Ballet in Sky: Online Satellite Downlink with Joint ISL Balancing and GS Selection under Uncertainty

Jianping Huang[1,*], Xiang Liu[1,2,*], Feng Shan[1,†]
[1] Southeast University, China [2] The Chinese University of Hong Kong, Hong Kong
Email: {jphuang, xiangliu, shanfeng}@seu.edu.cn

***Abstract*—Low Earth Orbit (LEO) Earth Observation (EO) satellites generate massive data but encounter significant downlink challenges due to a fundamental imbalance between uneven data distributions and fluctuating downlink opportunities. This issue is critically exacerbated by uncertain Satellite-to-GS Link (SGL) quality. To tackle the complex scheduling under uncertainty, we first propose *FlowDance*, a performance-guaranteed approximation algorithm for the oracle case with known SGL quality, by jointly optimizing ISL data balancing and GS group selection. We achieve this through a novel network flow model that captures the trade-off between balancing duration and downlink opportunity. Building on this, we design *Ballet*, a novel online learning framework. It develops a partition-batch-based exploration policy to estimate SGL qualities, which serve as the input for the exploitation process, *i.e.,* leveraging *FlowDance* to guide the downlink scheduling. We provide a theoretical analysis for *Ballet* and verify its performance through extensive simulations using real-world satellites and GS deployment data. The results show that *Ballet* significantly enhances data downlink throughput compared to state-of-the-art methods, with an average improvement of at least 29.8%.**

I. Introduction

Advancements in aerospace technology have propelled Low Earth Orbit (LEO) satellites as a promising solution for various fields, driven by their low-latency, wide coverage, and high-throughput communication capabilities [1], [2]. As reported by the open UCS Satellite Database [3], nearly half of the LEO satellites currently in orbit are dedicated to Earth Observation (EO) missions such as disaster monitoring [4] and precision agriculture [5]. They are equipped with high-resolution optical and multi-spectral sensors that generate terabytes of data daily [6], which must be downlinked to Ground Stations (GSs) [7] for timely analysis and storage.

However, the primary obstacle to efficient data downlink is a fundamental systemic *imbalance*, where data collected by satellites is severely mismatched with their available downlink opportunities. This imbalance stems from two main sources. First, *uneven data distribution* across the constellation arises from different EO tasks generating vastly different data volumes; for instance, a disaster monitoring mission might generate hundreds of gigabytes [8], while routine ocean tracking produces only tens of megabytes [9], [10]. Moreover, emergency EO missions often need satellites with optimal regional coverage, leading to inconsistent data acquisition across different satellites. Second, *there are unbalanced downlink opportunities* due to sparse GS locations and short satellite-to-GS contact windows [6], [11], [12]. This problem is critically exacerbated by the *uncertainty* of the Satellite-to-GS Links (SGLs), where SGLs face communication fluctuations due to atmospheric factors such as extreme weather. Such unpredictable fluctuations can result in significant signal attenuation and unstable downlink performance, which can only be observed upon establishing a connection [13], [14]. As a result, data-heavy satellites may struggle with poor-quality links, while lightly loaded satellites inefficiently utilize high-capacity connections, creating a severe performance bottleneck. Although deploying more global GSs is a straightforward solution, practical implementation is often constrained by the prohibitive costs of building GSs in remote, desirable locations like oceans and deserts.

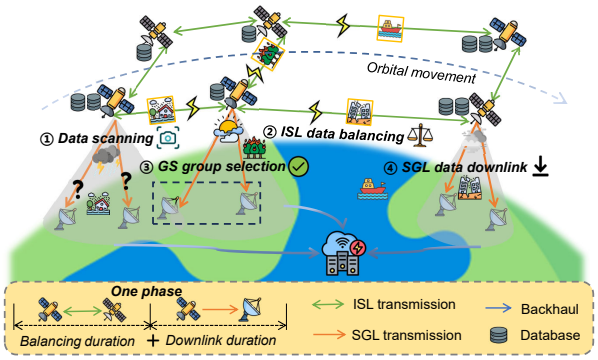

Fig. 1. Multi-satellite data downlink faces challenges due to *imbalance*, stemming from uneven data distribution and unbalanced downlink opportunities, particularly under SGL uncertainty. Our system workflow consists of data scanning, ISL data balancing, GS group selection, and SGL data downlink.

To counteract this imbalance, a promising direction is utilizing high-throughput Inter-satellite Links (ISLs) for *ISL data balancing* [11], [15], as depicted in Fig. 1. Modern constellations are increasingly equipped with laser-based ISLs capable of gigabit-per-second transfer rates [16]. Operationally, this involves scheduling a dedicated *balancing duration* for satellites to exchange data via ISLs before the *downlink duration*, *i.e.,* the remaining time allocated for data downlink to GSs. Therefore, heavily loaded satellites can offload data to their less congested peers, leading to a more balanced and efficient utilization of the subsequent downlink resources. Furthermore, by employing multi-access techniques [17], satellites can downlink data to multiple GSs simultaneously through

* Equal contribution. † Corresponding author.

a *GS group selection* process, maximizing throughput. However, effectively orchestrating ISL balancing and multi-GS downlinking in a dynamic and uncertain network environment presents significant optimization challenges:

- Jointly optimizing ISL data balancing and GS group selection is complicated by SGL uncertainty. Whether an ISL transfer is beneficial depends on the recipient satellite's unknown future downlink capacity, while allocating a high-quality SGL to an under-loaded satellite can waste scarce downlink opportunities. This decision coupling under uncertainty calls for an adaptive co-optimization strategy to maximize overall throughput.
- Designing an effective ISL data balancing strategy is non-trivial. Satellites must choose transfer amounts and targets without overwhelming the recipient's downlink capacity and must balance a timing trade-off: a shorter balancing duration may not mitigate imbalance, whereas a longer one can lower total downlink throughput.
- Selecting GS groups is NP-hard even with perfect SGL information. Each GS can serve only one satellite at a time, and per-satellite access budgets induce competition for GSs; short satellite-GS contact windows further constrain the timely identification of effective GS groups.

Existing research fails to address these challenges, often tackling them either separately or under simplifying assumptions. Most studies presuppose known and stable SGL communication quality [18]–[21], making their non-adaptive methods unsuitable for dynamic space environments. While some works acknowledge SGL dynamics [6], [12], they often rely on unrealistic assumptions of perfect communication quality forecasting. Studies on ISL data balancing either require costly additional relay satellites [22], or assume ideal SGL conditions, failing to effectively address the core problem of uneven data distribution [23]. Moreover, the literature on multi-GS downlinking [18], [24] typically overlooks the combinatorial complexities of GS group selection under practical constraints such as budget limitations and access competition.

To bridge this gap, we design *Ballet*, a novel online learning-based framework for multi-satellite data downlink scheduling. First, we present *FlowDance*, a provable approximation method for scenarios with known SGL quality. *FlowDance* employs a modified network flow to jointly optimize ISL data balancing and GS group selection. Next, *Ballet* incorporates a partition-batch exploration policy to learn SGL qualities effectively, which informs *FlowDance* to enhance downlink scheduling. This paper makes the following contributions:

- We are the first to formulate the problem of maximizing multi-satellite downlink throughput by jointly optimizing ISL data balancing and GS group selection under network uncertainty. We highlight that this holistic approach is critical to avoiding the severe performance bottlenecks found in existing research.
- We propose *FlowDance*, an approximation algorithm with a performance guarantee for the deterministic scheduling problem (*i.e.,* oracle case with known SGL quality). It effectively solves the joint optimization via a modified network flow formulation.
- We design *Ballet*, a novel online learning framework to tackle SGL uncertainty. It develops a partition-batch-based exploration policy to learn SGL qualities and uses *FlowDance* for exploitation. We provide a rigorous theoretical regret analysis for *Ballet*.
- We conduct extensive simulations using real-world satellites. The results show that *Ballet* significantly outperforms baseline methods, achieving at least a 29.8% average improvement in data downlink throughput.

The paper is organized as follows. Section II presents the system model, Sections IV and V present the algorithms and evaluation, followed by the related work and the conclusion.

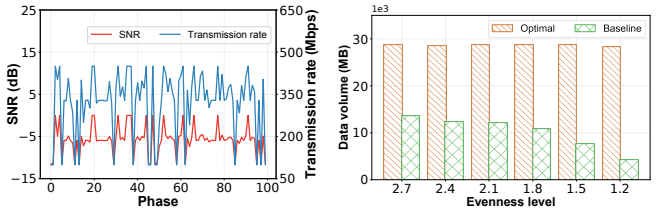

(a) Fluctuating but bounded communication quality (b) Data downlink throughput under uneven data distribution

Fig. 2. Motivating case studies for multi-satellite data downlink scheduling.

TABLE I
COMPARISON OF RELATED WORKS ON SATELLITE DATA DOWNLINKING

Existing Works	Uncertainty	Uneven Data	GS Group
Ballet (ours)	√	√	√
LBCO [23]	×	√	×
Umbra [6]	×	×	×
L2D2 [12]	√	×	×
FALCON [24]	×	×	√
AEROPATH [20]	×	×	√

II. Motivation and System Model

A. Motivating Case Studies for Satellite Downlink

Dynamic and Learnable SGL Communication. As shown in Fig. 2(a), a case study based on ITU-R P.838 [25] confirms time-varying Signal-to-Noise Ratios (SNRs) and transmission rates due to atmospheric factors like weather. Critically, we observe that these fluctuations, while unpredictable at any instant, are bounded and exhibit statistical patterns (*e.g.,* log-normal distribution [26]). This suggests that the expected SGL performance is learnable over time, motivating an online exploration-exploitation learning approach.

Inefficiency from Uneven Data Distribution. Fig. 2(b) uses a power-law workload to emulate uneven data distribution, where a small fraction of satellites carries most of the cached data. The baseline policy without ISL balancing can misallocate scarce high-rate SGL opportunities to lightly loaded satellites, while heavily loaded satellites remain bottlenecked, resulting in substantial throughput loss compared

to the brute-force optimum. This motivates the proactive balancing of cached data before downlink.

Existing studies, as shown in Table I, separately tackle SGL uncertainty, GS group selection, and uneven data, causing suboptimal results and motivating our integrated framework.

B. System Settings and Models

Multi-Satellite Data Downlink Network. As depicted in Fig. 1, this network comprises N satellites, $\mathcal{S} = \{s_1, \ldots, s_N\}$, M fixed GSs, $\mathcal{G} = \{g_1, \ldots, g_M\}$, and a central cloud data center, $\mathcal{D}$. These satellites conduct diverse Earth observation tasks, resulting in varying data volumes cached onboard. Satellites frequently experience uneven data distribution due to varying tasks and inconsistent observation opportunities. The GS network receives this satellite data and relays it to the cloud center for processing and storage.

System Workflow. This system operates over a time horizon discretized into T consecutive phases, $\mathcal{T} = \{1, 2, \ldots, T\}$, each of duration δ. Within each phase t, the workflow unfolds through four sequential steps: ① *Data Acquisition*: Each satellite performs EO tasks, acquiring a specific amount of downlink data at the start of this phase. ② *ISL Data Balancing*: To counteract data imbalance and leverage downlink opportunities efficiently, satellites exchange data via ISLs for a specified balancing duration, $\tau_t \leq \delta$. ③ *GS Group Selection*: To enhance downlink throughput, each satellite selects a group of GSs for simultaneous data transmission, while adhering to resource constraints. ④ *SGL Data Downlink*: Satellites downlink cached data (potentially updated) to selected GSs during the remaining downlink duration of $\delta - \tau_t$. Our phase-wise optimization naturally assumes that the network topology, *i.e.,* ISL and SGL connectivity, remains static within each phase, as δ is sufficiently short compared to orbital periods. Intra-phase dynamics (*e.g.,* switching overheads or congestion) can be modeled via effective capacity or duration adjustments, leaving full time-unrolled optimization for future work.

ISL Data Balancing. At phase t, via antenna transceivers on both sides [16], each satellite s_i can establish ISLs with neighbors, denoted by $\mathcal{A}_{i,t}$. For any pair $s_i, s_{i'} \in \mathcal{A}_{i,t}$, the ISL rate is modeled as $\gamma_{i,i',t} = Z_0 P_i G_i G_{i'} \rho^f_{i,i',t}$ [22], where Z_0 is a system constant, P_i is the transmission power, G_i and $G_{i'}$ are antenna gains, and $\rho^f_{i,i',t}$ represents free-space path loss. If $s_{i'} \notin \mathcal{A}_{i,t}$, then $\gamma_{i,i',t} = 0$. Let the minimum ISL rate be $\underline{\gamma} = \min_{\forall s_i, s_{i'} \in \mathcal{A}_{i,t}, t \in \mathcal{T}} \gamma_{i,i',t}$. Theoretically, satellites can form a connected graph with multi-hop paths, determined by a shortest-path algorithm. Thus, we model the balancing decision as a set of single-hop flows $\Phi_t = \{\phi_{i,i',t} | s_{i'} \in \mathcal{A}_{i,t}\}$ over direct links. These flows are constrained by the link capacity and the balancing duration τ_t:

$$\phi_{i,i',t} \leq \gamma_{i,i',t} \tau_t, \forall s_i \in \mathcal{S}, s_{i'} \in \mathcal{A}_{i,t}, t \in \mathcal{T}. \tag{1}$$

After balancing, let the updated data volume in s_i be $Q_{i,t}$:

$$Q_{i,t}(\Phi_t, \tau_t) = q_{i,t} + \sum_{s_{i'} \in \mathcal{S}} \phi_{i',i,t} - \sum_{s_{i'} \in \mathcal{S}} \phi_{i,i',t} \geq 0, \quad \forall s_i \in \mathcal{S}, t \in \mathcal{T}. \tag{2}$$

GS Group Selection. Due to orbital dynamics, at phase t, each satellite s_i has LoS visibility to a subset of GSs $\mathcal{M}_{i,t}$. Using advanced multi-access capabilities (*e.g.,* phased array antennas enabling σ_i simultaneous spot beams), s_i can simultaneously connect to σ_i GSs. Let the GS group selection decision be $\Pi_t = \{\pi_{i,t} | \forall s_i \in \mathcal{S}\}$, which is given by:

$$|\pi_{i,t}| \leq \sigma_i, \forall s_i \in \mathcal{S}, t \in \mathcal{T}. \tag{3}$$

Furthermore, to ensure signal integrity and prevent cross-satellite interference, a single-channel GS receiver typically locks onto and processes the downlink from only one satellite at a time [12], imposing a mutual exclusivity constraint:

$$\pi_{i,t} \cap \pi_{j,t} = \emptyset, \forall s_i, s_{j,j \neq i} \in \mathcal{S}, t \in \mathcal{T}. \tag{4}$$

We can address GS with multiple receivers by treating them as several co-located virtual GS nodes. Then we define $\mathcal{N}_{j,t}$ as the set of satellites visible to GS g_j at phase t.

Uncertain SGLs. SGL transmissions are susceptible to atmospheric uncertainties. Following [22], the achievable SGL rate $r_{i,j,t}$ between satellite s_i and GS g_j is a function of the SNR, *e.g.,* $SNR_{i,j,t} = P_i G_i G_j \rho^f_{i,j,t} \rho^a_{i,j,t} / N_0$. Here, $\rho^a_{i,j,t}$ is time-varying attenuation [25] and N_0 denotes noise power. $r_{i,j,t}$ is then determined by $r_{i,j,t} = B \log_2(1 + SNR_{i,j,t})$, where B is the channel bandwidth. In our online setting, we model $r_{i,j,t}$ as a random variable that is independent and identically distributed (i.i.d.) across phases, with $\mathbb{E}[r_{i,j,t}] = \mu_{i,j}$, $r_{i,j,t} \in [\underline{r}, \overline{r}]$, and $\forall s_i \in \mathcal{S}, g_j \in \mathcal{G}$. Here, $\underline{r}$ and $\overline{r}$ are the upper and lower bounds, as illustrated in Fig. 6(a).

Remark. Empirical evidence [27] highlights a significant performance gap between modern laser-based ISLs and conventional SGLs. We thus assume this system is SGL-bottlenecked, meaning the minimum ISL rate exceeds the maximum aggregate downlink rate: $\underline{\gamma} > \overline{R} = \max_{s_i,t} R_{i,t}$, where $R_{i,t} = \sum_{g_j \in \pi_{i,t}} r_{i,j,t}$. Then, ISL data transfers are limited mainly by the allocated balancing duration τ_t and data availability, not by link capacity.

C. Problem Formulation

The goal is to maximize the total expected data throughput from satellites to the ground over T phases. This requires making optimal decisions at each $t \in \mathcal{T}$ for: (1) the balancing duration, $\tau_t \in [0, \delta]$; (2) the ISL data flows $\Phi_t = \{\phi_{i,i',t} | s_i, s_{i'} \in \mathcal{S}\}$ during τ_t; and (3) the GS group selection $\Pi_t = \{\pi_{i,t} | s_i \in \mathcal{S}\}$. We define data downlinked from satellite s_i at phase t as $\chi_i(\Pi_t, \Phi_t, \tau_t)$, so its expected value is: $\mathbb{E}[\chi_i(\Pi_t, \Phi_t, \tau_t)] = \min\{Q_{i,t}(\Phi_t, \tau_t), (\delta - \tau_t) \sum_{g_j \in \pi_{i,t}} \mu_{i,j}\}$, which is limited by the available cached data and the expected SGL capacity. With these system settings and models, the overall optimization problem, P1, is formulated as follows:

$$\textbf{(P1)} \quad \max_{\{(\Pi_t, \Phi_t, \tau_t) | t \in \mathcal{T}\}} \sum_{t \in \mathcal{T}} \sum_{s_i \in \mathcal{S}} \mathbb{E}[\chi_i(\Pi_t, \Phi_t, \tau_t)]$$
$$s.t. \quad Eqs.\ (1), (2), (3), (4), \quad 0 \leq \tau_t \leq \delta, \forall t \in \mathcal{T}.$$

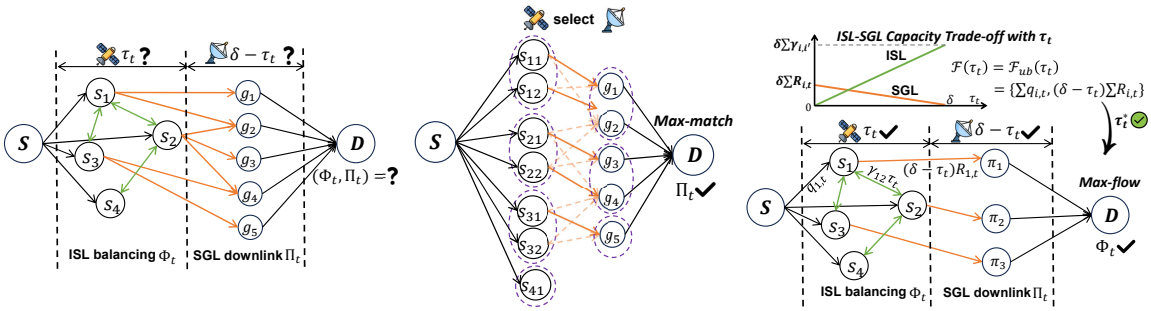

(a) Network flow graph construction (b) Strategic GS group selection (c) Adaptive ISL flow-optimization balancing

Fig. 3. Key steps of *FlowDance*. (a) Construct a τ_t-dependent flow graph for P2; (b) select GS groups via maximum matching; (c) compute ISL balancing via max-flow. The top plot shows the ISL–SGL capacity trade-off controlled by τ_t.

Algorithm 1: *FlowDance* Algorithm

Input: $\mathcal{S},\mathcal{G},\{r_{i,j,t}\},\delta,\eta,\Pi_t,\Phi_t,t\in\mathcal{T}$

// **Network Graph Flow Construction**

$\mathcal{V}=\mathcal{S}\cup\mathcal{G}\cup\{S,D\}$, $\mathcal{E}=\emptyset$;

for $s_i\in\mathcal{S},g_j\in\mathcal{M}_i$ **do**

  $\mathcal{E}=\mathcal{E}\cup\{e_{S,s_i},e_{s_i,g_j},e_{s_i,s_{i'}},e_{g_j,D}\}$, $c_{S,s_i}=q_{i,t}$, $c_{s_i,g_j}=(\delta-\tau_t)r_{i,j,t}$, $c_{s_i,s_{i'}}=\gamma_{i,i',t}\tau_t$, $c_{g_j,D}=\infty$;

// **Strategic GS Group Selection**

Split each $s_i\in\mathcal{S}$ into σ_i copies cp_i, forming $\mathcal{B}(\{cp_i\},\mathcal{G})$;

Find a maximum edge-weighted matching $\Pi_t=\{\pi_{i,t}\}$ for $\mathcal{B}$, where $\pi_{i,t}$ includes GSs matching to copies of s_i;

// **Adaptive ISL Flow-Optimized Balancing**

Merge $\pi_{i,t}$ to a super GS with $c_{s_i,\pi_{i,t}}=(\delta-\tau_t)\sum_{g_j\in\pi_{i,t}}r_{i,j,t}$;

Find τ_t via binary search over $[0,\delta]$ to solve $\Delta H(\tau_t)=0$;

Let $\mathcal{F}(\tau_t)$ be the maximum flow and Φ_t be the schedule;

Output: $\Pi_t,\Phi_t,\tau,\mathcal{F}(\tau)$

III. Algorithm for the Oracle Case

In this section, we first address the oracle case of Problem P1, referred to as Problem P2, where $\{r_{i,j,t}\}$ are known in advance for each phase. We then analyze P2's hardness and propose a provable approximation method, named *FlowDance*.

A. Hardness of the oracle problem

Formally, the definition of Problem P2 is as follows:

$$\textbf{(P2)}\quad \max_{(\Pi_t,\Phi_t,\tau_t)}\sum_{s_i\in\mathcal{S}}\chi_i(\Pi_t,\Phi_t,\tau_t)$$
$$s.t.\quad Eqs.\ (1),(2),(3),(4),\quad 0\le\tau_t\le\delta,\forall t\in\mathcal{T}.$$

Theorem 1. *Problem P2 is strongly NP-hard.*

Proof. We reduce the strongly NP-hard Partition problem to a special case of P2. Given a Partition instance with set $\{a_1,\dots,a_m\}$ where $\sum_{i=1}^m a_i = 2B$, we construct a P2 instance with: two satellites s_1,s_2, each with initial data $q_{1,t}=q_{2,t}=B$; m GSs $\mathcal{G}=\{g_1,\dots,g_m\}$, where $r_{i,j,t}=a_j\delta$ for all s_i; $\sigma_i\ge m$, $\mathcal{M}_{i,t}=\mathcal{G}$, and $\tau_t=0$. Since $\tau_t=0$, we have $Q_{i,t}=q_{i,t}=B$ and $\chi_i=\min\{B,\delta\sum_{g_j\in\pi_{i,t}}r_{i,j,t}/\delta\}=\min\{B,\sum_{g_j\in\pi_{i,t}}a_j\}$. The maximum total downlink is $2B$ if and only if there exists a partition: each satellite selects disjoint GS sets $\pi_{1,t}$ and $\pi_{2,t}$ such that $\sum_{g_j\in\pi_{1,t}}a_j=\sum_{g_j\in\pi_{2,t}}a_j=B$. This completes the reduction. □

B. FlowDance Algorithm

Solving Problem P2 is non-trivial due to the coupling of GS groups (Π_t) and ISL data flows (Φ_t), along with balancing duration (τ_t). Thus, we propose *FlowDance*, detailed in Algorithm 1, which consists of three core steps: (1) constructing a network flow graph that embeds the balancing duration trade-off; (2) determining the GS group selection via maximum weight matching; and (3) optimizing the balancing duration and ISL flows by solving a parameterized max-flow problem.

*1) **Network Flow Graph Construction**:* To address the joint optimization, we first model the system as a network graph $G(\mathcal{V},\mathcal{E},\tau_t)$, where $\mathcal{V}$ and $\mathcal{E}$ are the vertex and edge sets, as illustrated in Fig. 3(a). The core is embedding the critical trade-off governed by the balancing duration τ_t directly into the graph's edge capacities. Initially, $\mathcal{V}=\mathcal{S}\cup\mathcal{G}\cup\{S,\mathcal{D}\}$, with S and $\mathcal{D}$ as the source and the destination. $\mathcal{E}$ includes SGL edges e_{s_i,g_j} for all $g_j\in\mathcal{M}_{i,t}$, with capacity $c_{s_i,g_j}=(\delta-\tau_t)r_{i,j,t}$, and ISL edges $e_{s_i,s_{i'}}$ for all $s_{i'}\in\mathcal{A}_{i,t}$, with capacity $c_{s_i,s_{i'}}=\gamma_{i,i',t}\tau_t$. Then, S connects to each s_i via edges e_{S,s_i}, having capacity $c_{S,s_i}=q_{i,t}$, representing the initial cached data. All GSs connect to $\mathcal{D}$ with edges of infinite capacity, reflecting that backhaul links are not the bottleneck.

*2) **Strategic GS Group Selection**:* The GS group selection Π_t is a combinatorial problem constrained by satellite multi-access capabilities ($|\pi_{i,t}|\le\sigma_i$) and GS exclusivity ($\pi_{i,t}\cap\pi_{j,t}=\emptyset$). Directly applying maximum flow algorithms on the graph G cannot enforce these constraints. Thus, we first consider the subproblem of maximizing the total downlink rate, effectively setting the balance duration $\tau_t=0$. Then, we construct a bipartite graph $\mathcal{B}(\{cp_i\},\mathcal{G})$ by virtually splitting each satellite s_i into σ_i copies, as illustrated in Fig. 3(b). Each copy represents a potential SGL for satellite s_i. The weight of an edge between a satellite copy and a GS is set to the SGL rate $r_{i,j,t}$, prioritizing higher-capacity links to maximize total throughput. We solve the maximum weight bipartite matching problem on $\mathcal{B}$ using standard max-flow algorithms. The resulting selected GS groups $\Pi_t=\{\pi_{i,t}\}$ combine the copies cp_i into the original node s_i, with each $\pi_{i,t}$ containing GSs matched to copies of s_i.

*3) **Adaptive ISL Flow-Optimized Balancing**:* With Π_t, we merge all GSs $g_j\in\pi_{i,t}$ into a super node, denoted as π_i.

This creates a new graph, as shown in Fig. 3(c), featuring edge e_{s_i,π_i} whose capacity is defined as $c_{s_i,\pi_i} = (\delta - \tau_t)R_{i,t}$, where $R_{i,t} = \sum_{g_j \in \pi_{i,t}} r_{i,j,t}$. The remaining challenge is to find the optimal $\tau_t \in [0, \delta]$ that maximizes the total network flow $\mathcal{F}(\tau_t)$. As illustrated in the conceptual plot of Fig. 3(c), τ_t governs a fundamental trade-off. Increasing τ_t expands ISL capacity for data balancing but shrinks SGL capacity available for downlink. A naive search for τ_t is inefficient. Instead, we leverage a key insight: the maximum total throughput is achieved at a unique point τ_t^* where $\mathcal{F}(\tau_t^*)$ equals its theoretical upper bound, $\mathcal{F}_{ub}(\tau_t^*)$, as proven in Lemma 2. This upper bound is determined by the total cached data, $\sum_{s_i} q_{i,t}$, and the total potential downlink capacity, $\sum(\delta - \tau_t)R_{i,t}$. Formally, $\mathcal{F}_{ub}(\tau_t) = \min\{\sum_{s_i} q_{i,t}, \sum_{s_i}(\delta - \tau_t)R_{i,t}\}$. To find such τ_t^* efficiently, we use the non-increasing gap function $\Delta H(\tau_t) = \mathcal{F}_{ub}(\tau_t) - \mathcal{F}(\tau_t)$, as shown in Lemma 3, to pinpoint the τ_t^* in binary search. Once τ_t^* is found (*i.e.,* $\Delta H(\tau_t^*) = 0$), the max-flow process outputs the optimal ISL flows Φ_t and the total flow $\mathcal{F}(\tau_t^*)$ that serves as the solution for Problem P2.

C. Analysis of FlowDance Algorithm

Lemma 1. *The optimal solution for Problem P2 does not exceed* $\min\{\sum_{s_i \in \mathcal{S}} q_{i,t}, \sum_{s_i \in \mathcal{S}} \delta R_{i,t}\}$*, where* $\Pi_t = \{\pi_{i,t}\}$ *is obtained from the FlowDance Algorithm.*

Proof Sketch. The total downlinked data is limited by the sum data available across all satellites, $\sum_{s_i} q_{i,t}$, and the maximum data transmitted through the selected SGLs by maximum weighted matching in duration δ, $\sum_{s_i} \delta R_{i,t}$. The optimal throughput is capped by these two aggregate quantities. □

Lemma 2. *Assume that* τ_t^* *is the minimum value such that* $\mathcal{F}(\tau_t^*) = \mathcal{F}_{ub}(\tau_t^*)$*; then* $\mathcal{F}(\tau_t^*) \geq \max_{\tau_t' \in [0,\delta]} \mathcal{F}(\tau_t')$.

Proof. Let τ_t^* be the minimum value where $\mathcal{F}(\tau_t^*) = \mathcal{F}_{ub}(\tau_t^*)$. We show $\mathcal{F}(\tau_t^*) \geq \mathcal{F}(\tau_t')$ for any $\tau_t' \in [0, \delta]$. If $\tau_t' \geq \tau_t^*$, the claim holds because $\mathcal{F}(\tau_t') \leq \mathcal{F}_{ub}(\tau_t') \leq \mathcal{F}_{ub}(\tau_t^*) = \mathcal{F}(\tau_t^*)$, since $\mathcal{F}_{ub}(\tau)$ is non-increasing in τ. If $\tau_t' < \tau_t^*$, we prove by contradiction. Assume $\mathcal{F}(\tau_t') > \mathcal{F}(\tau_t^*)$. By the minimality of τ_t^*, we know $\mathcal{F}(\tau_t') < \mathcal{F}_{ub}(\tau_t')$. **Case (1):** $\mathcal{F}_{ub}(\tau_t^*) = \sum q_{i,t}$. Since $\tau_t' < \tau_t^*$, it follows that $\mathcal{F}_{ub}(\tau_t') = \sum q_{i,t} = \mathcal{F}(\tau_t^*)$. Then $\mathcal{F}(\tau_t') < \mathcal{F}_{ub}(\tau_t') = \mathcal{F}(\tau_t^*)$. **Case (2):** $\mathcal{F}_{ub}(\tau_t^*) = \sum(\delta - \tau_t^*)R_{i,t}$. Let $\Delta\tau = \tau_t^* - \tau_t' > 0$. We partition satellites into $\mathcal{S}_1 = \{s_i \mid Q'_{i,t}(\tau_t') \leq (\delta - \tau_t')R_{i,t}\}$ (data-limited) and $\mathcal{S}_2 = \mathcal{S} \setminus \mathcal{S}_1$ (downlink-limited). We have:

$$\begin{aligned}
\mathcal{F}(\tau_t') &= \sum_{s_i \in \mathcal{S}_1} Q'_{i,t}(\tau_t') + \sum_{s_i \in \mathcal{S}_2} (\delta - \tau_t')R_{i,t} \\
&= \sum_{s_i \in \mathcal{S}_1} Q'_{i,t}(\tau_t') + \sum_{s_i \in \mathcal{S}_2} (\delta - \tau_t^*)R_{i,t} + \sum_{s_i \in \mathcal{S}_2} \Delta\tau R_{i,t} \\
&\overset{(a)}{<} \sum_{s_i \in \mathcal{S}_1} Q'_{i,t}(\tau_t') + \sum_{s_i \in \mathcal{S}_2} (\delta - \tau_t^*)R_{i,t} + \sum_{s_i \in \mathcal{S}_2} \underline{\gamma}\Delta\tau \\
&\overset{(b)}{\leq} \sum_{s_i \in \mathcal{S}} (\delta - \tau_t^*)R_{i,t} = \mathcal{F}_{ub}(\tau_t^*) = \mathcal{F}(\tau_t^*). \qquad (5)
\end{aligned}$$

Inequality (a) holds since $R_{i,t} \leq \overline{R} < \underline{\gamma}$, which implies $\Delta\tau R_{i,t} < \underline{\gamma}\Delta\tau$ for each $s_i \in \mathcal{S}_2$. Inequality (b) must hold. Otherwise, there exists $\theta \in (0, \Delta\tau)$ such that for $\tau_t'' = \tau_t^* - \theta$,

$$\sum_{s_i \in \mathcal{S}_1} Q'_{i,t}(\tau_t') + \sum_{s_i \in \mathcal{S}_2} \underline{\gamma}\theta = \sum_{s_i \in \mathcal{S}_1} (\delta - \tau_t'')R_{i,t}, \qquad (6)$$

which implies $\mathcal{F}(\tau_t'') = \mathcal{F}_{ub}(\tau_t'')$ with $\tau_t'' < \tau_t^*$, contradicting the minimality of τ_t^*. Both cases lead to a contradiction, so we conclude that $\mathcal{F}(\tau_t^*) \geq \max_{\tau_t' \in [0,\delta]} \mathcal{F}(\tau_t')$. □

Lemma 3. $H(\tau_t)$ *is a non-increasing function.*

Proof Sketch. To prove $H(\tau_t)$ is non-increasing, we show $H(\tau_t^1) \geq H(\tau_t^2)$ for any $\tau_t^1 < \tau_t^2$ in the following cases. Let τ_t^* be minimal such that $H(\tau_t^*) = 0$. (1) $\tau_t^1 < \tau_t^2 \leq \tau_t^*$: $\mathcal{F}(\tau_t)$ is non-decreasing as longer balancing aids data distribution before downlink saturation. Since $\mathcal{F}_{ub}(\tau_t)$ is non-increasing, $H(\tau_t)$ is non-increasing. (2) $\tau_t^1 < \tau_t^* \leq \tau_t^2$: $H(\tau_t^2) = 0$ by definition of τ_t^*. As $H(\tau_t^1) \geq 0$, $H(\tau_t^1) - H(\tau_t^2) \geq 0$. (3) $\tau_t^* \leq \tau_t^1 < \tau_t^2$: $H(\tau_t^1) = H(\tau_t^2) = 0$. □

Lemma 4. *The optimal* τ_t^* *satisfies* $\tau_t^* \leq \frac{\overline{R}}{\underline{\gamma}+\overline{R}}\delta$.

Proof Sketch. If no data balancing is needed, $\tau_t^* = 0$, and the lemma holds. Otherwise, we consider a bottleneck satellite s_l that receives data. To achieve the optimal flow, the total data on s_l after balancing, $Q'_{l,t}(\tau_t^*)$, should not exceed its downlink capacity, $(\delta - \tau_t^*)R_{l,t}$. In the worst-case scenario, where only one satellite offloads data to s_l, this implies $q_{l,t} + \underline{\gamma}\tau_t^* \leq (\delta - \tau_t^*)R_{l,t}$. Solving this for τ_t^* and applying the bounds $q_{l,t} \geq 0$ and $R_{l,t} \leq \overline{R}$, we get the stated result. □

Theorem 2. *For Problem P2 with known SGL rates in phase* t*, FlowDance returns a feasible solution* (Π_t, Φ_t, τ_t) *such that* $X \geq \frac{\underline{\gamma}}{\underline{\gamma}+\overline{R}} X^*$*, where* X^* *is the optimal throughput.*

Proof. Let X be the throughput of our *FlowDance* Algorithm and X^* be that of the optimal solution. By leveraging Lemmas 1 and 2, the approximation ratio is:

$$\begin{aligned}
\frac{X}{X^*} &\geq \frac{\min\{\sum_{s_i \in \mathcal{S}} q_{i,t}, \sum_{s_i \in \mathcal{S}}(\delta - \tau_t^*)R_{i,t}\}}{\min\{\sum_{s_i \in \mathcal{S}} q_{i,t}, \sum_{s_i \in \mathcal{S}} \delta R_{i,t}\}} \\
&\overset{(a)}{\geq} \frac{\sum_{s_i \in \mathcal{S}}(\delta - \tau_t^*)R_{i,t}}{\sum_{s_i \in \mathcal{S}} \delta R_{i,t}} \overset{(b)}{\geq} 1 - \frac{\overline{R}}{\underline{\gamma}+\overline{R}} = \frac{\underline{\gamma}}{\underline{\gamma}+\overline{R}}, \qquad (7)
\end{aligned}$$

where inequality (a) holds since $\sum_{s_i \in \mathcal{S}}(\delta - \tau_t^*)R_{i,t} \leq \sum_{s_i \in \mathcal{S}} \delta R_{i,t}$ and $\min\{Q, B\}/\min\{Q, C\} \geq B/C$ for any $Q \geq 0$ and $0 \leq B \leq C$. Inequality (b) is due to Lemma 4.

Furthermore, *FlowDance* employs the network flow algorithms [28] for GS group selection, with a running time of $O(n^3)$ where $n = \max\{\sum_{s_i \in \mathcal{S}} \sigma_i, M\}$, and binary search for finding τ_t^* with a logarithmic factor. This leads to the polynomial time complexity. In practice, the effective n is determined by the active satellites/GSs in that phase rather than the total constellation size, making the per-phase runtime compatible with short operational phase windows. □

Algorithm 2: *Ballet* framework

Input: $V = ((\overline{r} - \underline{r})^2 T / P\overline{MN})^{\frac{2}{3}} \ln T^{\frac{1}{3}}$

// **Exploration process**

$\tilde{r}_{i,j} = 0, v_{i,j} = 0, \forall s_i \in \mathcal{S}, \forall g_j \in \mathcal{G}$;

for *phase* $t = 1$ *to* P, *period* $p = 1$ *to* $V\overline{MN}$ **do**

  $t' = (t + p \times P), l = (p+1)\%\overline{N}$;

  for *batch* $b = 1$ *to* $\overline{M}$ **do**

    $p = p + 1$;

    if $l \leq |\mathcal{N}_{j,t'}|$ **then**

      Let s_i be the l-th element, $v_{i,j} = v_{i,j} + 1$;

      $\tilde{r}_{i,j} = \left(\tilde{r}_{i,j} \times (v_{i,j} - 1) + r_{i,j,t'}\right) / v_{i,j}$;

// **Exploitation process**

for *phase* $t = PV\overline{MN} + 1$ *to* T **do**

  Using *FlowDance* with $\{\tilde{r}_{i,j}\}, \Pi_t = \Phi_t = \emptyset, \tau_t = 0$;

  Downlink data transmission upon (Π_t, Φ_t, τ_t);

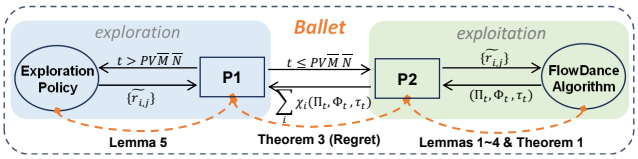

Fig. 4. Roadmap of online scheduling algorithm analysis

IV. Online Scheduling Framework Design

This section proposes an online framework, *Ballet* , for Problem P1, facing two key challenges: (1) How to effectively learn the underlying transmission rate of each SGL in a short-lived and intermittent satellite-ground contact topology with limited feedback and observation opportunities? (2) Within a finite time horizon, how to strategically manage the trade-off between exploring SGLs to refine rate estimates for long-term performance improvement and exploiting current knowledge of SGL rates to maximize immediate data throughput?

Ballet tackles these challenges via two core components: an exploration policy for systematic SGL rate estimation and the *FlowDance* Algorithm for exploiting these estimates to optimize data balancing and GS group selection. The online data downlink scheduling process of *Ballet* is outlined in Algorithm 2. Specifically, the first T_1 phases are an exploration stage, where V samples for each SGL are collected to produce transmission rate estimates $\{\tilde{r}_{i,j}\}$, serving as crucial inputs for the exploitation process. The specific duration T_1 and the sample count V are functions of the system's periodicity P and exploration parameters $\overline{M}, \overline{N}$, as detailed in Section IV-A. In each phase t of the remaining $T - T_1$ phases, *Ballet* employs *FlowDance*, which uses the learned rates $\{\tilde{r}_{i,j}\}$ to determine decisions (Π_t, Φ_t, τ_t), then guides the actual data downlink. Finally, we present the theoretical regret analysis of *Ballet* in Theorem 3, with the overall roadmap illustrated in Fig. 4.

A. Exploration Policy

Obtaining reliable estimates of SGL quality is challenging due to the intermittent and periodic satellite-to-GS contacts, the GS access budget of each satellite, and the GS constraint of serving only one satellite at a time. A generic exploration policy would fail to respect these constraints. Thus, we design a tailored exploration policy, including a partition-batch round-robin mechanism, as detailed in Lines 3-8 of Algorithm 2. This exploration policy operates on the observation that the overall satellite-to-GS contact topology reveals a shared periodicity [29], repeating every P phases, where each potential SGL gets an opportunity to be active within one period. The exploration stage spans $T_1 = PV\overline{MN}$ phases, designed to ensure that each SGL is sampled V times systematically, where $\overline{N} = \max_{g_j,t} |\mathcal{N}_{j,t}|, \overline{M} = \lceil (\max_{s_i,t} |\mathcal{M}_{i,t}|)/(\min_{s_i} \sigma_i) \rceil$.

To handle the per-GS constraint, we logically organize the SGLs into $\overline{N}$ partitions to account for each GS communicating with only one satellite. At a given phase, $\overline{N}$ is the maximum number of satellites simultaneously visible to any single GS. The l-th partition allows the selection of SGLs involving the l-th satellite (in a consistent order) within the set $\mathcal{N}_{j,t}$ of satellites visible to g_j. Moreover, to comply with the GS access budget σ_i of each satellite s_i, the visible GSs $\mathcal{M}_{i,t}$ are divided into $\overline{M}$ batches, ensuring that only one batch is explored at any time. During each phase t of T_1, the exploration policy iterates through partitions l and batches b.

Afterward, the observed rate $r_{i,j,t}$ is used to update the average estimate as $\tilde{r}_{i,j} = \left(\tilde{r}_{i,j} \times (v_{i,j} - 1) + r_{i,j,t}\right) / v_{i,j}$. Here, $v_{i,j}$ tracks the number of observations collected for the SGL from satellite s_i to GS g_j.

B. Theoretical Analysis of Ballet

Lemma 5. *Define a clear event* $\mathcal{A} := \{\forall s_i \in \mathcal{S}, g_j \in \mathcal{G}, |\tilde{r}_{i,j} - \mu_{i,j}| \leq \Delta\}$, *where* $\Delta = \sqrt{2(\overline{r} - \underline{r})^2 \ln T / V}$. *With* $T_1 = PV\overline{MN}$, *we have* $Pr(\mathcal{A}) = 1 - \frac{2NM}{T^4}$ *based on [30].*

Proof. We consider the complementary event $\mathcal{A}^c$ as follows:

$$\begin{aligned} Pr(\mathcal{A}^c) &\leq Pr(\exists s_i \in \mathcal{S}, g_j \in \mathcal{G}, |\tilde{r}_{i,j} - \mu_{i,j}| \geq \Delta) \\ &\leq NM \max_{s_i \in \mathcal{S}, g_j \in \mathcal{G}} Pr(|\tilde{r}_{i,j} - \mu_{i,j}| \geq \Delta) \\ &\leq 2NMe^{-\frac{2\Delta^2}{(\overline{r}-\underline{r})^2} V} = \frac{2NM}{T^4}. \end{aligned} \tag{8}$$

Hence, it always holds that $Pr(\mathcal{A}) = 1 - \frac{2NM}{T^4}$. □

Theorem 3. *Let* $T_1 = PV\overline{MN}$, $V = (\frac{(\overline{r}-\underline{r})^2 T}{P\overline{MN}})^{\frac{2}{3}} \ln T^{\frac{1}{3}}$, *and* $\alpha = \frac{\gamma}{\gamma + \overline{R}}$. *The total expected regret of Algorithm 2 satisfies:*

$$\widetilde{R}(T) = \mathbb{E}\Big[\sum_{t \in \mathcal{T}} (\alpha X_t^* - \widetilde{X}_t)\Big] \leq O(MT^{\frac{2}{3}} (\ln T)^{\frac{1}{3}}). \tag{9}$$

Here, $X_t^* = \sum_{s_i \in \mathcal{S}} \chi_i(\Pi_t^*, \Phi_t^*, \tau_t^*)$ *is the optimal throughput, and* $\widetilde{X}_t = \sum_{s_i \in \mathcal{S}} \chi_i(\widetilde{\Pi}_t, \widetilde{\Phi}_t, \widetilde{\tau}_t)$ *is our algorithm's throughput under the true rates* $\{\mu_{i,j}\}$.

Proof. $\widetilde{R}(T)$ is the sum of regrets from the exploration phase ($t \leq T_1$) and the exploitation phase ($t > T_1$). The exploration regret is at most $T_1 M \overline{r} \delta$. For exploitation, our algorithm computes decisions $(\widetilde{\Pi}_t, \widetilde{\Phi}_t, \widetilde{\tau}_t)$ using the estimated rates $\{\tilde{r}_{i,j}\}$.
Case 1: Under event $\mathcal{A}$. Let $X_t^{\mu}(\Pi, \Phi, \tau)$ and $X_t^{\tilde{r}}(\Pi, \Phi, \tau)$ be the throughput of a fixed decision (Π, Φ, τ) under $\{\mu_{i,j}\}$ and $\{\tilde{r}_{i,j}\}$, respectively. Under $\mathcal{A}$, for any feasible decision, we have $|X_t^{\mu}(\Pi, \Phi, \tau) - X_t^{\tilde{r}}(\Pi, \Phi, \tau)| \leq M\Delta\delta$. Let $(\widetilde{\Pi}_t^*, \widetilde{\Phi}_t^*, \widetilde{\tau}_t^*)$

be the optimal under $\{\tilde{r}_{i,j}\}$. By Theorem 2, $X_t^{\tilde{r}}(\widetilde{\Pi}_t, \widetilde{\Phi}_t, \widetilde{\tau}_t) \geq \alpha X_t^{\tilde{r}}(\widetilde{\Pi}_t^*, \widetilde{\Phi}_t^*, \widetilde{\tau}_t^*) \geq \alpha X_t^{\tilde{r}}(\Pi_t^*, \Phi_t^*, \tau_t^*)$. Therefore,

$$\begin{aligned}\widetilde{X}_t &= X_t^{\mu}(\widetilde{\Pi}_t, \widetilde{\Phi}_t, \widetilde{\tau}_t) \\ &\geq X_t^{\tilde{r}}(\widetilde{\Pi}_t, \widetilde{\Phi}_t, \widetilde{\tau}_t) - M\Delta\delta \\ &\geq \alpha X_t^{\tilde{r}}(\Pi_t^*, \Phi_t^*, \tau_t^*) - M\Delta\delta \\ &\geq \alpha X_t^{\mu}(\Pi_t^*, \Phi_t^*, \tau_t^*) - (\alpha+1)M\Delta\delta \\ &= \alpha X_t^* - (\alpha+1)M\Delta\delta. \end{aligned} \tag{10}$$

Since $\alpha \leq 1$, it follows that $\alpha X_t^* - \widetilde{X}_t \leq 2M\Delta\delta$.

Case 2: Under event $\mathcal{A}^c$. We employ a trivial bound $\widetilde{R}(T) \leq TM\overline{r}\delta$ due to its small probability by Eq. (8).

Summing the regrets over the horizon yields:

$$\begin{aligned}\widetilde{R}(T) &\leq T_1 M\overline{r}\delta + 2(T-T_1)M\Delta\delta + TM\overline{r}\delta\frac{2NM}{T^4} \\ &\leq T_1 M\overline{r}\delta + 2TM\Delta\delta + \frac{2NM^2\overline{r}\delta}{T^3} \\ &= O(MT^{\frac{2}{3}}(\ln T)^{\frac{1}{3}}). \end{aligned} \tag{11}$$

□

V. Simulation

In this section, we demonstrate the superiority of our proposed algorithm compared with other alternative algorithms.

A. Simulation setup

Simulation parameters. Our dataset includes 12 SkySat EO satellites [31], which provide high-frequency surveillance through detailed geospatial data, and 30 GSs from the real-world SatNOGS database [7]. We calculate satellite-to-GS contacts over time using TLEs [32]. Following [6], this simulation operates in 1-minute phases within a shared period P of 600 phases. Each satellite has a GS access budget of up to 6, a bandwidth of 1 GHz, and an expected transmission rate of 200 Mbps for each SGL. We employ a log-normal rainfall intensity model [26] to simulate random shadowing effects, resulting in transmission rates ranging from 100 Mbps to 450 Mbps. Laser-based ISLs operate at 2 Gbps. To account for uneven data distribution across satellites, data sizes follow a uniform distribution between 50 MB and 10 GB.

Baseline algorithms. To show the effectiveness of *Ballet*, we introduce the following algorithms for comparison. (1) *M-Umbra* [6]: Satellites select GSs based on the last-phase transmission rate using a weighted-edge matching algorithm, but they do not balance inter-satellite data. (2) *ISLCO* [23]: Satellites employ the ISLCO strategy for data balancing, and each satellite independently selects GSs using the Upper Confidence Bound (UCB). (3) *Hurry* [15]: Each satellite chooses a single GS and constructs a min-cost max-flow graph that includes all ISLs and SGLs to schedule downlink data. (4) *GreedyISL*: Satellites select the least loaded satellite for data balancing and match GSs like the *M-Umbra* strategy. (5) *RAN*: Data-overloaded satellites randomly offload data to under-loaded satellites and choose available GSs. (6) *BFS*: Satellites use brute-force search for GS selection, serving as a benchmark for the oracle problem.

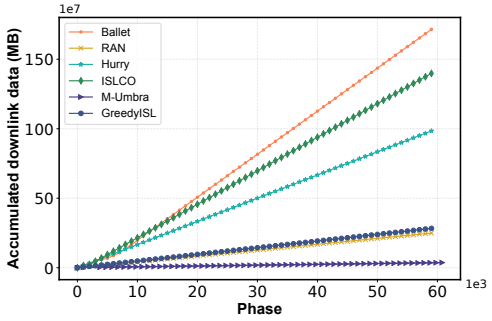

Fig. 5. Accumulated downlink data

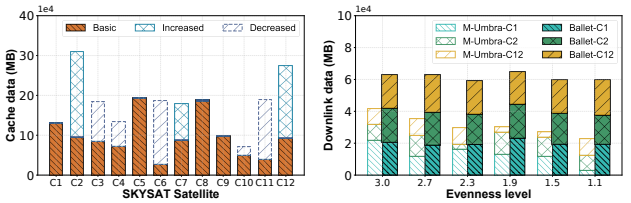

(a) Before-and-after cached data (b) Performance with uneven data

Fig. 6. Case on ISL data balancing of FlowDance Policy.

B. Simulation results

Algorithm performance over time. Fig. 5 shows the overall performance of different algorithms over time. Our *Ballet* achieves the highest accumulated downlink data, outperforming *M-Umbra* by up to 39× and other baselines by at least 29.8%. This lead, particularly pronounced after 20,000 phases, stems from increasingly accurate rate estimates via its bandit framework. In contrast, *M-Umbra* performs the worst by neglecting ISL data balancing. This oversight causes some satellites to accumulate substantial cached data without matching GSs to downlink. ISLCO *combines* UCB to learn rates by exploring less frequent and inefficient SGLs, but it suffers from GS selection conflicts, hindering data balancing and reducing throughput. *Hurry* also shows suboptimal performance as its min-cost max-flow scheduling neither accounts for SGL rate fluctuations nor considers ISL balancing. Both *GreedyISL* and *RAN* share similar limitations with *Hurry*, which lacks learning and employs random or naive data balancing strategies, resulting in lower downlink data. Nevertheless, they still outperform *M-Umbra* by using some form of ISL balancing.

Impact of Adaptive Data Balancing. To further verify the effectiveness of inter-satellite data balancing in our proposed *FlowDance*, we present a case study on the balancing process in Figs. 6(a) and 6(b). Fig. 6(a) illustrates the comparison of cached data before and after the flow scheduling process, where ISL data balancing is determined by the optimal maximum flow process. SkySat C2, C7, and C12, having under-loaded data and superior downlink quality, accept data offloaded from SkySat C3, C4, C6, C10, and C11, experiencing data overload or lower downlink quality. The increased data resulting from this ISL data balancing is highlighted. In

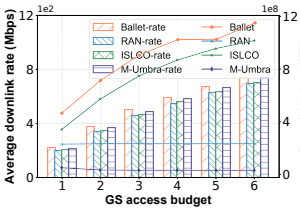

Fig. 7. Downlink rate and data under varying GS access budget.

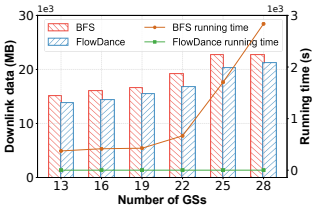

Fig. 8. Efficiency and Effectiveness of *FlowDance* Policy

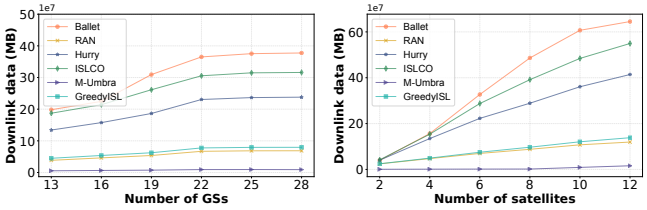

(a) (b)

Fig. 9. Scalability on number of GSs/satellites.

Fig. 6(b), Unlike *M-Umbra*, *FlowDance* effectively mitigates bottlenecks and stabilizes cached data on SkySat C1, C2, and C12, regardless of varying data imbalance. Quantitatively, in this scenario, *FlowDance* enhances total throughput by $1.98\times$ compared to *M-Umbra*.

Impact of strategic GS selection. Fig. 7 evaluates the sensitivity of *FlowDance* to the GS access budgets, comparing it with other GS group selection strategies of *ISLCO*, *RAN*, and *M-Umbra*, respectively. Overall, the aggregated rate $R_{i,t}$ at each phase t shows that both *FlowDance* and *M-Umbra* utilize a weighted-edge matching algorithm, achieving comparable and better performance than *ISLCO* and *RAN*. This is because *ISLCO* and *RAN* employ greedy and random strategies for GS group identification, whereas *FlowDance* and *M-Umbra* optimize the matching between satellites and GSs based on downlink rates or cached data, alleviating competition for SGLs. With the support of our adaptive data balancing mechanism, *Ballet* incorporating *FlowDance* significantly exceeds *M-Umbra* in total average downlink data by $800TB$.

Effectiveness and Efficiency of *FlowDance*. We compare *FlowDance* with *BFS*, which solves the GS selection in the oracle problem through a brute-force search of all possible GS groups, as illustrated in Fig. 8. Simulation results indicate that as the number of GSs increases, *FlowDance* achieves 91.1% of the optimal *BFS* performance. This significantly exceeds the approximation ratio bound of $\underline{\gamma}/(\underline{\gamma}+\overline{R})$ specified in Theorem 2, calculated as 73.5% in this simulation setting. Besides, we analyze the time cost of *BFS* and *FlowDance*. As the number of GSs rises, the computational time of *BFS* escalates exponentially, while *FlowDance*, which operates on a polynomial-time algorithm with respect to the number of GSs, grows much more slowly, highlighting its efficiency. The effectiveness and efficiency of *FlowDance* provide a solid foundation for the learning-based *Ballet*.

Scalability of GS/satellite Numbers. Finally, we investigate the scalability of *Ballet* in comparison to other algorithms as the number of GSs and satellites increases. As illustrated in Fig. 9(a), *Ballet* consistently outperforms all other algorithms across varying GS counts, with an increasing performance gap as the number of GSs rises, revealing its ability to leverage additional GSs to enhance system throughput. Fig. 9(b) further shows that *Ballet* exhibits the fastest growth trend when the number of satellites exceeds 4, indicating its robustness in resolving conflicts associated with GS selection for larger satellite numbers. In short, *Ballet* maintains stable performance across different scales of GSs and satellites, achieving an average improvement of at least 12.8% in accumulated downlink data compared to these baseline algorithms.

VI. Related Work

Satellite Data Downlink Scheduling. The increase in EO services has led to significant research on optimizing satellite data downlink scheduling to manage the growing data volume, including heuristic multipath routing strategies [24], collaborative ground station selection and scheduling strategies [20], and joint optimization approaches for observation routing and computation node selection [21]. However, these studies often assume ideal communication conditions, while in practice, research indicates that SGLs often encounter weather-related attenuation, such as atmospheric interference, rain, and cloud/fog, which can diminish downlink signals by up to 20 dB, negatively affecting transmission quality [12], [26].

Inter-Satellite Collaboration and Load Balancing. The emergence of ISLs has transformed satellite networks from independent nodes into collaborative constellations, with laser-based ISLs achieving throughputs of up to hundreds of Gbps. Recent studies have examined ISL topology design and routing mechanisms [19], [24]. While these works focus mainly on optimizing routing efficiency, they overlook uneven data distribution across satellites. A few works [22], [23] attempt to balance data distribution but involve additional relay satellites for downlinking, which increases deployment and management costs [22] or assume known and stable SGLs [23].

Satellite-to-Multiple-GS Transmissions. Advanced techniques such as Orthogonal Frequency Division Multiple Access (OFDMA) and Multiple-Input Multiple-Output (MIMO) have been explored to enhance satellite downlink capacity and mitigate user quality of experience issues [6]. However, existing studies mainly focus on point-to-point setups in satellite networks [12], [15]. Research on satellite-to-multiple-GS transmissions [22], [24] often assumes uniform channel conditions and overlooks the combinatorial complexities of optimal GS group selection under resource constraints, including budget limitations and GS competition. Learning-based approaches such as multi-armed bandits have also been used in satellite networks [18], [33], but they do not jointly capture ISL data balancing, GS group selection, and uncertain SGLs in multi-satellite downlink scheduling.

VII. Conclusion

This paper investigates multi-satellite data downlink scheduling under uncertain satellite-to-GS link (SGL) quality. We develop *FlowDance*, a provably efficient approximation algorithm for the deterministic oracle setting with known SGL quality, and *Ballet*, an online learning framework that estimates SGL qualities via a partition-batch exploration policy and exploits *FlowDance* for scheduling. We provide a regret bound for *Ballet* and an approximation ratio for *FlowDance*. Extensive simulations using real-world satellite and GS deployment data show that *Ballet* improves average downlink throughput by 29.8% over baselines. Future work includes incorporating more accurate SGL prediction models, handling dynamic ISL topologies, and extending the framework to heterogeneous satellite capabilities and mission priorities.

VIII. Acknowledge

This work is supported by the National Natural Science Foundation of China Grant 62232004, 62472090, 62132009, 62402102, the National Natural Science Foundation of Jiangsu Grants BK20242026, BK20241275, the Jiangsu Provincial Key Laboratory of Network and Information Security Grant BM2003201, the Key Laboratory of Computer Network and Information Integration of the Ministry of Education of China under Grant 93K-9, the Collaborative Innovation Center of Novel Software Technology.

References

[1] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff *et al.*, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2020.

[2] H. Pan, L. Qiu, B. Ouyang, S. Zheng, Y. Zhang, Y.-C. Chen, and G. Xue, "Pmsat: Optimizing passive metasurface for low earth orbit satellite communication," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023, pp. 1–15.

[3] "UCS satellite database," https://www.ucsusa. org/resources/satellite-database, accessed Jul. 18, 2024.

[4] Z. Wang, X. Wang, W. Wu, and G. Li, "Continuous change detection of flood extents with multisource heterogeneous satellite image time series," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–18, 2023.

[5] A. Brook, V. De Micco, G. Battipaglia, A. Erbaggio, G. Ludeno, I. Catapano, and A. Bonfante, "A smart multiple spatial and temporal resolution system to support precision agriculture from satellite images: Proof of concept on Aglianico vineyard," *Remote Sensing of Environment*, vol. 240, p. 111679, 2020.

[6] B. Tao, M. Masood, I. Gupta, and D. Vasisht, "Transmitting, fast and slow: Scheduling satellite traffic through space and time," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023, pp. 1–15.

[7] "SatNOGS network," https://network.satnogs.org/, accessed Jun. 3, 2025.

[8] B. Tao, O. Chabra, I. Janveja, I. Gupta, and D. Vasisht, "Known knowns and unknowns: Near-realtime earth observation via query bifurcation in serval," in *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, 2024, pp. 809–824.

[9] I. A. Csiszar, J. T. Morisette, and L. Giglio, "Validation of active fire detection from moderate-resolution satellite sensors: The modis example in northern eurasia," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 7, pp. 1757–1764, 2006.

[10] K.-F. Dagestad and J. Röhrs, "Prediction of ocean surface trajectories using satellite derived vs. modeled ocean currents," *Remote sensing of environment*, vol. 223, pp. 130–142, 2019.

[11] Z. M. Bakhsh, Y. Omid, G. Chen, F. Kayhan, Y. Ma, and R. Tafazolli, "Multi-satellite MIMO systems for direct satellite-to-device communications: A survey," *IEEE Communications Surveys & Tutorials*, 2024.

[12] D. Vasisht, J. Shenoy, and R. Chandra, "L2D2: Low latency distributed downlink for LEO satellites," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 151–164.

[13] A. D. Panagopoulos, P.-D. M. Arapoglou, and P. G. Cottis, "Satellite communications at Ku, Ka, and V bands: Propagation impairments and mitigation techniques," *IEEE communications surveys & tutorials*, vol. 6, no. 3, pp. 2–14, 2004.

[14] C. Kourogiorgas and A. D. Panagopoulos, "A rain-attenuation stochastic dynamic model for LEO satellite systems above 10 GHz," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 2, pp. 829–834, 2014.

[15] H. Luo, W. Liu, Q. Zhang, Z. Yang, Q. Lin, W. Zhu, K. Qiu, Z. Chen, and Y. Gao, "Hurry: Dynamic collaborative framework for low-orbit mega-constellation data downloading," in *European Conference on Parallel Processing*. Springer, 2024, pp. 269–282.

[16] "Starlink," https://www.starlink.com/technology, accessed Jul. 18, 2024.

[17] B.-H. Yeh, J.-M. Wu, and R. Y. Chang, "Efficient doppler compensation for LEO satellite downlink OFDMA systems," *IEEE Transactions on Vehicular Technology*, 2024.

[18] Z. Xu, G. Xu, H. Wang, W. Liang, Q. Xia, and S. Wang, "Enabling streaming analytics in satellite edge computing via timely evaluation of big data queries," *IEEE Transactions on Parallel and Distributed Systems*, vol. 35, no. 1, pp. 105–122, 2023.

[19] Z. Lai, H. Li, Y. Wang, Q. Wu, Y. Deng, J. Liu, Y. Li, and J. Wu, "Achieving resilient and performance-guaranteed routing in space-terrestrial integrated networks," in *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 2023, pp. 1–10.

[20] W. Liu, Q. Wu, Z. Lai, H. Li, Y. Li, and J. Liu, "Enabling ubiquitous and efficient data delivery by leo satellites and ground station networks," in *GLOBECOM 2022-2022 IEEE Global Communications Conference*. IEEE, 2022, pp. 687–692.

[21] Z. Zhai, L. Zeng, T. Ouyang, S. Yu, Q. Huang, and X. Chen, "SECO: Multi-satellite edge computing enabled wide-area and real-time earth observation missions," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*. IEEE, 2024, pp. 2548–2557.

[22] P. Wang, H. Li, B. Chen, and S. Zhang, "Enhancing earth observation throughput using inter-satellite communication," *IEEE Transactions on Wireless Communications*, vol. 21, no. 10, pp. 7990–8006, 2022.

[23] P. He, J. Hu, X. Fan, D. Wu, R. Wang, and Y. Cui, "Load-balanced collaborative offloading for LEO satellite networks," *IEEE Internet of Things Journal*, vol. 10, no. 21, pp. 19 075–19 086, 2023.

[24] M. Lyu, Q. Wu, Z. Lai, H. Li, Y. Li, and J. Liu, "Falcon: Towards fast and scalable data delivery for emerging earth observation constellations," in *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 2023, pp. 1–10.

[25] "ITU-R P.838," https://www.itu.int/rec/R-REC-P.838, accessed Jun. 3, 2025.

[26] S. Shrestha and D.-Y. Choi, "Characterization of rain specific attenuation and frequency scaling method for satellite communication in South Korea," *International Journal of Antennas and Propagation*, vol. 2017, no. 1, p. 8694748, 2017.

[27] I. Del Portillo, B. G. Cameron, and E. F. Crawley, "A technical comparison of three low earth orbit satellite constellation systems to provide global broadband," *Acta astronautica*, vol. 159, pp. 123–135, 2019.

[28] B. V. Cherkassky and A. V. Goldberg, "On implementing the push—relabel method for the maximum flow problem," *Algorithmica*, vol. 19, pp. 390–410, 1997.

[29] M. Lv, X. Peng, W. Xie, and N. Guan, "Task allocation for real-time earth observation service with leo satellites," in *2022 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2022, pp. 14–26.

[30] "Probability bounds," https://stanford.edu/~jduchi/projects/probability_bounds.pdf, accessed Jun. 3, 2025.

[31] "SkySat," https://www.eoportal.org/satellite-missions/skysat, accessed Jul. 18, 2024.

[32] "Celestrak," https://celestrak.com/, accessed Jun. 3, 2025.

[33] X. Gao, J. Wang, X. Huang, Q. Leng, Z. Shao, and Y. Yang, "Energy-constrained online scheduling for satellite-terrestrial integrated networks," *IEEE Transactions on Mobile Computing*, vol. 22, no. 4, pp. 2163–2176, 2021.