
Research Article

Intelligent Orchestration of Distributed Large Foundation Model Inference at the Edge

Fernando Koch^{1*}, Aladin Djuhera², Alecio Binotto³

¹ Department of Electrical Engineering and Computer Sciences, Florida Atlantic University, Boca Raton, FL, USA

² Chair of Theoretical Information Technology, Technical University Munich, Munich, Germany

³ Zeiss Digital Partners, Carl Zeiss AG, Munich, Germany

E-mail: kochf@fau.edu

Received: 20 March 2025; Revised: 18 June 2025; Accepted: 27 June 2025

Abstract: Large Foundation Models (LFMs), including multi-modal and generative models, promise to unlock new capabilities for next-generation Edge AI applications. However, performing inference with LFM in resource-constrained and heterogeneous edge environments, such as Multi-access Edge Computing (MEC), presents significant challenges for workload orchestration due to time-varying network, compute, and storage conditions. In particular, current split inference strategies, which partition LFM layers across nodes, are not designed to adapt to fluctuating workloads, dynamic bandwidth conditions, or evolving privacy constraints in high-utilization MEC environments. In this work, we propose a novel *adaptive split inference orchestration framework* that elevates both the placement and partitioning of LFM layers to runtime-tunable variables. Specifically, our framework enables real-time, quality-of-service (QoS)-aware management of inference workloads by extending conventional orchestrators with three key services: (1) *Capacity-aware workload distribution*, which continuously profiles node resources and selects an optimal subset of MEC nodes; (2) *Dynamic partition migration*, which transparently relocates pre-cut LFM segments in response to changes in utilization or network conditions; (3) *Real-time reconfiguration*, which dynamically re-splits LFM layers to balance latency, throughput, and privacy. We formalize the joint placement-partitioning problem, outline a reference architecture and algorithmic workflow, and discuss applicability in representative smart city, V2X, and industrial edge scenarios.

Keywords: large language models, split inference, edge artificial intelligence, edge computing

1. Introduction

Inference, the real-time forward execution of a trained model, is typically less compute-intensive than model training and may require substantial resources for layer activation, memory, and model storage [1]. This holds particularly for Large Foundation Models (LFMs) such as transformer-based large language models (LLMs) [2], e.g., Llama [3] and Qwen [4] models. Performing inference efficiently on the edge remains challenging, especially in Multi-Access Edge Computing (MEC) environments with limited and heterogeneous compute resources [5].

Distributed Split Inference (DSI) [6] has emerged as a promising approach to mitigate this problem. This strategy aims to partition an LFM into multiple segments that are executed sequentially across different nodes.

DSI strategies have typically employed *static splits* of inference workloads, where some computation is executed locally, while heavier computation tasks are outsourced, alleviating the computational burden on the client device. Such *splits* are mostly predetermined before execution and therefore lack adaptability to dynamic and heterogeneous operational conditions, such as fluctuating network reliability, changing node utilization, or intermittent connectivity. Consequently, these approaches produce suboptimal performance, compromising latency, resource utilization, and quality of service (QoS) guarantees, especially in mission-critical or latency-sensitive applications such as those found in financial services, industrial manufacturing, retail operations, and logistics [7]. This problem becomes even more acute in *resource-constrained* and *heterogeneous* edge environments, where multiple users rely on accessing shared edge compute resources, and where data privacy regulations (e.g., GDPR [8]) often restrict offloading computations to the cloud. For instance, executing a 7B parameter LLM may incur 25 ms on an NVIDIA RTX A6000 edge node but over 250 ms on an NVIDIA Jetson accelerator. Meanwhile, backhaul latency can oscillate between sub-1 ms (mmWave) and 30 ms (congested Wi-Fi) [9, 10]. Such volatility renders any *a priori* static split untenable.

Contemporary orchestration frameworks such as Kubernetes, Ray Serve, InferLine, and KubeEdge excel at container or micro-batch scheduling but treat AI models as black boxes, lacking mechanisms for runtime layer re-partitioning or privacy-aware placement [11].

Further, recent works, e.g., on jamming resilience for Split Federated Learning with LLMs [12], Federated Split Learning for satellite-terrestrial networks [13], and on Split Federated Mutual Learning (SFML) for traffic classification [14], address training-time collaboration, yet still assume fixed split points. Consequently, the key problem of *joint, model-aware partition and placement* under dynamic edge conditions remains open.

In this work, we address this problem by introducing an *adaptive split inference orchestration framework*, extending existing workload orchestration systems with domain-specific capabilities that are specifically tailored for LFMs, such as (multi-modal) LLMs. We introduce the following capabilities by leveraging the modular architecture of these models:

1. **Distribution of workloads to edge nodes** that offer better performance or operational capacity than the original source node.
2. **Redistribution of split LFM partitions** across connected edge nodes to dynamically optimize resources under changing conditions.
3. **Adaptive reconfiguration of model splitting** (e.g., re-splitting) to further improve performance and resource utilization when required.

Our framework operates on the *computational graph of the LFM itself*, allowing decisions at the granularity of individual transformer blocks. This enables QoS-driven re-splitting that commodity orchestrators cannot express. Furthermore, since our solution emphasizes split inference, *privacy* can be implemented as an additional feature at no cost if sensitive LFM layers can be executed locally, which makes reverse engineering data from model weights significantly more challenging for attackers [15]. Through this approach, we establish a foundation for *privacy-preserving, real-time*, and *QoS-aware* AI inference in edge networks, aligning with key 6G objectives of seamless connectivity, low inference latency, and intelligent edge resource management [16]. Importantly, our framework elevates privacy guarantees by localizing sensitive computations, thereby reducing regulatory exposure and aligning with emerging compliance standards.

Our contributions to the Edge AI landscape are multifold:

- **A reference architecture** for adaptive orchestration of distributed inference workloads;
- **A method for dynamic redistribution** of model segments to accommodate fluctuating compute resources and connectivity, and;
- **Establish mechanisms** for real-time, Service-Level Agreement (SLA)-compliant partitioning that align inference execution with QoS targets.

2. Background and Related Work

2.1 Large Foundation Models at the Edge

Next-generation 6G-enabled services, for example in dense urban environments, will need to support a multitude of AI-driven applications underpinned by LFM [17, 18]. However, deploying such large models typically requires significant computational resources and raises privacy concerns when handling sensitive data, making their adoption particularly challenging for inference in edge environments [19, 20]. Further, as organizations strive to keep data on-premise (e.g., for regulatory compliance such as GDPR [8]), optimizing *distributed inference* becomes a critical enabler of low-latency and privacy-preserving AI services [1, 21].

For that, next-generation networks make use of MEC infrastructures [22], which embed computing resources directly within the network. While this brings compute closer to end-users, a single MEC-enabled base station can quickly become saturated by various inference workloads such as from smart city and crowd management, personalized user applications, or industrial applications [23]. Looking ahead, 6G is envisaged to evolve beyond a mere network infrastructure upgrade into a *trustworthy* [24] and *intelligent workload orchestration system*, enabling distributed, LFM-based AI services that seamlessly shift computation across user devices, edge, and cloud as network and compute conditions change [25, 26].

2.2 Challenges in AI Workload Orchestration

The current industry norm has been to integrate general-purpose orchestration platforms (e.g., Kubernetes or proprietary MEC orchestrators), which facilitate application deployment and scaling but were not designed for challenges in inference scaling of modern LFM architectures, including fine-grained model partitioning and real-time, potentially hardware-accelerated inference optimization [27]. Existing solutions thus primarily target stateless services or relatively simple microservices, neglecting the unique requirements of edge-based AI pipelines, particularly for dynamic model splitting, QoS-driven scheduling, and adaptive resource reallocation [28].

In particular, in Edge AI scenarios involving LFM, traditional workload orchestration platforms fall short of meeting Edge AI-centric inference demands as they lack mechanisms to dynamically redistribute or reconfigure large model partitions based on real-time changes in *network conditions*, *node utilization*, or *connectivity* [29]. While model-serving stacks such as Triton, InferLine, Ray Serve, and MLC-Serve introduce batching or replica autoscaling, they still treat neural networks as opaque binaries and cannot *re-partition* a model graph at runtime. Edge extensions like KubeEdge and OpenYurt inherit the same pod-level abstraction, leaving the joint problem of *runtime layer splitting and placement*, especially under privacy and QoS constraints, unaddressed (see Table 1).

Table 1. Comparison of **capability coverage** across mainstream workload orchestrators / serving stacks and our proposed *Adaptive Split Orchestrator*. ✓ indicates native support, ✗ indicates the capability is absent, and entries in parentheses denote the mechanism (e.g., *pods*, *replicas*). Our framework simultaneously (i) reasons over the *layer graph*, (ii) adapts *split points* and *placements at runtime*, (iii) schedules under explicit *QoS/SLA* targets, and (iv) enforces *privacy-aware* layer placement.

Orchestrator	Layer-Aware Graph	Runtime Split	Placement Adapt.	QoS-Aware	Privacy Support
Kubernetes	✗	✗	✓ (pods)	✗	✗
KubeEdge	✗	✗	✓ (edge pods)	✗	✗
Ray Serve	✗	✗	✓ (replicas)	✓ (latency tiers)	✗
NVIDIA Triton	✗	✗	✓ (inst. groups)	✓ (batch/queue)	✗
InferLine	✗	✗	✓ (replicas)	✓ (tail-latency)	✗
EdgeShard [30]	✓ (manual)	✓ (manual)	✗	✗	N/A
Ours	✓	✓ (runtime)	✓ (runtime)	✓ (SLA-aware)	✓ (layer scope)

As a result:

- **Latency spikes** occur when critical links become congested, delaying real-time applications.
- **Straggler problems** arise when tasks are bottlenecked on overloaded or slower nodes, degrading overall QoS.
- **Resource utilization** becomes imbalanced, either overloading certain nodes or leaving others underutilized, leading to missed SLAs.
- **Privacy risks** escalate when large volumes of sensitive data must be offloaded to remote servers due to inadequate local processing.

While current research predominantly focuses on efficient AI model *training* (e.g., hardware-efficient training [31], quantization [32], and federated learning [33]), the practical challenges of *inference scaling and efficiency* at the edge remain relatively overlooked [1]. Yet, these inference-related challenges are increasingly critical for the widespread adoption of LFM in industrial and commercial scenarios, particularly in future *AI as a Service* (AIaaS)-driven 6G networks [17].

2.3 Distributed and Adaptive Split Inference

To alleviate computational demands, *distributed split inference* (DSI) [34] has emerged as an approach within Edge AI, which partitions a model across different compute locations (e.g., client device, MEC node, cloud) to balance local processing with remote offloading. Here, early, lightweight, or privacy-sensitive model layers are often chosen to be executed locally on-device or on trusted MEC nodes, extracting compact feature maps to reduce data transmission overhead. Subsequent layers are then offloaded to remote cloud environments and processed by more powerful servers [34, 35].

There are different techniques to distribute workloads on heterogeneous computing environments. For example, the authors of [36, 37] proposed heuristics to distribute workloads among CPUs and GPUs (or any computing units), monitoring their execution and dynamically adapting the overall scheduling giving changing conditions to, e.g., maximize performance or minimize energy consumption.

Although DSI enables larger AI models to operate closer to data sources, current implementations predominantly employ static splits defined *a priori* based on expected conditions, without runtime adaptation. While some studies, such as EdgeShard [30], explore collaborative inference setups where a model is shared across edge nodes, these approaches continue to lack dynamic orchestration of model splits and thus cannot effectively respond to real-world changes in edge environments, such as fluctuating node workloads, intermittent connectivity, variable network reliability, or dynamically changing service demands. As a result, traditional solutions frequently lead to suboptimal latency, inefficient resource utilization, degraded service quality, and decreased compliance with QoS guarantees and SLAs [38]. Consequently, recent research highlights the benefits of *adaptive split inference*, wherein partition points or even partition strategies (e.g., layer reordering) can be reconfigured at runtime to maintain QoS under shifting conditions [28]. This approach, combined with optimal orchestration policies, has the potential to cater to the increasingly demanding AI inference workloads in future AIaaS 6G-enabled networks and edge environments.

In addition to inference, recent split learning approaches illustrate the state of today's limitations from the model training point of view. R-SFLLM [12] freezes an LLM split immediately after the embedding layer of a transformer so that raw tokens never leave the device, and then analyses jamming resilience for that single layout. FedSL-LSTM [13] for satellite-terrestrial anomaly detection pre-divides each LSTM into client- and server-side subnetworks before training begins. The interface remains unchanged to keep gateway compute low and ground timing predictable. SFML [14] adopts a fixed three-segment CNN in which only lightweight "head+tail" layers stay on edge routers while the compute-heavy core runs in the cloud, simplifying privacy guarantees but eliminating runtime flexibility. Across all three studies, split points are determined *a priori* and never migrated, leaving questions of dynamic re-splitting, QoS-driven placement, and privacy-aware adaptation unanswered, precisely the gaps our adaptive orchestrator addresses. Table 2 summarizes these representative frameworks, highlighting their fixed partitioning schemes and the lack of adaptive capabilities.

Thus, in standard implementations, orchestrators cannot dynamically decide to offload additional LFM layers to another edge node or re-split the network, leaving significant performance and reliability gains unrealized. Note that this problem *does not* originate from within the network itself as current 5G and future 6G architectures natively implement adaptive resource management strategies for various services and network slices [39]. In current deployments, effectiveness thus remains limited by the absence of real-time dynamic orchestration policies tailored explicitly to modern AI workloads and complex LFM deployment scenarios. This is especially problematic for heterogeneous compute nodes, which complicates uniform deployment strategies. Thus, a *one-size-fits-all* static partitioning rarely works, as local workloads, performance constraints, and available resources differ from one site to another [19].

Table 2. Comparison of representative split learning frameworks with respect to their application context, model partitioning strategies, and (lack of) adaptivity. While each approach adopts fixed split points to simplify deployment or privacy guarantees, none supports dynamic re-partitioning or runtime migration, highlighting the need for adaptive orchestration in real-world, QoS-sensitive Edge AI scenarios.

	R-SFLLM	FedSL-LSTM	SFML
Application	Jamming-resilient FL with LLMs over 6G wireless links	LSTM anomaly detection in satellite-terrestrial integrated networks	Encrypted-traffic classification on edge routers + cloud
Split Layer	After the <i>embedding</i> block; attention layers off-device	First LSTM layers on gateways; remainder on ground server	Lightweight head + tail on router; CNN core in cloud
Adaptivity	Static (single two-cut split, never moved during training)	Static (single two-cut split, never moved during training)	Static (single three-cut split, never moved during training)

2.4 Key Design Goals for Adaptive LFM Split Inference

Despite evidence that splitting models can significantly improve efficiency and privacy, practical deployments remain constrained by *static* or *coarse-grained* orchestration mechanisms [7, 40]. Today’s solutions thus seldom adapt to shifting network or compute conditions in real time, leading to latency spikes, resource imbalances, and potential SLA and QoS violations [41]. Meanwhile, next-generation 6G network architectures will further exacerbate the complexity of distributing large-scale inference workloads across heterogeneous edge topologies to support various commercial and operational AIaaS applications [18].

Hence, an *adaptive split inference* framework will be required that:

1. Dynamically reconfigures the partition of LFM layers among edge and cloud compute nodes,
2. Exploits real-time profiling of resource availability and network conditions,
3. Preserves data privacy by keeping sensitive computations locally, and
4. Ensures consistent, QoS-compliant performance under fluctuating workloads.

In the next section, we propose a novel orchestration method that closes this gap by intelligently managing LFM inference across edge compute infrastructures.

3. Proposal

Recall the dense urban environment scenario from earlier where multiple users split their inference workloads between their local devices and an edge node (e.g., 5G-MEC). To this end, the LFMs must be partitioned accordingly, resulting in several different split configurations depending on local compute capacity, privacy requirements, and edge node capacity. In this scenario, we need to address the following three problems:

1. How to ensure that split inference *can indeed* take place on the assigned edge node given QoS and/or SLA requirements?

2. How to *redistribute* the split inference request to other candidate nodes in a connected region in case inference on the originally assigned node is not possible?
3. How to *dynamically revise* suboptimal LFM splits to obtain the best possible configuration given the local- and wide-area edge compute capacity?

To address these challenges, we propose an *adaptive split inference orchestration* framework that dynamically manages LFM partitions across heterogeneous edge nodes. Figure 1 depicts a possible realization of this framework in a 5G/6G-MEC deployment, including key components for monitoring, decision-making, model partitioning, and reconfiguration. We outline a detailed reference architecture as follows.

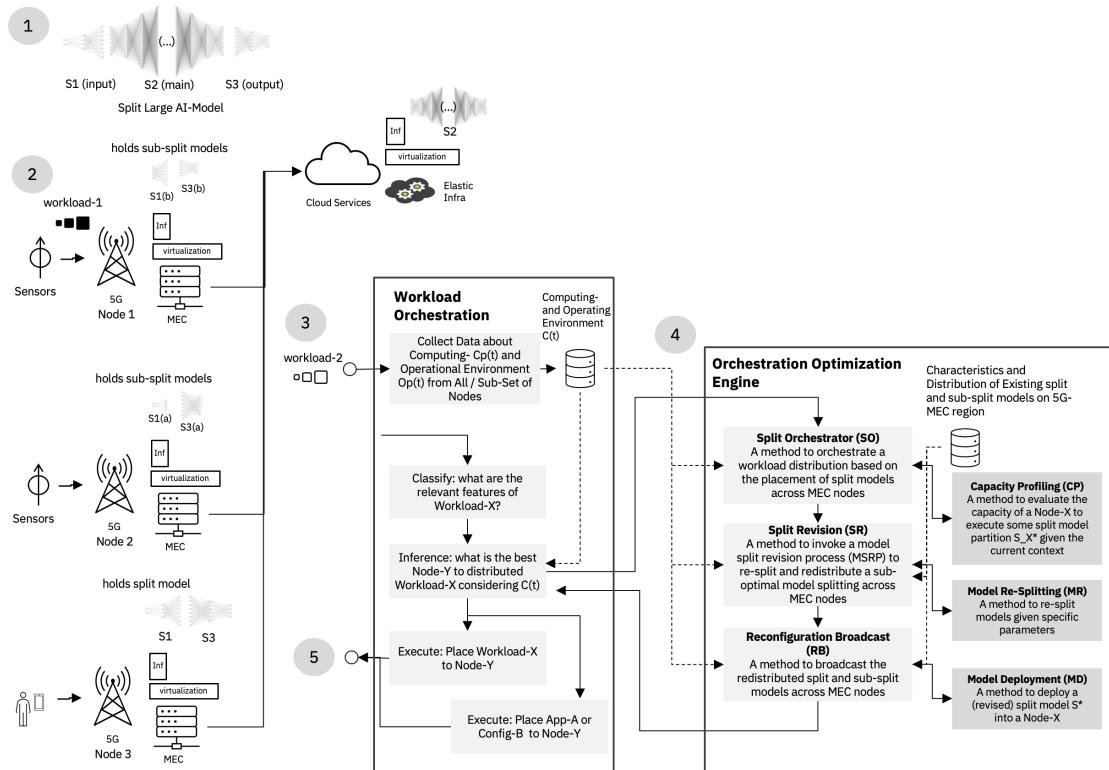


Figure 1. Reference architecture of the proposed adaptive split inference orchestration. Sub-split models (S1, S2, S3) are deployed across edge/cloud nodes, while a central orchestrator, guided by real-time capacity profiling, re-splits and reconfigures workloads on demand to meet QoS and privacy constraints. A corresponding workflow diagram of our proposed Algorithm 1 is given in Figure 2.

3.1 Reference Architecture

Our framework orchestrates on-demand allocation and reallocation of LFM partitions under evolving operational conditions via the following core modules:

1. **Monitoring & Capacity Profiling (CP):** Collects real-time metrics from edge nodes and the network environment, such as CPU/GPU utilization, memory usage, bandwidth, and latency. These metrics guide the orchestrator in partition placement and potential re-splitting decisions.
2. **Adaptive Orchestrator (AO):** Acts as the decision-making engine by evaluating whether to:
 - *Keep* the current split (no changes).

- *Redistribute* sub-splits across underutilized or more capable nodes.
- *Fully re-split* the model to find an updated partition configuration.

These decisions are informed by constraints like node capacity, privacy requirements, and expected QoS.

3. **Split Revision (SR):** Implements the logic to re-partition the LFM at different layer boundaries or blocks. This module may use heuristic, rule-based, or learning-based strategies to identify improved splits, respecting constraints such as local privacy boundaries.
4. **Reconfiguration Broadcast (RB):** Propagates new model partitions or sub-partitions to the selected nodes and updates local or remote orchestrators, ensuring future inference requests follow the revised configuration.

Our approach dynamically adapts split inference to fluctuating conditions while maintaining strict QoS and privacy requirements by combining these modules. Although **Kubernetes** already offers node-level telemetry, a pod scheduler, and ReplicaSet roll-outs, it treats the AI model as an *opaque container*. Instead, our *Monitoring & CP* module augments the standard metrics stream with *layer-granular* latency, activation size, and privacy tags. The *Adaptive Orchestrator* then optimizes both placement and split boundaries, whereas Kubernetes can only relocate whole pods. Finally, the *Reconfiguration Broadcast* disseminates on-the-fly-generated weight shards and graph rewiring commands—operations that cannot be expressed through a deployment update or ConfigMap patch. Collectively, these extensions lift orchestration from container granularity to *LFM-graph granularity*, which is essential for runtime re-splitting under QoS and privacy constraints.

The next subsections formalize the system model, define constraints, and describe the orchestration workflow for LFMs in detail.

3.2 Notation and System Model

We define key terminologies and orchestration concepts that underlie our adaptive split inference framework as follows.

Edge Nodes and Cloud. Let $\mathcal{N} = \{1, 2, \dots, n\}$ denote the set of n edge nodes, and let c refer to a (potential) cloud node (more capacity but increased latency). Each node $j \in \mathcal{N} \cup \{c\}$ has resource capacities for inference at time t , captured via capacity profiling (CP) as:

$$\text{CP}(n_j, t) = \{\text{CPU}_j(t), \text{GPU}_j(t), \text{Mem}_j(t), \text{NetCap}_j(t)\}, \quad (1)$$

which vary with concurrent workloads, hardware, and network conditions.

Model Partitioning. Consider an LFM \mathcal{M} segmented into k partitions or layers (e.g., from Transformer or neural networks architectures):

$$S = \{S_1, S_2, \dots, S_k\}. \quad (2)$$

Typically, S_1 handles raw (potentially private) data, and S_k generates the final outputs. Intermediate segments S_2, \dots, S_{k-1} often encompass the bulk of computation (e.g., multi-head self-attention in Transformers). A three-split example (S_1, S_2, S_3) might place S_1, S_3 on a local edge node (for privacy where user data is translated into/from vector embeddings) and offload the compute-intensive S_2 to a more capable edge or cloud node. Depending on the specific LFM architecture, splits may either be configured as self-contained building blocks (e.g., embeddings, self-attentions) or individual layers (e.g., from deep or convolutional neural networks) [7, 12].

Inference Requests. Inference tasks arrive as requests $\{r_1, r_2, \dots\}$, each with an associated workload \mathcal{W}_r . At a high level, each request utilizes the same model partitions $\{S_1, \dots, S_k\}$, but may require separate scheduling decisions depending on QoS constraints or real-time node capacities. We can treat each request as an instance of the partition assignment problem or, if simultaneous requests must be handled, sum over their respective costs when formulating a corresponding objective function.

Decision Variables. For computational convenience, we define a binary placement matrix $\mathbf{x} = [x_{i,j}]$, where $x_{i,j} = 1$ indicates partition S_j is assigned to node n_i , and $x_{c,j} = 1$ indicates assignment to the cloud node c . Each column corresponds to a partition, and each row to a node in $\mathcal{N} \cup \{c\}$. When multiple requests are considered, either the same \mathbf{x} can be reused if the system enforces a single partition layout, or a time/index extension can be introduced (e.g., \mathbf{x}_r for each request r).

With these concepts and terminologies in place, we may define an appropriate optimization objective as follows.

Objective Function. We aim to minimize the high-level cost:

$$\Phi(\mathbf{x}, \mathcal{C}(t)) = \alpha \mathcal{L}(\mathbf{x}, \mathcal{C}(t)) + \beta \mathcal{U}(\mathbf{x}, \mathcal{C}(t)) + \gamma \mathcal{P}(\mathbf{x}, \mathcal{C}(t)), \quad (3)$$

where:

- \mathcal{L} measures inference latency, including data transfer.
- \mathcal{U} captures resource usage imbalance or node overload.
- \mathcal{P} penalizes privacy violations (e.g., placing sensitive partitions on untrusted nodes).
- $\alpha, \beta, \gamma \geq 0$ weight the relative importance of latency, resource usage, and privacy, respectively.

Here, $\mathcal{C}(t)$ encapsulates the system state at time t , including node capacities, network bandwidths, and any QoS or SLA requirements. In scenarios with multiple concurrent requests, Φ can be extended to represent the sum or average cost across all active requests. In addition, to ensure valid assignments, we impose the following constraints:

1. **Unique Assignment.** Each partition S_j must be placed on exactly one node:

$$\sum_{i \in \mathcal{N}} x_{i,j} + x_{c,j} = 1, \quad \forall j \in \{1, \dots, k\}. \quad (4)$$

2. **Capacity Limits.** For each node $n_i \in \mathcal{N}$, the sum of resource loads from its assigned partitions cannot exceed that node's capacity:

$$\sum_{j=1}^k \text{load}(S_j) x_{i,j} \leq \text{capacity}(n_i, t). \quad (5)$$

An analogous constraint applies to the cloud node c if cloud resources are finite.

3. **Privacy Constraints.** Partitions handling sensitive data (e.g., S_1) must remain on trusted nodes:

$$x_{i,j} = 0, \quad \text{if } n_i \notin \text{trustedSet} \wedge (S_j \text{ is privacy-critical}). \quad (6)$$

Further, if LFM layer boundaries can be modified (e.g., subdividing S_2 into $\{S_{2a}, S_{2b}\}$, as for example in neural network layers, Transformer embeddings and attentions, etc.), we may treat the set of partitions S itself as part of the optimization. Herewith, we define the split revision as follows.

Split Revision. Let Ω denote the set of all valid splitting schemes. The orchestrator aims to solve:

$$\min_{S \in \Omega, \mathbf{x}} \Phi(\mathbf{x}, S, \mathcal{C}(t)), \quad (7)$$

to find an optimal split S^* and assignment \mathbf{x}^* that minimize the overall cost subject to the constraints above. This allows partitions and assignments to adapt dynamically to shifts in resource availability, privacy requirements, or workload demands, initiated and managed by the adaptive orchestrator.

3.3 Orchestration Workflow

Algorithm 1 outlines the main orchestration steps. The workflow begins by deploying a *baseline* partitioning (e.g., (S_1, S_2, S_3)) among a set of nodes. The system then continuously monitors resource usage and performance metrics to trigger dynamic adjustments.

1. **Initial Deployment.** Perform a static partitioning of the model based on coarse performance estimates (e.g., place S_1, S_3 locally for privacy, and put S_2 on a more powerful node or cloud instance c).
2. **Continuous Monitoring.** The Monitoring & CP module collects real-time metrics $CP(n_j, t)$ and calculates an *environment state* $\mathbf{E}(t)$ that captures fluctuations in node utilization, network throughput, or latency.

Algorithm 1: Adaptive Split Orchestration Workflow

Input: (i) Initial partitioning $\{S_1, \dots, S_P\}$, (ii) baseline mapping d_0 , (iii) monitoring intervals Δt , (iv) trigger-threshold vector $\Theta = \{L_{\max}, U_{\max}, B_{\min}, T_{\text{cool}}\}$

Initialize: Deploy baseline split (S_1, \dots, S_P) across nodes as per d_0 .

Set $t_{\text{last}} \leftarrow -\infty$

for each monitoring cycle $t \leftarrow 0, \Delta t, 2\Delta t, \dots$ **do**

 Collect environment metrics $\mathbf{E}(t)$ via Monitoring & CP.

$\text{reconf} \leftarrow \text{ShouldReconfigure}(\mathbf{E}(t), \Theta)$.

if (trigger condition is met, e.g., high latency, node overload, etc.) and reconf **then**

 Evaluate feasible mappings $\{d'\}$ given current partitions.

 Optionally call Model Re-Splitting to produce new partitions $\{S_i^*\}$.

 Determine best mapping $\hat{d} = \arg \min_{d'} C(d')$.

if $\hat{d} \neq d_t$ and $t - t_{\text{last}} \geq T_{\text{cool}}$ **then**

 Broadcast reconfiguration to all affected nodes via RB.

$t_{\text{last}} \leftarrow t$; $d_{t+\Delta t} \leftarrow \hat{d}$.

 Resume inference under current assignment $d_{t+\Delta t}$.

3. **Adaptive Decisions.** Based on the updated system states $\mathcal{C}(t), \mathbf{E}(t)$, the adaptive orchestrator continuously evaluates whether to *keep the current split* (if performance remains within SLA targets), *redistribute sub-splits* (reassigning some partitions S_j from node n_i to $n_{i'}$ by adjusting \mathbf{x} without altering the partition boundaries), or *perform full re-splitting* (to obtain a better partition set S^* via the SR module if incremental changes are insufficient or new privacy constraints arise). More formally,

- The adaptive orchestrator evaluates whether the current partition mapping d_t remains optimal under $\mathbf{E}(t)$. For each request r , the orchestrator checks:

$$\mathcal{C}(d_t) \stackrel{?}{\leq} \mathcal{C}(d') \quad \forall \text{ feasible } d'. \quad (8)$$

If a lower-cost (or higher-utility) mapping d' is found, a reconfiguration is triggered.

- If needed, the SR module modifies the set of partitions $\{S_1, \dots, S_P\}$ (e.g., subdividing a large block S_2 into new split configurations $\{S_{2a}, S_{2b}\}$), i.e.

$$\hat{d} = \underset{d \in \mathcal{D}(\text{new splits})}{\operatorname{argmin}} \mathcal{C}(d), \quad (9)$$

subject to constraints (e.g., compute, network, privacy).

4. **Reconfiguration Broadcast (RB).** Once a decision is made, the *RB* module disseminates the updated assignment \mathbf{x}^* or partition set S^* to relevant nodes, ensuring the new configuration is deployed consistently.

5. **Execution.** Inference resumes with the updated partition assignment \hat{d} . The orchestrator continues to monitor performance, forming a feedback loop, allowing the system to adapt further as conditions evolve.

Additional Trigger Conditions and Decision Logic. Table 3 summarizes the runtime metrics that feed the function `ShouldReconfigure(E(t), Θ)` in Algorithm 1. The orchestrator invokes a reconfiguration if *any* of the following holds for a monitoring window of length Δt :

1. **Latency threshold.** The exponentially weighted moving average (EWMA) of end-to-end inference latency $\bar{L}(t, \Delta t)$ exceeds L_{\max} , i.e. $\bar{L}(t, \Delta t) > L_{\max}$.
2. **Utilisation threshold.** The maximum GPU or CPU utilization across all nodes exceeds U_{\max} , i.e., $\max_{n \in \mathcal{N}} U_n(t) > U_{\max}$.
3. **Bandwidth drop.** The minimum available bandwidth across any active edge link drops below B_{\min} , i.e., $\min_{(i,j) \in \mathcal{L}} B_{ij}(t) < B_{\min}$.
4. **Privacy policy violation.** A new inference request is tagged with the identifier `privacy=high`, but the current partitioning would route raw features through an untrusted node.

Table 3. Monitored metrics and default trigger thresholds.

Metric	Symbol	Default value
EWMA latency	L_{\max}	150 ms
GPU/CPU utilisation	U_{\max}	0.85
Available link bandwidth (edge → edge)	B_{\min}	50 Mbps
Time-to-reconfigure cool-down	T_{cool}	30 s

If multiple triggers fire simultaneously, the system first attempts *placement* migration. If that cannot meet all constraints, the *split-revision* module is invoked. Reconfigurations are rate-limited by T_{cool} to prevent thrashing. The full control loop of our adaptive orchestration algorithm is illustrated in Figure 2, detailing its monitoring, decision, and reconfiguration stages.

Above outlined system model and orchestration workflow provide possible entry points for optimizations in real-world deployments. In practice, such an orchestration loop can be integrated into existing container platforms (e.g., extending Kubernetes with a custom controller that triggers model re-splitting when monitoring thresholds are exceeded). Partitioning decisions may rely on traditional heuristics (e.g., rule-based or greedy approaches) or adopt learning-based schemes (e.g., reinforcement learning) to continuously refine splitting strategies [42–44]. Alternatively, Python-based pipelines could invoke layer-partitioning heuristics based on state-of-the-art open-source frameworks, such as Huggingface, and broadcast updates via RESTful APIs. These approaches enable straightforward adoption of adaptive split inference within both on-premise and cloud-based edge deployments.

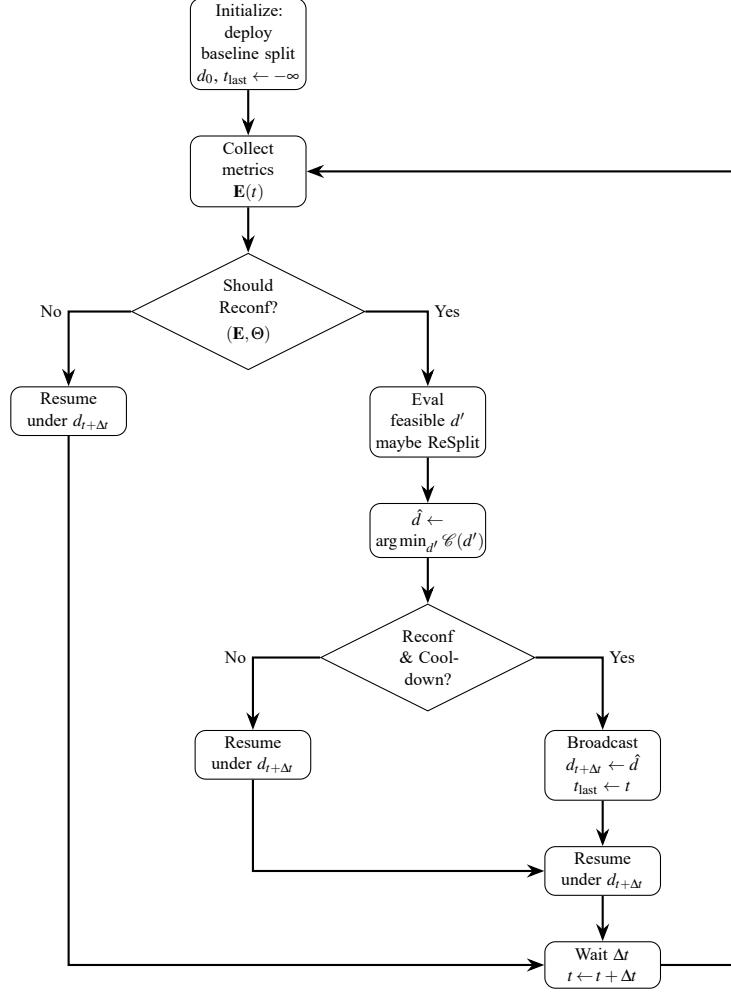


Figure 2. Control-flow diagram of the adaptive split orchestration loop described in Algorithm 1. The orchestrator periodically monitors environment metrics and triggers reconfiguration decisions when QoS thresholds or privacy constraints are violated. Feasible placements are evaluated, and, if no cool-down limit is active, a new mapping is broadcast to all nodes.

3.4 Privacy and Security Considerations

A core feature behind split inference is the preservation of data privacy by ensuring critical or sensitive operations remain on a trusted device or node. Thus, our framework permits:

1. **Selective Local Execution:** Some LFM blocks, especially those close to the input layer, may handle raw personal or private data. By design, these partitions can be configured to remain on the user's device or a trusted edge node (e.g., for compliance with GDPR). Formally, if S_i handles privacy-critical data, we require that

$$d_t(i) \in \mathcal{N}_{\text{trusted}} \quad \forall t, \quad (10)$$

where $\mathcal{N}_{\text{trusted}} \subseteq \mathcal{N} \cup \{c\}$ is the set of *trusted* nodes. Corresponding LFM splits can be obtained according to the model architecture, compute resources and privacy requirements (e.g., measured as layer depth) [42].

2. **Secure Communication Channels:** Intermediate activations (e.g., outputs from S_1 that serve as inputs to S_2) can be additionally encrypted and transmitted securely to the next node in the chain. This ensures that eavesdropping or tampering with partial model data (e.g., due to jamming wireless transmissions [12]) is substantially harder. Further,

the *RB* component may include additional cryptographic signatures so that only valid reconfiguration commands from the orchestrator are honored.

3. **Partition Metadata Obfuscation:** To further reduce risk, the orchestrator can store only references to partial model weights or encrypted partitions in a registry accessible to each node, such that no single node (other than the one hosting a given partition) stores the raw weights.

Our orchestration framework thus extends standard orchestration platforms but adds specialized components for real-time capacity profiling, model splitting, and reconfiguration in response to varying network and compute conditions. By leveraging partial splits of LFM layers, the framework also inherently supports privacy-preserving inference at the edge, ensuring that sensitive data never leaves a trusted domain.

4. Use Cases

4.1 Emergency Coordination in Smart Cities

In a highly connected smart city environment, where autonomous AI agents are responsible for managing infrastructure, monitoring public safety, and responding to critical incidents [17], large-scale AI inference is crucial for maintaining operational efficiency. During emergency scenarios, such as regional blackouts, cyber-attacks on urban infrastructure, or natural disasters, adaptive split inference ensures real-time decision-making despite fluctuating resource availability.

Consider a scenario where a massive earthquake disrupts transportation networks, damages critical infrastructure, and impairs traditional cloud connectivity. Smart city AI agents deployed across distributed MEC nodes can play a pivotal role in orchestrating emergency response through *AI Agents for Autonomous Coordination*, i.e. specialized AI agents trained for disaster response. This can be part of infrastructure monitoring bots, autonomous drones, and emergency service assistants, which rely on continuous, high-throughput AI inference. For instance, these agents need to process high-dimensional multi-modal data, including real-time video, LiDAR scans, and sensor data from IoT devices to provide continuous responses.

In this environment, *Adaptive Model Deployments* are initially instantiated as foundation model partitions, which are distributed across MEC nodes based on predefined computational capabilities and expected workloads. As infrastructure degradation leads to unstable connectivity and hardware failures, the system dynamically adjusts model partitions across available MEC nodes. For instance, if an AI agent controlling autonomous emergency drones detects a surge in demand for real-time object detection (e.g., identifying survivors in debris), the system triggers split revision *SR* to redistribute workloads efficiently. Consequently, when an MEC node reaches its computational threshold due to a high influx of emergency data streams, the Reconfiguration Broadcast *RB* mechanism ensures that AI agents can offload inference tasks to alternative nodes with idle capacity. The system dynamically revises the model split $S = (S_1, S_2, S_3)$ into a more optimized configuration $S^* = (S_1^*, S_2^*, S_3^*)$, continuously adapting to the dynamic conditions to maintain operational efficiency and robust performance.

4.2 Industry 4.0 Manufacturing Lines

Modern manufacturing floors increasingly integrate edge AI for tasks like predictive maintenance, anomaly detection, and quality control, often under tight latency requirements. In such environments, multiple MEC nodes (e.g., private 5G/6G-MEC base stations) may handle continuous data streams from high-speed sensors and robotic arms. When production ramps up unexpectedly, compute workloads spike and nodes near their capacity limits. The orchestrator then responds accordingly by reassigning or re-splitting the inference model across less-loaded nodes, preventing bottlenecks. Privacy-constrained segments, such as those inspecting proprietary designs, remain on trusted hardware, while more generic modules can be offloaded seamlessly to boost throughput.

4.3 Autonomous Vehicles and Intelligent Transport Systems

Vehicle-to-infrastructure (V2I) services increasingly rely on advanced AI models for collision avoidance, route planning, and traffic flow optimization. Edge nodes at roadside units (RSUs) offer local compute to complement on-board vehicle processors, reducing latency while offloading computationally heavy layers. In busy urban corridors, traffic sensors and autonomous cars generate significant inference workloads. If congestion surges or a particular RSU becomes overloaded, the orchestrator redistributes model partitions among neighboring edge nodes, ensuring split inference scales effectively. Such real-time adaptivity allows vehicles to maintain continuous, low-latency awareness, meeting strict safety and efficiency standards that are critical in next-generation transport systems.

5. Expected Results and Quantitative Benefits

Table 4. Quantitative comparison of static and adaptive split inference across key performance dimensions in edge environments. Results are sourced from prior studies and testbed evaluations of LLM deployments over 5G-MEC infrastructures [28, 43, 45–49]

Performance Dimension	Static Split Inference	Adaptive Split Inference
Latency (End-to-End)	High variance; $\sim 500\text{--}1000$ ms typical in 5G-MEC scenarios. Cannot guarantee low latency under network fluctuation (may spike beyond 1 s).	Low and stable; $\sim 100\text{--}300$ ms under same conditions. Dynamically maintains latency below target (often 50%+ faster than static).
Throughput	Limited by weakest link (device or network); lower overall. < 1 req/s on a baseline edge device alone.	Higher via parallelism and load balancing; utilizes multiple nodes. Achieved > 5 req/s in adaptive multi-node setup ($\geq 5\times$ higher).
Resource Utilization	$\sim 50\text{--}60\%$ of available edge / cloud resources used. Static partition leads to idle resources on one side and potential overload on the other.	$\sim 80\text{--}95\%$ resource utilization. Orchestrator keeps both edge and cloud busy, scaling across nodes; more efficient use of CPU, GPU, and bandwidth.
QoS Compliance (SLA HitRate)	Often poor under variability. QoS deadlines met in $\sim 60\text{--}70\%$ of cases; frequent SLA violations when conditions deviate from design point.	Near-guaranteed QoS even as conditions change. $\sim 95\text{--}99\%$ of requests meet latency/SLA targets due to real-time adaptation (much fewer deadline misses).
Reliability & Failover	Rigid deployment; single-point bottlenecks. $\sim 5\text{--}10$ inference failures or timeouts per hour observed under load/network issues.	Resilient, with dynamic re-routing and re-partitioning. Downtime incidents reduced to $\sim 0\text{--}2$ per hour. Maintains service continuity by avoiding overloads.
Privacy Adherence	Moderate. Fixed layer offload may send sensitive features off device. No ability to alter behavior for sensitive data—potential compliance issues.	High. Can execute sensitive layers locally and limit data exposure. Adaptive policy balances performance with data confidentiality (privacy preserved without sacrificing QoS).

Table 4 presents a comparative analysis derived from prior studies, experimental testbeds, and empirical estimates for a representative 5G-MEC-enabled edge environment serving text-generation LLMs ranging from 7B to 13B parameters. For further technical depth and context, we refer to the comprehensive evaluations in [28, 43, 45–49].

In particular, it becomes evident that adaptive split inference offers significant improvements across multiple operational dimensions. First, it consistently achieves lower and more stable end-to-end latency, typically maintaining inference response times within 100–300 ms, even under fluctuating 5G-MEC conditions, whereas static configurations frequently exceed 500ms with high variance. This translates to a latency reduction of over 50% in representative scenarios. Second, throughput is substantially improved: adaptive orchestration enables parallel execution across distributed nodes, reaching over 5 req/s compared to sub-1 req/s throughput observed in static, single-device deployments.

Resource utilization also benefits considerably. Static splits often lead to underutilized compute on one end (e.g., idle cloud GPUs) and overloading on the other, capping effective usage at 50–60%. In contrast, adaptive orchestration dynamically balances workloads, yielding sustained utilization rates of 80–95% across edge and cloud components. This directly contributes to higher QoS adherence: while static inference often fails to meet SLA constraints during network or

load fluctuations, achieving target latency in only 60–70% of cases, adaptive inference maintains compliance for 95–99% of requests through real-time partition and placement adjustments.

Further, reliability under stress is markedly improved. Static deployments are prone to bottlenecks and failure cascades, with up to 10 inference errors or timeouts per hour observed under load. Adaptive strategies mitigate this via failover-aware routing and reconfiguration, reducing service interruptions to near-zero levels. In addition, privacy adherence is strengthened. While static splits may expose sensitive intermediate representations to untrusted infrastructure which can be targeted by malicious attackers, the adaptive approach can selectively retain privacy-critical layers on-device, enabling compliance with confidentiality requirements without degrading system performance.

Table 5 summarizes the discussed median figures reported across four public 5G-MEC studies [30, 43, 49, 50] and corresponding latency distributions are visualized in Figure 3 (static vs. adaptive CDF). The adaptive curve reaches the 95 % completion mark below 300 ms, while the static curve stretches beyond 1s, highlighting the practical QoS benefit of adaptive orchestration at runtime.

Table 5. Static vs. adaptive split inference for text-generation LLMs in typical 5G-MEC deployments (median values across the studies in [30, 43, 49, 50]).

Metric	Static Split	Adaptive Split
End-to-end latency	500–1000 ms	100–300 ms
Throughput (requests s^{-1})	~1	~5
GPU/CPU utilisation	50–60 %	80–95 %
SLA hit-rate (400 ms budget)	60–70 %	95–99 %
Downtime incidents (per h)	5–10	0–2
Privacy compliance	Moderate	High

Across all evaluative dimensions (latency, throughput, utilisation, QoS robustness, and privacy), *adaptive split inference* consistently outperforms any static configuration. The small overhead of monitoring (≤ 10 ms per cycle) and graph rewiring is amortised by hundreds of ms saved per request, yielding a net performance gain an order of magnitude larger than the orchestration cost. Consequently, adaptive split inference emerges as the decisive design choice for LLM deployment in heterogeneous, bandwidth-variable 5G/6G edge networks.

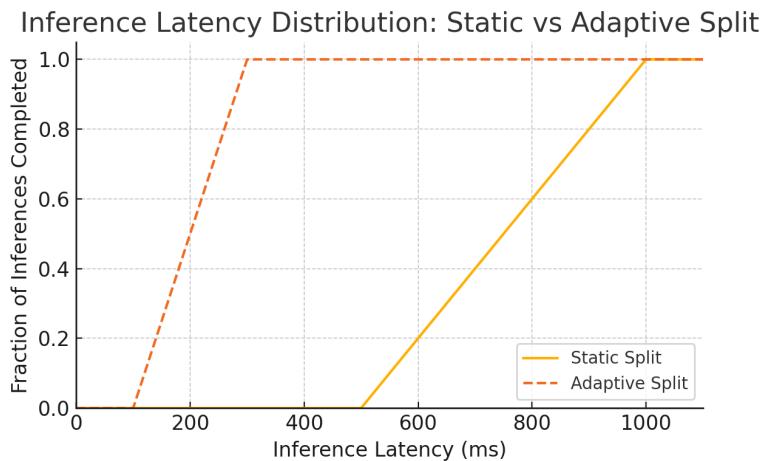


Figure 3. CDF of end-to-end inference latency for static (solid) vs. adaptive (dashed) split inference in a 5G-MEC scenario. 95 % of adaptive requests finish within 300 ms, while static requests may take up to 1s [30, 43, 49, 50].

6. Conclusions

This study has introduced an adaptive split inference orchestration framework designed to dynamically manage LFM partitions across heterogeneous edge nodes, addressing the inefficiencies inherent in static split inference methodologies. Our framework establishes a foundation for real-time, QoS-aware, and privacy-preserving AI inference in edge computing environments, which is particularly crucial for latency-sensitive and resource-constrained applications. The proposed approach optimizes performance, enhances resource efficiency, and fortifies privacy preservation by leveraging real-time monitoring, workload redistribution, and dynamic reconfiguration. In contrast to existing orchestration frameworks, we do not treat AI models as opaque containers, but instead support runtime model re-partitioning and privacy-aware layer placement, capabilities essential for deploying LFMs in dynamic edge environments. Our proposed orchestration model thus aligns with emerging objectives in the development of 6G networks by enabling intelligent, distributed AI processing at the edge. Its modular architecture facilitates integration with existing edge orchestrators while maintaining extensibility for future AI-driven optimizations.

To validate the practical relevance of our framework, we provided a comparative analysis grounded in prior studies, testbed results, and empirical evaluations. These results demonstrate consistent gains across latency, throughput, resource utilization, reliability, and privacy adherence, reinforcing the value of adaptive orchestration for real-world edge inference given the complexity of future foundation model deployments.

Nevertheless, several research directions emerge to further advance adaptive inference orchestration. Proposed and necessary items for a research agenda on future work include:

1. Investigating the deeper integration of AI-driven decision-making mechanisms, including reinforcement learning-based optimizations, to enhance inference orchestration.
2. Developing advanced privacy-preserving techniques, such as secure multi-party computation and homomorphic encryption, to ensure robust data security in distributed inference environments.
3. Designing adaptive network-aware partitioning strategies that dynamically adjust inference workload distribution based on real-time network conditions to optimize resource utilization and latency minimization.
4. Establishing standardized benchmarks and datasets for evaluating the performance of split inference frameworks in edge computing environments.

Beyond technical extensions, domain-specific deployments in autonomous systems, healthcare, and industrial automation present high-impact opportunities. For instance, we are working on an extension of our technology around augment reality in field operations [51] to improve local reasoning through distributed inference. We are also revising our previous idea about mesh-computing (grid computing) in management of telecommunication networks [52], applying this technology in distributed management for MEC environments. Likewise, examining the role of edge-driven inference in reducing energy consumption and enabling sustainable AI operations across smart cities and IoT ecosystems opens new interdisciplinary research avenues. Addressing these challenges will be essential to realizing scalable, trustworthy, and privacy-preserving AIaaS architectures in the forthcoming 6G edge landscape.

Acknowledgment

This paper provides a scholarly description of the published Patent *US20250071069A1* [53], along with references to published Patent *US18/449811* [54], which have been conceived and authored while the authors were employed at International Business Machines (IBM). IBM is the current assignee of both patent applications. Mr. Djuhera was supported by the German Federal Ministry of Research, Technology and Space (BMFTR) under *Grant 16KISK002*, as part of the *6G – life* research hub for the preparation of this manuscript and the supplementary experiments being presented.

Conflict of Interest

The authors declare that they have no conflicts of interest.

References

- [1] Zhou Z, Ning X, Hong K, Fu T, Xu J, Li S, et al. A survey on efficient inference for large language models. 2024. Available from: <https://arxiv.org/abs/2404.14294>. [Accessed 8 Mar 2025].
- [2] Vaswani A, Shazeer N, Parmar J, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Advances in Neural Information Processing Systems*. 2017; 30: 5998-6008.
- [3] Grattafiori A, Dubey A, Jauhri A, Pandey A, Kadian A, Al-Dahle A, et al. The Llama 3 herd of models. 2024. Available from: <https://arxiv.org/abs/2407.21783>. [Accessed 8 Mar 2025].
- [4] Yang A, Li A, Yang B, Zhang B, Hui B, Zheng B, et al. Qwen3 technical report. 2025. Available from: <https://arxiv.org/abs/2505.09388>.
- [5] Zhang X, Nie J, Huang Y, Xie G, Xiong Z, Liu J, et al. Beyond the cloud: edge inference for generative large language models in wireless networks. *IEEE Transactions on Wireless Communications*. 2024; 23(8): 7894-7908.
- [6] Karjee J, Naik SP, Anand K, Bhargav VN. Split computing: DNN inference partition with load balancing in IoT-edge platform for beyond 5G. *Measurement: Sensors*. 2022; 23: 100409.
- [7] Karjee J, Anand K, Bhargav VN, Naik PS, Dabbiru RBV, Srinidhi N. Split computing: dynamic partitioning and reliable communications in IoT-edge for 6G vision. In: *2021 8th International Conference on Future Internet of Things and Cloud (FiCloud)*; 2021 Aug 23-25; Virtual Conference. IEEE; 2021. p. 233-240.
- [8] European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council. Available from: <https://data.europa.eu/eli/reg/2016/679/oj>. [Accessed 8 Mar 2025].
- [9] Struye J, Lemic F, Famaey J. Towards ultra-low-latency mmWave Wi-Fi for multi-user interactive virtual reality. In: *GLOBECOM 2020 IEEE Global Communications Conference*; 2020 Dec 7-11; Taipei, Taiwan. IEEE; 2020. p. 1-6.
- [10] Lau I, Ekpo S, Zafar M, Ijaz M, Gibson A. Hybrid mmWave-LiFi 5G architecture for reconfigurable variable latency and data rate communications. *IEEE Access*. 2023; 11: 42850-42861.
- [11] Vasireddy I, Ramya G, Kandi P. Kubernetes and Docker load balancing: state-of-the-art techniques and challenges. *International Journal of Innovative Research in Engineering and Management*. 2023; 10(6): 49-54.
- [12] Djuhera A, Andrei VC, Li X, Mönich UJ, Boche H, Saad W. R-SFLLM: jamming resilient framework for split federated learning with large language models. 2024. Available from: <https://arxiv.org/abs/2407.11654>. [Accessed 8 Mar 2025].
- [13] Jiang W, Han H, Zhang Y, Mu J. Federated split learning for sequential data in satellite-terrestrial integrated networks. *Information Fusion*. 2024; 103: 102141.
- [14] Xia J, Wu M, Li P. SFML: a personalized, efficient, and privacy-preserving collaborative traffic classification architecture based on split learning and mutual learning. *Future Generation Computer Systems*. 2025; 162: 107487.
- [15] Xu R, Baracaldo N, Joshi J. Privacy-preserving machine learning: methods, challenges and directions. 2021. Available from: <https://arxiv.org/abs/2108.04417>. [Accessed 8 Mar 2025].
- [16] Letaief KB, Shi Y, Lu J, Lu J. Edge artificial intelligence for 6G: vision, enabling technologies, and applications. *IEEE Journal on Selected Areas in Communications*. 2021; 40(1): 5-36.
- [17] Saad W, Hashash O, Thomas CK, Chaccour C, Debbah M, Mandayam N, et al. Artificial general intelligence (AGI)-native wireless systems: a journey beyond 6G. 2024. Available from: <https://arxiv.org/abs/2405.02336>.
- [18] Zhou H, Hu C, Yuan Y, Cui Y, Jin Y, Chen C, et al. Large language model (LLM) for telecommunications: a comprehensive survey on principles, key techniques, and opportunities. *IEEE Communications Surveys & Tutorials*. 2024; 26(3): 1568-1602.
- [19] Li B, Jiang Y, Gadepally V, Tiwari D. LLM inference serving: survey of recent advances and opportunities. 2024. Available from: <https://arxiv.org/abs/2407.12391>. [Accessed 8 Mar 2025].
- [20] Yao Y, Duan J, Xu K, Cai Y, Sun Z, Zhang Y. A survey on large language model (LLM) security and privacy: the good, the bad, and the ugly. *High-Confidence Computing*. 2024; 4: 100211.
- [21] Li E, Zeng L, Zhou Z, Chen X. Edge AI: on-demand accelerating deep neural network inference via edge computing. *IEEE Transactions on Wireless Communications*. 2019; 19(1): 447-457.

[22] Filali A, Abouaomar A, Cherkaoui S, Kobbane A, Guizani M. Multi-access edge computing: a survey. *IEEE Access*. 2020; 8: 197017-197046.

[23] Lin Z, Qu G, Chen Q, Chen X, Chen Z, Huang K. Pushing large language models to the 6G edge: vision, challenges, and opportunities. 2023. Available from: <https://arxiv.org/abs/2309.16739>. [Accessed 8 Mar 2025].

[24] Fettweis GP, Boche H. On 6G and trustworthiness. *Communications of the ACM*. 2022; 65(4): 48-49.

[25] Camelo M, Cominardi L, Gramaglia M, Fiore M, Garcia-Saavedra A, Fuentes L, et al. Requirements and specifications for the orchestration of network intelligence in 6G. In: *2022 IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*; 2022 Jan 8-11; Las Vegas, NV, USA. IEEE; 2022. p. 1-9.

[26] Zeb S, Rathore MA, Hassan SA, Raza S, Dev K, Fortino G. Toward AI-enabled NextG networks with edge intelligence-assisted microservice orchestration. *IEEE Wireless Communications*. 2023; 30(3): 148-156.

[27] Chen W, Zhu Y, Liu J, Chen Y. Enhancing mobile edge computing with efficient load balancing using load estimation in ultra-dense network. *Sensors*. 2021; 21(9): 3135.

[28] Chen Y, Li R, Yu X, Zhao Z, Zhang H. Adaptive layer splitting for wireless LLM inference in edge computing: a model-based reinforcement learning approach. 2024. Available from: <https://arxiv.org/abs/2406.02616>. [Accessed 8 Mar 2025].

[29] Carrión C. Kubernetes scheduling: taxonomy, ongoing issues and challenges. *ACM Computing Surveys*. 2022; 55(7): 1-37.

[30] Zhang M, Shen X, Cao J, Cui Z, Jiang S. EdgeShard: efficient LLM inference via collaborative edge computing. *IEEE Internet of Things Journal*. 2024; 11(5): 7894-7908.

[31] Duan J, Zhang S, Wang Z, Jiang L, Qu W, Hu Q, et al. Efficient training of large language models on distributed infrastructures: a survey. 2024. Available from: <https://arxiv.org/abs/2407.20018>.

[32] Lang J, Guo Z, Huang S. A comprehensive study on quantization techniques for large language models. In: *2024 4th International Conference on Artificial Intelligence, Robotics, and Communication (ICAIRC)*; 2024 Jun 15-17; Singapore. IEEE; 2024. p. 224-231.

[33] Yao Y, Zhang J, Wu J, Huang C, Xia Y, Yu T, et al. Federated large language models: current progress and future directions. 2024. Available from: <https://arxiv.org/abs/2409.15723>. [Accessed 8 Mar 2025].

[34] Mohammed T, Joe-Wong C, Babbar R, Francesco MD. Distributed inference acceleration with adaptive DNN partitioning and offloading. In: *IEEE INFOCOM 2020 IEEE Conference on Computer Communications*; 2020 Jul 6-9; Virtual Conference. IEEE; 2020. p. 854-863.

[35] Timor N, Mamou J, Korat D, Berchansky M, Pereg O, Wasserblat M, et al. Distributed speculative inference (DSI): speculation parallelism for provably faster lossless language model inference. 2025. Available from: <https://arxiv.org/abs/2405.14105>. [Accessed 8 Mar 2025].

[36] Binotto APD, Pereira CE, Kuijper A, Stork A, Fellner DL. An Effective Dynamic Scheduling Runtime and Tuning System for Heterogeneous Multi and Many-Core Desktop Platforms. In: *IEEE International Conference on High Performance Computing and Communications*. 2011. p. 78-85.

[37] Binotto APD, Wehrmeister MA, Kuijper A, Pereira CE. Sm@rtConfig: A Context-Aware Runtime and Tuning System using an Aspect-Oriented Approach for Data Intensive Engineering Applications. *Control Engineering Practice*. 2013; 21(2): 204-217.

[38] Hudson N, Khamfroush H, Baughman M, Lucani DE, Chard K, Foster I. QoS-aware edge AI placement and scheduling with multiple implementations in FaaS-based edge computing. *Future Generation Computer Systems*. 2024; 157: 250-263.

[39] Thantharate A, Beard C. ADAPTIVE6G: adaptive resource management for network slicing architectures in current 5G and future 6G systems. *Journal of Network and Systems Management*. 2023; 31(1): 9.

[40] Zhou L, Wen H, Teodorescu R, Du DH. Distributing deep neural networks with containerized partitions at the edge. In: *2nd USENIX Workshop on Hot Topics in Edge Computing (HotEdge 19)*; 2019 Jul 11; Renton, WA, USA. USENIX; 2019.

[41] Li Y, Wang X, Gan X, Jin H, Fu L, Wang X. Learning-aided computation offloading for trusted collaborative mobile edge computing. *IEEE Transactions on Mobile Computing*. 2019; 19(12): 2833-2849.

[42] Li X, Bi S. Optimal AI model splitting and resource allocation for device-edge co-inference in multi-user wireless sensing systems. *IEEE Transactions on Wireless Communications*. 2024; 23(9): 11094-11108.

[43] Tuli S, Casale G, Jennings N R. SplitPlace: AI augmented splitting and placement of large-scale neural networks in mobile edge environments. 2022. Available from: <https://arxiv.org/abs/2205.10635>.

- [44] Lien S-Y, Yeh C-H, Deng D-J. Optimum splitting computing for DNN training through next generation smart networks: a multi-tier deep reinforcement learning approach. *Wireless Networks*. 2024; 30(3): 1737-1751.
- [45] International Telecommunication Union (ITU). IMT-2020 (5G) standard: minimum requirements related to technical performance for IMT-2020 radio interface(s). International Telecommunication Union; 2020. Available from: <https://www.itu.int/pub/R-REP-M.2410-2020>. [Accessed 8 Mar 2025].
- [46] Sarah A, Nencioni G, Khan MMI. Resource allocation in multi-access edge computing for 5G-and-beyond networks. *Computer Networks*. 2023; 227: 109720.
- [47] Xu Y, Gui G, Gacanin H, Adachi F. A survey on resource allocation for 5G heterogeneous networks: current research, future trends, and challenges. *IEEE Communications Surveys & Tutorials*. 2021; 23(2): 668-695.
- [48] Karjee J, Naik SP, Srinidhi N. Energy profiling based load-balancing approach in IoT-edge for split computing. In: *2021 IEEE 18th India Council International Conference (INDICON)*; 2021 Dec 19-21; Guwahati, India. IEEE; 2021. p. 1-6.
- [49] Zhang G, Guo W, Tan Z, Jiang H. AMP4EC: adaptive model partitioning framework for efficient deep learning inference in edge computing environments. 2025. Available from: <https://arxiv.org/abs/2504.00407>. [Accessed 8 Mar 2025].
- [50] Mudvari A, Vainio A, Ofeidis I, Tarkoma S, Tassiulas L. Adaptive compression-aware split learning and inference for enhanced network efficiency. 2024. Available from: <https://arxiv.org/abs/2311.05739>. [Accessed 8 Mar 2025].
- [51] Beloglazov A, Koch FL, Richter J, Steer K. *Augmented reality recommendations in emergency situations*. US Patent 14976512. 2017 Jun 22.
- [52] Assuncao MD, Koch FL, Westphall CB. Grids of agents for computer and telecommunication network management. *Concurrency and Computation: Practice and Experience*. 2004; 16(5). 413-424.
- [53] Djuhera A, Binotto APD, Koch FL, High R. *Orchestration of workloads involving an AI model*. US Patent 20250071069A1. 2025 Feb 27.
- [54] Djuhera A, Binotto APD, Koch FL, Angel NB. *Distributed execution of an artificial intelligence model*. US Patent 18449811. 2024 Dec 19.