to its simplicity and centralized resource utilization. Recent research primarily focuses on resource efficiency and latency optimization. Kwon et al.[14] proposed vLLM with a PagedAttention allocator, significantly reducing KV-cache overhead and fragmentation, achieving up to 4× throughput improvement. Wu et al. [26] introduced FastServe, leveraging multi-level feedback queue scheduling and proactive KV-cache management to reduce tail latency by up to 31× at the 99th percentile. Rajbhandari et al. [21] developed DeepSpeed Inference, combining multiple parallelism strategies with NVMe and CPU off-loading, allowing inference of substantially larger models and reducing latency by up to 7.3×. Crucially, the standard decoding process in remote inference is often autoregressive, generating one token at a time, which can be memory-intensive due to large key-value caches and lead to resource underutilization on powerful servers. Moreover, the cost associated with cloud GPU instances for continuous, often under-utilized, inference also presents a substantial economic burden, especially for high-throughput scenarios.

Table 2 compares the *SLED* and related works that deliver LLM inference service or propose model variants for edge devices, among which *SLED* stands out as the only approach that enables lossless LLM inference for heterogeneous edge devices, while maintaining a collaborative design that can flexibly accommodate increasingly large models. *SLED* directly addresses these limitations by fundamentally altering the inference paradigm. Instead of autoregressive token generation on the powerful central server, *SLED* offloads the preliminary token drafting to lightweight edge devices. This allows the central server to focus its considerable resources primarily on the more efficient and batchable task of verifying multiple drafted tokens. By doing so, *SLED* significantly improves the utilization of expensive server-side GPU resources without sacrificing model accuracy, leading to a more cost-effective and scalable distributed LLM inference system.

3 SLED Design

3.1 Speculative Decoding

SLED adopts speculative decoding[15], a recently proposed paradigm that leverages a lightweight draft model to generate multiple tokens. Instead of waiting for a single token from the large target model at each step, the draft model speculatively predicts a candidate sequence, which is then selectively verified or corrected by the more accurate but slower target model with one single forwarding. This design decouples the fast-path token generation from the slow-path verification, enabling improved throughput and reduced end-to-end latency without compromising output quality. In *SLED*,

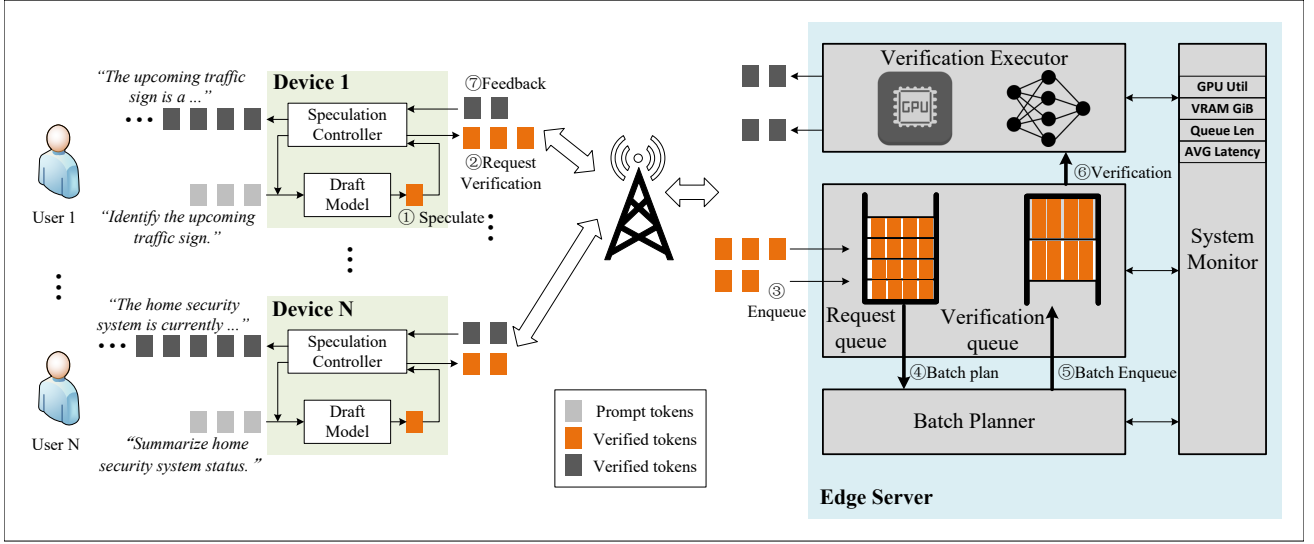


Figure 2: SLED architecture and request/verification data flow

this speculative execution is further optimized for edge-server deployment under variable connectivity and resource constraints.

According to the original work on speculative decoding[15], a draft token \tilde{x}_{n+1} is sampled from the draft model’s output distribution $p(x | x_1, x_2, \dots, x_n)$. This draft token is then accepted with probability:

$$\alpha = \min\left(1, \frac{p_{\text{target}}(\tilde{x}_{n+1} | x_1, \dots, x_n)}{p_{\text{draft}}(\tilde{x}_{n+1} | x_1, \dots, x_n)}\right), \quad (1)$$

where p_{target} is the distribution given by the large target model, and p_{draft} is that of the smaller draft model. Tokens that fail this acceptance criterion are resampled directly from $p_{\text{target}} - p_{\text{draft}}$. As a result, the generated sequence strictly follows the target model’s distribution, ensuring no degradation in accuracy compared with standard autoregressive decoding.

3.2 System Overview

Fig. 2 shows the detailed structure and data flow of the *SLED*. In close proximity to N edge devices, typically located at facilities such as base stations, the edge server provides substantial computational capabilities, leveraging specialized hardware like Graphics Processing Units (GPUs) or Neural Processing Units (NPUs). On this edge server, a single, comprehensive target model is deployed, optimized for efficiently verifying the draft tokens generated by the distributed edge devices.

Operationally, user-generated prompts, encompassing a wide array of task-specific requests, are initially received and tokenized locally by each edge device. Subsequently, the tokenized prompts, denoted as input sequences p^n where $n \in \{1, 2, \dots, N\}$, are processed by local draft models to generate speculative tokens. These draft tokens are then transmitted to the edge server for verification. Upon completion of the verification step, the edge server communicates the results back to the respective edge devices, specifically

identifying rejected token positions along with any necessary corrective tokens.

This drafting-verification workflow iteratively progresses, alternating between local speculation at the edge devices and centralized validation at the edge server, until the generated output reaches the predetermined desired length or the end-of-response token is encountered. This collaborative mechanism not only optimizes resource utilization by distributing computational tasks according to device capabilities but also significantly reduces latency and enhances overall efficiency by transmitting tokens rather than huge activations.

3.3 Dynamic Drafting on Edge Devices

On edge devices, each verification cycle is preceded by the generation of multiple draft tokens. The acceptance rate of these draft tokens serves as a crucial indicator of their quality, directly influenced by the capabilities of the draft models utilized. A higher acceptance rate is desirable as it signifies fewer verification iterations and consequently reduces the computational burden on the costly target model, thereby mitigating communication overhead inherent in edge computing scenarios.

Previous studies [11], validated by our preliminary experimental results, have established a correlation between the acceptance rate of draft tokens and their associated confidence scores derived from the output logits. As illustrated in Fig.3, draft tokens with higher confidence scores exhibit a significantly increased likelihood of acceptance by the target model.

Building upon this insight, we propose and implement a dynamic drafting mechanism on edge devices. This adaptive strategy modulates the speculative decoding length based on the real-time evaluation of token confidence scores. Formally, we introduce a threshold parameter, c_{th} , derived empirically, and define the decision-making process for triggering server verification for the draft tokens as follows:




Figure 3: Acceptance rate vs. confidence of draft tokens

$$c_s^i \begin{cases} < c_{th}, \text{request verification} \\ \geq c_{th}, \text{generate another token} \end{cases}, \quad (2)$$

where c_s^i represents the confidence score associated with token t_s^i .

Considering the inherent unreliability and fluctuating nature of network conditions in edge computing environments, such as variable round-trip time (RTT) and intermittent connectivity, we further enhance our system with an asynchronous decoding mechanism accompanied by a timeout protocol. Specifically, edge devices continue generating additional draft tokens using local lightweight LLM concurrently while awaiting verification responses from the edge server. If a verification response confirms acceptance of all previously sent draft tokens, these locally generated tokens seamlessly transition into the draft token queue for subsequent verification cycles, thus significantly reducing idle wait times.

Additionally, each verification request initiates a timer on the device side. If the verification response exceeds the timer due to server failures or network disruptions, the most recently-produced draft tokens are concatenated with existing draft tokens for subsequent verification attempts. To maintain continuity of user experience, the draft tokens generated during this period are released to users as a fallback when consecutive verification failures exceed the threshold.

3.4 Batched Verification on Edge Server

The edge server aggregates verification requests from multiple edge devices into batches to optimize computational efficiency and throughput. Our current implementation within *SLED* employs a static batching strategy. Under this scheme, incoming verification requests are temporarily queued until reaching a fixed batch size. Subsequently, a batch planner retrieves the queued requests, applies appropriate padding to equalize token lengths, and forwards the consolidated batch to the target LLM for verification.

A critical advantage of *SLED* lies in the target model’s ability to accept and verify draft tokens generated by diverse draft LLMs across heterogeneous edge devices provided that all edge devices share the same tokenizer. This compatibility effectively mitigates device heterogeneity, enabling each device to select a draft model suited to its computational constraints while ensuring scalable and efficient inference across a wide range of edge hardware.




Figure 4: WSTGR comparison between *SLED* and centralized LLM serving systems, highlighting improved scalability of the *SLED* framework.

4 Evaluation

In this section, we evaluate the performance and efficiency of the proposed *SLED* framework through extensive simulations and measurements. We assess *SLED*’s efficacy compared to a centralized LLM serving system which serves the decoding requests from edge device directly, and the edge-only inference system which generates all tokens locally on the devices across various metrics including throughput, system capacity, cost efficiency, and impact of speculative length on system capacity and the throughput.

To accurately simulate verification request workloads from edge devices utilizing speculative decoding, we adopt a Poisson-based modeling approach. Each edge device is considered an independent source of verification requests, with inter-arrival times following an exponential distribution. This modeling choice effectively captures the asynchronous and inherently stochastic nature of real-world device interactions, ensuring that the simulated workload closely mirrors realistic operational conditions. The device-specific request rate is derived directly from realistic device speculative decoding throughput, ensuring that the simulation’s temporal patterns closely align with practical speculative decoding workloads.

As for the device setting, we tested Raspberry Pi 4b, 5 and NVIDIA Jetson Orin Nano on the *SLED* system supported by 4 A100 GPUs. We deployed LLaMA family, including model sizes of 1B, 3B, 11B and 70B on the edge devices and verification server, supporting up to context length of 128k.

4.1 Whole System Token Generation Rate (WSTGR)

We first evaluate the Whole System Token Generation Rate (WSTGR), which is defined as the total number of tokens generated and verified by the entire inference system per second, and serves as a metric for the system’s overall productive output[18]. Given a certain time period we measure the total number of tokens generated by the *SLED* and centralized inference system. The verification workload model is derived from a Raspberry Pi 5 device running a LLaMA 3B model. Additionally, we evaluate two different target models (11B and 70B) on both the *SLED* system and a centralized inference system. As shown in Fig. 4, for both the 11B and 70B models, the WSTGR increases rapidly in the initial stages as batch size grows,

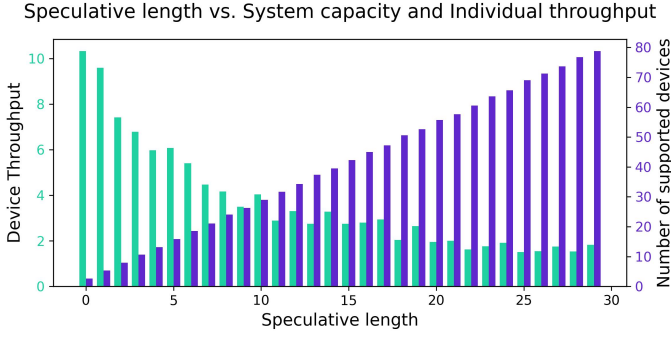


Figure 5: The Impact of speculative length on system-level capacity and device-level throughput, showing the speculative length should be considered to balance the tradeoff.

due to the amortization of fixed GPU launch and driver overhead, and improved utilization of GPU cores. The proposed *SLED* system achieves higher overall throughput than the centralized serving system under identical conditions, including the same number of devices and target model. This observed scalability demonstrates that *SLED* effectively utilizes distributed edge resources to enhance the system’s token generation capacity.

The more than twofold improvement in WSTGR over the centralized serving system stems from the efficient and balanced distribution of computational tasks in the *SLED* system. Specifically, in *SLED*, simple token generation tasks can be handled by relatively small models [15], such as those deployed on edge devices, while more complex tasks are offloaded to larger models on the edge server. This architectural separation allows computation to be distributed across edge devices and the edge server in a more resource-aligned and efficient manner. As a result, the computational capacity of the edge server is reserved for challenging verification tasks, rather than being consumed by processing simpler tokens from edge devices, unlike in a centralized LLM serving setup.

4.2 Speculative Length vs. throughput and capacity

In speculative decoding, the length of the draft sequence used for verification on the target model is defined as the speculative length, and it affects both token generation throughput and the capacity of *SLED*, that is, the number of edge devices supported by *SLED* simultaneously. In this experiment, we manually adjust the speculative length for drafting using LLaMA 1B model on a Raspberry Pi 5 device, and measure both the device throughput and overall system capacity. As shown in Fig. 5, increasing the speculative length results in lower device throughput but higher system capacity. This inverse relationship between per-device and system-level metrics highlights the importance of selecting an appropriate speculative length to balance the performance of individual edge devices and the system as a whole. On one hand, a longer speculative length reduces token generation on each device, since the drafting throughput remains stable, and a longer speculative length leads to a longer verification period, thereby reducing the response update rate. On the other hand, a longer verification period for individual devices




Figure 6: Pareto front showing optimal trade-offs between energy consumption per token and latency, highlighting the efficiency of *SLED*.

reduces the verification workload on the edge server, allowing it to support more devices concurrently.

4.3 Cost Efficiency and System Throughput

The cost efficiency of token generation in edge computing scenarios is a critical factor and has been considered in various edge inference system designs [12, 13]. In this paper, we compare the proposed *SLED* system, a centralized serving system, and an all-edge decoding system in terms of both cost efficiency and WSTGR. To systematically analyze the trade-off between cost and performance, we construct Pareto front visualizations, which highlight the non-dominated configurations that achieve the best balance between monetary cost and system throughput.

In our experiments, the cost and throughput metrics for edge inference scenarios were carefully computed based on a comprehensive capital expenditure (CAPEX) and operational expenditure (OPEX) model[25]. Specifically, we adopt the widely recognized CPU-hour cost model described by Walker[25] and the edge-compute modeling approach proposed by Eriksson[6]. The CAPEX component was determined by amortizing the purchase price of each edge device (Raspberry Pi 5 priced at \$80 [24]) over a three-year lifetime, assuming an average device utilization rate of 70%. The OPEX component included electrical consumption calculated from experimentally measured average power draw (8 W for Pi 5) [20] and industrial electricity rates (0.083 \$/kWh) [5]. Combining these costs, we obtained a unified hourly expense for each device, subsequently normalized by the experimentally measured token generation rates (tokens per second), as shown in Eq. 3. The resulting metric, expressed clearly as dollars per one thousand generated tokens (\$/1K tokens), enabled direct and transparent comparison across different experimental configurations and devices.

$$\text{Cost} = \frac{1000}{3600 R} \left(\frac{P_{\text{device}}}{3 \times 8760 \times 0.70} + \frac{P_{\text{avg}}}{1000} \times 0.083 \right) \quad (3)$$

Figure 6 compares the following three deployment strategies along a common cost–performance plane. Specifically, the strategies are: 1) all-Server executes every token-generation step on a bank of four NVIDIA A100–80 GB GPUs. 2) All-Edge places the same LLaMA draft model on each Raspberry Pi 5, with no server involvement. 3) *SLED* lets the Raspberry Pi 5 generate draft tokens, which are batch-verified on the A100 cluster with the same configuration of the centralized scenario. For every strategy, we sweep two orthogonal factors: quantization precision (16-, 8-, and 4-bit) and edge-device count $N \in \{1, 2, 4, 8, 16\}$. Cost is monetised as dollars per one-thousand verified tokens.

We observe that *SLED*'s skyline consistently dominates the Pareto frontier, achieving lower cost per 1K verified tokens while sustaining higher overall throughput. For instance, with the same system capacity and quantization level, *SLED* achieves a throughput of 137 tokens/s—representing a 3.5 \times improvement over the centralized baseline—while simultaneously reducing cost to just 29% of that. This advantage becomes more pronounced as the number of edge devices increases. Furthermore, quantization universally improves cost efficiency across all schemes by simultaneously reducing energy demand and increasing per-device generation rate. Notably, the 4-bit *SLED* configuration with 16 devices achieves 137 tokens/s at \$0.13 / 1K tokens, representing a 65% improvement in throughput over the best-performing All-Edge setup, with acceptable additional cost. These results substantiate the claim that *SLED* enables a superior cost–throughput trade-off, combining local cost efficiency with global throughput.

4.4 Impacts of Network Connectivity

We built a *SLED* prototype to quantify how network loss affects edge-side throughput. The edge runs on a Raspberry Pi 5B hosting meta-llama/Llama-3.2-1B-Instruct for drafting, while the server uses A100 GPU running meta-llama/Meta-Llama-3.1-70B-Instruct for batched verification. We instrumented the system to sweep packet-loss ratios from 0% to 100% under a fixed RTT budget and measured the committed token throughput at the edge. As shown in Fig.7, across commonly used settings in practice ($\leq 5\text{--}10\%$ loss), *SLED*'s edge throughput degraded only slightly (typically $<2\text{--}3\%$ versus the no-loss baseline), confirming that the proposed retry-then-fallback policy effectively masks moderate loss. Even under complete network unavailability (responses never arrive), the edge maintained ≥ 5.24 tokens/s, comfortably above typical human reading rates, ensuring responsive user-perceived progress. These results indicate that, in both engineering practice and research deployments, *SLED* preserves acceptable edge throughput under realistic packet-loss regimes, while providing graceful performance under extreme conditions.

Additionally, we evaluated the response quality of the *SLED* under varying packet loss rates in the prototype system. By sweeping the packet loss ratio from 0% to 100%, we assessed performance on GSM8K [3], a curated benchmark of linguistically diverse, grade-school math word problems designed to evaluate multi-step quantitative reasoning in language models. As shown in Fig.8, the response quality remains consistent with that of the target model when the verification packet loss rate is below 10%, which covers the most common real-world conditions. The quality gradually




Figure 7: The impact of packet loss on edge device-side throughput.




Figure 8: The impact of packet loss on the quality of generated response.

declines to the draft model level as the packet loss rate increases from 10% to 100%, indicating a loss of connection to the verification server. Based on the proposed retry-then-fallback policy, persistent packet loss triggers a fallback to the local draft model for token generation, leading to a degradation in output quality. These results demonstrate that under typical connectivity conditions, *SLED* can effectively serve edge devices without compromising accuracy.

Overall, the experimental evaluations underscore the significant advantages of *SLED* in distributed inference scenarios, including throughput, capacity, and cost efficiency, showcasing insightful findings in the *SLED* system to motivate more explorations in future work.

5 Conclusion and Future Work

This position paper presented the *SLED*, a novel distributed decoding framework designed for LLM deployment at the edge. Our extensive evaluation demonstrated that *SLED* significantly improves system throughput, capacity, and cost efficiency compared to traditional centralized approaches. The integration of speculative local

drafting and centralized verification establishes a balance of computational workload, making *SLED* particularly suitable for bringing LLMs towards the edge of the network. The position paper highlights that the *SLED* is more than a decoding enhancement—it opens the door to a more foundational and elastic approach to resource-aware LLM serving at the edge.

In future work, we will explore optimizing key–value caching strategies, potentially leveraging recent advances in PagedAttention, to further reduce server-side verification cost. Additionally, enhancing the adaptive capabilities of *SLED* for dynamic environments, such as extending *SLED*'s applicability to multi-modal scenarios, will be another interesting topic to focus on. Lastly, network conditions and resource-aware verification strategy could further broaden its practical impact in complex edge computing landscapes.

Acknowledgments

This material is based on work supported by the National Science Foundation under Grants No. 2315851 and 2106634.

References

- [1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., Red Hook, NY, USA, 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bf8ac142f64a-Paper.pdf
- [2] Jiasi Chen and Xukan Ran. 2019. Deep Learning With Edge Computing: A Review. *Proc. IEEE* 107, 8 (2019), 1655–1674. doi:10.1109/JPROC.2019.2921977
- [3] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. <https://arxiv.org/abs/2110.14168>
- [4] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale.
- [5] eia2025electric. 2025. *Electric Power Monthly: Average Price of Electricity to Ultimate Customers by End-Use Sector, March 2025*. <https://www.eia.gov/electricity/monthly/> accessed 17 Jun 2025.
- [6] Mats Eriksson. 2020. Cost modelling of edge compute.
- [7] Luyao Gao, Jianchun Liu, Hongli Xu, Xichong Zhang, Yunming Liao, and Liusheng Huang. 2025. Collaborative Speculative Inference for Efficient LLM Inference Serving.
- [8] Albert Gu and Tri Dao. 2024. Mamba: Linear-Time Sequence Modeling with Selective State Spaces.
- [9] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network.
- [11] Kaiyu Huang, Hao Wu, Zubo Shi, Han Zou, Minchen Yu, and Qingjiang Shi. 2025. SpecServe: Efficient and SLO-Aware Large Language Model Serving with Adaptive Speculative Decoding.
- [12] Erik Johannes Husom, Arda Goknil, Merve Astekin, Lwin Khin Shar, Andre Käsen, Sagar Sen, Benedikt Andreas Mithassel, and Ahmet Soylu. 2025. Sustainable LLM Inference for Edge AI: Evaluating Quantized LLMs for Energy Efficiency, Output Accuracy, and Inference Latency. arXiv:2504.03360 doi:10.48550/arXiv.2504.03360 Abstract & full paper report energy-/cost-efficiency benchmarks on Raspberry Pi 4.
- [13] SiYoung Jang and Roberto Morabito. 2025. Edge-First Language Model Inference: Models, Metrics, and Trade-offs. arXiv:2505.16508 doi:10.48550/arXiv.2505.16508
- [14] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. <https://arxiv.org/abs/2309.06180>
- [15] Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*. PMLR, Honolulu, Hawaii, USA, 19274–19286.
- [16] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. AWQ: Activation-aware Weight Quantization for On-Device LLM Compression and Acceleration. In *Proceedings of Machine Learning and Systems*, P. Gibbons, G. Pekhimenko, and C. De Sa (Eds.), Vol. 6. MLSys, Santa Clara, CA, USA, 87–100. https://proceedings.mlsys.org/paper_files/paper/2024/file/42a452cbafa9dd64e9ba4aa95cc1ef21-Paper-Conference.pdf
- [17] Shih-yang Liu, Zechun Liu, Xijie Huang, Pingcheng Dong, and Kwang-Ting Cheng. 2023. LLM-FP4: 4-Bit Floating-Point Quantized Transformers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Singapore, 592–605. doi:10.18653/v1/2023.emnlp-main.39
- [18] Xiaoxuan Liu, Cade Daniel, Langxiang Hu, Woosuk Kwon, Zhuohan Li, Xiangxi Mo, Alvin Cheung, Zhijie Deng, Ion Stoica, and Hao Zhang. 2024. Optimizing speculative decoding for serving large language models using goodput.
- [19] Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. LLM-Pruner: On the Structural Pruning of Large Language Models.
- [20] Les O'Donnell. 2023. *Raspberry Pi 5 Review: A New Standard for Makers*. <https://www.tomshardware.com/reviews/raspberry-pi-5> Tom's Hardware, accessed 17 Jun 2025.
- [21] Samyam Rajbhandari, Reza Yazdani Aminabadi, Minjia Zhang, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Jeff Rasley, Shaden Smith, Olatunji Ruwase, and Yuxiong He. 2022. DeepSpeed Inference: Enabling Efficient Inference of Transformer Models at Unprecedented Scale. <https://arxiv.org/abs/2207.00032>
- [22] Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. 2024. A Simple and Effective Pruning Approach for Large Language Models.
- [23] Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. MobileBERT: a Compact Task-Agnostic BERT for Resource-Limited Devices.
- [24] Eben Upton. 2023. *Introducing Raspberry Pi 5*. <https://www.raspberrypi.com/news/introducing-raspberry-pi-5/> Raspberry Pi Foundation blog, accessed 17 Jun 2025.
- [25] Edward Walker. 2009. The Real Cost of a CPU Hour. *IEEE Computer* 42, 4 (April 2009), 35–41. doi:10.1109/MC.2009.127
- [26] Bingyang Wu, Yinmin Zhong, Zili Zhang, Shengyu Liu, Fangyue Liu, Yuanhang Sun, Gang Huang, Xuanzhe Liu, and Xin Jin. 2024. Fast Distributed Inference Serving for Large Language Models. <https://arxiv.org/abs/2305.05920> arXiv:2305.05920
- [27] Daliang Xu, Wangsong Yin, Hao Zhang, Xin Jin, Ying Zhang, Shiyun Wei, Mengwei Xu, and Xuanzhe Liu. 2025. EdgeLLM: Fast On-Device LLM Inference With Speculative Decoding. *IEEE Transactions on Mobile Computing* 24, 4 (2025), 3256–3273. doi:10.1109/TMC.2024.3513457
- [28] Shengyuan Ye, Jiangsu Du, Liekang Zeng, Wenzhong Ou, Xiaowen Chu, Yutong Lu, and Xu Chen. 2024. Galaxy: A Resource-Efficient Collaborative Edge AI System for In-situ Transformer Inference.
- [29] Shengyuan Ye, Bei Ouyang, Liekang Zeng, Tianyi Qian, Xiaowen Chu, Jian Tang, and Xu Chen. 2025. Jupiter: Fast and Resource-Efficient Collaborative Inference of Generative LLMs on Edge Devices.
- [30] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. 2022. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. OSDI, Carlsbad, CA, USA, 521–538.
- [31] Zhongzhi Yu, Zheng Wang, Yuhan Li, Haoran You, Ruijie Gao, Xiaoya Zhou, Sreenidhi Reedy Bommu, Yang Katie Zhao, and Yingyan Celine Lin. 2024. EDGE-LLM: Enabling Efficient Large Language Model Adaptation on Edge Devices via Layerwise Unified Compression and Adaptive Layer Tuning and Voting.
- [32] Chao Zeng, Songwei Liu, Yusheng Xie, Hong Liu, Xiaojiang Wang, Miao Wei, Shu Yang, Fangmin Chen, and Xing Mei. 2024. ABQ-LLM: Arbitrary-Bit Quantized Inference Acceleration for Large Language Models.
- [33] Mingjin Zhang, Xiaoming Shen, Jiannong Cao, Zeyang Cui, and Shan Jiang. 2025. EdgeShare: Efficient LLM Inference via Collaborative Edge Computing. *IEEE Internet of Things Journal* 12, 10 (2025), 13119–13131. doi:10.1109/JIOT.2024.3524255