Journal of Intelligent Computing and Networking
https://www.ffspub.com/index.php/jicn/index
ISSN 3079-9228 (print)
ISSN 3107-2798 (online)
E-mail: jicn.office@ffspub.com

Article

Aggregation-Guided Buffering and Orchestration for Wide-Area Mixture-of-Agents in the Optical Internet of Agents

Xiao Lin, Zuqi Teng, Wang Li, Ziang Feng, Weiling Lu, Enguo Chen†
College of Physics and Information Engineering, Fuzhou University, Fuzhou 350108, China
†*E-mail: ceg@fzu.edu.cn*

Received: August 2, 2026 / Revised: November 2, 2026 / Accepted: November 6, 2026 / Published online: November 13, 2026

Abstract: Wide-area Mixture-of-Agents (MoA) collaboration requires the exchange of large payloads over the Optical Internet of Agents (O-IoA) under fluctuating traffic conditions. To increase agent participation within collaboration deadlines under network resource contention, we propose the Aggregation-Guided Buffering and Orchestration (AGBO) framework. AGBO coordinates hierarchical update aggregation, intermediate buffering, and optical transmission using a capacitated time-expanded resource graph (TERG). We develop analytical models showing that temporal buffering reduces collaboration failure probability at the cost of higher candidate-path search complexity. We also develop an AGBO algorithm combining two-stage maximum-flow orchestration with Meta Agent (MA) selection. Simulation results show that AGBO retains more Work Agents (WAs) across different wavelength-utilization levels, payload sizes, and collaboration deadlines. At a wavelength utilization of 60%, it improves the WA participation ratio by at least 57 percentage points over the two baselines. In a convolutional neural network (CNN)-based collaborative-learning proxy, AGBO also achieves faster accuracy improvement and higher global test accuracy alongside higher WA participation.

Keywords: Optical Internet of Agents; Mixture-of-Agents; Spatio-Temporal Orchestration; Hierarchical Aggregation; Store-and-Forward
https://doi.org/10.64509/jicn.24.140

1 Introduction

1.1 Background

Recent advances in large language models (LLMs) have accelerated the development of autonomous agents equipped with perception, memory, reasoning, and tool-use capabilities [1]. Nevertheless, the performance of an individual agent remains constrained by its finite model capacity, local observations, and limited access to task-relevant information. The Internet of Agents (IoA) has therefore been proposed to support interactions among heterogeneous agents distributed across edge devices, computing infrastructures, and administrative domains [2]. Within the IoA, the Mixture-of-Agents (MoA) paradigm enables multiple agents to complete tasks collaboratively through iterative information exchange and aggregation [3]. Although MoA was originally introduced for collaborative reasoning among LLM-based agents, its hierarchical interaction pattern can also support distributed learning and perception workloads [4].

Deploying such collaborative systems over wide-area optical infrastructure gives rise to the Optical Internet of Agents (O-IoA). In a typical hierarchical MoA architecture [5], geographically distributed Work Agents (WAs) generate task-related local update payloads and transmit them to selected Meta Agents (MAs). Each MA performs hierarchical update aggregation to produce an Aggregated Trace Block (ATB), which is then delivered to a Global Agent (GA) for global aggregation. Depending on the model scale and the amount of intermediate information exchanged during each collaboration round, both local updates and the resulting ATBs can impose substantial communication demands [6]. Concurrent transmissions from multiple WAs and MAs therefore generate heavy fan-in traffic toward their aggregation nodes [7].

† Corresponding author: Enguo Chen
*Academic Editor: Chunxiao Jiang
© 2026 The authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Wide-area optical networks, however, are subject to time-varying background traffic and fragmented wavelength availability [8]. As traffic loads shift among commercial, residential, campus, and data-center regions, the wavelengths available on an optical link may vary considerably over time [9, 10]. Conventional MoA transmission schemes generally require an end-to-end (E2E) lightpath to be immediately available when a transmission request arrives. In the direct-transmission baseline (Vanilla) [3], a WA directly requests an E2E lightpath to the GA. In the layered-aggregation baseline (Layered) [5], a WA first requests an E2E lightpath to its assigned MA, after which the MA requests another E2E lightpath to the GA. Because these schemes do not support temporal storage at intermediate nodes, a WA is excluded from the current collaboration round if the required lightpath is unavailable when the request arrives. The resulting loss of WA participation reduces the amount and diversity of local information available for global aggregation.

1.2 Related Work

Optical Computing Power Networks (OCPNs) have been studied as a means of interconnecting geographically distributed computing facilities through high-capacity, low-latency optical channels [11]. Existing studies have explored computation-aware routing, wavelength assignment, reconfigurable optical networking, and joint computing-network resource allocation to support distributed AI workloads [8–13]. These approaches provide an important foundation for integrating computing services with optical transport. However, many existing formulations optimize resource allocation in a static network setting or assume that a request must be served using only the resources available at its arrival time. They therefore provide limited temporal flexibility when a transient background-traffic surge prevents the establishment of an E2E lightpath. Moreover, application-layer aggregation dependencies in multi-agent collaboration are generally not jointly considered with optical transmission and node-storage decisions.

LLM ensemble methods such as LLM-Blender and LLM-TOPLA combine candidate responses from multiple language models into a single output [14, 15]. When this aggregation pattern is mapped onto the O-IoA, each WA requires an E2E transmission path to the GA, resulting in a substantial concurrent fan-in load. Recent studies have explored layered aggregation in MoA [5] and related forms of hierarchical collaboration in LLM-based multi-agent systems [16, 17]. In the layered O-IoA architecture considered here, each MA aggregates the updates from its associated WAs into a single ATB and forwards it to the GA. When an MA combines multiple WA updates into an ATB of the same size as an individual update, this reduces both the number of payloads and the total data volume delivered to the GA. At the agent layer, dynamic agent-routing and model-selection methods have also been proposed to select logical agent paths or models based on task requirements and inference costs [7, 17–19]. These methods make decisions about the logical workflow, while physical optical routing determines how the resulting updates are carried over the network. Existing agent-layer methods typically assume that the underlying network can deliver the selected updates, whereas conventional optical resource-allocation methods do not explicitly consider aggregation dependencies, collaboration deadlines, or the impact of network blocking on WA participation.

Store-and-Forward (SnF) routing provides temporal flexibility for wide-area data transport by allowing payloads to be temporarily stored at intermediate nodes and forwarded when subsequent transmission resources become available [20]. By linking transmission opportunities across successive time slots, SnF can avoid discarding a request solely because a complete E2E lightpath is unavailable at its arrival time. Recent studies have further explored the joint use of intermediate storage and in-network computing for delay-tolerant distributed tasks [8, 20, 21].

However, existing hierarchical aggregation and SnF-based transport mainly address different aspects of the problem. Hierarchical aggregation reduces the traffic forwarded toward the global aggregation stage, but typically does not model how time-varying optical-resource availability and intermediate storage affect the delivery of individual updates. SnF, in contrast, provides temporal flexibility for data delivery, but is generally formulated for transport requests whose transmission process does not involve the release dependencies introduced by intermediate aggregation. When temporal buffering is introduced into a hierarchical WA–MA–GA collaboration process, these two mechanisms become tightly coupled: the spatio-temporal routes of WA updates determine when an MA can complete aggregation and release its ATB, while the resources consumed in the WA–MA stage affect the feasible transmission opportunities available to the subsequent MA–GA stage. Moreover, the selection of MA-hosting nodes changes both the aggregation structure and the resulting spatio-temporal traffic demands.

Therefore, the key issue considered in this work is not hierarchical aggregation, SnF routing, or MA selection in isolation, but their deadline-constrained coordination over time-varying optical resources. This coupling motivates a unified orchestration framework that jointly accounts for aggregation dependencies, MA-hosting decisions, spatio-temporal optical transmission, and intermediate storage.

1.3 Motivations and Contributions

The above analysis shows that simply combining hierarchical aggregation with temporal buffering does not directly yield an effective solution for wide-area MoA. Under conventional E2E transmission, WA–MA and MA–GA lightpaths must be available when the corresponding transmission requests arrive. Incorporating SnF creates additional temporal dependencies into this process because WA updates may arrive at their selected MAs at different future time slots. The release time of an MA-generated ATB therefore depends on the latest arrival among its associated WA updates, and the remaining time for the subsequent MA–GA transmission is constrained by the common collaboration deadline. Temporal flexibility can consequently improve transmission feasibility, but excessive waiting may also delay aggregation and reduce the time available for completing the remaining collaboration stages.

This coupling also substantially increases the orchestration complexity. Each WA update is no longer associated with

only a spatial route in the current network snapshot, but may follow a spatio-temporal route containing both optical transmission segments and intermediate storage transitions. The WA–MA routing decisions determine not only the consumed wavelength and storage resources, but also the aggregation associations and ATB release times. These decisions, in turn, affect the residual resources and feasible spatio-temporal routes available for the MA–GA stage. In addition, changing the MA-hosting configuration alters both the hierarchical aggregation structure and the resulting traffic distribution. Consequently, routing, temporal storage, aggregation timing, and MA-hosting decisions cannot be treated independently.

To address these coupled challenges, we propose the Aggregation-Guided Buffering and Orchestration (AGBO) framework for deadline-constrained wide-area MoA. Here, aggregation-guided orchestration accounts for how WA arrivals determine ATB release times and the resources available to subsequent transmission. A capacitated time-expanded resource graph (TERG) represents time-varying wavelength availability and intermediate storage opportunities. For a given MA configuration, AGBO decomposes the WA–MA and MA–GA orchestration into two sequential maximum-flow stages, coupled through WA–MA associations, ATB release layers, and residual resource capacities. An outer MA-selection search evaluates alternative hosting configurations according to the final number of WAs whose contributions reach the GA within the deadline.

The main contributions of this paper are summarized as follows:

- We propose the AGBO framework for deadline-constrained wide-area MoA collaboration over the O-IoA. AGBO coordinates MA-hosting decisions and WA–MA–GA aggregation dependencies with time-varying wavelength availability and intermediate storage, allowing buffering opportunities to be exploited within the collaboration deadline.
- Building on the bulk-transfer SnF analysis in our previous work [22], which did not consider MoA, we develop analytical models for Vanilla, Layered, and the transmission architecture underlying AGBO. These models characterize the effects of hierarchical aggregation and temporal buffering on candidate-path search complexity and round failure probability.
- We develop the AGBO algorithm by combining AGBO Graph Extension (AGBO-GE), AGBO Two-Stage Orchestration (AGBO-TO) for a given MA set, and AGBO MA Selection (AGBO-MS). The two transmission stages are coupled through WA–MA associations, ATB release layers, and residual resources, while the outer search explores MA configurations that increase final WA participation.
- We compare Vanilla, Layered, and AGBO across wavelength-utilization levels, payload sizes, and collaboration deadlines. AGBO retains more WAs by exploiting wavelength and storage resources over time, with a participation–latency trade-off under some settings. A CNN-based collaborative-learning proxy using MNIST and FEMNIST illustrates the application-level outcomes associated with different participation patterns under the adopted workload.

1.4 Notations

The principal notations used in the system model, analytical comparison, orchestration algorithm, and performance evaluation are summarized in Table 1. Only the principal notations used across multiple sections are listed; local variables are defined when they first appear.

Table 1: Summary of Key Notations

Notation	Description
$G=(V,E)$	Physical optical network graph.
$V_{\mathrm{W}}, V_{\mathrm{M}}^{C}, V_{\mathrm{M}}, v_{\mathrm{G}}$	WA nodes, candidate/selected MA nodes, and GA node.
F, t_{ddl}, L	Payload size, collaboration deadline, and number of TERG layers.
t_{loc}	Common local processing time before network transmission.
T_{net}	Network scheduling window remaining after local processing.
Δt	Duration of one temporal layer.
$G^{L}, E_{\mathrm{s}}^{L}, E_{\mathrm{t}}^{L}, c(e)$	Capacitated TERG with spatial edges, temporal edges, and payload-unit capacities.
p_b, p_s	Spatial-edge and temporal-edge unavailability probabilities.
$C(N,L), P_f(N,L)$	Candidate-path count and probability that all candidates fail.
$N_{\mathrm{W}}, N_{\mathrm{M}}, N_{\mathrm{WpM}}^{j}$	WA count, MA count, and WAs assigned to MA j.
$N_{\mathrm{W2G}}^{j}, N_{\mathrm{W2M}}^{j}, N_{\mathrm{M2G}}^{j}$	Node counts along WA–GA, WA–MA, and MA–GA routes.
γ_j	Effective aggregation factor of subtree j.
$C_{\mathrm{V}}, C_{\mathrm{L}}, C_{\mathrm{AGBO}}$ $P_{\mathrm{V}}, P_{\mathrm{L}}, P_{\mathrm{AGBO}}$	Scheme-level search complexities and round failure probabilities.
$C_{\mathrm{stree}}, P_{\mathrm{stree}}$	Subtree search complexity and failure probability.
$A_{\mathrm{W}}, A_{\mathrm{M}}, R_{\mathrm{W2M}}$ $P_{\mathrm{W2M}}, P_{\mathrm{M2G}}, \Pi$	Admitted sets, associations, selected paths, and final orchestration result.

1.5 Organization

Section 2 presents the system model and AGBO framework, including the wide-area O-IoA architecture, hierarchical aggregation, and TERG resource representation. Section 3 compares simplified transmission models for Vanilla, Layered, and the architecture underlying AGBO. Section 4 develops the AGBO algorithm, including aggregation-aware orchestration, TERG-based aggregator selection, and complexity analysis. Section 5 evaluates the network-level performance and collaborative-learning proxy. Section 6 concludes the paper.

2 System Model and AGBO Framework

2.1 Wide-Area O-IoA Architecture

We model the physical substrate of the wide-area O-IoA as a directed graph $G=(V,E)$, where V and E denote the sets

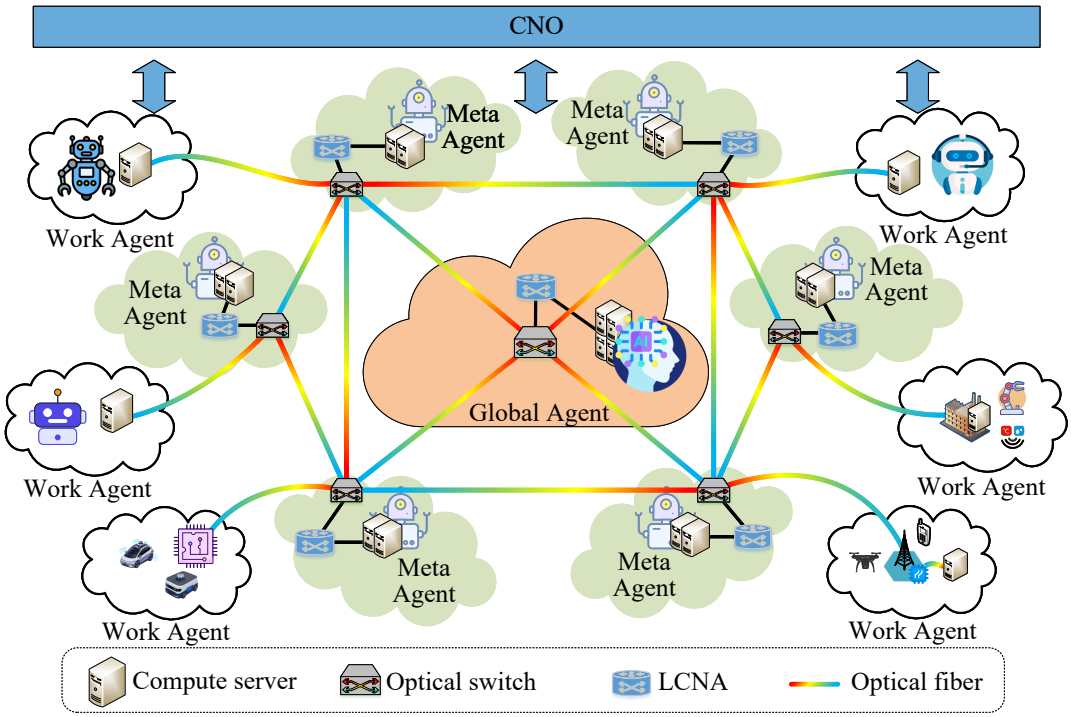


Figure 1: Illustration of the wide-area O-IoA architecture.

of geographically distributed infrastructure nodes and optical fiber links, respectively, as illustrated in Figure 1. Each node $v \in V$ provides communication and storage resources, while nodes satisfying the MA-hosting requirements form the candidate pool V_{M}^{C}. For each collaboration round, AGBO selects the active MA set $V_{\mathrm{M}} \subseteq V_{\mathrm{M}}^{C}$. The candidate pool specifies eligible locations; the active set is determined by the orchestration process rather than supplied as part of the collaboration request.

A Local Computing and Networking Agent (LCNA) is deployed at each infrastructure node to monitor local resource states, report them to the Centralized Network Orchestrator (CNO), and execute the resource configurations issued by the CNO. The CNO maintains a global view of the current network state and the reservations established for subsequent time slots. In particular, the CNO records the residual wavelength resources on each optical link and the residual storage resources at each infrastructure node after accounting for existing reservations. This reservation-aware resource view enables the CNO to jointly orchestrate optical transmission and intermediate storage over time.

The application-layer WA–MA–GA architecture is mapped onto the physical O-IoA substrate. The set $V_{\mathrm{W}} \subseteq V$ identifies the physical nodes hosting the WAs listed in the collaboration request, while $v_{\mathrm{G}} \in V$ identifies the physical node hosting the GA. Each WA performs its local task and generates a local update payload.

A single-round collaboration request is represented by the four-tuple $r = (V_{\mathrm{W}}, v_{\mathrm{G}}, F, t_{\mathrm{ddl}})$, where V_{W} and v_{G} identify the physical nodes hosting the WAs listed in the request and the GA, respectively. Here, F denotes the size of one local update payload, and t_{ddl} denotes the end-to-end collaboration deadline measured from the beginning of local processing. A local update payload and the ATB generated after MA aggregation are assumed to have the same size F.

Let t_{loc} denote the common local processing time before WA updates become available for transmission. The remaining network scheduling window is $T_{\mathrm{net}} = t_{\mathrm{ddl}} - t_{\mathrm{loc}}$.

AGBO coordinates two distinct information flows. The control flow is established between the CNO and the LCNAs. The LCNAs report local resource states to the CNO, while the CNO issues configurations for MA hosting, transmission, and storage to the corresponding nodes. The application-layer data flow follows the hierarchical WA–MA–GA structure. WAs transmit their local updates to the selected MAs, and each MA performs hierarchical update aggregation to generate an ATB before forwarding it to the GA.

For an admitted request, the CNO first identifies the infrastructure nodes satisfying the MA-hosting requirements. It then selects the active MA-hosting nodes, assigns WAs to the selected MAs, and reserves the required optical-link and node-storage resources within the collaboration deadline. The LCNAs subsequently execute the issued resource configurations. After the scheduled WA updates assigned to an MA arrive at its hosting node, the MA performs update aggregation and generates one ATB for the subsequent MA–GA transmission stage.

2.2 Hierarchical Update Aggregation

Storage-aided transmission enables a payload to wait for resources that become available in a subsequent time slot. However, temporal storage alone does not reduce the number of WA updates that must eventually be forwarded toward the GA. If every WA independently transmits its local update over a long-distance path to the GA, a large number of concurrent flows may traverse the links surrounding the GA.

AGBO introduces intermediate MAs to reduce the communication demand of the MA–GA transmission stage. Each selected MA is assigned a subset of the participating WAs.

The WAs independently deliver their local updates to the corresponding MA-hosting node along the paths orchestrated by the CNO. The MA aggregates the local updates successfully received by the time aggregation is triggered and generates one ATB.

Both an individual local update and the resulting ATB are assumed to have the same size F. Therefore, hierarchical update aggregation does not reduce the dimensionality or size of an individual payload. Instead, it replaces multiple WA update payloads with one ATB in the MA–GA transmission stage. For an MA assigned N^j_{WpM} WAs, hierarchical aggregation reduces the number of payloads entering the corresponding MA–GA stage from N^j_{WpM} local updates to one ATB.

The aggregation process is performed by the logical MA application hosted on the selected infrastructure node. The LCNA neither performs update aggregation nor interprets the contents of the local updates. Instead, it allocates and configures local computing and storage resources according to the decisions issued by the CNO. Similarly, intermediate storage can be provided by any infrastructure node along an orchestrated spatio-temporal path and is not restricted to MA-hosting nodes.

The specific update-aggregation rule depends on the collaborative workload and is introduced together with the simulation configuration in Section 5. At the framework level, AGBO requires only that multiple local updates be combined into one ATB of the same size F.

2.3 Time-Expanded Resource Graph

To represent time-varying optical-link resources and intermediate storage, the CNO constructs a capacitated time-expanded resource graph (TERG) over the complete physical topology. This resource representation supports the AGBO algorithm in Section 4. Section 3 uses a regular fixed-route restriction of the TERG for analytical comparison.

After local processing, the CNO divides the network scheduling window into slots of duration Δt. Let $t_\ell = (\ell - 1)\Delta t$ denote the start of slot ℓ, measured from the beginning of this network window. Slot ℓ spans $[t_\ell, t_\ell + \Delta t)$ and corresponds to temporal layer ℓ. The first temporal layer represents the currently available communication and storage resources, whereas the subsequent layers represent the residual resources derived from existing reservation records. The physical graph G is then expanded along the temporal dimension to construct the TERG $G^L = (V^L, E^L_s \cup E^L_t, c)$. The time-expanded node set is $V^L = \{v^{(\ell)} : v \in V, 1 \le \ell \le L\}$, where $v^{(\ell)}$ is the virtual copy of physical node v in temporal layer ℓ. Each temporal layer therefore contains one copy of the complete physical node set.

For each physical optical link $(u,v) \in E$ and temporal layer ℓ, the spatial edge $(u^{(\ell)}, v^{(\ell)}) \in E^L_{\mathrm{s}}$ represents transmission over the physical link (u,v) during time slot ℓ. A spatial edge therefore corresponds to one physical optical link in one time slot rather than an entire E2E lightpath. A multi-hop transmission path within one time slot consists of a sequence of consecutive spatial edges.

For each physical node $v \in V$ and each $1 \le \ell < L$, the temporal edge $(v^{(\ell)}, v^{(\ell+1)}) \in E^L_{\mathrm{t}}$ represents the storage of a payload at node v from time slot ℓ to time slot $\ell + 1$. Because all nodes in V are assumed to provide storage capability, a temporal edge is constructed between each pair of consecutive temporal copies of every physical node. Its usable capacity depends on the residual storage available at the node after accounting for existing reservations.

A spatio-temporal transmission path may contain both spatial and temporal edges. Traversing a spatial edge represents optical transmission over a physical link during the corresponding time slot, whereas traversing a temporal edge represents retaining the payload at the same physical node until the next time slot. Consequently, AGBO can orchestrate a payload through a sequence of optical transmission segments and intermediate storage transitions without requiring an immediately available E2E lightpath in a single network snapshot.

The edge-capacity function c converts the residual link-transmission and node-storage capacities into integer payload units of size F. For a spatial edge $e = (u^{(\ell)}, v^{(\ell)})$, let F_e denote the residual data volume that physical link (u,v) can transmit during time slot ℓ after accounting for existing reservations. For a temporal edge $e = (v^{(\ell)}, v^{(\ell+1)})$, F_e denotes the residual storage capacity available at node v during the corresponding inter-slot interval. The payload-unit capacity of either edge type is defined as $c(e) = \lfloor F_e / F \rfloor$.

When $c(e) = 0$, the corresponding spatial or temporal edge cannot accommodate an additional payload and is unavailable to the current request. When $c(e) > 0$, up to $c(e)$ payload units can use the edge without violating existing reservations.

The TERG represents the network scheduling window after local processing, whose duration is $T_{\mathrm{net}} = t_{\mathrm{ddl}} - t_{\mathrm{loc}}$. With a uniform scheduling interval Δt, the maximum number of complete temporal layers is $L_{\mathrm{ddl}} = \lfloor (t_{\mathrm{ddl}} - t_{\mathrm{loc}})/\Delta t \rfloor$. The CNO constructs $L \le L_{\mathrm{ddl}}$ layers and sets $L = L_{\mathrm{ddl}}$ when the full available window is used.

The TERG encodes communication and storage resources but does not explicitly represent computing capacity as an edge- or node-capacity dimension. Instead, node computing capability is used to identify the infrastructure nodes eligible to host an MA. The selection of active MA-hosting nodes is jointly considered with communication and storage orchestration in the subsequent algorithm, but no divisible computing-capacity allocation is performed over the TERG.

The general TERG established in this section provides a common resource representation for the subsequent analytical comparison and AGBO algorithm. Section 3 derives simplified analytical models by restricting the general TERG to fixed routes with regular temporal-layer structures, while Section 4 uses the general TERG to formulate and solve the aggregation-aware spatio-temporal orchestration problem.

3 Analysis of MoA Transmission Schemes

3.1 Analytical Assumptions and Regular TERG Model

Starting from the TERG established in Section 2.3, we construct simplified analytical models for Vanilla, Layered,

and the hierarchical, storage-aided transmission architecture underlying AGBO. These models compare candidate-path search complexity and approximate round failure probability under the stated assumptions; they do not represent the execution complexity or exact operational failure probability of the AGBO algorithm.

To focus on network transmission and buffering, the analytical model neglects local processing time by setting $t_{\mathrm{loc}} = 0$. The collaboration deadline t_{ddl} therefore coincides with the available network scheduling window.

For analytical tractability, each source–destination pair is assigned a fixed physical route. Each temporal layer contains a copy of this route, and temporal copies of the same physical node in adjacent layers are connected by temporal edges. A fixed route contains N nodes, including its source and destination, and the corresponding regular TERG contains L temporal layers.

Following the analytical model established in our previous work [22], $C(N,L)$ denotes the total number of candidate spatio-temporal paths from the source node in the first layer to the destination node in any admissible layer. In this section, routing search complexity is measured by the number of candidate paths that must be examined in the worst case. The function $P_f(N,L)$ denotes the probability that all $C(N,L)$ candidate paths are unavailable.

The functions $C(N,L)$ and $P_f(N,L)$ are defined for $N \geq 2$ and $L \geq 1$. For the boundary values required by the subsequent expressions, we define $C(N,0) = 0$ and $P_f(N,0) = 1$.

The spatial-edge unavailability probability p_b represents the probability that a spatial edge cannot provide the wavelength resource required by a request. The temporal-edge unavailability probability p_s represents the probability that a temporal edge cannot accommodate the payload because of insufficient residual storage. The effects of payload size, link rate, slot duration, and residual resource capacities are therefore abstracted into p_b and p_s in the analytical model.

To compare the effects of hierarchical aggregation and temporal buffering, we assume independent resource states on different edges and in different time slots, with link-disjoint fixed routes for different WAs. Correlations among overlapping candidate spatio-temporal paths are neglected. These assumptions yield tractable approximations for candidate-path search complexity and round failure probability.

For the analytical comparison, a collaboration round succeeds when all WA updates and MA-generated ATBs required by the corresponding scheme reach their destinations within the deadline. The round failure probability is therefore the probability that at least one required transmission does not complete.

3.2 TERG Models for MoA Transmission

Vanilla: We define Vanilla [3] as a two-tier WA–GA transmission architecture in which each WA directly delivers its local update to the GA through an E2E path. Vanilla supports neither intermediate update aggregation nor temporal buffering and therefore operates on a single network snapshot. Its TERG model is depicted in Figure 2.

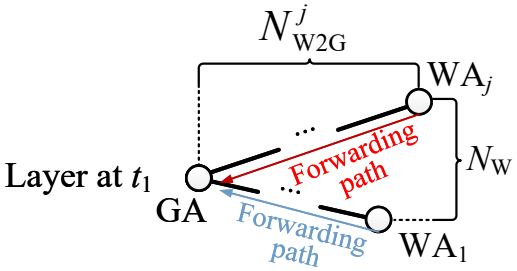


Figure 2: TERG representation of Vanilla.

Let N_{W} denote the number of WAs listed in the collaboration request, and let N^j_{W2G} denote the number of nodes on the fixed route from the j-th WA to the GA. Because one fixed WA–GA route is examined for each WA in the current time slot, the routing search complexity of Vanilla is

$$C_{\mathrm{V}}(N_{\mathrm{W}}) = N_{\mathrm{W}}. \tag{1}$$

Under the assumptions of link-disjoint routes and independent edge unavailability, the collaboration round succeeds only if all WA–GA transmissions succeed. The round failure probability of Vanilla is

$$P_{\mathrm{V}}(N_{\mathrm{W}}) = 1 - \prod_{j=1}^{N_{\mathrm{W}}} \left[1 - P_f\left(N^j_{\mathrm{W2G}}, 1\right)\right]. \tag{2}$$

Layered: We define Layered [5] as a three-tier WA–MA–GA transmission architecture. Each WA delivers its local update to an assigned MA through an E2E path. After receiving all assigned WA updates, the MA performs hierarchical update aggregation and generates one ATB, which is then delivered to the GA through another E2E path. Layered supports update aggregation but not temporal buffering and therefore also operates on a single network snapshot. Its TERG model is depicted in Figure 3.

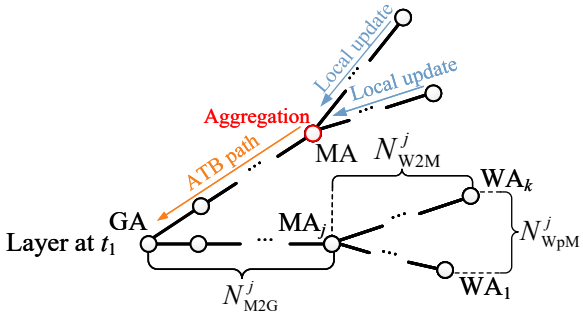


Figure 3: TERG representation of Layered.

Let N_{M} denote the number of selected MA-hosting nodes, and let N^j_{WpM} denote the number of WAs assigned to the j-th MA. The aggregation subtree rooted at the j-th MA contains N^j_{WpM} WA–MA routes and one MA–GA route. The routing search complexity of Layered is therefore

$$C_{\mathrm{L}}(N_{\mathrm{M}}) = \sum_{j=1}^{N_{\mathrm{M}}} \left(N^j_{\mathrm{WpM}} + 1\right). \tag{3}$$

Let N^j_{W2M} denote the number of nodes on each WA–MA route in the j-th aggregation subtree, N^j_{M2G} denote the number of nodes on its MA–GA route, and N^j_{W2G} denote the number of nodes on the corresponding reference WA–GA route used for direct transmission.

We introduce the aggregation factor $\gamma_j = N^j_{\mathrm{W2M}}/(N^j_{\mathrm{WpM}}N^j_{\mathrm{W2G}})$ to approximate the effective resource availability gain associated with hierarchical update aggregation. Either assigning more WAs to an MA or shortening the WA–MA route produces a smaller γ_j. We consider deployment configurations satisfying $0 < \gamma_j \leq 1$. The effective spatial-edge and temporal-edge unavailability probabilities of the j-th aggregation subtree are approximated by $p'_{b,j} = \gamma_j p_b$ and $p'_{s,j} = \gamma_j p_s$, respectively. This scaling captures the reduction in effective resource demand at the analytical level and does not alter the physical state of any individual link or infrastructure node.

Let $P_{\mathrm{stree}}(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, L, \gamma_j)$ denote the failure probability of the j-th aggregation subtree over an L-layer TERG. The Layered aggregation subtree corresponds to the special case $L = 1$, in which temporal buffering is unavailable. Its approximate expression and derivation are provided in Appendix A.2.

Because a Layered round fails if any aggregation subtree fails, its round failure probability is

$$P_{\mathrm{L}}(N_{\mathrm{M}}) = 1 - \prod_{j=1}^{N_{\mathrm{M}}} \left[1 - P_{\mathrm{stree}}\left(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, 1, \gamma_j\right)\right]. \tag{4}$$

Layered reduces the number of long-distance payloads entering the GA from N_{W} local updates to N_{M} ATBs. However, all WA–MA and MA–GA transmissions must still be completed using resources available in the current network snapshot.

AGBO: We consider the three-tier WA–MA–GA transmission architecture supported by the AGBO framework. It combines hierarchical aggregation with storage-aided transmission across multiple TERG layers. WA updates and MA-generated ATBs may be buffered at intermediate nodes before transmission resumes in a later slot. Figure 4 illustrates the simplified TERG representation used in this analysis.

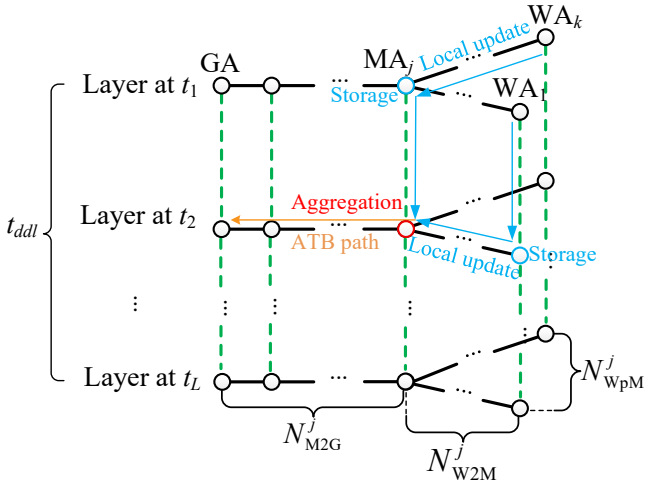


Figure 4: Simplified TERG representation of the hierarchical transmission architecture underlying AGBO.

In the theoretical model considered in this section, an MA generates its ATB only after all WAs assigned to its aggregation subtree have successfully delivered their local updates. The entire collaboration round is considered to have failed if any required WA–MA or MA–GA transmission cannot be completed within the admissible L-layer time window.

Let $C_{\mathrm{stree}}(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, L)$ denote the routing search complexity of the j-th AGBO aggregation subtree. It depends on the WA–MA route length, the MA–GA route length, the number of WAs assigned to the MA, and the number of temporal layers. Its approximate expression and derivation are provided in Appendix A.1. The candidate-path search complexity of the simplified AGBO transmission model is

$$C_{\mathrm{AGBO}}(N_{\mathrm{M}}, L) = \sum_{j=1}^{N_{\mathrm{M}}} C_{\mathrm{stree}}\left(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, L\right). \tag{5}$$

The aggregation factor γ_j and the effective unavailability probabilities $p'_{b,j}$ and $p'_{s,j}$ follow the same definitions as those used for Layered. Unlike Layered, however, AGBO is affected by temporal-edge unavailability because payloads may traverse temporal edges between consecutive time slots.

The failure probability of the j-th AGBO aggregation subtree is denoted by $P_{\mathrm{stree}}(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, L, \gamma_j)$. Its approximate expression and derivation are provided in Appendix A.2. Assuming that the failure events of different aggregation subtrees are independent, the approximate round failure probability of the simplified AGBO transmission model is

$$P_{\mathrm{AGBO}}(N_{\mathrm{M}}, L) = 1 - \prod_{j=1}^{N_{\mathrm{M}}} \left[1 - P_{\mathrm{stree}}\left(N^j_{\mathrm{W2M}}, N^j_{\mathrm{M2G}}, N^j_{\mathrm{WpM}}, L, \gamma_j\right)\right]. \tag{6}$$

Layered can therefore be interpreted as the $L = 1$ special case of the aggregation-subtree failure model, whereas AGBO exploits the additional spatial and temporal paths available when $L > 1$.

3.3 Numerical Results

We numerically compare candidate-path search complexity and approximate round failure probability under the simplified transmission models for Vanilla, Layered, and AGBO. We focus on the operating regime $p_s \ll p_b$, in which storage resources are substantially more available than optical transmission resources. The spatial-edge unavailability probability p_b varies over $[0.01, 0.1]$, the temporal-edge unavailability probability is set to $p_s = 10^{-4}$, the maximum number of temporal layers is set to $L = 4$, and the number of WAs listed in the collaboration request, N_{W}, ranges from 2 to 20.

For Layered and AGBO, each aggregation subtree is assumed to contain the same number of WAs, with $N^j_{\mathrm{WpM}} = 2$. Consequently, the number of active MAs is $N_{\mathrm{M}} = N_{\mathrm{W}}/2$ for the considered even values of N_{W}. The route-length parameters are fixed within each model configuration, and the aggregation factors are calculated as $\gamma_j = N^j_{\mathrm{W2M}}/(N^j_{\mathrm{WpM}}N^j_{\mathrm{W2G}})$. Under the symmetric-subtree setting, all aggregation subtrees use the same aggregation factor. Specifically, we set $N^j_{\mathrm{W2M}} = 3$, $N^j_{\mathrm{M2G}} = 3$, and $N^j_{\mathrm{W2G}} = 5$.

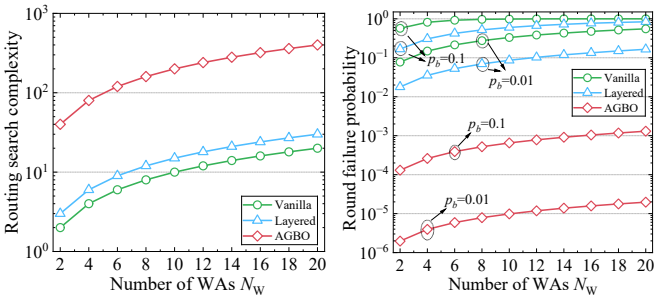
As N_{W} increases, all three schemes require more WA-related transmissions. Routing search complexity increases

in Figure 5(a), and round failure probability increases in Figure 5(b). Vanilla has the lowest routing search complexity because it examines one fixed WA–GA route for each WA within the current network snapshot. Layered additionally considers the WA–MA routes and one MA–GA route in each aggregation subtree. AGBO further expands these fixed physical routes across multiple temporal layers and consequently examines a larger number of candidate spatio-temporal paths.

The temporal expansion used by AGBO provides additional opportunities to bypass temporary spatial-edge unavailability, although examining these additional paths increases routing search complexity. Under the considered parameter settings, AGBO therefore achieves a lower round failure probability than Vanilla and Layered, as shown in Figure 5(b). This improvement is achieved at the expense of higher routing search complexity, revealing a trade-off between the number of candidate paths examined and the probability of completing the entire collaboration round.

Figure 5(b) also shows that, as p_b decreases from 0.1 to 0.01, the round failure probabilities of all three schemes decrease because the required spatial transmissions are more likely to find available resources. AGBO maintains a lower round failure probability by exploiting both spatial paths available in the current time slot and transmission opportunities available in subsequent temporal layers.

The collaboration deadline t_{ddl} determines the maximum number of temporal layers available to AGBO. Specifically, a larger t_{ddl} allows the CNO to incorporate more network snapshots into the TERG, resulting in a larger admissible value of L. Under a fixed time-slot duration, t_{ddl} and L have a monotonic relationship; hence, increasing t_{ddl} expands the temporal scheduling window considered by AGBO.

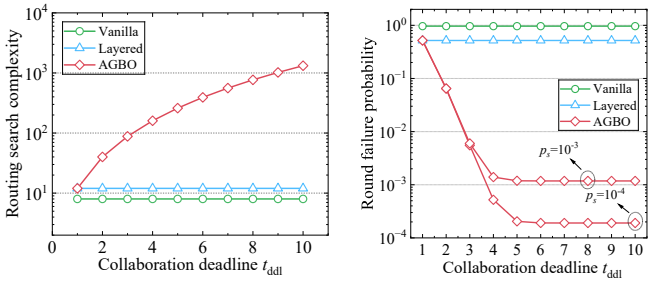


(a) Routing search complexity. (b) Round failure probability.

Figure 5: Routing search complexity and round failure probability under varying N_{W}.

Figure 6 investigates the impact of t_{ddl} on routing search complexity and round failure probability. In this experiment, we set $p_b = 0.1$, $p_s \in \{10^{-4}, 10^{-3}\}$, $N_{\mathrm{M}} = 4$, and $N_{\mathrm{W}} = 8$, while keeping the remaining parameters unchanged from the preceding experiment. For Figure 6, time is normalized by the temporal-layer duration. With local processing time neglected and integer-valued deadlines, $L = t_{\mathrm{ddl}}$.

Vanilla and Layered operate only on the network state available when a transmission request arrives. Their routing search complexities and round failure probabilities are therefore unaffected by t_{ddl}, because neither scheme can exploit resources represented in subsequent temporal layers. In contrast, a longer t_{ddl} increases the number of admissible TERG layers for AGBO and thereby enlarges the set of candidate spatio-temporal paths. Although the expanded path set increases routing search complexity, it can reduce the round failure probability by allowing payloads to be temporarily stored at intermediate nodes and forwarded when transmission resources become available in later temporal layers.



(a) Routing search complexity. (b) Round failure probability.

Figure 6: Routing search complexity and round failure probability under varying t_{ddl}.

Overall, the analytical and numerical results indicate that hierarchical update aggregation reduces the number of long-distance payloads delivered to the GA, while storage-aided spatio-temporal transmission enables AGBO to exploit resources available in subsequent time slots. AGBO consequently reduces the round failure probability under the considered assumptions, although it incurs higher routing search complexity. These observations motivate the aggregation-aware spatio-temporal orchestration formulation and AGBO algorithm developed in Section 4.

4 AGBO Algorithm

4.1 Overview

Building on the system model and TERG established in Section 2, this section presents the AGBO algorithm. In this study, AGBO selects suitable MA-hosting nodes and constructs feasible WA–MA and MA–GA transmission paths to maximize the number of participating WAs. This objective reflects the service preference considered here, while alternative objectives may be formulated according to the network operator's requirements and management policies. The resulting transmission and storage schedules must satisfy the collaboration deadline and the wavelength and storage capacity constraints. AGBO jointly considers MA-hosting selection, WA–MA association, and TERG-based transmission and storage scheduling.

The AGBO algorithm combines an outer MA-selection search with an inner two-stage orchestration process. AGBO-MS explores alternative MA configurations, and AGBO-TO evaluates each configuration through the WA–MA and MA–GA transmission stages. AGBO-GE constructs the extended TERGs required by the two maximum-flow computations. Only WAs associated with MAs whose ATBs reach the GA are retained in the final result.

The algorithm uses computing capability to determine MA-hosting eligibility; computing capacity, processing load, and computation delay are not explicitly optimized. All optical nodes support full wavelength conversion, so wavelength

continuity is not required and one available wavelength is reserved independently on each traversed spatial edge. Local updates and MA-generated ATBs are modeled as equal-sized payload units. Each spatial or temporal edge has a capacity equal to the number of additional payload units it can accommodate after existing reservations are accounted for.

4.2 Aggregation-Aware Orchestration

For a given MA set, AGBO-TO coordinates the WA–MA and MA–GA stages through aggregation-dependent release layers and residual resources. AGBO-GE first converts each routing stage into a single-source, single-sink maximum-flow problem while preserving the source release-layer constraints.

Given a source set S, a destination set D, and the source release-layer mapping $\boldsymbol{\tau} = \{\tau_u : u \in S\}$, AGBO-GE creates a copy of G^L and adds a virtual source s_0 and a virtual sink d_0, as described in Algorithm 1. For each source $u \in S$, the edge $(s_0, u^{(\tau_u)})$ has unit capacity, ensuring that the source contributes at most one payload. Each temporal replica $v^{(\ell)}$ of a destination $v \in D$ is connected to d_0 with capacity $|S|$, allowing each payload to terminate at any destination replica in an admissible layer.

All capacities in the extended graph are integers. The resulting maximum flow can therefore be decomposed into unit-flow paths, each of which identifies a source–destination association and its spatio-temporal route.

Algorithm 1 AGBO-GE: Release-Aware TERG Extension

Input: Capacitated TERG G^L, source set S, destination set D, and source release-layer mapping $\boldsymbol{\tau} = \{\tau_u : u \in S\}$
Output: Extended single-source, single-sink graph G_{Ext}

Create G_{Ext} as a copy of G^L
Add a virtual source s_0 and a virtual sink d_0 to G_{Ext}
for all $u \in S$ **do**
    Add edge $(s_0, u^{(\tau_u)})$ with capacity 1
end for
for all $v \in D$ **do**
    for $\ell = 1$ to L **do**
        Add edge $(v^{(\ell)}, d_0)$ with capacity $|S|$
    end for
end for
return G_{Ext}

The same graph-extension module is used in both stages, with different source sets, destination sets, and release layers. Algorithm 2 presents AGBO-TO, which evaluates the complete WA–MA–GA delivery process for a fixed MA-hosting set (V_{M}) through two sequential maximum-flow stages. The first stage jointly determines the WA–MA associations and routes the provisionally admitted WA updates to the selected MAs. The second stage routes the resulting MA-generated ATBs to the GA.

In the WA–MA stage, all WA updates are available in the first temporal layer. AGBO-GE is therefore invoked with V_{W} as the source set, V_{M} as the destination set, and $\tau_w = 1$ for every WA w. The resulting integral maximum flow determines the provisionally admitted WAs, their serving MAs, and the corresponding WA–MA spatio-temporal paths.

The residual capacities of the spatial and temporal edges used by these paths are then decremented in a local copy of the TERG. For each reached MA m, the algorithm calculates its ATB release layer as $\tau_m = \max_{w \in W_m} \ell_w$, where W_m is the set of WAs associated with m and ℓ_w is the temporal layer in which the update of WA w reaches its assigned MA. Because the aggregation delay is assumed to be negligible, the MA–GA transmission can start in layer τ_m.

In the MA–GA stage, AGBO-GE is applied to the residual TERG with the reached MAs as the source set and the GA as the destination. The second maximum flow identifies the MAs whose ATBs can reach the GA before the collaboration deadline. Any MA whose ATB cannot reach the GA in this stage, together with its associated WAs, is removed from the final collaboration result.

The tentative reservations of the discarded WA–MA paths are released before the final schedule is committed, but the released resources are not reused during the current algorithm execution. The algorithm returns $\Pi = (A_{\text{W}}, A_{\text{M}}, R_{\text{W2M}}, P_{\text{W2M}}, P_{\text{M2G}})$, which represents the final orchestration result for the given MA set.

Algorithm 2 AGBO-TO: Aggregation-Aware Two-Stage Orchestration

Input: G^L, V_{W}, V_{M}, and v_{G}
Output: Orchestration result Π

Create a working copy $\overline{G}^L \leftarrow G^L$ and set $\tau_w \leftarrow 1$ for all $w \in V_{\text{W}}$
$G_{\text{W2M}} \leftarrow \text{AGBO-GE}(\overline{G}^L, V_{\text{W}}, V_{\text{M}}, \boldsymbol{\tau}_{\text{W}})$
$f_{\text{W2M}} \leftarrow \textsc{DinicMaxFlow}(G_{\text{W2M}}, s_0, d_0)$
Extract $\widetilde{A}_{\text{W}}$, A_{M}^{in}, $\widetilde{R}_{\text{W2M}}$, $\widetilde{P}_{\text{W2M}}$, and $\{\ell_w\}$ from the unit-flow decomposition of f_{W2M}
if $\widetilde{A}_{\text{W}} = \emptyset$ **then**
    return an empty orchestration result
end if
Decrease by one the residual capacity of every edge traversed by each path in $\widetilde{P}_{\text{W2M}}$
Remove zero-capacity edges from $\overline{G}^L$ to obtain G^L_{res}
for all $m \in A_{\text{M}}^{\text{in}}$ **do**
    $W_m \leftarrow \{w : (w,m) \in \widetilde{R}_{\text{W2M}}\}$, $\tau_m \leftarrow \max_{w \in W_m} \ell_w$
end for
$G_{\text{M2G}} \leftarrow \text{AGBO-GE}(G^L_{\text{res}}, A_{\text{M}}^{\text{in}}, \{v_{\text{G}}\}, \boldsymbol{\tau}_{\text{M}})$
$f_{\text{M2G}} \leftarrow \textsc{DinicMaxFlow}(G_{\text{M2G}}, s_0, d_0)$
Extract $A_{\text{M}}^{\text{out}}$ and P_{M2G} from the unit-flow decomposition of f_{M2G}
Retain the MAs whose ATBs reach the GA: $A_{\text{M}} \leftarrow A_{\text{M}}^{\text{out}}$
Retain their WA–MA associations: $R_{\text{W2M}} \leftarrow \{(w,m) \in \widetilde{R}_{\text{W2M}} : m \in A_{\text{M}}\}$
Retain the WAs associated with the successful MAs: $A_{\text{W}} \leftarrow \{w : \exists m, (w,m) \in R_{\text{W2M}}\}$
Retain their WA–MA paths: $P_{\text{W2M}} \leftarrow \{p_w \in \widetilde{P}_{\text{W2M}} : w \in A_{\text{W}}\}$
Release the tentative reservations of $\widetilde{P}_{\text{W2M}} \setminus P_{\text{W2M}}$
$\Pi \leftarrow (A_{\text{W}}, A_{\text{M}}, R_{\text{W2M}}, P_{\text{W2M}}, P_{\text{M2G}})$
return Π

4.3 TERG-Based Aggregator Selection

The intermediate MAs serve as aggregators. Algorithm 3 presents AGBO-MS, which searches over subsets of the eligible candidate pool V_{M}^{C}, using AGBO-TO to evaluate each configuration by the final number of admitted WAs.

AGBO-MS first constructs an initial MA set from eligible MA-hosting nodes located in the intermediate topological region between the peripheral WAs and the GA. Let $d_G(u,v)$ denote the shortest-hop distance in the physical topology, and let $H_{\max} = \max_{w \in V_{\mathrm{W}}} d_G(w, v_{\mathrm{G}})$. The target distance H_{tar} is typically set close to $H_{\max}/2$, and the eligible MA-hosting nodes whose distances to the GA are closest to H_{tar} form the initial MA set. The target-hop rule provides a topology-informed starting point in an intermediate region for the consecutive WA–MA and MA–GA transmission stages. It is an initialization heuristic rather than an optimality condition. The subsequent expansion, substitution, and contraction operations adjust the active MA configuration within the same candidate pool. Section 5.8 evaluates sensitivity to the starting configuration.

Algorithm 3 AGBO-MS: TERG-Based Aggregator Selection

Input: TERG G^L, WA set V_{W}, candidate MA set V_{M}^{C}, GA node v_{G}, and target hop count H_{tar}
Output: Selected MA set V_{M}^{*} and orchestration result Π^{*}

Compute $d_G(v, v_{\mathrm{G}})$ for all $v \in V_{\mathrm{M}}^{C}$
Select all nodes in V_{M}^{C} whose hop distances to v_{G} are closest to H_{tar}, and denote them by V_{M}^{*}
$\Pi^{*} = (A_{\mathrm{W}}^{*}, \ldots) \leftarrow$ AGBO-TO$(G^L, V_{\mathrm{W}}, V_{\mathrm{M}}^{*}, v_{\mathrm{G}})$
while true do
    Generate the expansion, substitution, and contraction neighborhood $\mathcal{N}(V_{\mathrm{M}}^{*})$
    Set the current solution as the iteration-best candidate: $(V_{\mathrm{M}}^{\mathrm{best}}, \Pi^{\mathrm{best}}, A_{\mathrm{W}}^{\mathrm{best}}) \leftarrow (V_{\mathrm{M}}^{*}, \Pi^{*}, A_{\mathrm{W}}^{*})$
    for all $V' \in \mathcal{N}(V_{\mathrm{M}}^{*})$ **do**
        $\Pi' = (A'_{\mathrm{W}}, \ldots) \leftarrow$ AGBO-TO$(G^L, V_{\mathrm{W}}, V', v_{\mathrm{G}})$
        if $|A'_{\mathrm{W}}| > |A_{\mathrm{W}}^{\mathrm{best}}|$ **then**
            Update the iteration-best candidate: $(V_{\mathrm{M}}^{\mathrm{best}}, \Pi^{\mathrm{best}}, A_{\mathrm{W}}^{\mathrm{best}}) \leftarrow (V', \Pi', A'_{\mathrm{W}})$
        end if
    end for
    if $|A_{\mathrm{W}}^{\mathrm{best}}| \le |A_{\mathrm{W}}^{*}|$ **then**
        break
    end if
    Accept the best improving candidate as the new current solution: $(V_{\mathrm{M}}^{*}, \Pi^{*}, A_{\mathrm{W}}^{*}) \leftarrow (V_{\mathrm{M}}^{\mathrm{best}}, \Pi^{\mathrm{best}}, A_{\mathrm{W}}^{\mathrm{best}})$
end while
return $V_{\mathrm{M}}^{*}, \Pi^{*}$

Starting from the initial solution, AGBO-MS explores three types of MA-configuration transformations: MA-set expansion, MA-host substitution, and MA-set contraction. MA-set expansion activates one previously inactive candidate MA, MA-host substitution replaces one active MA-hosting node with an inactive candidate, and MA-set contraction deactivates one active MA. Each transformed configuration is evaluated by AGBO-TO on a fresh copy of the original TERG, using the coupled WA–MA and MA–GA orchestration process.

At each iteration, AGBO-MS selects the neighboring configuration with the largest number of finally admitted WAs. The current MA configuration is updated only if this neighbor admits more WAs; otherwise, the search terminates. The returned solution is therefore locally optimal with respect to the considered MA-set expansion, MA-host substitution, and MA-set contraction neighborhoods, although global optimality is not guaranteed.

4.4 Complexity Analysis

AGBO-MS evaluates each candidate MA set by invoking the two-stage orchestration algorithm. We therefore analyze the computational complexity in two steps: the cost of evaluating one fixed MA set and the number of such evaluations performed by AGBO-MS.

For an L-layer TERG, AGBO-GE adds one virtual source and one virtual sink. The resulting extended graph contains $|V'| = L|V| + 2$ nodes. For a source set S and a destination set D, the number of edges is bounded by $|E'| \le L|E| + (L-1)|V| + |S| + L|D|$, where the four terms correspond to spatial edges, temporal edges, virtual-source edges, and virtual-sink edges, respectively. Because edges with zero residual capacity can be removed, this expression provides an upper bound on the size of the extended graph.

For a fixed MA set, the two-stage orchestration algorithm solves one WA–MA maximum-flow problem and one MA–GA maximum-flow problem. Using Dinic's algorithm, each maximum-flow computation has a worst-case complexity of $O(|E'||V'|^2)$. The graph-construction, flow-decomposition, and residual-capacity-update steps do not change this asymptotic upper bound. Therefore, the complexity of evaluating one fixed MA set is $T_{\mathrm{fixed}} = O(|E'||V'|^2)$, where $|V'|$ and $|E'|$ denote the larger of the corresponding graph-size bounds for the two orchestration stages.

Let $n_{\mathrm{M}} = |V_{\mathrm{M}}^{C}|$. In one AGBO-MS iteration, the expansion and contraction neighborhoods each contain at most $O(n_{\mathrm{M}})$ candidate configurations, whereas the substitution neighborhood contains at most $O(n_{\mathrm{M}}^2)$ configurations. Consequently, one complete neighborhood search evaluates at most $O(n_{\mathrm{M}}^2)$ MA configurations.

AGBO-MS accepts a new configuration only when the integer-valued objective $|A_{\mathrm{W}}|$ strictly increases. Therefore, the number of accepted improving iterations is at most $|V_{\mathrm{W}}|$. Combining the cost of one fixed-MA evaluation with the number of candidate evaluations gives an overall worst-case complexity of $O(|V_{\mathrm{W}}||V_{\mathrm{M}}^{C}|^2|E'||V'|^2)$. Under the additional scaling assumption $|V_{\mathrm{W}}| = O(|V_{\mathrm{M}}^{C}|)$, this bound simplifies to $O(|V_{\mathrm{M}}^{C}|^3|E'||V'|^2)$.

5 Evaluation and Discussion

5.1 Experimental Setup

We use a MATLAB-based simulator to compare Vanilla, Layered, and AGBO. The primary network contains 38 nodes and 59 optical links. One central node hosts the GA, while the remaining 37 nodes host WAs. When a WA and its selected MA

share a physical node, the local update is directly available to that MA without consuming wavelength or storage resources.

Each optical link provides five wavelengths, each operating at $B_\lambda = 1$ Gbps. Full wavelength conversion is supported, and each payload occupies one wavelength on every traversed link. Each temporal layer of the TERG represents a scheduling interval of $\Delta t = 16$ s. For a payload of size F GB, one wavelength on each traversed link is reserved for $q(F)$ consecutive temporal layers, where $q(F)$ is calculated by rounding up as $q(F) = \lceil 8F/(B_\lambda \Delta t) \rceil$. We use a default temporal interval of 16 s and evaluate sensitivity to this interval in Section 5.7. Each node provides 50 GB of storage, representing the storage-rich operating regime considered in the current evaluation. Limited and heterogeneous storage settings are left for future study.

Unless otherwise specified, the local processing time on each WA is fixed at 100 s and is treated as an exogenous application-side parameter rather than an optimization variable. The collaboration deadline is 230 s, leaving a 130-s network transmission window. With $\Delta t = 16$ s per temporal layer, the default TERG contains $L = 8$ temporal layers, while the remaining 2 s is treated as a discretization margin. The background wavelength utilization U is defined as the fraction of occupied wavelength units across all spatial edges of the TERG. For the network-level evaluation, each experimental point is obtained from 100 independent background-reservation instances. The same instance is applied to all three schemes for a fair comparison, and the reported value is the arithmetic mean over the 100 instances. The three schemes are configured as follows:

- **Vanilla** [3] directly transmits WA updates to the GA without intermediate aggregation or temporal buffering; an update is excluded if an immediate E2E route cannot be established.
- **Layered** [5] uses hierarchical WA–MA–GA aggregation and the same candidate MA pool and AGBO-MS search as AGBO, but restricts the TERG to one temporal layer and excludes temporal buffering.
- **AGBO** uses the proposed framework and algorithm to select MA-hosting nodes and schedule wavelength and storage resources across multiple temporal layers.

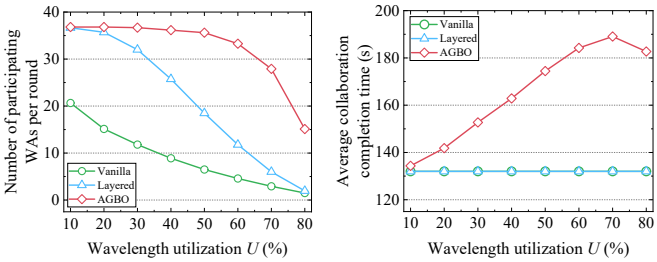
The network-level metrics are the number of participating WAs and the average collaboration completion time. Collaboration completion time is measured from the beginning of local processing until the last payload retained in the final collaboration result reaches the GA. It includes the fixed 100 s of local processing.

5.2 Network-Level Performance

Impact of wavelength utilization:

Figure 7 compares the three schemes as U increases from 10% to 80%. Higher utilization leaves fewer wavelength resources available and reduces WA participation.

AGBO retains more WAs by using transmission opportunities in later temporal layers, although some admitted payloads may reach the GA later. The WA participation ratio is the number of admitted WAs divided by the 37 available WAs. At $U = 60\%$, the WA participation ratios of AGBO, Layered, and Vanilla are approximately 90%, 32%, and 12%, respectively. AGBO therefore improves participation by at least 57 percentage points over the two baselines. Even at $U = 80\%$, it retains approximately 15 WAs, corresponding to about 41% participation.



(a) Number of participating WAs (b) Average collaboration completion time

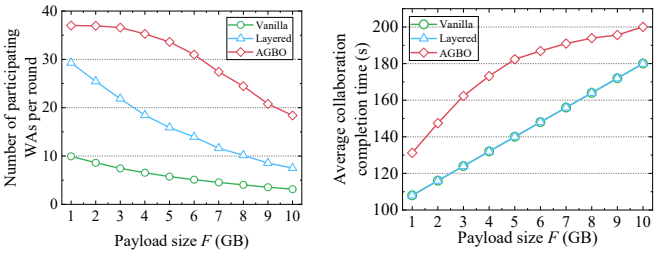
Figure 7: Impact of background wavelength utilization U on (a) the number of participating WAs and (b) the average collaboration completion time.

The completion time of AGBO increases until approximately $U = 70\%$ and then decreases at $U = 80\%$, where the admitted WA set becomes substantially smaller. The average completion times of Vanilla and Layered remain constant across the tested utilization levels. These results illustrate a participation–latency trade-off, so completion time should be evaluated together with WA participation.

Impact of payload size:

Figure 8 evaluates payload sizes from 1 to 10 GB. A larger payload requires one wavelength on each traversed link to be reserved over more consecutive temporal layers, thereby increasing resource contention and reducing the number of admitted WAs.

At $F = 10$ GB, AGBO admits approximately 18 WAs, whereas Layered and Vanilla admit approximately 8 and 3 WAs, respectively. These results indicate that storage-aided spatio-temporal routing maintains higher WA participation when payload transmission requires longer consecutive wavelength reservations across multiple temporal layers.



(a) Number of participating WAs (b) Average collaboration completion time

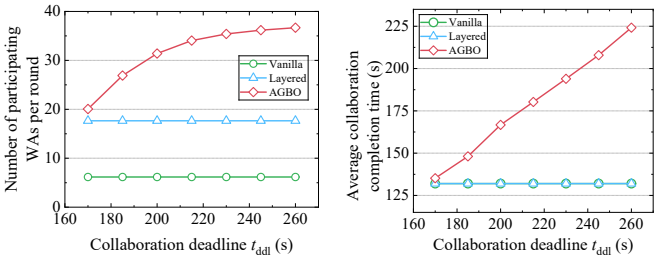
Figure 8: Impact of payload size F on (a) the number of participating WAs and (b) the average collaboration completion time.

Figure 8(b) shows that the average collaboration completion time increases with payload size for all three schemes. Vanilla and Layered have identical completion times across the tested payload sizes. AGBO admits more WAs than either baseline, as shown in Figure 8(a), and has a longer average completion time.

5.3 Impact of the Collaboration Deadline

Figure 9 evaluates collaboration deadlines from 170 to 260 s under $U = 50\%$ and $F = 4$ GB. A longer deadline increases the admissible number of TERG layers, allowing AGBO to exploit additional transmission opportunities. The number of participating WAs under AGBO increases from approximately 20 at 170 s to approximately 37 at 260 s. Its completion time also increases because some newly admitted payloads are transmitted through later temporal layers. The deadline can therefore be adjusted to balance participation and completion time according to application requirements.

Vanilla and Layered use only the current resource snapshot and do not exploit resources in later temporal layers. Their participation and completion times remain unchanged under the adopted no-waiting baseline model.



(a) Number of participating WAs (b) Average collaboration completion time

Figure 9: Impact of the collaboration deadline t_{ddl} on (a) the number of participating WAs and (b) the average collaboration completion time under $U = 50\%$ and $F = 4$ GB.

5.4 Collaborative-Learning Proxy Evaluation

The collaborative-learning evaluation is conducted separately from the preceding network-level experiments and is used as an application-level proxy for examining how network-induced differences in WA participation may be reflected in the outcome of a concrete collaborative workload. It is not intended to emulate the semantic aggregation or reasoning process of a typical LLM-based MoA system. In particular, the arithmetic model averaging adopted in this experiment is a workload-specific rule for the convolutional neural network (CNN)-based proxy rather than an aggregation requirement imposed by the AGBO framework.

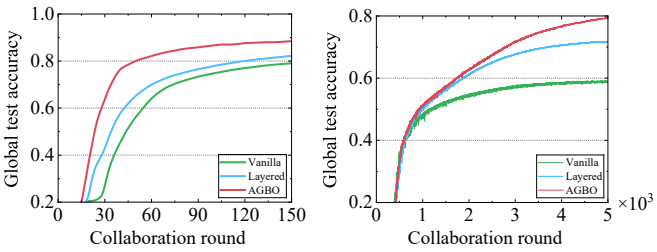
A fixed payload size is used for all schemes and collaboration rounds so that the experiment focuses on the application-level behavior associated with different network-induced participation patterns. Each WA trains a CNN with the same architecture on its local dataset and generates a local model update. Both MNIST and FEMNIST are partitioned among the WAs using a Dirichlet distribution with concentration parameter $\alpha = 0.5$. The same model initialization, data partitions, and training configuration are used for all three schemes.

For Layered and AGBO, each MA applies ordinary averaging to the successfully received WA updates and generates one ATB. The GA then applies equal-weight averaging to the received ATBs. For Vanilla, the GA directly averages the WA updates successfully delivered to it. This two-level averaging rule is specific to the adopted collaborative-learning proxy. When different MAs aggregate different numbers of WAs, equal weighting of the resulting ATBs may change the effective contribution of individual WAs. Therefore, the results shown below should be interpreted as the overall application-level outcome under the adopted hierarchical averaging configuration, rather than as a controlled estimate of the causal effect of WA participation.

For both hierarchical schemes, AGBO-MS searches over different MA configurations with the objective of increasing the final number of participating WAs under the considered configuration neighborhoods. The learning evaluation uses $U = 50\%$, $t_{ddl} = 230$ s, and $F = 4$ GB for all three schemes.

As shown in Figure 10, AGBO achieves faster accuracy improvement and higher final global test accuracy than the two baselines under the adopted collaborative-learning proxy. By the final reported round, AGBO reaches approximately 0.88 global test accuracy on MNIST, compared with approximately 0.82 for Layered and 0.79 for Vanilla. On FEMNIST, the corresponding accuracies are approximately 0.79, 0.72, and 0.59.

These results show that, under the adopted CNN-based proxy and two-level aggregation rule, the different network orchestration schemes lead to different application-level learning outcomes together with their different WA participation levels. The experiment is therefore intended as an illustrative downstream workload rather than a direct validation of LLM-based semantic MoA reasoning. Application-specific semantic aggregation mechanisms can be incorporated at the MA and GA layers without changing the network-side objective of transporting the required WA and ATB payloads within the collaboration deadline, provided that the resulting aggregated information can be represented by the ATB abstraction used in the present framework.



(a) MNIST (b) FEMNIST

Figure 10: Global test accuracy over consecutive collaboration rounds on (a) MNIST and (b) FEMNIST.

5.5 Evaluation Across Network Topologies

To examine performance across different network structures and sizes, we extend the evaluation from the original Metro topology (38 nodes and 59 links) to USNET (24 nodes and 43 links) and Germany50 (50 nodes and 88 links). The comparison uses $U = 50\%$ and the remaining default settings in Section 5.1. Within each topology, all three schemes are evaluated using the same background-resource instances.

Because the numbers of available WAs differ across the evaluated topologies, we use the WA participation ratio to enable a normalized cross-topology comparison. This ratio is defined as the final number of admitted WAs divided by the number of available WAs in the corresponding topology.

Figure 11 shows that AGBO maintains approximately 97%–98% participation and achieves the highest participation ratio on all three topologies. These results indicate that its participation advantage is maintained across the evaluated network structures and sizes, including the larger 50-node network.

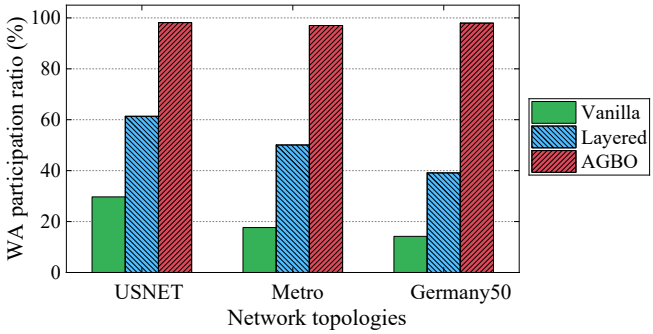


Figure 11: WA participation ratios of Vanilla, Layered, and AGBO across three network topologies at $U = 50\%$.

5.6 Ablation Study of AGBO Mechanisms

To assess the contributions of temporal SnF, hierarchical aggregation, and MA selection, we compare Vanilla, Layered, AGBO, and two ablation baselines at $U = 40\%$ and $U = 60\%$. SnF-only delivers WA updates directly to the GA through spatio-temporal paths with intermediate storage, without hierarchical aggregation. AGBO (fixed MA) retains aggregation and two-stage orchestration but disables the subsequent AGBO-MS search, keeping the same initially selected MA set across both utilization levels and all instances. All schemes use the default settings in Section 5.1 and the same 100 background-resource instances at each utilization level.

Figure 12 shows the average number of participating WAs per round. At both utilization levels, SnF-only admits more WAs than Vanilla, illustrating the benefit of temporal transmission opportunities. With SnF enabled, AGBO (fixed MA) further increases WA participation relative to SnF-only under the selected MA placement. AGBO achieves the highest WA participation at both utilization levels. Its gain over AGBO (fixed MA) is larger at $U = 60\%$ than at $U = 40\%$, suggesting that the MA-configuration search can be more beneficial under heavier contention in the tested scenarios. Together, these comparisons indicate complementary benefits from temporal SnF, hierarchical aggregation, and MA selection under the evaluated configurations.

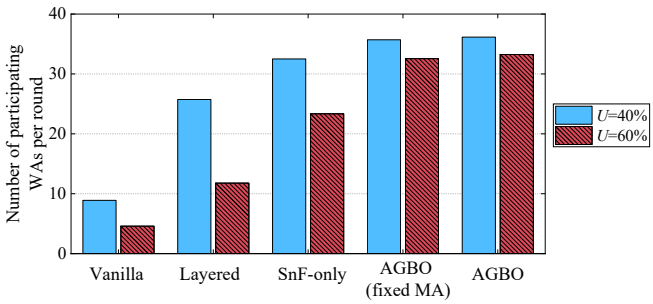


Figure 12: WA participation in the mechanism-ablation comparison at $U = 40\%$ and $U = 60\%$.

5.7 Sensitivity to Temporal Granularity

We compare temporal intervals of 8, 16, and 32 s at $U = 40\%$, with the physical payload transmission duration fixed at 32 s. Table 2 shows that AGBO maintains nearly full WA participation across the tested intervals and consistently exceeds the two baselines. Refining the interval from 16 to 8 s yields little additional participation benefit.

Table 2: WA participation under different temporal intervals at $U = 40\%$.

Scheme	Δt		
	8 s	16 s	32 s
Vanilla	9.29	9.29	9.29
Layered	27.50	27.50	27.50
AGBO	36.99	36.99	36.93

For a fixed scheduling horizon, finer discretization increases the number of TERG layers and the graph size. The interval should therefore be selected according to resource-variation timescales, payload transmission requirements, and the controller's computational budget. The 16-s default is a representative setting rather than a universally optimal choice.

5.8 Sensitivity to MA Initialization

To assess initialization sensitivity, we compare three AGBO initializations on the 38-node Metro topology at $U = 40\%$ and $U = 60\%$. The All candidates initialization activates the entire candidate MA-hosting pool, whereas the 2-hop subset and 3-hop subset initializations activate candidates two and three hops from the GA, respectively. All three use the same candidate pool, background-resource instances at each utilization level, and subsequent expansion, substitution, and contraction search.

Table 3 shows nearly identical WA participation across the three initializations. The reported means are equal at $U = 40\%$ and differ by at most 0.02 WAs at $U = 60\%$, indicating limited sensitivity to the tested initial configurations.

Table 3: WA participation under different initial MA configurations.

Initial MA configuration	Utilization (U)	
	40%	60%
All candidates	36.21	34.09
2-hop subset	36.21	34.11
3-hop subset	36.21	34.09

6 Conclusion

This paper proposes the AGBO framework for wide-area MoA collaboration over the O-IoA. A TERG represents time-varying wavelength availability and intermediate storage opportunities. Under the adopted analytical assumptions, simplified models indicate that temporal expansion can reduce

the modeled collaboration failure probability while enlarging the candidate-path search space. This structural trade-off motivates the AGBO algorithm, which combines two-stage maximum-flow orchestration with MA-configuration search.

Simulation results show that AGBO retains more WAs across different background wavelength-utilization levels, payload sizes, and collaboration deadlines. At a background wavelength utilization of $U = 60\%$, AGBO improves the WA participation ratio by at least 57 percentage points over the considered baselines. Under the adopted CNN-based collaborative-learning proxy and two-level aggregation rule, AGBO also achieves faster accuracy improvement and higher global test accuracy together with its higher WA participation.

These results demonstrate the potential of AGBO for supporting geo-distributed collaborative applications over optical infrastructure. By coordinating hierarchical aggregation with wavelength and temporal storage allocation, AGBO improves agent participation under resource contention and collaboration deadlines. The application-specific aggregation rule can be selected according to the corresponding workload, while the AGBO framework focuses on the network-side orchestration of WA and ATB payloads. The same spatio-temporal orchestration principle may also be extended to other deadline-constrained multi-agent services. Future work will examine correlation-aware modeling and validate the analytical approximations under shared-link and correlated-resource conditions.

Funding

This work was supported by the Key Technology Project of Fujian Province under Grant 2024HZ022005.

Author Contributions

Conceptualization, X.L. and W.L.; methodology, X.L. and Z.T.; software, X.L., Z.T. and W.L.; validation, X.L. and Z.T.; formal analysis, X.L., Z.T. and Z.F.; investigation, W.L., Z.T. and Weiling Lu; resources, X.L., Z.T. and Z.F.; data curation, Z.T., Z.F. and Weiling Lu; writing–original draft, X.L., W.L. and Z.T.; writing–review and editing, X.L., E.C. and Z.T.; visualization, X.L., Z.T. and Z.F.; supervision, X.L. and E.C.; project administration, X.L. and E.C.; funding acquisition, X.L. and E.C.. All authors have read and agreed to the published version of the manuscript.

Conflict of Interest

All the authors declare that they have no conflict of interest.

References

[1] Boateng, G.O., Sami, H., Alagha, A., Elmekki, H., Hammoud, A., Mizouni, R., Mourad, A., Otrok, H., Bentahar, J., Muhaidat, S., *et al.*: A Survey on Large Language Models for Communication, Network, and Service Management: Application Insights, Challenges, and Future Directions. IEEE Communications Surveys & Tutorials **28**(1), 527–566 (2026). https://doi.org/10.1109/COMST.2025.3564333

[2] Wang, Y., Guo, S., Pan, Y., Su, Z., Chen, F., Luan, T.H., Li, P., Kang, J., Niyato, D.: Internet of Agents: Fundamentals, Applications, and Challenges. IEEE Transactions on Cognitive Communications and Networking **12**, 4476–4501 (2026). https://doi.org/10.1109/TCCN.2025.3623369

[3] Qian, C., Xie, Z., Wang, Y., Liu, W., Zhu, K., Xia, H., Dang, Y., Du, Z., Chen, W., Yang, C., *et al.*: Scaling Large Language Model-based Multi-Agent Collaboration. In Proceedings of The Thirteenth International Conference on Learning Representations (ICLR 2025), pp. 1–18 (2025)

[4] Yao, D., Rabbani, T., Li, T.: Federation over Text: Insight Sharing for Multi-Agent Reasoning. In Proceedings of The Fourteenth International Conference on Learning Representations (ICLR 2026), pp. 1–46 (2026).

[5] Wang, J., Wang, J., Athiwaratkun, B., Zhang, C., Zou, J.Y.: Mixture-of-Agents Enhances Large Language Model Capabilities. In Proceedings of The Thirteenth International Conference on Learning Representations (ICLR 2025), pp. 1–20 (2025).

[6] Fu, T., Min, Z., Zhang, H., Yan, J., Dai, G., Ouyang, W., Wang, Y.: Cache-to-Cache: Direct Semantic Communication Between Large Language Models. In Proceedings of The Fourteenth International Conference on Learning Representations (ICLR 2026), pp. 1–29 (2026).

[7] Zhang, C., Li, C., Shen, H., Zhang, J., Zou, J., Xie, N., Wang, P., Sun, Q., Bae, S.-H., Wang, X., *et al.*: Efficient and Interpretable Multi-Agent LLM Routing via Ant Colony Optimization. arXiv preprint arXiv:2603.12933 (2026). https://doi.org/10.48550/arXiv.2603.12933

[8] Lin, X., Yang, H., Huang, Y., Chen, E.: Design of a Bulk Data Transfer Approach for Deadline-Critical Tasks in Optical Computing Power Networks. In Proceedings of 2025 13th International Conference on Information and Communication Networks (ICICN), pp. 239–244 (2025). https://doi.org/10.1109/ICICN67355.2025.11430506

[9] Zarei, M.H., Hadi, M., Pakravan, M.R.: Dynamic Routing and Spectrum Assignment in Elastic Optical Networks Using Graph Signal Correlation. Journal of Optical Communications and Networking **18**(8), 848–861 (2026). https://doi.org/10.1364/JOCN.592663

[10] Fan, Y., Li, Y., Natalino, C., Guo, J., Wu, W., Ruan, R., Wang, W., Zhao, Y., Zhang, J.: Distributed model training task migration for hotspot management in intelligent computing center interconnection with tidal characteristics. IEEE Transactions on Network and Service Management **22**(6), 5489–5502 (2025). https://doi.org/10.1109/TNSM.2025.3590011

[11] Chen, L., Li, Y., Natalino, C., Li, Y., Zhang, B., Fan, Y., Wang, W., Zhao, Y., Zhang, J.: Reliable and efficient RAR-based distributed model training in computing power network. Journal of Optical Communications and Networking **16**(5), 527–540 (2024). https://doi.org/10.1364/JOCN.511165

[12] Lian, M., Zhao, Y., Li, X., Liu, W., Li, Y., Tornatore, M., Zhang, J.: Resource allocation in flexible-bandwidth fine-grained optical transport networks for geo-distributed machine learning. IEEE Internet of Things Journal **12**(13), 25601–25619 (2025). https://doi.org/10.1109/JIOT.2025.3558933

[13] Hu, Q., Wang, W., Huang, C., Wang, X., Li, Y., Zhao, Y., Zheng, Y., Tan, Y., Zhang, J.: Task placement and traffic interleaving for cross-datacenter LLM training over optical networks. Journal of Optical Communications and Networking **18**(2), 137–149 (2026). https://doi.org/10.1364/JOCN.579324

[14] Jiang, D., Ren, X., Lin, B.Y.: LLM-Blender: Ensembling Large Language Models with Pairwise Ranking and Generative Fusion. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, pp. 14165–14178 (2023). https://doi.org/10.18653/v1/2023.acl-long.792

[15] Tekin, S. F., Ilhan, F., Huang, T., Hu, S., Liu, L.: LLM-TOPLA: Efficient LLM Ensemble by Maximising Diversity. In Findings of the Association for Computational Linguistics: EMNLP 2024, pp. 11951–11966 (2024). https://doi.org/10.18653/v1/2024.findings-emnlp.698

[16] Wang, Z., Moriyama, S., Wang, W.-Y., Gangopadhyay, B., Takamatsu, S.: Talk Structurally, Act Hierarchically: A Collaborative Refinement Framework for LLM Multi-Agent Systems. IEEE Transactions on Artificial Intelligence **7**(9), 4993–5003 (2026). https://doi.org/10.1109/TAI.2026.3676025

[17] Ma, X., Ma, Y., Lin, C., Yan, S., Bi, J., Cao, Z., Tian, Y., Tresp, V., Schuetze, H.: Self-Evolving Multi-Agent Systems via Textual Backpropagation. In Findings of the Association for Computational Linguistics: ACL 2026, pp. 9918–9951 (2026). https://doi.org/10.18653/v1/2026.findings-acl.483

[18] Liu, Z., Zhang, Y., Li, P., Liu, Y., Yang, D.: A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration. In Proceedings of Conference on Language Modeling (COLM), pp. 1–30 (2024)

[19] Ong, I., Almahairi, A., Wu, V., Chiang, W.-L., Wu, T., Gonzalez, J.E., Kadous, M.W., Stoica, I.: RouteLLM: Learning to Route LLMs with Preference Data. In Proceedings of The Thirteenth International Conference on Learning Representations (ICLR 2025), pp. 1–16 (2025).

[20] Wang, Y., Kong, L., Zhu, M., Gu, J., Cai, Y., Zhang, J.: Dynamic SFC mapping and time scheduling with store-and-forward scheme based on DRL in inter-datacenter elastic optical networks. IEEE Transactions on Networking **34**, 4110–4125 (2026). https://doi.org/10.1109/TON.2026.3674479

[21] Lin, X., Cai, Z., Zeng, C., Feng, Z., Lu, W., Chen, Y., Chen, E.: Fluctuation-Aware Routing and Scheduling for Bulk Transfers in Computing Power Optical Networks. In Proceedings of 2026 11th International Conference on Computer and Communication Systems (ICCCS), pp. 119–124 (2026). https://doi.org/10.1109/ICCCS69761.2026.11613512

[22] Lin, X., Sun, W., Wang, X., Yue, S., Veeraraghavan, M., Hu, W.: Time-space decoupled SnF scheduling of bulk transfers across inter-datacenter optical networks. IEEE Access **8**, 24829–24846 (2020). https://doi.org/10.1109/ACCESS.2020.2970542

Appendix

A Approximate Derivation of Aggregation-Subtree Expressions

This appendix presents the approximate derivation of the aggregation-subtree expressions used in Section 3. To keep the analysis tractable, we adopt a segment-wise relaxation. The WA–MA and MA–GA stages are evaluated separately in terms of their candidate-path search contributions and failure probabilities, while their temporal ordering is retained through the MA release layer. We further assume that all WA updates associated with a given release-layer case arrive at the MA in the same layer and treat the cases associated with different release layers as approximately independent. Consequently, the following expressions provide analytical approximations rather than exact combinatorial or probabilistic characterizations.

Consider the j-th aggregation subtree T_j, which contains N^j_{WpM} WAs associated with one MA. Under the regular fixed-route model, each WA–MA route contains N^j_{W2M} nodes, while the MA–GA route contains N^j_{M2G} nodes. The WAs first transmit their local updates to the MA, after which the MA generates one ATB and forwards it to the GA.

The analysis considers L release-layer cases according to the temporal layer in which the MA becomes ready to transmit its ATB. In the l-th case, all WA updates are assumed to arrive at the MA in layer l. The WA–MA stage can therefore use layers 1 through l, whereas the MA–GA stage can use layers l through L, as illustrated in Figure A1.

For one WA branch, $C(N^j_{\mathrm{W2M}}, l)$ denotes the number of candidate paths that reach the MA no later than layer l. These candidates include the paths already available within the first $l-1$ layers. Therefore, the number of additional candidate

paths that become available when the temporal range expands from $l-1$ to l layers is approximated by $C(N^j_{\mathrm{W2M}},l)-C(N^j_{\mathrm{W2M}},l-1)$. Multiplying this quantity by the N^j_{WpM} WA branches gives the approximate WA–MA search contribution of the l-th release-layer case.

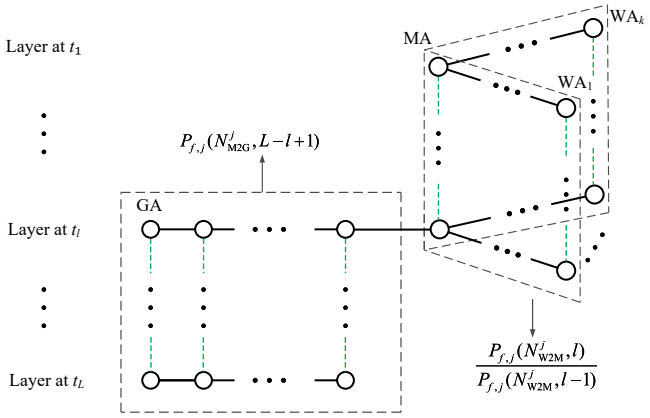


Figure A1: Illustration of the l-th release-layer case for aggregation subtree T_j. The WA–MA stage uses candidate paths available within the first l temporal layers, while the MA–GA stage starts from the MA release layer and uses the remaining $L-l+1$ layers.

A.1 Candidate-Path Search Complexity

Once the MA becomes ready in layer l, the remaining temporal range contains $L-l+1$ layers. The corresponding number of candidate MA–GA paths is $C(N^j_{\mathrm{M2G}},L-l+1)$. Under the segment-wise relaxation, the two stage-specific search contributions are added instead of explicitly enumerating their cross-stage Cartesian product. The approximate search complexity of the l-th release-layer case is therefore

$$C^j_{\mathrm{case}}(l)=N^j_{\mathrm{WpM}}\left[C(N^j_{\mathrm{W2M}},l)-C(N^j_{\mathrm{W2M}},l-1)\right]+C(N^j_{\mathrm{M2G}},L-l+1). \tag{A1}$$

Summing the search contributions over all possible MA release layers gives

$$\begin{aligned}&C_{\mathrm{stree}}\left(N^j_{\mathrm{W2M}},N^j_{\mathrm{M2G}},N^j_{\mathrm{WpM}},L\right)\\&=\sum_{l=1}^{L}N^j_{\mathrm{WpM}}\left[C(N^j_{\mathrm{W2M}},l)-C(N^j_{\mathrm{W2M}},l-1)\right]\\&\quad+\sum_{l=1}^{L}C(N^j_{\mathrm{M2G}},L-l+1).\end{aligned} \tag{A2}$$

The boundary condition $C(N,0)=0$ is adopted. Equation (A2) retains the temporal ordering between the two transmission stages while avoiding the combinatorial growth caused by explicitly enumerating all cross-stage path combinations.

A.2 Aggregation-Subtree Failure Probability

We next derive the approximate failure probability of subtree T_j using the same release-layer decomposition. As defined in Section 3, the aggregation factor γ_j defines the effective spatial-edge and temporal-edge unavailability probabilities as $p'_{b,j}=\gamma_j p_b$ and $p'_{s,j}=\gamma_j p_s$. For brevity, let $P_{f,j}(N,l)$ denote the probability that all candidate paths are unavailable, evaluated by substituting these effective probabilities into the basic function $P_f(N,l)$.

For one WA–MA branch, the candidate paths available within the first $l-1$ layers form a subset of those available within the first l layers. Under the nested-path and independence approximations, the ratio

$$\frac{P_{f,j}(N^j_{\mathrm{W2M}},l)}{P_{f,j}(N^j_{\mathrm{W2M}},l-1)} \tag{A3}$$

approximates the probability that all additional candidate paths made available by extending the temporal range from $l-1$ to l layers are unavailable. Hence, the probability that at least one of these additional candidate paths is available for a WA to reach the MA is

$$1-\frac{P_{f,j}(N^j_{\mathrm{W2M}},l)}{P_{f,j}(N^j_{\mathrm{W2M}},l-1)}. \tag{A4}$$

The subtree model adopts an all-or-nothing assumption. Therefore, all N^j_{WpM} WAs must successfully reach the MA in the considered release-layer case. Assuming that the WA branches are approximately independent, the success probability of the WA–MA stage is

$$\left[1-\frac{P_{f,j}(N^j_{\mathrm{W2M}},l)}{P_{f,j}(N^j_{\mathrm{W2M}},l-1)}\right]^{N^j_{\mathrm{WpM}}}. \tag{A5}$$

After aggregation, the MA has $L-l+1$ temporal layers available for forwarding the ATB to the GA. The success probability of the MA–GA stage is

$$1-P_{f,j}(N^j_{\mathrm{M2G}},L-l+1). \tag{A6}$$

Under the approximation that the two stages are independent, the success probability of the l-th release-layer case is

$$\begin{aligned}S_j(l)=&\left[1-P_{f,j}(N^j_{\mathrm{M2G}},L-l+1)\right]\\&\times\left[1-\frac{P_{f,j}(N^j_{\mathrm{W2M}},l)}{P_{f,j}(N^j_{\mathrm{W2M}},l-1)}\right]^{N^j_{\mathrm{WpM}}}.\end{aligned} \tag{A7}$$

The failure probability of this release-layer case is therefore $1-S_j(l)$. Under the additional approximation that the cases associated with different MA release layers are independent, the subtree is approximated as failed when all L release-layer cases fail. Its approximate failure probability is therefore

$$
\begin{aligned}
&P_{\text{stree}}\left(N^j_{\text{W2M}}, N^j_{\text{M2G}}, N^j_{\text{WpM}}, L, \gamma_j\right) \\
&\quad = \prod_{l=1}^{L} \Bigg\{ 1 - \left[1 - P_{f,j}(N^j_{\text{M2G}}, L-l+1)\right] \\
&\qquad \times \left[1 - \frac{P_{f,j}(N^j_{\text{W2M}}, l)}{P_{f,j}(N^j_{\text{W2M}}, l-1)}\right]^{N^j_{\text{WpM}}} \Bigg\}.
\end{aligned} \tag{A8}
$$

The boundary condition $P_{f,j}(N,0) = 1$ is used for the case $l = 1$. When $L = 1$, the expression reduces to the single-layer aggregation subtree considered by Layered. For $L > 1$, the expression accounts for the additional transmission opportunities created by extending the TERG over multiple temporal layers.