

Progressive Replacement Sequence Planning for Laser-Enhanced BeiDou Navigation Constellations

Huan Yan*, Juan A. Fraire^{†‡}, Ziqi Yang*, Kanglian Zhao*,
Wenfeng Li*, Jinjun Zheng[§], Ziyu Wang[§], Changgang Zheng[§],
* Nanjing University, Nanjing, China

[†] Inria, INSA Lyon, CITI, UR3720, 69621 Villeurbanne, France

[‡] CONICET – Universidad Nacional de Córdoba, Córdoba, Argentina

[§] China Academy of Space Technology, Beijing, China

Abstract—The early satellites of the third-generation BeiDou Navigation Satellite System (BDS-3) are approaching the end of their design lifetime, making progressive constellation replacement an inevitable engineering task. Meanwhile, laser inter-satellite links (LISLs) provide high-precision time transfer, sub-millimeter-level ranging, and high-rate data forwarding capabilities, offering a promising upgrade path for future BeiDou satellites. This paper investigates the replacement-sequence planning problem for progressively replacing legacy microwave satellites with laser-enabled satellites from a networking perspective. To address this problem, three main contributions are made: (i) BeiDou progressive replacement is formulated as a network-aware launch sequence planning problem over evolving laser-microwave hybrid constellation states; (ii) an integer linear programming (ILP) model is developed to evaluate the networking gain of each candidate launch action under practical engineering constraints; and (iii) a priority-aware search heuristic is introduced to reduce the action evaluation space for efficient round-wise decision making. Simulations based on the actual BDS-3 constellation demonstrate that the proposed replacement sequence achieves a more favorable hybrid-network evolution than the historical BDS-3 launch order, including earlier formation of a ground-accessible laser-enhanced network structure, better utilization of laser-link resources, fewer residual non-anchor satellites, and lower microwave-layer waiting delay. The proposed framework can serve as a network-aware decision-support tool for future BeiDou evolution, and the resulting replacement sequence can provide direct engineering guidance for BeiDou replacement planning.

Index Terms—BeiDou Navigation System, constellation replacement, laser inter-satellite link, laser-microwave hybrid network, launch sequence planning, integer linear programming.

I. INTRODUCTION

In 2017, China launched the first satellite of the third-generation BeiDou Navigation Satellite System (BDS-3) [1]. Since the design lifetime of a BDS-3 satellite is approximately ten years [2], the early satellites in the current constellation are now approaching the end of their service life and must be replaced by newly launched satellites.

Meanwhile, a growing body of research has investigated the application of laser inter-satellite links (LISLs) to future navigation constellations. Compared with the Ka-band phased-array microwave ISLs currently adopted in BeiDou [2], [3],

LISLs offer several prominent advantages. First, laser links can achieve picosecond-level time transfer accuracy [4], [5], thereby significantly improving the timing performance of the system. Second, the continuous sub-millimeter ranging capability of laser links can provide a large amount of high-precision inter-satellite measurement data [6], [7], which in turn improves autonomous orbit determination and enables better GNSS positioning services for users [8]. Third, the Gbps-level transmission rate of laser links enables high-speed inter-satellite networking and large-volume data forwarding and exchange, thus substantially enhancing the communication and network management capability of the system [7], [9].

Motivated by these facts, the next major evolution of the third-generation BeiDou constellation is to progressively replace the legacy microwave satellites reaching the end of their lifetime with newly launched satellites carrying laser payloads, hereafter referred to as *laser satellites* [10], [11]. It should be noted, however, that laser payloads do not eliminate the need for microwave payloads. In order to support diversified inter-satellite measurement tasks, navigation satellites must establish links and perform ranging with different objects within a short time. Because laser links require sophisticated pointing, acquisition, and tracking (PAT), their link switching process is slow and incurs a high reconfiguration cost [10]. Therefore, even the newly launched laser satellites still need to be equipped with fast-switching microwave phased-array antennas. In this sense, the laser payload is an incremental enhancement rather than a complete replacement of the microwave payload.

Most existing studies focus on topology construction for laser inter-satellite links in navigation constellations and typically optimize global metrics such as Position Dilution of Precision (PDOP) [11], end-to-end network delay [12], and network-wide connectivity [10]. However, these studies usually assume that the laser-enabled satellites have already been fully or partially deployed as a given set, while overlooking the gradual replacement process itself. In the early replacement stage, the number of deployed laser satellites is still limited, the resulting laser topology may be disconnected, and many launched laser satellites may remain outside the connected laser part of the constellation. Under such conditions, conventional metrics such as PDOP and network-wide connectivity are no longer suitable for guiding launch planning.

(Corresponding authors: Kanglian Zhao; Jinjun Zheng.)




Fig. 1. Intermediate laser–microwave hybrid constellation during progressive replacement. A ground-accessible local laser subnetwork coexists with residual microwave satellites, while some launched laser satellites may remain isolated.

Moreover, replacing an existing navigation constellation is fundamentally different from deploying a new constellation from scratch. During the initial deployment of BDS-3, the primary objective was to ensure basic Radio Navigation Satellite Service (RNSS) capability, including global coverage and geometric indicators such as PDOP [3]. In contrast, in the replacement problem studied in this paper, the original microwave constellation is already operational. The system therefore evolves from a functionally complete but capability-limited architecture to a stronger laser–microwave hybrid architecture, rather than from no service to full service. Consequently, the key planning problem is no longer how to establish the basic functionality of the constellation, but how to optimize the intermediate states generated during a multi-year replacement process.

The historical deployment of BDS-3 itself lasted from 2017 to 2020 [1], [2]. Likewise, the progressive replacement process considered here also spans multiple rounds separated by production scheduling, launch-site coordination, and mission preparation. The replacement order therefore directly affects the networking quality of the intermediate laser–microwave hybrid constellation over each inter-launch interval.

Fig. 1 illustrates the intermediate hybrid constellation state during progressive replacement. In this state, a ground-accessible local laser subnetwork coexists with residual microwave satellites, while some launched laser satellites may remain outside the connected laser subnetwork because of visibility limitations and terminal-resource constraints. Such laser–microwave coexistence defines the core networking scenario to be optimized during replacement.

Based on the above observations, this paper investigates the launch planning problem for progressively replacing legacy microwave satellites with laser-enabled satellites in the BeiDou navigation constellation. A round-wise launch planning framework is developed to determine the replacement sequence from

a networking perspective. For each launch round, a candidate action denotes a feasible set of satellites to be launched under practical payload constraints. Each action is evaluated by an integer linear programming (ILP) model based on a single-commodity-flow formulation, which characterizes the networking gain of the post-launch laser–microwave hybrid constellation. To improve computational efficiency, a priority-aware action search heuristic is further introduced to reduce the number of candidate actions requiring full ILP evaluation by exploiting visibility-based prior knowledge. This pruning mechanism is particularly relevant for future enlarged navigation architectures, where BeiDou may be integrated with low-orbit augmentation satellites [13], [14] or Earth–Moon libration-point navigation satellites [15], leading to a rapidly increasing action space.

Accordingly, the main contributions of this paper are summarized as follows:

- The progressive replacement problem for the BeiDou navigation constellation is formulated from a networking perspective. A system model is established for the laser–microwave hybrid replacement scenario, explicitly characterizing the interaction between the ground-accessible local laser subnetwork and the residual microwave satellites under practical engineering constraints.
- A round-wise launch planning framework is proposed. The framework combines a single-commodity-flow ILP model for network-performance evaluation with a priority-aware action search heuristic for reducing the action evaluation space, thereby enabling efficient replacement-sequence planning.
- Simulations based on the actual BDS-3 constellation are conducted to compare the proposed sequence with the historical BDS-3 launch order. The results show that the proposed sequence yields more favorable hybrid-network evolution and can serve as a direct engineering reference for replacing legacy microwave satellites from a networking perspective.

The remainder of this paper is organized as follows. Section II establishes the system model and engineering assumptions for the progressive replacement scenario. Section III presents the proposed launch planning method, including the state-level ILP model and the priority-aware action search heuristic. Section IV evaluates the computational feasibility and network-performance gains of the proposed framework. Finally, Section V concludes this paper.

II. SYSTEM MODEL

This section establishes the system model for the progressive replacement problem. The round-wise replacement process and its associated planning objective are first formalized at the constellation level. The topology abstraction underlying the action gain is then introduced, followed by the practical engineering constraints that govern the replacement process.

A. Progressive Replacement Problem Statement

We consider the progressive replacement of the BDS-3 constellation. Let $\mathcal{S}$ denote the set of all satellites in the

original constellation. During the replacement process, legacy microwave satellites are gradually replaced by newly launched satellites carrying laser payloads. The replacement proceeds over a sequence of launch rounds indexed by $r = 1, 2, \dots, R$.

Before round r , let $\mathcal{H}_{r-1} \subseteq \mathcal{S}$ denote the set of satellites that have already been launched as laser satellites in previous rounds. In round r , the planner selects one launch action from a candidate action set $\mathcal{A}_r$. Each action $a \in \mathcal{A}_r$ corresponds to launching a feasible subset of satellites under practical launch constraints. Once action a is executed, the set of currently available laser satellites becomes

$$\mathcal{L}_r(a) = \mathcal{H}_{r-1} \cup a, \quad (1)$$

while the remaining satellites are still treated as microwave satellites:

$$\mathcal{R}_r(a) = \mathcal{S} \setminus \mathcal{L}_r(a). \quad (2)$$

Accordingly, each candidate action a induces a post-launch hybrid network state, denoted by

$$\mathcal{N}_r(a) = (\mathcal{L}_r(a), \mathcal{R}_r(a)), \quad (3)$$

in which the launched laser satellites and the residual microwave satellites coexist and interact with each other.

The quality of this post-launch hybrid state is characterized by a gain function

$$G_r(a) = G(\mathcal{N}_r(a)), \quad (4)$$

which measures the networking benefit brought by launching action a at round r .

In practice, different launch rounds are usually separated by non-negligible time intervals due to satellite production, launch-site scheduling, mission preparation, and other engineering factors. For this reason, the replacement process is formulated as a round-wise greedy planning problem: at each round, the selected action should maximize the gain of the *immediate* post-launch hybrid network, so that each intermediate replacement stage remains as favorable as possible from the networking perspective. The launch decision at round r is therefore written as

$$a_r^* \in \arg \max_{a \in \mathcal{A}_r} G_r(a). \quad (5)$$

After the optimal action is selected and executed, the historical launched set is updated as

$$\mathcal{H}_r = \mathcal{H}_{r-1} \cup a_r^*. \quad (6)$$

Hence, the progressive replacement problem can be interpreted as a sequence of round-wise decisions, each aimed at selecting the launch action that yields the largest gain for the resulting laser-microwave hybrid network state. It should be emphasized that this formulation pursues per-round optimality of the immediate post-launch state rather than global optimality over the entire replacement horizon. Accordingly, the optimality-preserving property established in the appendix is to be understood within a single launch round, not as a guarantee of a globally optimal replacement sequence.

B. Topology-Based Action Gain Definition

This paper studies the progressive replacement process from a networking perspective, where the resulting network performance is fundamentally determined by the underlying laser and microwave topologies. Accordingly, this subsection first characterizes the visibility models and topology evolution mechanisms of the two networks, and then defines the launch-action gain function G based on these topology abstractions.

1) *Laser and Microwave Visibility Models*: Two satellites are said to be *visible* to each other if there is no celestial-body blockage between them and they are located within each other's beam coverage range. Visibility is the prerequisite for establishing a link. Since satellites are constantly moving, the visibility relation between any pair of satellites also varies continuously with time.

To handle this dynamic visibility, a finite-state automaton (FSA) abstraction is adopted [16]. Specifically, the system period is partitioned into a sequence of equal-length time segments, referred to as *FSA states*. If two satellites remain visible throughout the entire duration of one FSA state, they are regarded as visible in that state; otherwise, they are regarded as invisible in that state. In this way, the continuously varying visibility process is discretized into a sequence of static visibility snapshots.

The FSA durations are chosen differently for laser and microwave links because of their distinct physical characteristics.

Laser visibility FSA states: LISL switching relies on the PAT process. Because PAT is complex and incurs a high switching cost, the corresponding laser topology should not change too frequently. Otherwise, the links would be repeatedly torn down and re-established across consecutive states, causing a substantial portion of each state duration to be consumed by PAT. For this reason, previous studies commonly adopt a relatively long FSA duration for laser-topology evolution, typically on the order of one hour [10].

Microwave visibility FSA states: By contrast, phased-array microwave links are intended to switch rapidly among multiple observable objects so as to collect abundant inter-satellite measurement data. If the microwave FSA duration is chosen excessively long, the visibility criterion becomes overly restrictive, because two satellites must remain continuously visible over the entire FSA interval in order to be regarded as visible in that state. This reduces the effective visibility and correspondingly decreases the number of feasible link partners. Consequently, previous studies usually adopt a much shorter FSA duration for microwave-topology evolution, typically on the order of five minutes [13].

Let $\mathcal{K}$ denote the set of laser visibility states. For each laser state $k \in \mathcal{K}$, let $\mathcal{Q}_k$ denote the associated set of microwave visibility states. This two-layer abstraction reflects the fact that the performance of the laser topology should be evaluated under the coarser-grained laser FSA states, whereas the performance of the microwave topology must be assessed under the finer-grained microwave visibility states. Since laser satellites are also equipped with phased-array terminals, they participate in the construction of the microwave topology as well. By contrast, microwave satellites do not participate in the construction of the laser topology.

2) *Laser and Microwave Topology Evolution Under FSA States*: Based on the FSA abstraction above, the laser and microwave topologies are modeled hierarchically, as illustrated in Fig. 2.

Laser topology: Within each laser FSA state, the laser topology is fixed. Therefore, the local laser subnetwork formed by laser satellites interconnected through laser links in that state can be modeled as a static graph. Since a laser satellite typically carries multiple laser terminals, the resulting laser subnetwork can often behave as a connected backbone over the deployed laser satellites. Accordingly, the end-to-end transmission performance inside the laser subnetwork is mainly dominated by physical propagation delay.

Microwave topology: Within each microwave FSA state, the microwave topology is further refined in time. Specifically, each microwave state is divided into a sequence of superframes, typically of one minute, and each superframe is further divided into a sequence of time slots, typically of three seconds. A slot is the basic time unit for establishing phased-array microwave links between satellites, whereas a superframe is the basic planning unit of the microwave topology scheduler, i.e., the scheduler typically determines the phased-array link schedule for all slots contained in one superframe at a time. Accordingly, the microwave topology evolves dynamically over time according to the slot-level scheduling decisions.

The phased-array terminal currently carried by a BeiDou satellite can communicate with at most one target within a single slot [1], [2], [17]. As a result, the microwave network usually cannot provide an end-to-end transmission path within one slot. Data delivery through phased-array links therefore relies on the evolution of a dynamic topology across multiple slots, which is characteristic of a delay-tolerant networking (DTN) regime [18]. Under such a store-carry-forward mechanism, the dominant delay component is usually the waiting delay incurred while data are buffered at intermediate nodes until the next scheduled contact becomes available, which is quantized in multiples of the slot duration. By contrast, the physical signal propagation delay is typically below one second in Earth orbit and is therefore of secondary importance.

The laser topology and the microwave topology therefore exhibit fundamentally different characteristics. Within each laser FSA state, the laser subnetwork can be viewed as a high-speed connected backbone. By contrast, the microwave subnetwork is a slot-evolving DTN-style topology, whose connectivity should be interpreted at the superframe scale rather than at the scale of an individual slot.

3) *Definition of Action Gain*: Based on the hybrid-topology abstraction above, the gain of a launch action is defined from two complementary perspectives.

1) *Gain in local laser-subnetwork size*: Executing action a at round r enlarges the set of currently available laser satellites from the historical set $\mathcal{H}_{r-1}$ to $\mathcal{L}_r(a) = \mathcal{H}_{r-1} \cup a$. Under each laser FSA state, these currently available laser satellites may form a new connected laser subnetwork. Let

$$\mathcal{T}_{r,k}(a) \quad (7)$$




Fig. 2. Hierarchical topology abstraction of the laser-microwave hybrid constellation. Laser topology changes across laser FSA states, whereas microwave topology evolves across microwave FSA states, superframes, and time slots.

denote the connected local laser subnetwork formed under round r , action a , and laser state k , and let

$$S_{r,k}(a) = |\mathcal{T}_{r,k}(a)| \quad (8)$$

denote its size, measured by the number of mutually connected laser satellites.

As illustrated in Fig. 1, launched laser satellites may either be incorporated into the connected local laser subnetwork or remain isolated. From the networking perspective, an isolated laser satellite cannot benefit from the high-rate laser backbone and, in that state, functionally degenerates into a microwave-only node. Therefore, the planning objective naturally favors the largest possible connected local laser subnetwork.

A larger local laser subnetwork brings multiple benefits. First, all satellites inside the subnetwork can return telemetry data through the laser network, which is substantially faster than DTN-style return through phased-array links, where even waiting for one slot already incurs 3 seconds of delay. Second, because laser links provide significantly higher synchronization accuracy, the satellites inside the laser subnetwork can form a local high-precision timing network and can continuously obtain high-quality inter-satellite measurements. Third, a larger local laser subnetwork means that a larger fraction of users can directly exploit the high-quality laser-enabled part of the constellation. For example, once a short-message user [19] accesses a laser satellite inside the laser subnetwork, the corresponding message can be forwarded through the high-rate laser backbone with much lower overall latency.

To characterize the effect of action a across all laser states, the average local-laser-subnetwork size induced by executing action a in round r is defined as

$$\bar{S}_r(a) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} S_{r,k}(a). \quad (9)$$

2) *Gain in visibility support to the remaining microwave satellites*: The second aspect of the action gain concerns the microwave-state visibility relationship between the connected local laser subnetwork formed after executing action a and the residual microwave satellites.

Whenever a residual microwave satellite is visible (based on microwave FSA states) to at least one laser satellite contained in the local laser subnetwork, a phased-array link can be established between that residual microwave satellite and the subnetwork. Such a direct contact is beneficial because it enables one-hop access from the residual microwave satellite to the high-rate laser backbone, after which the data can continue to propagate through the connected laser subnetwork toward the ground segment. In addition, the time-synchronization accuracy is closely related to the number of timing-transfer hops: fewer hops generally lead to smaller accumulated synchronization errors. Since the laser subnetwork can be regarded as a high-precision timing network, one-hop access from a microwave satellite to the laser subnetwork is highly beneficial for reducing time synchronization error [20].

For each laser state k , the microwave-visibility support provided by the local laser subnetwork is defined as

$$C_{r,k}(a) = \sum_{q \in \mathcal{Q}_k} \sum_{m \in \mathcal{R}_r(a)} \sum_{i \in \mathcal{T}_{r,k}(a)} \delta_{im}^{(q)}, \quad (10)$$

where

$$\delta_{im}^{(q)} = \begin{cases} 1, & \text{if laser satellite } i \text{ and microwave satellite } m \text{ are visible in microwave FSA state } q, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

Averaging over all laser states yields

$$\bar{C}_r(a) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} C_{r,k}(a). \quad (12)$$

3) *Overall action gain*: Combining the two aspects above, the gain of action a in round r is defined as

$$G_r(a) = (\bar{S}_r(a), \bar{C}_r(a)), \quad (13)$$

where $\bar{S}_r(a)$ characterizes the average size of the connected local laser subnetwork after launching action a , and $\bar{C}_r(a)$ characterizes the average visibility support provided by that subnetwork to the residual microwave satellites. Therefore, $G_r(a)$ jointly quantifies, for the current launch round, both the connectivity scale of the laser-enabled part of the constellation and its support capability toward the still-unreplaced microwave part.

C. Engineering Constraints

Beyond the gain definition above, the replacement process is subject to several engineering constraints.

1) *Launch-mode constraint*: Launch actions are determined jointly by the adopted rocket type and the satellite platform to be deployed. In the BDS-3 setting, MEO satellites are usually launched in a two-satellite mode, whereas IGSO and GEO satellites are usually launched individually [21], [22]. In particular, when two MEO satellites are launched in one mission, the corresponding pair usually belongs to the same orbital plane; this characteristic is determined by the post-separation orbit-injection capability of the satellites rather than by the launcher alone. This reflects the diversity of feasible launch modes and implies that the candidate action set $\mathcal{A}_r$ is constrained by practical launcher configurations.

2) *Ground-segment connectivity and high-rate satellite-to-ground link constraint*: Telemetry data and service-related traffic, such as global short-message data, must eventually be delivered to ground stations. Meanwhile, the high-precision timing capability of the laser subnetwork depends on ground-referenced clock synchronization and the subsequent dissemination of timing information among laser satellites. Therefore, a local laser subnetwork is useful only when it remains connected to the ground segment during the replacement process.

However, the legacy S-band downlink [1]–[3] carried by current BeiDou satellites is insufficient to match the throughput of high-rate laser inter-satellite links and would therefore become the bottleneck of the overall system. Although satellite-to-ground laser transmission is a natural high-rate option, its performance is highly sensitive to atmospheric and weather conditions [23]. A more practical solution is therefore to equip part of the laser satellites with Ka-band large-aperture reflector antennas so as to construct high-rate satellite-to-ground links whose bandwidth can be matched with that of the inter-satellite laser backbone [24]. These links provide both the data-return path to the ground segment and the physical basis for satellite-to-ground time synchronization.

3) *Laser-terminal resource constraint*: The number of on-board laser terminals carried by each satellite also constrains the replacement problem. From the networking perspective, equipping satellites with more laser terminals generally improves connectivity and enhances performance. However, doing so also increases the overall satellite cost. Therefore, the number of onboard laser terminals must be explicitly modeled as a per-satellite resource constraint in the planning problem.

III. PROGRESSIVE LAUNCH PLANNING METHOD

This section presents the proposed method for evaluating candidate actions and selecting the optimal one at each launch round.

A. Round-Wise Planning Framework

Following the problem formulation in Section II-A, the progressive replacement process is solved in a round-wise manner. At each round, the planner starts from the current historical launched set and the corresponding candidate action set, and searches the action that maximizes the gain of the immediate post-launch hybrid constellation state.

The search process consists of two tightly coupled stages within each round. In the first stage, candidate actions that fail the necessary high-rate satellite-to-ground accessibility condition are removed, and the remaining actions are organized and screened by the priority-aware heuristic. In the second stage, the surviving actions are evaluated through the per-state ILP model, whose outputs are aggregated over all laser FSA states to obtain the round-level action score.

After the best action is identified, it is executed, the historical launched set is updated, and the planner proceeds to the next round. In this way, the overall replacement problem is solved as a sequence of round-wise decisions, while the detailed state-level evaluation of each action is carried out by the optimization model introduced in the following subsections.

B. Evaluation of Candidate Actions Across Laser States

1) *Evaluation rationale:* The networking gain of a candidate launch action is inherently state-dependent, since the feasible laser topology is determined by the current laser FSA state. Therefore, a launch action cannot be evaluated by a single static topology. Instead, for a fixed round r and a candidate action $a \in \mathcal{A}_r$, the local laser subnetwork must be evaluated separately under each laser state $k \in \mathcal{K}$.

For each tuple (r, a, k) , the planner determines which currently available laser satellites can be supported as the most favorable local laser subnetwork under the corresponding visibility condition, while satisfying ground-segment connectivity and engineering constraints. The resulting state-level quantities are then averaged over all laser states to obtain the round-level gain of the action.

2) *Evaluation graph under a given laser state:* Consider a fixed launch round r , a candidate action $a \in \mathcal{A}_r$, and a laser state $k \in \mathcal{K}$. After executing action a , the set of currently available laser satellites is denoted by $\mathcal{L}_r(a)$, and the set of ground stations is denoted by $\mathcal{G}$. To evaluate whether the selected laser satellites can support data backhaul and ground-referenced time synchronization, ground stations are explicitly included in the state-level evaluation graph. The node set is defined as

$$\mathcal{V}_{r,k}(a) = \mathcal{L}_r(a) \cup \mathcal{G}. \quad (14)$$

Under laser state k , the edge set contains two types of visibility-based candidate links. The first type is the inter-satellite laser link set $\mathcal{E}_{r,k}^{ss}(a)$ among satellites in $\mathcal{L}_r(a)$. The second type is the high-rate satellite-to-ground link set $\mathcal{E}_{r,k}^{sg}(a)$ between laser satellites and ground stations. A satellite-to-ground edge is included only when the satellite is equipped with a high-rate satellite-to-ground terminal and is visible to the corresponding ground station. Satellites without such terminals are treated as unavailable for high-rate ground access, even if a geometric line-of-sight to a ground station exists.

The resulting state-level laser-visibility graph is

$$\mathcal{G}_{r,k}^L(a) = (\mathcal{V}_{r,k}(a), \mathcal{E}_{r,k}^{ss}(a) \cup \mathcal{E}_{r,k}^{sg}(a)). \quad (15)$$

The graph in (15) provides the visibility-based support graph for evaluating action a under laser state k . It specifies the candidate laser and high-rate satellite-to-ground links available to the state-level ILP. Resource constraints, including laser-terminal limitations, ground-station capacity, and feeder-link thresholds, are not embedded in the graph itself; they are imposed later through the ILP constraints.

C. Single-State ILP for Local Laser-Subnetwork Evaluation

The evaluation graph $\mathcal{G}_{r,k}^L(a)$ in (15) specifies all visible laser inter-satellite links and high-rate satellite-to-ground links under that state. The task is to determine which currently available laser satellites should be selected into the local laser subnetwork and how they should be interconnected so that the resulting subnetwork is as beneficial as possible while remaining ground-accessible.

To this end, each node $v \in \mathcal{V}_{r,k}(a)$ is associated with a binary variable

$$z_v = \begin{cases} 1, & \text{if node } v \text{ is selected into the} \\ & \text{ground-connected laser subnetwork,} \\ 0, & \text{otherwise,} \end{cases} \quad (16)$$

where, in particular, $z_i = 1$ for a laser satellite i indicates that satellite i is selected into the local laser subnetwork. Likewise, each undirected edge $\{u, v\} \in \mathcal{E}_{r,k}^{ss}(a) \cup \mathcal{E}_{r,k}^{sg}(a)$ is associated with

$$y_{uv} = \begin{cases} 1, & \text{if edge } \{u, v\} \text{ is selected into the} \\ & \text{ground-connected laser subnetwork,} \\ 0, & \text{otherwise.} \end{cases} \quad (17)$$

A selected edge must connect two selected nodes:

$$y_{uv} \leq z_u, \quad \forall \{u, v\} \in \mathcal{E}_{r,k}(a), \quad (18)$$

$$y_{uv} \leq z_v, \quad \forall \{u, v\} \in \mathcal{E}_{r,k}(a), \quad (19)$$

where

$$\mathcal{E}_{r,k}(a) = \mathcal{E}_{r,k}^{ss}(a) \cup \mathcal{E}_{r,k}^{sg}(a).$$

Each satellite can carry only a limited number of laser terminals. Let d_i^{ss} denote the maximum number of inter-satellite laser links supported by laser satellite i . Then

$$\sum_{\{i,j\} \in \mathcal{E}_{r,k}^{ss}(a)} y_{ij} \leq d_i^{ss}, \quad \forall i \in \mathcal{L}_r(a). \quad (20)$$

Similarly, let d_v^{sg} denote the maximum number of high-rate satellite-to-ground links supported by node v . Then

$$\sum_{\{v,w\} \in \mathcal{E}_{r,k}^{sg}(a)} y_{vw} \leq d_v^{sg}, \quad \forall v \in \mathcal{V}_{r,k}(a). \quad (21)$$

In addition to terminal-resource limits, the selected local laser subnetwork must remain connected to the ground segment. To express this requirement, one ground station is selected as an auxiliary root. For each ground station $g \in \mathcal{G}$, introduce

$$\rho_g = \begin{cases} 1, & \text{if } g \text{ is selected as the auxiliary root,} \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

Exactly one ground station is selected as the root, and that ground station must itself be included in the ground-connected laser subnetwork:

$$\sum_{g \in \mathcal{G}} \rho_g = 1, \quad (23)$$

$$\rho_g \leq z_g, \quad \forall g \in \mathcal{G}. \quad (24)$$

Ground accessibility is a global connectivity requirement and does not admit a compact direct linear description. A single-commodity-flow construction is therefore adopted. Let $\mathcal{D}_{r,k}(a)$ denote the directed arc set obtained by replacing each undirected edge in $\mathcal{E}_{r,k}(a)$ with two opposite arcs. For each $(u, v) \in \mathcal{D}_{r,k}(a)$, let $f_{uv} \geq 0$ denote an auxiliary flow variable, and let $s_g \geq 0$ denote the source flow injected from ground station g . These flow variables are introduced only to enforce ground-rooted connectivity; they do not represent the actual physical traffic process in the hybrid network.

Flow is allowed on an arc only when the corresponding undirected edge is selected:

$$f_{uv} \leq (|\mathcal{V}_{r,k}(a)| - 1) y_{uv}, \quad \forall (u, v) \in \mathcal{D}_{r,k}(a). \quad (25)$$

Only the selected root can inject source flow:

$$0 \leq s_g \leq |\mathcal{V}_{r,k}(a)| \rho_g, \quad \forall g \in \mathcal{G}. \quad (26)$$

The total injected flow must equal the number of selected nodes:

$$\sum_{g \in \mathcal{G}} s_g = \sum_{v \in \mathcal{V}_{r,k}(a)} z_v. \quad (27)$$

For each selected laser satellite, one unit of flow is consumed:

$$\sum_{(u,i) \in \mathcal{D}_{r,k}(a)} f_{ui} - \sum_{(i,w) \in \mathcal{D}_{r,k}(a)} f_{iw} = z_i, \quad \forall i \in \mathcal{L}_r(a). \quad (28)$$

For each ground station, the corresponding conservation equation is

$$\sum_{(u,g) \in \mathcal{D}_{r,k}(a)} f_{ug} + s_g - \sum_{(g,w) \in \mathcal{D}_{r,k}(a)} f_{gw} = z_g, \quad \forall g \in \mathcal{G}. \quad (29)$$

Together, (25)–(29) guarantee that every selected laser satellite belongs to a connected local laser subnetwork and that this subnetwork remains connected to the ground segment through the selected root ground station.

In addition, the number of deployed high-rate satellite-to-ground links should scale with the size of the selected local laser subnetwork, so that the throughput of the laser backbone can be matched by the ground-return capability. Let

$$\mathcal{C} = \{(\eta, \beta)\}$$

denote a set of threshold pairs, where η is a local-laser-subnetwork-size threshold and β is the minimum required number of deployed high-rate satellite-to-ground links once that threshold is reached. For each $(\eta, \beta) \in \mathcal{C}$, introduce an auxiliary binary variable $b_{\eta,\beta}$ and impose

$$\sum_{i \in \mathcal{L}_r(a)} z_i \geq \eta b_{\eta,\beta}, \quad (30)$$

$$\sum_{i \in \mathcal{L}_r(a)} z_i \leq (\eta - 1) + |\mathcal{L}_r(a)| b_{\eta,\beta}, \quad (31)$$

$$\sum_{\{u,v\} \in \mathcal{E}_{r,k}^{sg}(a)} y_{uv} \geq \beta b_{\eta,\beta}. \quad (32)$$

After the support graph has been restricted by the above constraints, the state-level contribution of action a can be quantified by two gain components. The primary component is the size of the selected local laser subnetwork:

$$S_{r,k}(a) = \sum_{i \in \mathcal{L}_r(a)} z_i. \quad (33)$$

The secondary component is the microwave-visibility support provided by the selected local laser subnetwork to the residual microwave satellites. For each currently available laser satellite $i \in \mathcal{L}_r(a)$, define

$$c_i = \sum_{q \in \mathcal{Q}_k} \sum_{m \in \mathcal{R}_r(a)} \delta_{im}^{(q)}, \quad (34)$$

where $\delta_{im}^{(q)}$ is the visibility indicator defined in (11). The total microwave-visibility support is then given by

$$C_{r,k}(a) = \sum_{i \in \mathcal{L}_r(a)} c_i z_i. \quad (35)$$

The gain comparison follows a lexicographic priority: the local-laser-subnetwork size is maximized first, and the microwave-visibility support is used to distinguish solutions with the same subnetwork size. This lexicographic rule is implemented through the following weighted objective:

$$F_{r,k}(a) = M \sum_{i \in \mathcal{L}_r(a)} z_i + \sum_{i \in \mathcal{L}_r(a)} c_i z_i, \quad (36)$$

where M is a constant chosen larger than the maximum attainable value of the secondary term; since the secondary term is bounded by $\sum_{i \in \mathcal{L}_r(a)} c_i$, any $M > \sum_{i \in \mathcal{L}_r(a)} c_i$ is sufficient. With this choice, selecting one additional laser satellite always outweighs any possible difference in microwave-visibility support. Thus, the first term dominates the optimization, while the second term only refines the selection among solutions with the same number of selected laser satellites.

Finally, for each fixed tuple (r, a, k) , the single-state local-laser-subnetwork evaluation problem is written as

$$\begin{aligned} \mathcal{P}_{r,k}(a) : \quad & \max F_{r,k}(a) \\ \text{s.t.} \quad & (18) - (21), (23) - (32). \end{aligned} \quad (37)$$

Dense-topology refinement: The ILP in (37) is designed primarily for replacement-sequence evaluation rather than full topology construction. Its objective determines the best candidate launch action by prioritizing the size of the ground-connected laser subnetwork and its microwave-visibility support. However, when multiple feasible solutions achieve the same value of $F_{r,k}(a)$, their selected LISLs may still differ. In particular, the resulting edge set may use only a subset of available laser terminals, although a denser topology is preferable for topology-level operation because it can provide more inter-satellite measurement opportunities and reduce the hop distance within the local laser subnetwork.

When an explicit local laser topology is required, a secondary refinement can be applied after solving (37). Specifically, the optimal value of the primary action-evaluation objective is first fixed, and the number of selected LISLs is then maximized:

$$\max D_{r,k}(a) = \sum_{\{i,j\} \in \mathcal{E}_{r,k}^{ss}(a)} y_{ij}. \quad (38)$$

This refinement does not change the launch-action score, the selected subnetwork size, or the microwave-support capability. It only selects, among primary-optimal solutions, a denser ground-connected laser topology. Therefore, it can be used as a topology-construction option for post-processing and network-performance evaluation, while the main replacement-sequence planning logic remains governed by (37). This refinement is referred to as DTopo-ILP in the subsequent evaluation.

D. Multi-State Action Evaluation

For a candidate action $a \in \mathcal{A}_r$, the single-state ILP in (37) is solved under every laser state $k \in \mathcal{K}$. Let

$$F_{r,k}^*(a) \quad (39)$$

denote the optimal objective value obtained under state k .

Following the gain definition in Section II-B3, the round-level quality of action a is determined by aggregating its state-level performance over all laser states. For implementation convenience, this aggregation is represented by the scalar score

$$J_r(a) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} F_{r,k}^*(a), \quad (40)$$

which serves as the computational realization of the round-level gain comparison.

Accordingly, after all admissible candidate actions have been evaluated, the launch decision at round r is obtained as

$$a_r^* \in \arg \max_{a \in \mathcal{A}_r} J_r(a). \quad (41)$$

E. Priority-Aware Action Search Heuristic

The priority-aware search exploits visibility-based prior knowledge and necessary high-rate satellite-to-ground accessibility conditions to reduce the number of candidate actions requiring full ILP evaluation. The search procedure does not replace the state-level ILP model. Instead, it determines which actions should be evaluated first and which actions can be safely excluded once a dominant action has been found.

1) High-rate satellite-to-ground accessibility screening:

Not all satellites are equipped with high-rate satellite-to-ground links. Consequently, some candidate actions cannot satisfy the ground-segment accessibility requirement in the current round, even before detailed ILP evaluation. Such actions are filtered out in advance. The screening rule removes any candidate action that fails a necessary high-rate satellite-to-ground accessibility condition when combined with the already launched set $\mathcal{H}_{r-1}$. This preprocessing step reduces the action space while remaining consistent with the engineering constraints introduced in Section II-C2.

2) *Comparison classes and full local laser subnetwork:* To make the action comparison precise, the candidate action set is partitioned according to action cardinality. For each admissible integer h , define

$$\mathcal{A}_r^{(h)} = \{a \in \mathcal{A}_r : |a| = h\}, \quad (42)$$

where $|a|$ denotes the number of satellites contained in action a . Thus, $\mathcal{A}_r^{(h)}$ collects all candidate actions with the same action cardinality. These comparison classes are traversed in descending order of h , so that actions launching more satellites are considered before actions launching fewer satellites.

For a candidate action a at round r , a state-level solution is said to form a *full local laser subnetwork* under laser state k if all currently available laser satellites are included in the selected ground-connected local laser subnetwork, i.e.,

$$\sum_{i \in \mathcal{L}_r(a)} z_{i,r,k}^*(a) = |\mathcal{L}_r(a)|. \quad (43)$$

An action is referred to as *round-wise full-subnetwork feasible* if (43) holds for every laser state $k \in \mathcal{K}$. Such an action attains the maximum possible primary objective value within its comparison class.

3) *High-priority action subset and ordering rule:* Within each comparison class, the search first identifies a subset of actions that may achieve the primary objective, i.e., forming a full local laser subnetwork. The key observation is that a full local laser subnetwork can exist only if all currently available laser satellites are ground-reachable in the underlying laser-visibility graph before resource constraints are imposed. Therefore, this reachability condition can be used as a necessary screening rule before full ILP evaluation.

For a candidate action $a \in \mathcal{A}_r^{(h)}$, the state-level laser-visibility graph $\mathcal{G}_{r,k}^L(a)$ has already been defined in (15). Let $\text{Reach}_{r,k}^G(a)$ denote the set of laser satellites that are connected to at least one ground station in this graph. An action a is included in the high-priority subset $\mathcal{A}_{r,H}^{(h)}$ if all currently available laser satellites are ground-reachable for every laser state:

$$\mathcal{L}_r(a) \subseteq \text{Reach}_{r,k}^G(a), \quad \forall k \in \mathcal{K}. \quad (44)$$

Actions satisfying (44) are regarded as high-priority actions because they are the only actions that may form a round-wise full local laser subnetwork. If an action fails this reachability condition, then at least one available laser satellite is disconnected from the ground segment even in the relaxed visibility graph. After adding laser-terminal constraints, ground-station capacity constraints, and feeder-link threshold constraints, such an action still cannot form a full local laser subnetwork.

The actions in $\mathcal{A}_{r,H}^{(h)}$ are then ranked in descending order of net visibility-support gain. This ranking is computed under a hypothetical full-subnetwork condition: all satellites in $\mathcal{L}_r(a)$ are temporarily assumed to be included in the local laser subnetwork, and the corresponding microwave-visibility support to the residual microwave satellites is evaluated using the same visibility-support term defined in Section III-C. Equivalently, the ranking score can be written as

$$\Gamma_r(a) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} C_{r,k}^{\text{full}}(a), \quad (45)$$

where $C_{r,k}^{\text{full}}(a)$ denotes the microwave-visibility support obtained by setting $z_i = 1$ for all $i \in \mathcal{L}_r(a)$ in the secondary support term of the state-level objective. This score therefore measures the support capability that action a would provide if it were later verified by the ILP model to form a full local laser subnetwork.

The ranking in (45) does not replace ILP-based evaluation; it only determines the traversal order inside the high-priority subset. Once a ranked action is verified by the state-level ILP to form a round-wise full local laser subnetwork, its earlier position in the ranking implies that its microwave-visibility support is no smaller than that of any later full-subnetwork action in the same comparison class. This property enables the early-termination rule described next.

4) *Optimality-preserving early termination:* The early-termination rule is applied within each comparison class. Since the comparison classes are traversed in descending order of

action cardinality, once a round-wise full-subnetwork feasible action is found in class $\mathcal{A}_r^{(h)}$, no action in any lower-cardinality class $\mathcal{A}_r^{(h')}$ with $h' < h$ can outperform it in the primary objective. Therefore, all remaining lower-cardinality classes can be excluded from further search.

Within a fixed comparison class $\mathcal{A}_r^{(h)}$, the high-priority subset $\mathcal{A}_{r,H}^{(h)}$ contains all actions that may form a round-wise full local laser subnetwork. Actions outside this subset cannot become full-subnetwork feasible because they fail the relaxed ground-reachability condition before resource constraints are imposed. Therefore, once a full-subnetwork feasible action has been found in $\mathcal{A}_{r,H}^{(h)}$, actions outside the high-priority subset in the same class do not need to be further examined for the primary objective.

The actions in $\mathcal{A}_{r,H}^{(h)}$ are traversed in descending order of net visibility-support gain and evaluated by the state-level ILP model. Once the first round-wise full-subnetwork feasible action is found, the search within the current comparison class can be terminated. Any later full-subnetwork feasible action in the same ranked list has no larger secondary visibility-support gain, while any non-full-subnetwork action is dominated in the primary objective. Consequently, the remaining unevaluated actions in the current comparison class and all lower-cardinality classes can be safely excluded. A rigorous proof of this optimality-preserving property is provided in the appendix.

If no round-wise full-subnetwork feasible action is found in $\mathcal{A}_{r,H}^{(h)}$, the remaining actions in $\mathcal{A}_r^{(h)} \setminus \mathcal{A}_{r,H}^{(h)}$ are still evaluated by the state-level ILP model. This step is necessary because, when no full-subnetwork action exists in the current class, the optimal action may be a non-full-subnetwork action.

F. Overall Procedure

The complete planning process is an outer round-wise iterative procedure over the remaining unreplaced satellites. In each round, the procedure generates candidate actions, applies high-rate satellite-to-ground accessibility screening, constructs comparison classes, performs priority-aware action search, evaluates selected actions by the state-level ILP model, and updates the launched set. Algorithm 1 summarizes the core workflow.

IV. EVALUATION

This section evaluates the proposed progressive launch planning framework from the perspective of network performance. A launch sequence is not assessed as an isolated ordering result; instead, it is evaluated through the laser-microwave hybrid constellation states induced after successive replacement rounds.

A. Simulation Setup

1) *Baseline*: To the best of our knowledge, the network-aware replacement-sequence planning problem for an operational navigation constellation such as BeiDou has not been directly addressed in the literature. Launch-sequence

Algorithm 1 Priority-Aware Progressive Launch Planning

Input: Historical launched set $\mathcal{H}_0$, total satellite set $\mathcal{S}$, laser-state set $\mathcal{K}$, state-dependent visibility data
Output: Round-wise replacement sequence $\{a_r^*\}$

- 1: $r \leftarrow 1$
- 2: **while** $\mathcal{S} \setminus \mathcal{H}_{r-1} \neq \emptyset$ **do**
- 3: Generate the candidate action set $\mathcal{A}_r$
- 4: Remove actions violating the necessary high-rate satellite-to-ground accessibility condition
- 5: Partition the remaining actions into comparison classes $\mathcal{A}_r^{(h)}$ by action cardinality
- 6: Sort the comparison classes in descending order of h
- 7: Initialize the evaluated-action set $\mathcal{E}_r \leftarrow \emptyset$
- 8: **for all** comparison classes $\mathcal{A}_r^{(h)}$ **do**
- 9: Construct the high-priority subset $\mathcal{A}_{r,H}^{(h)}$ by relaxed ground-reachability screening
- 10: Rank actions in $\mathcal{A}_{r,H}^{(h)}$ by descending net visibility-support gain
- 11: Evaluate ranked actions in $\mathcal{A}_{r,H}^{(h)}$ by the state-level ILP model and add them to $\mathcal{E}_r$ until a round-wise full-subnetwork feasible action is found or all actions are evaluated
- 12: **if** a round-wise full-subnetwork feasible action is found **then**
- 13: Terminate the search within the current class and all lower-cardinality classes
- 14: **break**
- 15: **else**
- 16: Evaluate the remaining actions in $\mathcal{A}_r^{(h)} \setminus \mathcal{A}_{r,H}^{(h)}$ by the state-level ILP model and add them to $\mathcal{E}_r$
- 17: **end if**
- 18: **end for**
- 19: Select $a_r^* = \arg \max_{a \in \mathcal{E}_r} J_r(a)$, where $J_r(a)$ is obtained by aggregating the state-level ILP evaluation results of action a
- 20: Update $\mathcal{H}_r \leftarrow \mathcal{H}_{r-1} \cup a_r^*$ and set $r \leftarrow r + 1$
- 21: **end while**

optimization has been studied for LEO mega-communication-constellations [25], but that setting targets the initial deployment of a new constellation from scratch rather than the network-aware progressive replacement of an existing one. Consequently, no directly comparable method is available as a baseline. The historical initial deployment order of BDS-3 is adopted as the comparison baseline [21], [22]. Specifically, this baseline replaces satellites strictly following the original BDS-3 launch sequence. This choice is consistent with a lifetime-driven engineering strategy, because satellites launched earlier are naturally closer to the end of their design life and would be replaced earlier under a straightforward replacement policy.

2) *Constellation and temporal settings*: The simulated constellation is the actual BDS-3 constellation, consisting of 24 MEO satellites, 3 GEO satellites, and 3 IGSO satellites. After each launch action is executed, the resulting laser-microwave hybrid constellation is simulated over a 7-day horizon. The average network gain over these seven days is used as the evaluation score of that action.

TABLE I
MAIN SIMULATION PARAMETERS

Parameter	Value
Constellation	24 MEO + 3 GEO + 3 IGSO
Evaluation horizon per action	7 days
Laser FSA duration	1 h
Microwave FSA duration	5 min
Superframe duration	1 min
Time slot duration	3 s
MEO laser pointing range	70°
GEO/IGSO laser pointing range	80°
MEO microwave pointing range	60°
GEO/IGSO microwave pointing range	45°
Ground-station pointing range	85°
Ground stations (GSs)	Sanya (18.23°N, 109.02°E) Weinan (34.31°N, 109.28°E) Kashi (39.47°N, 75.989°E)

The temporal hierarchy of the hybrid network follows the system model in Section II. The laser FSA duration is set to 1 h, and the microwave FSA duration is set to 5 min. Each microwave state is further divided into superframes of 1 min, and each superframe is divided into time slots of 3 s.

The pointing ranges are configured according to representative satellite-link settings [13], [26]. For laser links, the pointing range is set to 70° for MEO satellites and 80° for GEO/IGSO satellites. For phased-array microwave links, the pointing range is set to 60° for MEO satellites and 45° for GEO/IGSO satellites. The pointing range of each ground station is set to 85°.

Three ground stations are considered, namely Sanya, Weinan, and Kashi. The main constellation, temporal, visibility, and ground-segment parameters are summarized in Table I.

3) Engineering constraints:

a) *Launch-mode setting:* The launch-mode setting follows the historical BDS-3 deployment mode. MEO satellites are launched in a “one-rocket two-satellite” mode, where the two satellites belong to the same orbital plane, whereas IGSO and GEO satellites are launched in a “one-rocket one-satellite” mode. In the present study, this setting determines the rule by which each candidate launch action a is generated.

b) *High-rate satellite-to-ground link setting:* Fig. 3 shows, for each satellite, the number of laser FSA states in which it is visible to at least one ground station over 30 days. A clear difference can be observed among different orbital types. The three GEO satellites are visible to the ground segment in all 744 laser FSA states. Two IGSO satellites are also visible in all 744 states, and the remaining IGSO satellite is visible in 672 states. In contrast, the 24 MEO satellites are visible only in 269–291 states. Therefore, GEO and IGSO satellites provide nearly continuous or highly persistent ground access, whereas MEO satellites can access the ground segment only during intermittent passes.

The persistent ground visibility of GEO/IGSO satellites makes them more suitable for hosting high-rate satellite-to-ground terminals. Such terminals generally require large-aperture antennas, sufficient onboard power, and stable ground-access opportunities to match the bandwidth of the inter-satellite laser backbone. In contrast, the intermittent ground-access windows and smaller platform capacity of MEO




Fig. 3. Number of laser FSA states in which each satellite is visible to at least one ground station over the considered 30-day observation window. GEO and IGSO satellites exhibit much more persistent ground visibility than MEO satellites.

satellites, approximately 1 ton [2], limit both antenna utilization efficiency and payload accommodation capability. GEO and IGSO satellites have larger platforms, approximately 5 ton [2], and much longer ground-access durations. Accordingly, high-rate satellite-to-ground links are deployed only on GEO and IGSO satellites in the simulation setting.

To match the bandwidth capability of the laser backbone and the satellite-to-ground return segment, the feeder-link thresholds are configured, in the notation of Section III-C, as

$$(\eta, \beta) \in \{(1, 1), (10, 2), (20, 3)\},$$

meaning that at least one high-rate satellite-to-ground link is required when the local laser subnetwork contains one laser satellite, at least two such links are required when the subnetwork size reaches 10, and at least three such links are required when the subnetwork size reaches 20. The first threshold prevents the first launched laser satellite from becoming disconnected from the ground segment, whereas the latter two thresholds approximate the increasing traffic load carried by a larger local laser subnetwork.

c) Laser-terminal and ground-station resource setting:

Following common assumptions in prior studies [11], each satellite is assumed to carry 3 laser terminals so as to support the formation of a highly connected laser network in the fully replaced stage. In addition, because ground-segment resources are also limited, each ground station is assumed to maintain at most two simultaneous high-rate satellite-to-ground links.

B. ILP Model Computational Feasibility and Search Efficiency

This part evaluates the computational feasibility of the ILP model and the efficiency of the priority-aware action search procedure. All ILP instances are solved by the commercial solver Gurobi 12.0. The experiments are conducted on a personal computer with an AMD Ryzen 7 7700X 8-Core Processor CPU and 32 GB memory.

TABLE II
COMPUTATIONAL FEASIBILITY AND SEARCH EFFICIENCY

Item	Value
Solver	Gurobi 12.0
CPU	AMD Ryzen 7 7700X 8-Core
Memory	32 GB
Number of launch rounds	18
Raw candidate-action count	567
Action count after ground-link screening	439
ILP fully evaluated action count	116
Reduction in full ILP evaluations	73.6%
Total evaluation time	503 s
Average time per evaluated action	4.33 s

TABLE III
NODE-TYPE DEFINITION IN THE HYBRID NETWORK

Type	Definition	Operational implication
A1	In laser subnet, GS-invisible	Anchor satellite, Ground access via laser subnetwork and high-rate sat-to-ground link
A2	In laser subnet, GS-visible	Anchor satellite, Direct (via legacy S-Band link) or laser-assisted ground access (via high-rate sat-to-ground link)
A3	Outside laser subnet, GS-visible	Anchor satellite, Direct ground access via legacy S-Band link
A4	Outside laser subnet, GS-invisible	Non-anchor satellite, Ground access by establishing microwave links with A1, A2, and A3 satellites

Table II summarizes the results. Across all replacement rounds, the raw candidate action space contains 567 actions. The complete priority-aware action search procedure first applies the necessary high-rate satellite-to-ground accessibility screening, after which 439 actions remain admissible. Among these admissible actions, only 116 actions are fully evaluated by the ILP-based network-performance model according to the priority-aware ranking strategy. Compared with exhaustive evaluation of all admissible actions, the number of full ILP evaluations is reduced by 73.6%. The total evaluation time is 503 s, and the average time per evaluated action is 4.33 s. These results indicate that the ILP model is lightweight enough for repeated round-wise evaluation, while the action-search procedure substantially reduces the number of full evaluations.

C. Network-Performance Results

The following results evaluate the hybrid-network states produced by the proposed replacement sequence from three aspects: the structural evolution of the hybrid constellation, the performance of the ground-connected laser subnetwork, and the residual microwave-layer performance.

1) *Evolution of the Hybrid Network Structure*: For subsequent analysis, each satellite in the hybrid constellation is classified into one of four disjoint types according to two attributes: whether it belongs to the ground-connected laser subnetwork and whether it is directly visible to the ground segment. The definitions and operational implications of the four types are summarized in Table III.

A1 and A2 jointly form the ground-connected laser subnetwork. Although A1 satellites do not have direct ground-station visibility, their telemetry return and time synchronization can




Fig. 4. Evolution of A1–A4 satellite types under the proposed ILP-based replacement sequence. A1 and A2 form the ground-connected laser subnetwork, A3 denotes satellites outside the laser subnetwork but visible to the ground segment, and A4 denotes non-anchor satellites.




Fig. 5. Evolution of A1–A4 satellite types under the historical BDS-3 launch sequence. The ground-connected laser subnetwork is not formed in the early rounds because high-rate satellite-to-ground links are introduced late.

be supported by multi-hop laser paths within the laser subnetwork. A2 satellites are inside the laser subnetwork and are also directly visible to the ground. A3 satellites are outside the laser subnetwork, but they are visible to the ground segment and can therefore communicate with the ground segment through the legacy S-band grounding link. The anchor-satellite set is defined as the union of A1, A2, and A3, covering all satellites with direct ground access or laser-assisted ground access. Satellites in A4 are defined as non-anchor satellites, since they are outside the laser subnetwork and invisible to the ground segment. A4 satellites face both data-return pressure and time-synchronization pressure, since they must rely on intermittent microwave contacts with anchor satellites to reach the ground or obtain timing information.

Fig. 4 and Fig. 5 compare the node-type evolution under the proposed ILP-based sequence and the historical BDS-3

deployment sequence. Under the proposed sequence, the size of the ground-connected laser subnetwork increases steadily from the first replacement round. This trend indicates that the proposed sequence continuously expands a laser subnetwork connected to the ground segment, rather than merely increasing the number of launched laser satellites. As the laser subnetwork expands, the number of A4 satellites decreases and eventually becomes zero. Equivalently, the fraction of anchor satellites, i.e., $A1+A2+A3$, increases throughout the replacement process. The increasing anchor-satellite fraction directly reduces the number of satellites that rely on intermittent microwave contacts for data return and time synchronization.

The A1 class represents an important structural gain introduced by the laser subnetwork. Without the laser subnetwork, these satellites would be non-anchor satellites because they are not directly visible to the ground segment. After being incorporated into the ground-connected laser subnetwork, they can reach the ground through laser paths and become indirectly ground-connected anchor satellites. This transformation accelerates telemetry return and places these satellites in a high-stability laser timing domain.

The historical BDS-3 launch order exhibits a different behavior. In the early rounds, it mainly launches MEO satellites and does not form a ground-connected laser subnetwork. This is consistent with the original BDS-3 deployment logic, whose primary objective was to establish global RNSS service coverage. However, this logic is not well matched to the replacement problem considered in this paper. In the replacement scenario, both legacy microwave satellites and newly launched laser satellites still retain their original RNSS service capability. The critical issue is therefore not to re-establish global service coverage from scratch, but to progressively construct a ground-connected laser backbone while maintaining the hybrid-network performance. Since the historical order delays the launch of GEO/IGSO satellites that can host high-rate satellite-to-ground terminals, the laser subnetwork remains disconnected from the ground in the early rounds. Consequently, the structural benefit of laser replacement is postponed.

The laser topology used in these results is obtained by the dense-topology refinement introduced in Section III-C, denoted as DTopo-ILP. DTopo-ILP preserves the primary replacement-sequence objective, namely maximizing the ground-connected laser-subnetwork size and its microwave-visibility support, and then selects a dense inter-satellite laser topology among the primary-optimal solutions. This topology construction method provides a representative realization compatible with the proposed launch sequence. The focus of this paper remains replacement-sequence planning rather than topology construction; therefore, other topology construction methods can also be applied under the same sequence if they satisfy the required connectivity and resource constraints.

2) Properties of the Ground-Connected Laser Subnetwork:

Fig. 6 and Fig. 7 show the evolution of high-rate satellite-to-ground links and laser-terminal utilization. Under the proposed sequence, the number of high-rate satellite-to-ground links increases gradually from one to three as the laser subnetwork expands. This behavior is consistent with the feeder-link




Fig. 6. Evolution of high-rate satellite-to-ground links and activated LISLs per satellite under the proposed ILP-based replacement sequence.




Fig. 7. Evolution of high-rate satellite-to-ground links and activated LISLs per satellite under the historical BDS-3 launch sequence.

threshold design in Section IV-A: a larger laser subnetwork carries more traffic and therefore requires more satellite-to-ground return capacity. The more balanced introduction of GEO/IGSO satellites enables a better match between the growth of the inter-satellite laser backbone and the growth of the satellite-to-ground return segment.

The average number of activated inter-satellite laser links per laser satellite also increases rapidly and approaches the terminal limit as the replacement process proceeds. Since each satellite carries three laser terminals, an average value close to three indicates that most available laser terminals are effectively used. This demonstrates the advantage of DTopo-ILP from the link-utilization perspective. A denser laser topology generally provides more diverse end-to-end paths, more inter-satellite measurement opportunities, and stronger robustness against temporary link unavailability.

By contrast, under the historical BDS-3 sequence, the first several rounds do not include GEO/IGSO satellites with high-rate satellite-to-ground terminals. As a result, the number




Fig. 8. Average and maximum satellite-to-ground delay of the ground-connected laser subnetwork under the proposed ILP-based replacement sequence.




Fig. 9. Average and maximum satellite-to-ground delay of the ground-connected laser subnetwork under the historical BDS-3 launch sequence.

of high-rate satellite-to-ground links remains zero, and no effective ground-connected laser topology can be constructed during those rounds. When the first GEO satellite is launched, the laser subnetwork size has already become large (17). Such delayed feeder-link deployment may lead to a bandwidth mismatch between the inter-satellite laser backbone and the satellite-to-ground return segment. Compared with this pattern, the proposed sequence introduces high-rate ground access earlier and more evenly, which is more suitable for progressive laser replacement.

Fig. 8 and Fig. 9 show the communication-delay performance of the ground-connected laser subnetwork. Under the proposed sequence, the average satellite-to-ground delay of laser-subnetwork satellites remains below 0.4 s, and the maximum delay remains below 1 s throughout the replacement process. This result reflects the advantage of the continuous laser transmission regime over the intermittent microwave DTN regime. In the microwave network, a missed contact

opportunity by one time slot immediately introduces an additional waiting delay of 3 s, which is already much larger than the propagation delay inside the laser subnetwork.

For the historical BDS-3 sequence, no laser-subnetwork delay is reported in the early rounds because no ground-connected laser subnetwork exists. After a GEO satellite is launched and the laser subnetwork becomes connected to the ground, its delay is also much lower than that of the microwave DTN regime. However, this benefit appears only after several replacement rounds. The proposed sequence is therefore advantageous not only in the delay level after the laser subnetwork is formed, but also in the earlier activation of the laser communication regime.

3) *Microwave-Layer Performance of Residual Non-Anchor Satellites (A4 Satellites)*: This subsection examines how the structural evolution of the hybrid constellation affects the residual microwave network. The microwave topology in this evaluation is constructed by DFCP [17], which is used as a representative microwave scheduling method. DFCP emphasizes fair contact opportunities, so that each satellite tends to establish links with as many distinct visible objects as possible over a superframe. Such a contact pattern is consistent with the measurement-diversity requirement of navigation satellites, for which diversified inter-satellite ranging geometries are important. Similar to DTopo-ILP, DFCP is not the unique possible microwave topology construction method. It is adopted to evaluate the network states induced by different replacement sequences under a representative and fairness-oriented microwave scheduling policy.

Among the four node classes, A4 satellites are the main performance bottleneck of the residual microwave network. A1 and A2 satellites are inside the laser subnetwork, which provides continuous high-rate data return and a high-stability timing domain. A3 satellites are outside the laser subnetwork, but they are directly visible to the ground segment and can therefore perform telemetry return and time synchronization through the legacy RF ground link. In contrast, A4 satellites are neither inside the laser subnetwork nor visible to the ground segment. They must rely on microwave contacts with anchor satellites to return data and synchronize time. Therefore, the ranging composition and waiting delay of A4 satellites are the most relevant indicators for the residual microwave-layer performance.

Fig. 10 and Fig. 11 show the ranging-link composition of A4 satellites. Under the proposed sequence, a non-negligible portion of A4 ranging links is established with other A4 satellites in the early rounds. These A4–A4 links provide measurement opportunities, but they are less useful for data return and time synchronization. Timing transfer through an “anchor–nonanchor–nonanchor” chain increases the number of microwave hops and may accumulate synchronization errors.

As the proposed replacement sequence progresses, the composition of A4 ranging links changes substantially. The A4–A4 component decreases, while contacts from A4 satellites to A1, A2, and A3 satellites become dominant. This trend has two implications. First, from the communication perspective, A4 satellites obtain more opportunities to access satellites that can eventually reach the ground. Second, from the timing per-




Fig. 10. Ranging-link composition of A4/non-anchor satellites under the proposed ILP-based replacement sequence. The components indicate the average number of ranging links from A4 satellites to A1, A2, A3, and A4 satellites.




Fig. 11. Ranging-link composition of A4/non-anchor satellites under the historical BDS-3 launch sequence.

spective, the synchronization path becomes shorter and more reliable. An A4 satellite connected to A3 can synchronize with the ground through a two-hop ground-related path, while an A4 satellite connected to A1 or A2 can access the high-stability laser timing domain in a single microwave hop. In particular, the A4-A1 ranging component directly reflects the gain introduced by the laser subnetwork: satellites that would otherwise be non-anchor satellites become anchor nodes after being connected to the ground through the laser subnetwork.

Under the historical BDS-3 sequence, the A4 ranging composition remains almost unchanged in the early rounds because no ground-connected laser subnetwork is formed. Consequently, A4 satellites continue to maintain a large frac-




Fig. 12. Average and maximum waiting delay for A4/non-anchor satellites to access anchor satellites under the proposed ILP-based replacement sequence. Only DTN waiting delay is considered, excluding propagation delay.




Fig. 13. Average and maximum waiting delay for A4/non-anchor satellites to access anchor satellites under the historical BDS-3 launch sequence. Only DTN waiting delay is considered, excluding propagation delay.

tion of A4-A4 contacts during the early replacement stage. Only after GEO/IGSO satellites are launched and a ground-connected laser subnetwork emerges does the ranging composition begin to shift toward anchor-related contacts. This result further indicates that the historical deployment logic delays the microwave-layer benefit of laser replacement.

Fig. 12 and Fig. 13 compare the waiting delay of A4 satellites before accessing anchor satellites. The reported delay is the DTN waiting delay in the microwave scheduling layer, excluding propagation delay. Under the DTN regime, the dominant delay component is the slot-level waiting time before a non-anchor satellite obtains a contact opportunity to an anchor satellite; the propagation delay is comparatively small.

Under the proposed sequence, both the average and maximum waiting delays decrease significantly as the laser subnetwork expands. The average waiting delay drops from the second-level range in the early rounds to nearly zero when A4 satellites almost disappear. The maximum waiting delay

also decreases from several tens of seconds to zero in the final rounds. This result indicates that the proposed sequence reduces not only the number of non-anchor satellites, but also the contact delay experienced by the remaining non-anchor satellites.

The historical BDS-3 sequence exhibits a delayed improvement. Since the early rounds do not form a ground-connected laser subnetwork, the anchor set does not expand through laser connectivity, and the A4 waiting delay remains almost unchanged. After GEO/IGSO satellites are launched and the laser subnetwork becomes connected to the ground, the waiting delay begins to decrease. Nevertheless, at comparable replacement rounds, the proposed ILP-based sequence generally provides lower A4 waiting delay. This confirms that the proposed sequence accelerates the transition from an intermittent microwave-dominated regime to a hybrid regime in which most satellites can directly or indirectly access the ground.

D. Discussion and Future Extensions

The above results demonstrate the effectiveness of the proposed ILP-based replacement-sequence planning framework from both structural and performance perspectives. Compared with the historical BDS-3 launch order, the proposed sequence forms a ground-connected laser subnetwork much earlier and expands it steadily throughout the replacement process. This leads to a continuous reduction in A4/non-anchor satellites, a more balanced growth of high-rate satellite-to-ground links, higher utilization of inter-satellite laser terminals, and lower access delay for the residual microwave network. These results indicate that the replacement order should not simply inherit the initial constellation deployment logic. Instead, the launch sequence should be optimized according to the evolving hybrid-network state, especially the ground connectivity of the laser subnetwork and its interaction with the remaining microwave satellites.

The proposed framework also has the potential to support more general multi-round planning. In the current implementation, the replacement decision is made in a round-wise greedy manner: at each round, the best launch action is selected according to the network gain after that action. This setting is suitable when launch actions are sufficiently separated in time. However, in practical deployment, two launch opportunities may be temporally close, and a short-term look-ahead decision may be preferable. In such cases, the action definition can be extended from a single launch action to a batch action. For example, when two consecutive MEO launches are considered jointly, the candidate action can be defined as selecting four satellites at once. The same evaluation framework can then be used to identify the best four-satellite batch according to the hybrid-network state after the batch replacement. The selected satellites can subsequently be assigned to two physical launches according to engineering constraints. This example shows that the proposed evaluation framework is not restricted to one-step greedy planning, but can be naturally extended to limited-horizon or batch-wise replacement decisions.

Several limitations should also be noted. The present study focuses on the replacement-sequence problem from the net-

work perspective. The resulting sequence therefore provides an engineering recommendation based on hybrid-network connectivity, laser-subnetwork construction, microwave-layer support, and satellite-to-ground access capability. In an actual constellation replacement program, additional factors may further constrain or modify the final replacement order. For example, certain aging satellites may have to be replaced with high priority due to degraded onboard performance, even if they are not the most beneficial choices from the network-formation perspective. Future BeiDou generations may also introduce new services, payload types, or mission requirements, which could change the relative importance of different orbital slots and satellite types. In addition, satellite and launch-vehicle production cycles, launch-site availability, mission integration schedules, and operational risk management may all affect the executable replacement plan. Therefore, the proposed method should be viewed as a network-aware decision-support framework rather than a complete engineering scheduling system. Integrating the proposed network-performance model with lifetime assessment, manufacturing constraints, launch scheduling, and service-evolution requirements constitutes an important direction for future work.

V. CONCLUSION

This paper investigated the progressive replacement-sequence planning problem for the BeiDou navigation constellation, where legacy microwave satellites are gradually replaced by laser-enabled satellites. A round-wise launch planning framework was proposed to determine the replacement sequence. Each candidate launch action is evaluated from a networking perspective by considering the formation of a ground-connected laser subnetwork and its support to residual microwave satellites. An ILP model was developed for network evaluation, and a priority-aware search heuristic was introduced to reduce the number of actions requiring full ILP evaluation.

Simulation results based on the actual BDS-3 constellation show that the proposed ILP-based sequence provides better hybrid-network evolution than the historical BDS-3 launch order. It enables earlier laser-network formation, more stable growth of anchor satellites, higher laser-link utilization, and lower microwave-layer waiting delay.

As one of the first studies to explicitly investigate the replacement-sequence planning problem for navigation constellations, this work highlights the importance of launch-order optimization in progressive constellation evolution. The proposed method provides both a network-aware decision-support tool and a direct engineering reference for future BeiDou replacement planning.

REFERENCES

- [1] J. Xie and G. Wang, "Innovation and technology characteristics of beidou-3," *Space Int.*, vol. 11, pp. 4–7, 2017.
- [2] Z. Chen and X. Wu, "General design of the third generation beidou navigation satellite system," *Journal of Nanjing University of Aeronautics & Astronautics*, vol. 52, no. 6, pp. 835–845, 2020.
- [3] J. Xie and C. Kang, "Engineering innovation and the development of the bds-3 navigation constellation," *Engineering*, 2021.

- [4] P. Exertier, E. Samain, N. Martin, C. Courde, M. Laas-Bourez, C. Fousard, and P. Guillemot, "Time transfer by laser link: data analysis and validation to the ps level," *Advances in Space Research*, vol. 54, no. 11, pp. 2371–2385, 2014.
- [5] N. Yang, X. Guo, L. Sun, S. Liu, G. Chen, and J. YANG, "Cooperative time synchronization via laser/radio inter-satellite links in navigation constellations," *SCIENTIA SINICA*, vol. 51, no. 1, p. 019516, 2021.
- [6] P. Henkel, "Precise point positioning with kepler," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*. IEEE, 2019, pp. 1–5.
- [7] L. Xiangnan, L. Yingfei, X. Chengyong, C. Ming, and L. Xiaoliang, "Study on integrated technique of laser ranging and communication and its applications in deep space," *Journal of Deep Space Exploration*, vol. 5, no. 2, pp. 147–153, 2018.
- [8] G. Giorgi, T. D. Schmidt, C. Trainotti, R. Mata-Calvo, C. Fuchs, M. M. Hoque, J. Berdermann, J. Furthner, C. Günther, T. Schuldt *et al.*, "Advanced technologies for satellite navigation and geodesy," *Advances in Space Research*, vol. 64, no. 6, pp. 1256–1273, 2019.
- [9] L. Liu, "Laser communications in space i optical link and terminal technology," *Chinese Journal of Lasers*, vol. 34, no. 1, p. 3, 2007.
- [10] S. Liu, J. Yang, X. Guo, and L. Sun, "Inter-satellite link assignment for the laser/radio hybrid network in navigation satellite systems," *GPS Solutions*, vol. 24, no. 2, p. 49, 2020.
- [11] A. Nardin, J. A. Fraire, and F. Dovis, "Contact plan design for gnss constellations: A case study with optical intersatellite links," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 3, pp. 1981–1995, 2021.
- [12] L. Zeng, X. Lu, Y. Bai, B. Liu, and G. Yang, "Topology design algorithm for optical inter-satellite links in future navigation satellite networks," *Gps Solutions*, vol. 26, no. 2, p. 57, 2022.
- [13] H. Yan, J. A. Fraire, Z. Yang, and K. Zhao, "Contact plan design for cross-linked gnss: An ilp approach for extended applications," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 62, pp. 6751–6770, 2026.
- [14] Y. Yang, Y. Mao, X. Ren, X. Jia, and B. Sun, "Demand and key technology for a leo constellation as augmentation of satellite navigation systems," *Satellite Navigation*, vol. 5, no. 1, p. 11, 2024.
- [15] H. Yan, J. A. Fraire, Z. Yang, and K. Zhao, "Joint contact planning for navigation and communication in gnss-libration point systems," *IEEE Transactions on Vehicular Technology*, vol. 75, no. 5, pp. 8565–8580, 2026.
- [16] H. S. Chang, B. W. Kim, C. G. Lee, S. L. Min, Y. Choi, H. S. Yang, D. N. Kim, and C. S. Kim, "Fsa-based link assignment and routing in low-earth orbit satellite networks," *IEEE transactions on vehicular technology*, vol. 47, no. 3, pp. 1037–1048, 1998.
- [17] Z. Yan, J. A. Fraire, K. Zhao, H. Yan, P. G. Madoery, W. Li, and H. Yang, "Distributed contact plan design for gnss," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 1, pp. 660–672, 2020.
- [18] J. A. Fraire and J. M. Finochietto, "Design challenges in contact plans for disruption-tolerant satellite networks," *IEEE Communications Magazine*, vol. 53, no. 5, pp. 163–169, 2015.
- [19] C. Hongliang, M. Yinan, G. Changjiang, G. Weiguang, Z. Tianqiao, L. Gang, S. Bo, X. Jie, L. Hongyang, M. Yue *et al.*, "Bds-3 performance assessment: Pnt, sbas, ppp, smc and sar," *Acta Geodaetica et Cartographica Sinica*, vol. 50, no. 4, p. 427, 2021.
- [20] L. Sun, J. Yang, W. Huang, L. Xu, S. Cao, and H. Shao, "Inter-satellite time synchronization and ranging link assignment for autonomous navigation satellite constellations," *Advances in Space Research*, vol. 69, no. 6, pp. 2421–2432, 2022.
- [21] China Satellite Navigation Office, "Launch list," 2026. [Online]. Available: <http://www.beidou.gov.cn/xt/fsgl/>
- [22] Satellite Encyclopedia (SatWiki), "Beidou-3," 2026. [Online]. Available: <https://sat.huijiwiki.com/wiki/BeiDou3> (in Chinese)
- [23] N. Maharjan, N. Devkota, and B. W. Kim, "Atmospheric effects on satellite-ground free space uplink and downlink optical transmissions," *Applied Sciences*, vol. 12, no. 21, p. 10944, 2022.
- [24] H. Yan, J. Fraire, Z. Yang, and K. Zhao, "Optimized contact plan design for reflector and phased array terminals in cislunar space networks," *Science China Information Sciences*, 2026.
- [25] T. Huang, Q. Li, C. Xu, M. Xu, S. Wang, G. Huang, and X. Liu, "From earth to orbit: Launch sequence optimization for LEO mega-constellations," *IEEE Transactions on Mobile Computing*, vol. 24, no. 9, pp. 8625–8641, 2025.
- [26] B. XU, K. HAN, R. DONG, W. GONG, and Q. REN, "A high-speed laser backbone node deployment approach for next-generation gnss," *Acta Aeronautica et Astronautica Sinica*, vol. 46, no. 9, 2025.

APPENDIX A

PROOF OF THE OPTIMALITY-PRESERVING PROPERTY OF THE PRIORITY-AWARE SEARCH

This appendix proves that the priority-aware search in Algorithm 1 is optimality-preserving with respect to exhaustive ILP-based evaluation over the admissible action set. The proof is given for a fixed launch round r . The actions removed by the high-rate satellite-to-ground accessibility screening are excluded from the admissible set, since they cannot satisfy the necessary ground-access condition.

A. Preliminaries

After the accessibility screening, the remaining candidate actions are partitioned into comparison classes according to action cardinality:

$$\mathcal{A}_r^{(h)} = \{a \in \mathcal{A}_r : |a| = h\}. \quad (46)$$

For all actions in the same class $\mathcal{A}_r^{(h)}$, the number of currently available laser satellites after executing the action is identical:

$$N_r^{(h)} = |\mathcal{L}_r(a)| = |\mathcal{H}_{r-1}| + h, \quad \forall a \in \mathcal{A}_r^{(h)}. \quad (47)$$

For a candidate action a , the round-level score is obtained by averaging the optimal state-level ILP objective values:

$$J_r(a) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} F_{r,k}^*(a), \quad (48)$$

where $F_{r,k}^*(a)$ is the optimal value of the state-level ILP in (37). According to (36), $F_{r,k}^*(a)$ consists of a primary term, namely the number of selected laser satellites, and a secondary term, namely the microwave-visibility support provided by the selected local laser subnetwork. The constant M is chosen sufficiently large so that the primary term is lexicographically prioritized over the secondary term.

For an action $a \in \mathcal{A}_r^{(h)}$, recall that it is *round-wise full-subnetwork feasible* if, for every laser state $k \in \mathcal{K}$,

$$\sum_{i \in \mathcal{L}_r(a)} z_{i,r,k}^*(a) = |\mathcal{L}_r(a)| = N_r^{(h)}. \quad (49)$$

That is, all currently available laser satellites are included in the ground-connected local laser subnetwork in every laser state.

B. Completeness of the High-Priority Subset

The high-priority subset $\mathcal{A}_{r,H}^{(h)}$ is constructed by the relaxed ground-reachability condition in (44). The following lemma states that this subset contains all actions that may form a round-wise full local laser subnetwork.

Lemma A1: If an action $a \in \mathcal{A}_r^{(h)}$ is round-wise full-subnetwork feasible, then

$$a \in \mathcal{A}_{r,H}^{(h)}. \quad (50)$$

Proof: If a is round-wise full-subnetwork feasible, then for every laser state $k \in \mathcal{K}$, the state-level ILP selects all satellites in $\mathcal{L}_r(a)$ into a ground-connected local laser subnetwork. Therefore, in the selected ILP solution, every satellite in $\mathcal{L}_r(a)$

is connected to the ground segment through feasible laser and high-rate satellite-to-ground links.

The high-priority screening graph $\mathcal{G}_{r,k}^L(a)$ is a relaxed visibility graph: it contains the visibility-based candidate links, while laser-terminal constraints, ground-station capacity constraints, and feeder-link threshold constraints are not imposed at this screening stage. Removing constraints cannot destroy an existing ground-connected path. Hence, every satellite in $\mathcal{L}_r(a)$ must be ground-reachable in $\mathcal{G}_{r,k}^L(a)$ for every $k \in \mathcal{K}$. Therefore, condition (44) holds, and $a \in \mathcal{A}_{r,H}^{(h)}$. $\square$

Lemma A1 implies that actions outside $\mathcal{A}_{r,H}^{(h)}$ cannot be round-wise full-subnetwork feasible.

C. Dominance of Full-Subnetwork Actions

Lemma A2: Consider two actions $a, b \in \mathcal{A}_r^{(h)}$. If a is round-wise full-subnetwork feasible and b is not, then

$$J_r(a) > J_r(b) \quad (51)$$

under the lexicographic weighting in (36).

Proof: Since a and b belong to the same comparison class, they have the same available laser-satellite count $N_r^{(h)}$. For action a , the primary term reaches its maximum in every laser state:

$$\sum_{i \in \mathcal{L}_r(a)} z_{i,r,k}^*(a) = N_r^{(h)}, \quad \forall k \in \mathcal{K}. \quad (52)$$

Since b is not round-wise full-subnetwork feasible, there exists at least one laser state $\bar{k}$ such that

$$\sum_{i \in \mathcal{L}_r(b)} z_{i,r,\bar{k}}^*(b) < N_r^{(h)}. \quad (53)$$

Thus, the average primary gain of a is strictly larger than that of b .

The secondary visibility-support term is finite because the number of satellites, laser states, microwave states, and candidate actions is finite. Since M is chosen to enforce lexicographic priority, the strictly larger primary term of a dominates any possible secondary-term advantage of b . Therefore, $J_r(a) > J_r(b)$. $\square$

Lemma A2 shows that, within a fixed comparison class, any round-wise full-subnetwork feasible action dominates any non-full-subnetwork action.

D. Ordering Among Full-Subnetwork Actions

Lemma A3: For two round-wise full-subnetwork feasible actions $a, b \in \mathcal{A}_{r,H}^{(h)}$, if

$$\Gamma_r(a) \geq \Gamma_r(b), \quad (54)$$

then

$$J_r(a) \geq J_r(b). \quad (55)$$

Proof: Since both a and b are round-wise full-subnetwork feasible, all satellites in $\mathcal{L}_r(a)$ and $\mathcal{L}_r(b)$ are selected in every laser state. Moreover, a and b belong to the same comparison class, so their available laser-satellite counts are identical and equal to $N_r^{(h)}$. Therefore, the primary term of the state-level objective is identical for the two actions in every laser state.

Under the full-subnetwork condition, setting $z_i = 1$ for all $i \in \mathcal{L}_r(a)$ makes the secondary support term equal to the full-subnetwork visibility-support quantity used in (45). Hence, the round-level secondary term of a full-subnetwork action is exactly represented by $\Gamma_r(\cdot)$. Therefore, if $\Gamma_r(a) \geq \Gamma_r(b)$, then the secondary contribution of a is no smaller than that of b , while the primary contribution is identical. This yields $J_r(a) \geq J_r(b)$. $\square$

Lemma A3 justifies the use of descending net visibility-support gain as the traversal order inside the high-priority subset.

E. Optimality-Preserving Early Termination

Theorem A1: The early-termination rule in Algorithm 1 is optimality-preserving. That is, the algorithm returns the same round-wise optimal action as exhaustive ILP-based evaluation over all admissible candidate actions.

Proof: The comparison classes are traversed in descending order of action cardinality. Consider the first comparison class $\mathcal{A}_r^{(h)}$ in which the algorithm finds a round-wise full-subnetwork feasible action. Let $\hat{a}$ denote the first such action encountered in the high-priority subset $\mathcal{A}_{r,H}^{(h)}$ under descending $\Gamma_r(\cdot)$ ordering.

First, consider the remaining unevaluated actions in $\mathcal{A}_{r,H}^{(h)}$. If such an action is not round-wise full-subnetwork feasible, then it is dominated by $\hat{a}$ according to Lemma A2. If it is round-wise full-subnetwork feasible, then it appears after $\hat{a}$ in the descending $\Gamma_r(\cdot)$ ordering. Therefore, its net visibility-support gain is no larger than that of $\hat{a}$, and Lemma A3 implies that its round-level score cannot exceed $J_r(\hat{a})$.

Second, consider actions in the same comparison class but outside the high-priority subset, i.e., actions in $\mathcal{A}_r^{(h)} \setminus \mathcal{A}_{r,H}^{(h)}$. By Lemma A1, these actions cannot be round-wise full-subnetwork feasible. Since $\hat{a}$ is round-wise full-subnetwork feasible, Lemma A2 gives

$$J_r(\hat{a}) > J_r(a), \quad \forall a \in \mathcal{A}_r^{(h)} \setminus \mathcal{A}_{r,H}^{(h)}. \quad (56)$$

Thus, the non-priority actions in the same class can be safely discarded once $\hat{a}$ has been found.

Third, consider any action in a lower-cardinality class $\mathcal{A}_r^{(h')}$ with $h' < h$. The maximum possible selected local-laser-subnetwork size of such an action is

$$N_r^{(h')} = |\mathcal{H}_{r-1}| + h' < |\mathcal{H}_{r-1}| + h = N_r^{(h)}. \quad (57)$$

Therefore, even if the lower-cardinality action forms a full local laser subnetwork, its primary objective value is strictly smaller than that of $\hat{a}$. Under the lexicographic weighting in (36), this primary-objective gap cannot be compensated by the secondary visibility-support term. Hence, no action in a lower-cardinality class can outperform $\hat{a}$.

Finally, all higher-cardinality classes have already been processed before $\mathcal{A}_r^{(h)}$. If no full-subnetwork action was found in those classes, Algorithm 1 evaluates all actions in each such class by the state-level ILP model before moving to the next class. Therefore, no action in a higher-cardinality class is skipped. Since all actions discarded after finding $\hat{a}$ are dominated by an already evaluated action, removing them

cannot change the maximizer of $J_r(a)$ over the admissible action set. The selected action

$$a_r^* = \arg \max_{a \in \mathcal{E}_r} J_r(a)$$

is therefore identical to the action that would be obtained by exhaustive ILP-based evaluation. $\square$

F. Remark

The proof shows that the priority-aware search does not approximate the state-level ILP evaluation. All evaluated actions are still scored by the exact ILP model in (37). The heuristic only reduces the number of actions requiring full ILP evaluation by exploiting three dominance relations: the action-cardinality dominance of the primary objective, the relaxed ground-reachability necessary condition for full-subnetwork feasibility, and the secondary visibility-support ordering among full-subnetwork actions.