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

Jianping Huang, Xiang Liu, Feng Shan

Southeast University

Email: shanfeng@seu.edu.cn

SECON 2026

Presenter: Jianping Huang

Outline

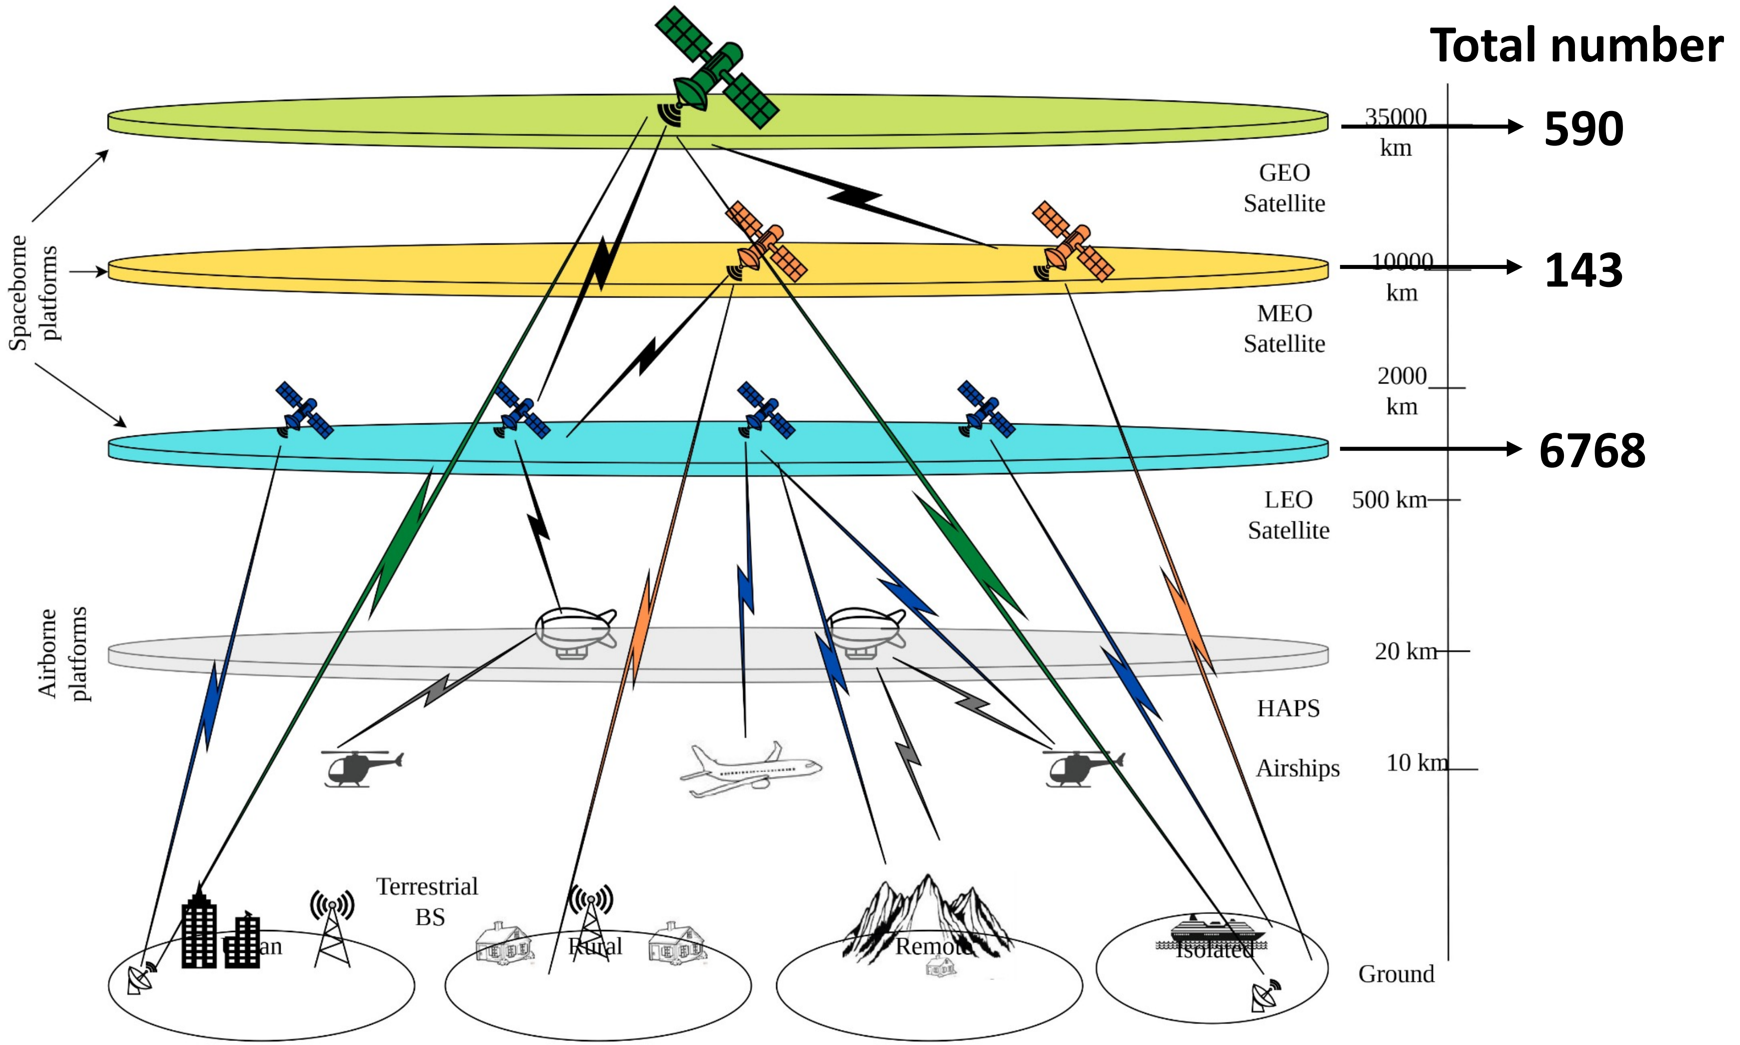

LEO Earth Observation (LEO) Satellites

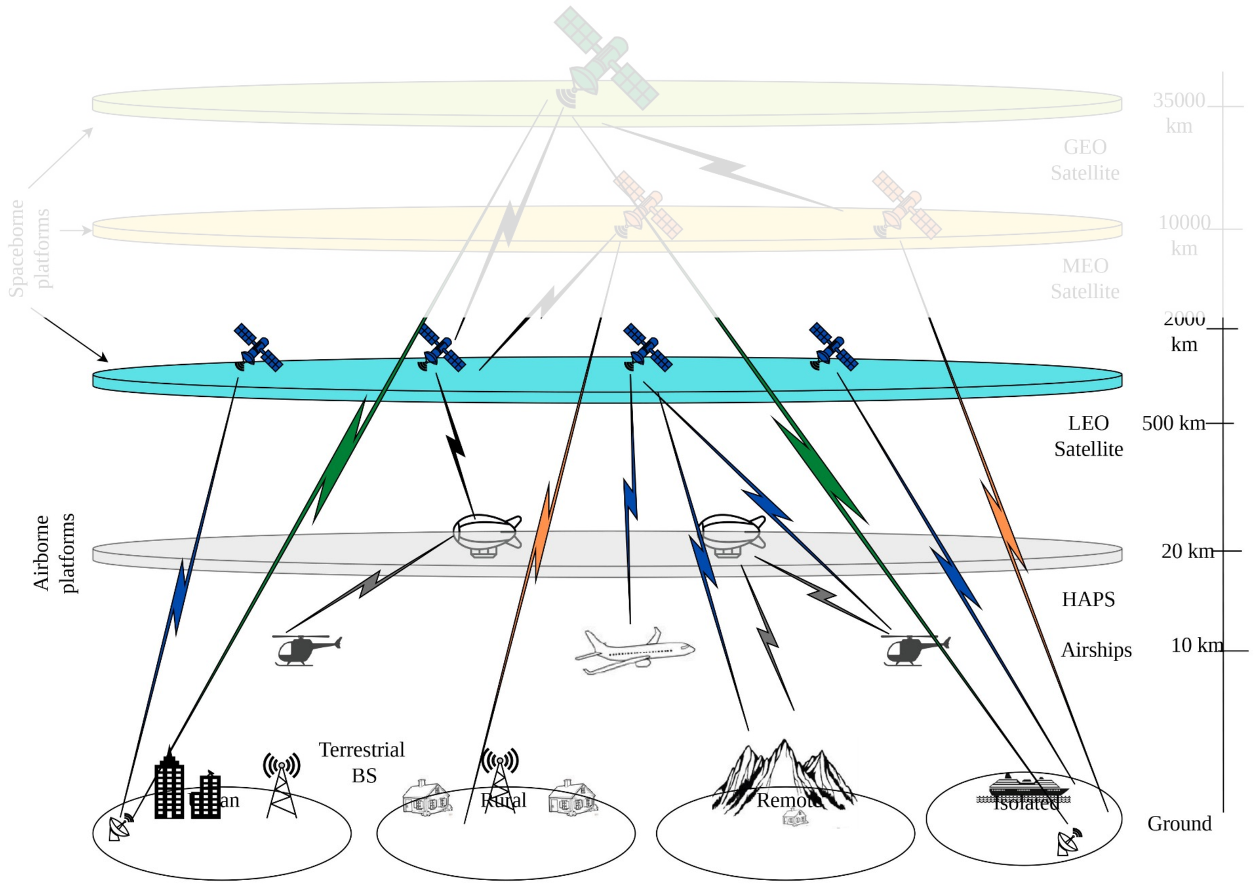

Characteristics of transfer data

- **Imaging data volume**: >1TB per day
- **Satellite orbits**: predictable location
- **Link**: contact is short-lived (≈10 min) and bandwidth is varying
- **Data download process**: satellite → ground station → cloud
- **Goal**: maximize the end-to-end throughput

✓ ***Uneven Data Distribution***

- Different missions generate vastly different data volumes

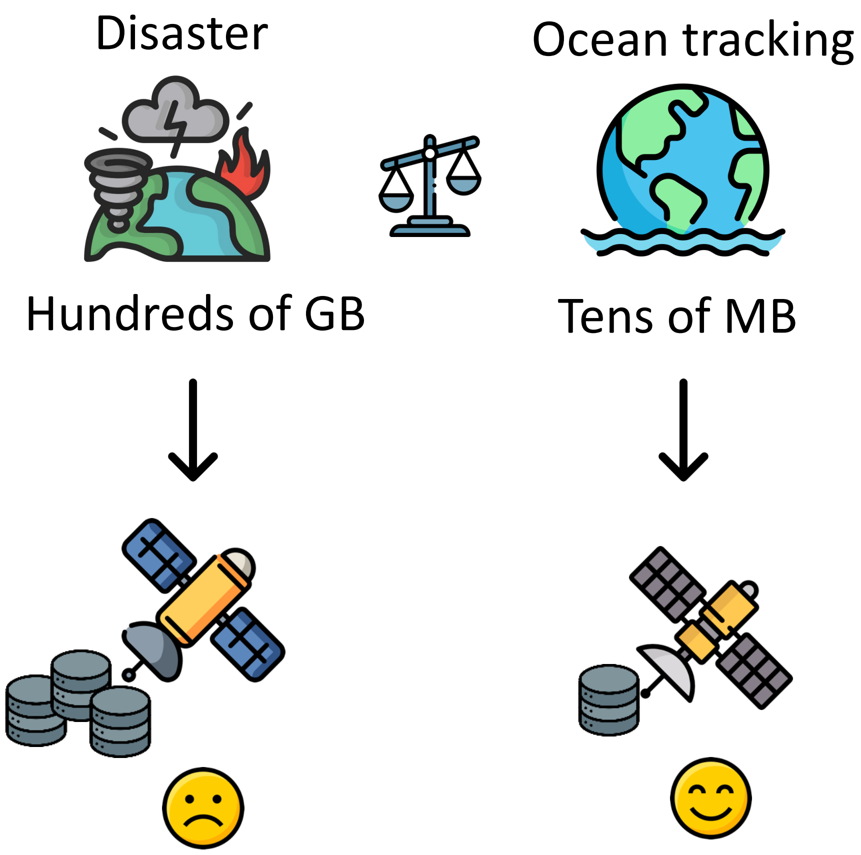

✓ ***Uneven Data Distribution***

✓ ***Unbalanced Downlink Opportunities***

- GS locations are sparse and often far from data-heavy regions
- Satellite-to-GS contact windows are very short

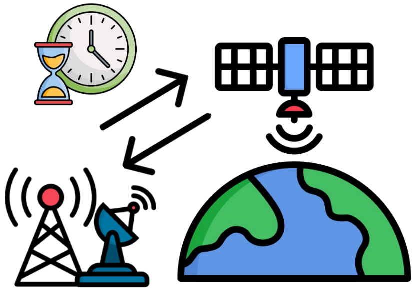

- ***Uneven Data Distribution***
- ***Unbalanced Downlink Opportunities***
- ***High-speed Inter-Satellite Links (ISLs)***

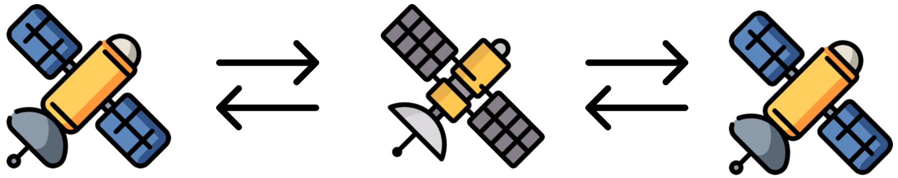

High-speed ISL →20Gbps

- ✓ ***Uneven Data Distribution***
- ✓ ***Unbalanced Downlink Opportunities***
- ✓ ***High-speed Inter-Satellite Links (ISLs)***

Imbalance across satellites need data balancing via ISL

Uncertain SGL Quality

- Satellite-to-GS Links (SGLs) are highly uncertain
- Atmospheric factors like rain, clouds, and fog cause significant signal attenuation
- The achievable transmission rate fluctuates unpredictably

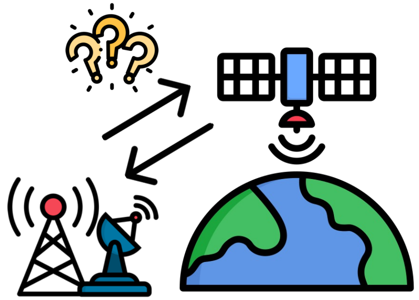

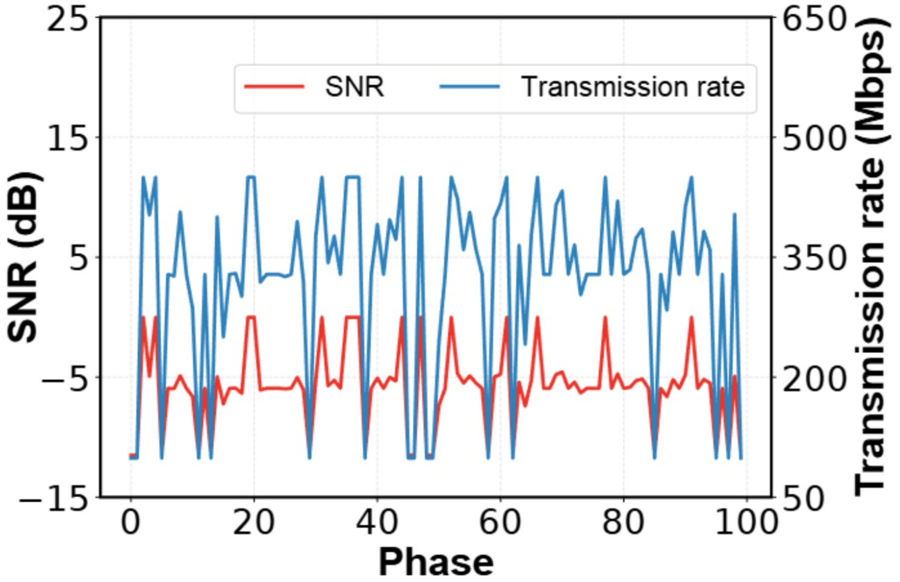

Uncertain SGL Quality

- Satellite-to-GS Links (SGLs) are highly uncertain
- Atmospheric factors like rain, clouds, and fog cause significant signal attenuation
- The achievable transmission rate fluctuates unpredictably

Fluctuations are bounded and follow statistical patterns

Outline

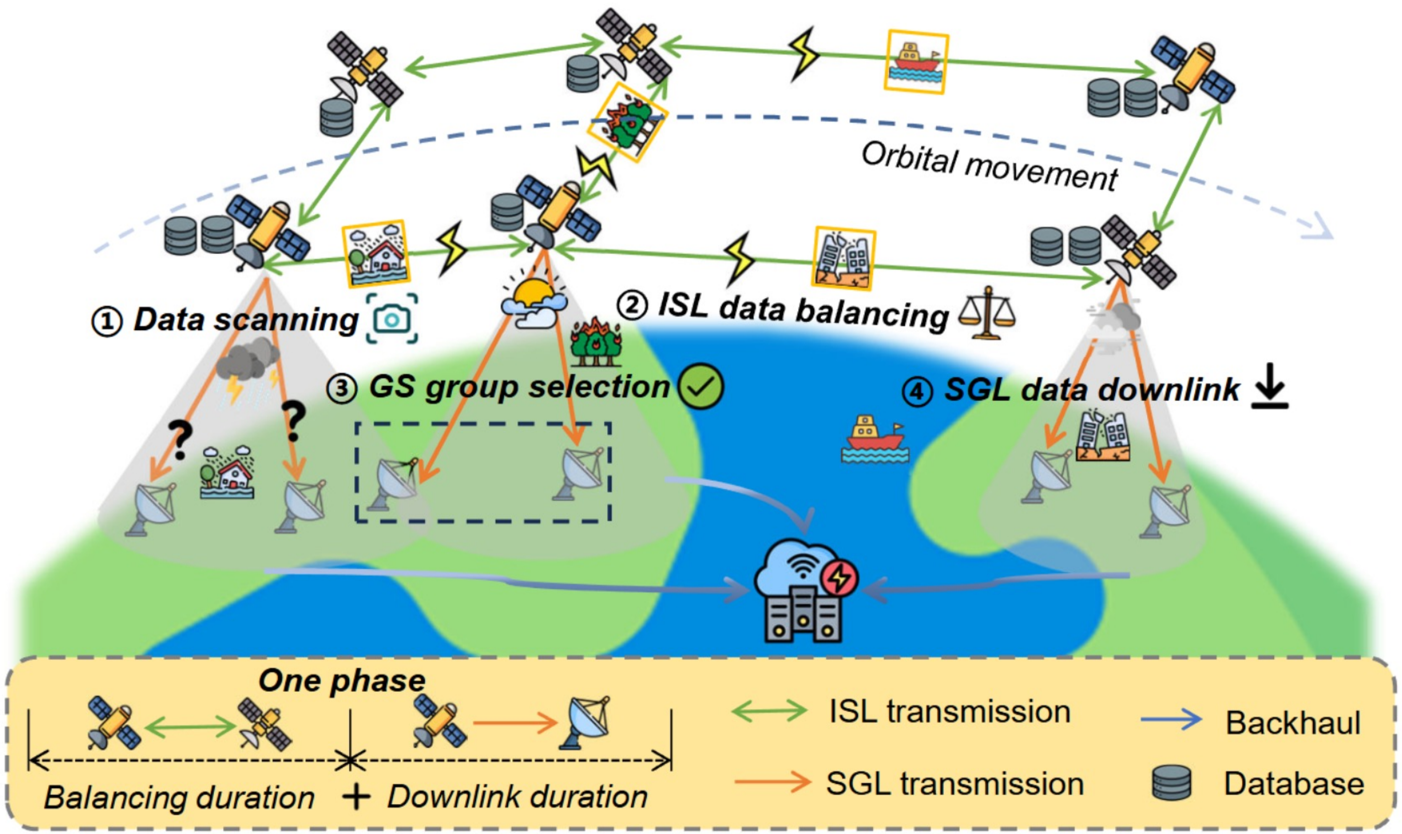

Multi-Satellite Data Downlink Network

Use high-throughput ISLs to balance data before downlink, and use multi-access techniques to select a group of GSs for each satellite.

Problem Objective

Maximizing multi-satellite downlink throughput by jointly optimizing ISL data balancing and GS group selection under network uncertainty

Constraints

- ISL capacity limits
- GS exclusivity
- Satellite access budget Data causality

Value Prop

✓ Resolve the Fundamental Imbalance in Satellite Data Downlink

✓ Enable Adaptive Decision-Making Under Real-World Uncertainty

✓ Unlock the Full Potential of Inter-Satellite Collaboration

Problem Objective

Maximizing multi-satellite downlink throughput by jointly optimizing ISL data balancing and GS group selection under network uncertainty

Constraints

- ISL capacity limits
- GS exclusivity
- Satellite access budget Data causality

Value Prop

✓ Resolve the Fundamental Imbalance in Satellite Data Downlink
✓ Enable Adaptive Decision-Making Under Real-World Uncertainty
✓ Unlock the Full Potential of Inter-Satellite Collaboration

Challenges

Joint Optimization under Uncertainty

- ISL balancing decisions depend on unknown future SGL qualities of receiving satellites.

Problem Objective

Maximizing multi-satellite downlink throughput by jointly optimizing ISL data balancing and GS group selection under network uncertainty

Constraints

- ISL capacity limits
- GS exclusivity
- Satellite access budget Data causality

Challenges

- **Joint Optimization under Uncertainty**
- **Effective ISL Balancing Strategy**
  - Increasing balancing duration τ improves data distribution but reduces time for downlink.

Value Prop

✓ Resolve the Fundamental Imbalance in Satellite Data Downlink
✓ Enable Adaptive Decision-Making Under Real-World Uncertainty
✓ Unlock the Full Potential of Inter-Satellite Collaboration

Problem Objective

Maximizing multi-satellite downlink throughput by jointly optimizing ISL data balancing and GS group selection under network uncertainty

Constraints

- ISL capacity limits
- GS exclusivity
- Satellite access budget Data causality

Challenges

- **Joint Optimization under Uncertainty**
- **Effective ISL Balancing Strategy**
- **Complex GS Group Selection**
  - Each GS can serve only one satellite, and each satellite has a limited access budget.

Value Prop

✓ Resolve the Fundamental Imbalance in Satellite Data Downlink
✓ Enable Adaptive Decision-Making Under Real-World Uncertainty
✓ Unlock the Full Potential of Inter-Satellite Collaboration

Workflow & Formulation

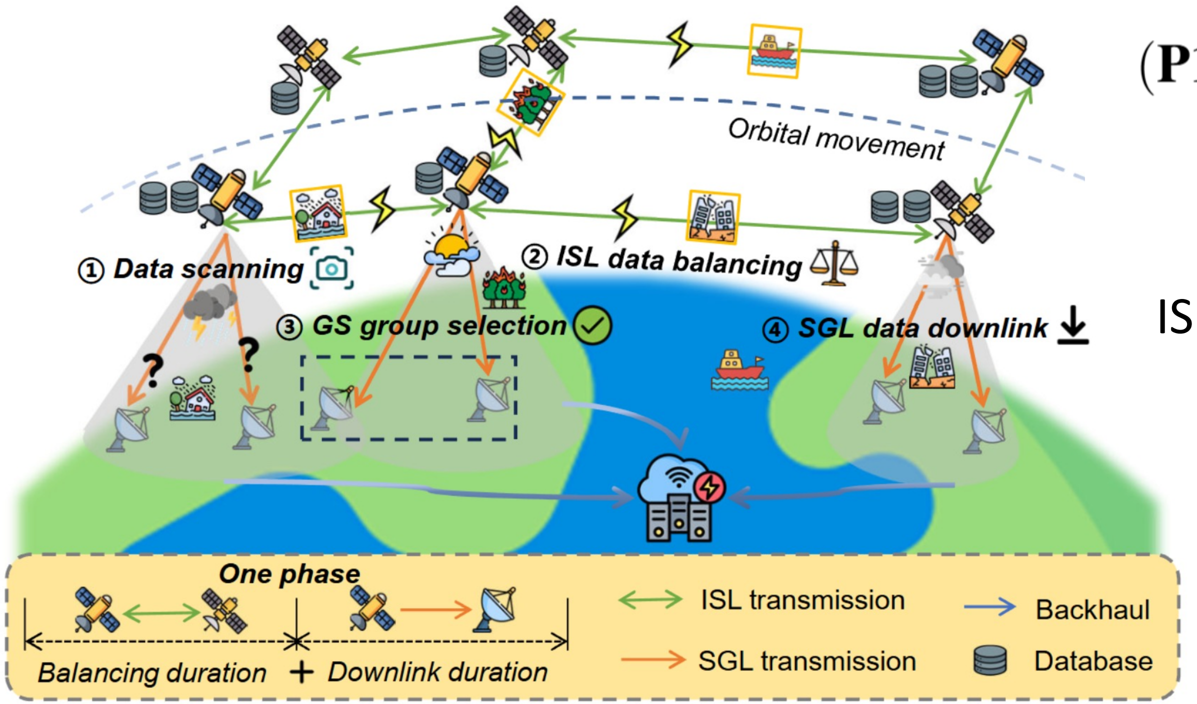

System Workflow

$$(\mathbf{P1}) \quad \max_{\{(\Pi_t,\Phi_t,\tau_t)|t\in\mathcal{T}\}} \sum_{t\in\mathcal{T}} \sum_{s_i\in\mathcal{S}} \mathbb{E}[\chi_i(\Pi_t,\Phi_t,\tau_t)]$$

$$s.t. \quad \phi_{i,i',t} \le \gamma_{i,i',t}\tau_t, \forall s_i \in \mathcal{S}, s_{i'} \in \mathcal{A}_{i,t}, t \in \mathcal{T}.$$

ISL flows $Q_{i,t}(\Phi_t,\tau_t) = q_{i,t} + \sum_{s_{i'}\in\mathcal{S}} \phi_{i',i,t} - \sum_{s_{i'}\in\mathcal{S}} \phi_{i,i',t} \ge 0,$

$$\forall s_i \in \mathcal{S}, t \in \mathcal{T}.$$

GS group assignment

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

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

Outline

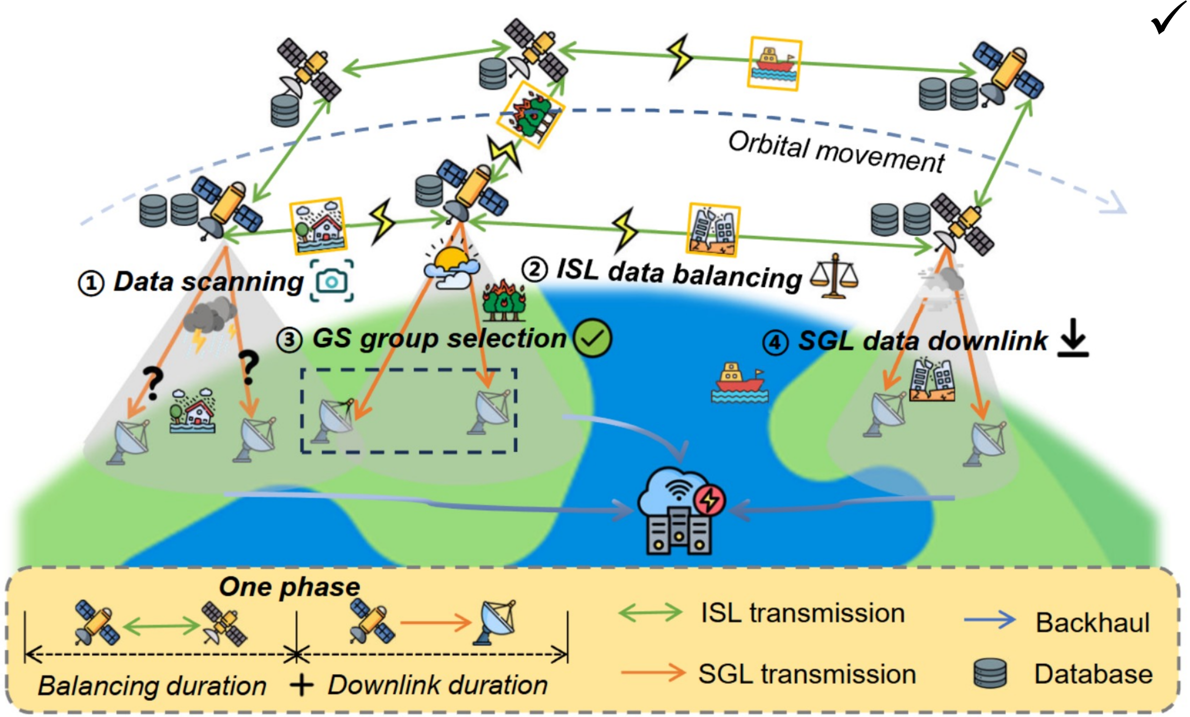

✓ ***Oracle setting: FlowDance***

- Maximize throughput in a single phase, assuming all **SGL rates are known**.
- Model the problem as a **network flow graph** where the balancing duration controls the capacity of ISL edges vs. SGL edges.

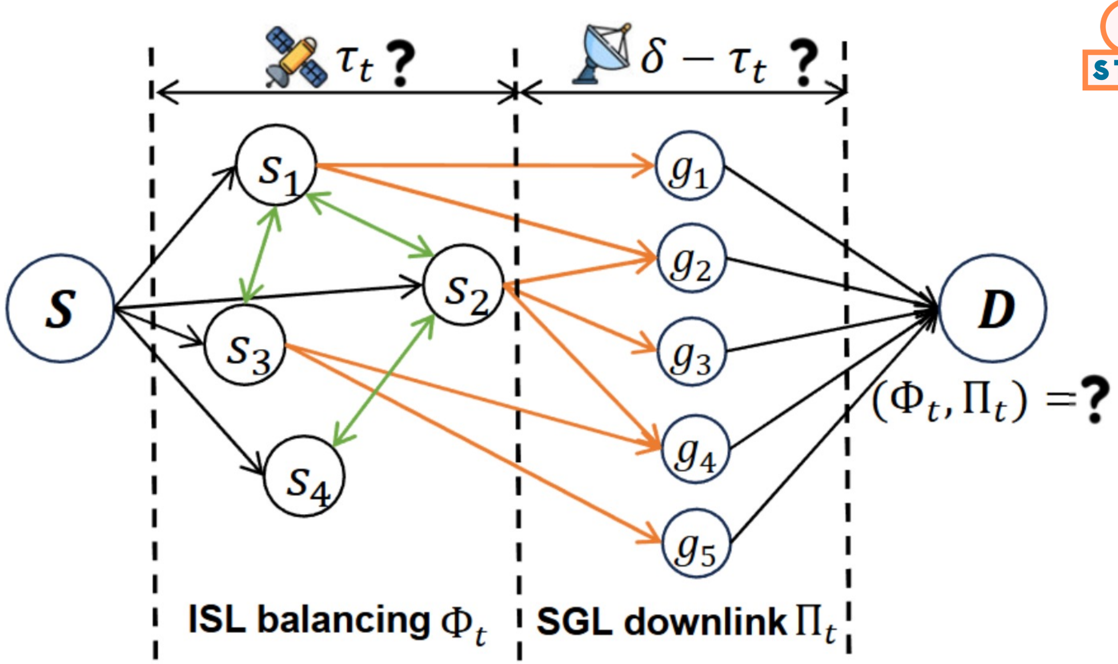

1 STEP ***Network Flow Graph Construction***

- Model the system as a **directed graph** $G(\mathcal{V}, \mathcal{E}, \tau_t)$ with source $\mathcal{S}$ and destination $\mathcal{D}$
- Embed the **ISL-SGL capacity** into edge capacities
  - SGL edges: $c_{s_i,g_j} = (\delta - \tau_t) \cdot r_{i,j,t}$
  - ISL edges: $c_{s_i,s_{i'}} = \gamma_{i,i',t} \cdot \tau_t$
  - Source-to-satellite edges: $c_{S,s_i} = q_{i,t}$
- The **balancing duration** τ_t governs the split between ISL capacity and SGL capacity

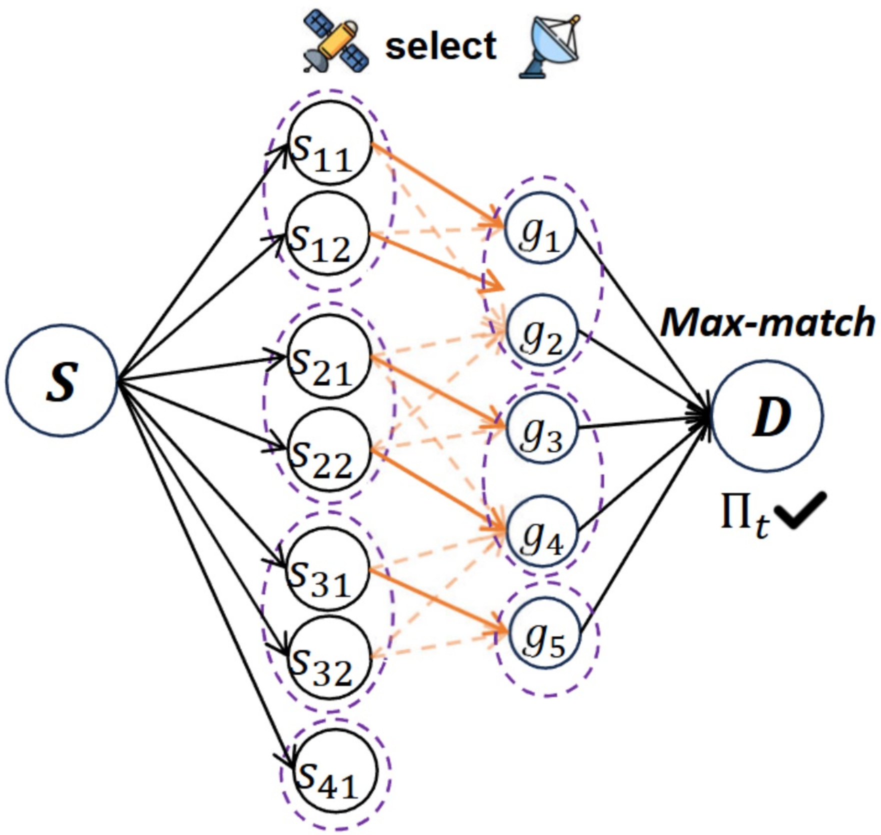

STEP 2 ***Strategic GS Group Selection***

- Construct a **bipartite graph** $\mathcal{B}(\{cp_i\}, \mathcal{G})$ by splitting each satellite s_i into σ_i copies
- Edge weight between a satellite copy and a GS is set to the SGL rate $r_{i,j,t}$
- Solve **maximum weight bipartite matching** $\Pi_t = \{\pi_{i,t}\}$
  - Per-satellite access budget: $|\pi_{i,t}| \leq \sigma_i$
  - GS exclusivity: $\pi_{i,t} \cap \pi_{j,t} = \emptyset$

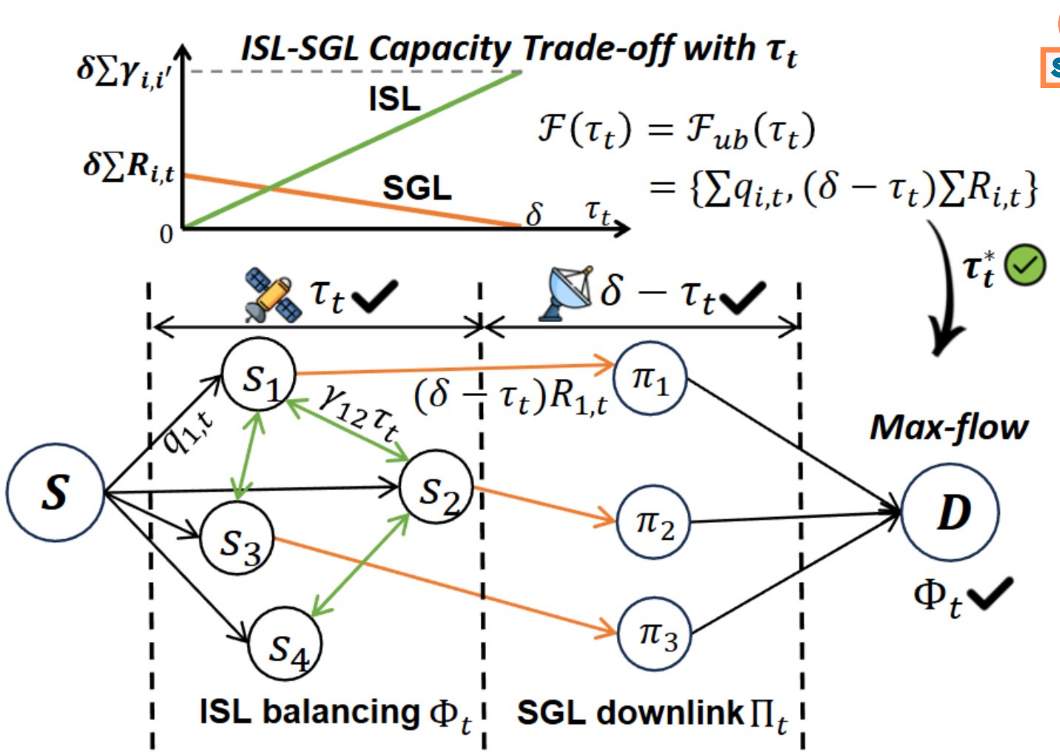

3 STEP *Adaptive ISL Flow-Optimized Balancing*

- Merge selected GSs $\pi_{i,t}$ into a super node with capacity $c_{s_i,\pi_i} = (\delta - \tau_t) \cdot R_{i,t}$
- Use binary search over $\tau_t \in [0, \delta]$ to find τ_t^* where $\Delta H(\tau_t) = \mathcal{F}_{ub}(\tau_t) - \mathcal{F}(\tau_t) = 0$
- The max-flow process outputs optimal ISL flows Φ_t and total throughput $\mathcal{F}(\tau_t^*)$

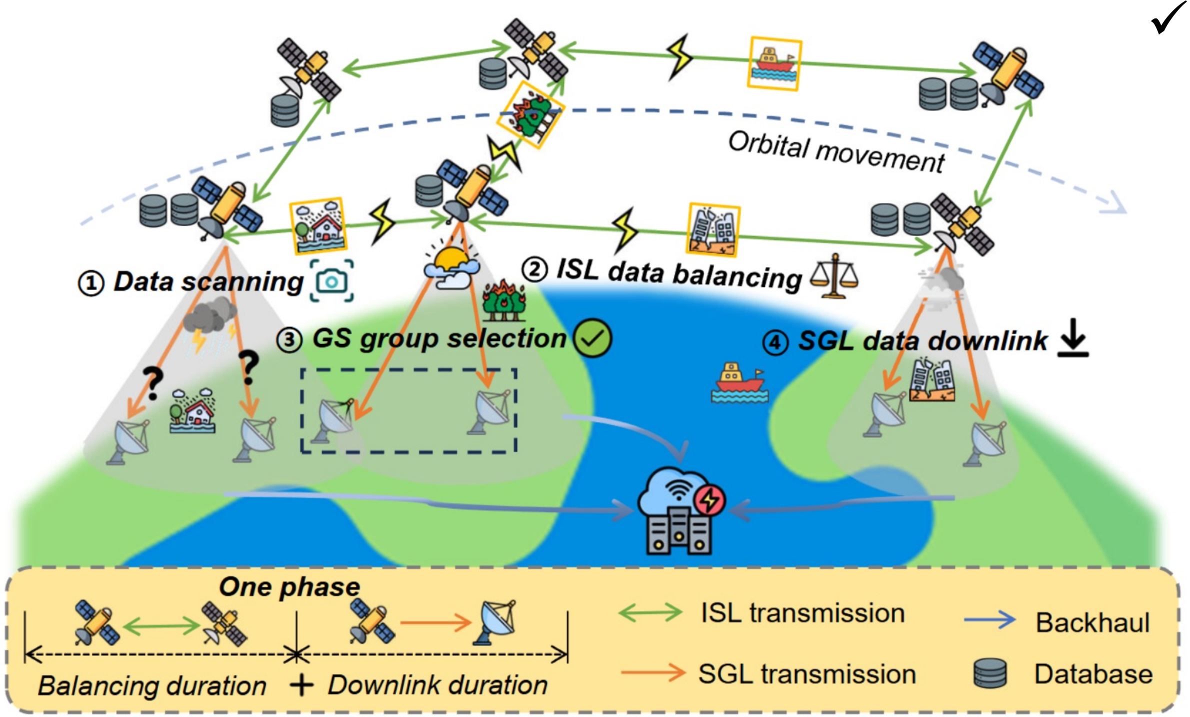

✓ ***Online setting: Ballet***

- How to learn SGL rates when contacts are short, periodic, and constrained?
- Handle GS exclusivity by cycling through the set of satellites visible to each GS, satellite access budget by dividing its visible GSs into groups.
- Use estimated rates as input to ***FlowDance.***

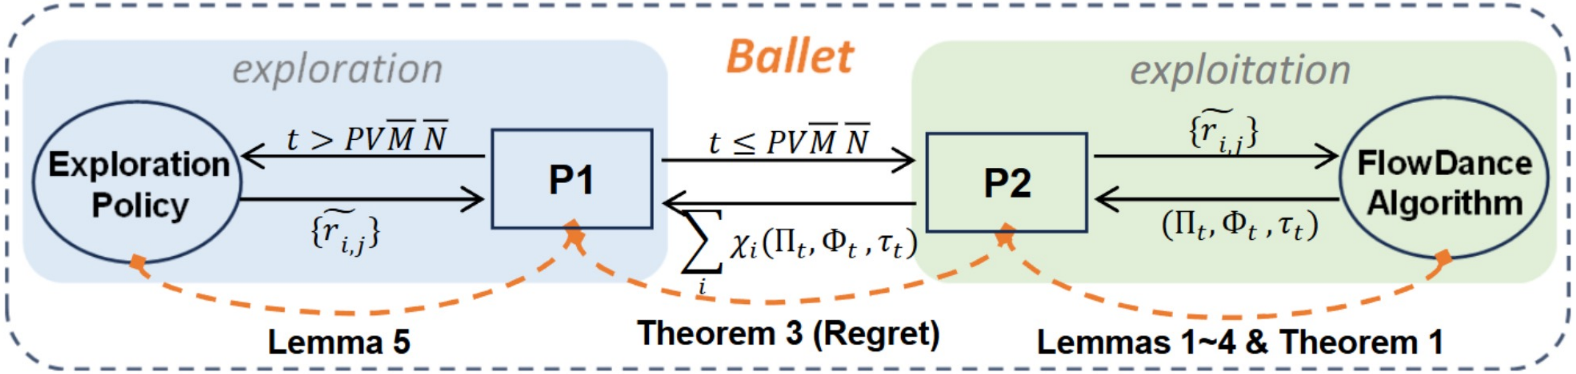

Theoretical Analysis

- ✓ Approximation Ratio of ***FlowDance*** $\alpha = \frac{X}{X^*} \ge \frac{\underline{\gamma}}{\underline{\gamma}+\bar{R}}$
- ✓ Regret of ***Ballet***

minimum ISL rate

maximum aggregate downlink rate

$$|\tilde{r}_{i,j} - \mu_{i,j}| \le \Delta, \quad \Delta = \sqrt{\frac{2(\bar{r}-\underline{r})^2 \ln T}{V}}$$

$$\widetilde{X}_t \ge \alpha X_t^{\tilde{r}}(\Pi_t^*, \Phi_t^*, \tau_t^*) - M\Delta\delta \ge \alpha X_t^* - (\alpha+1)M\Delta\delta$$

$$\alpha X_t^* - \widetilde{X}_t \le 2M\Delta\delta$$

$$T_1 = PV\overline{MN} \text{ and } V = \left(\frac{(\bar{r}-\underline{r})^2 T}{P\overline{MN}}\right)^{\frac{2}{3}} (\ln T)^{\frac{1}{3}}$$

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

Outline

Simulation Setting

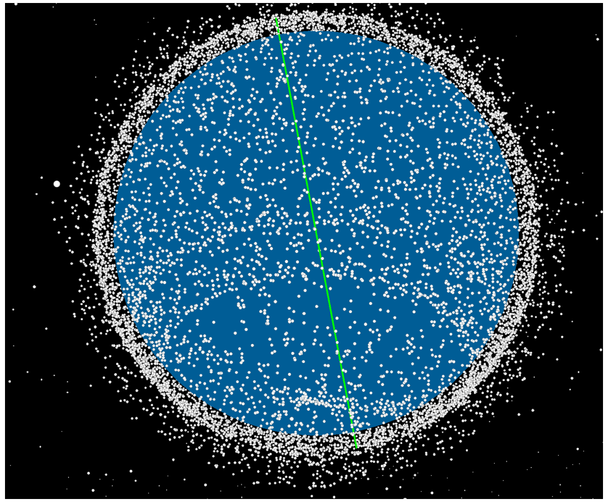

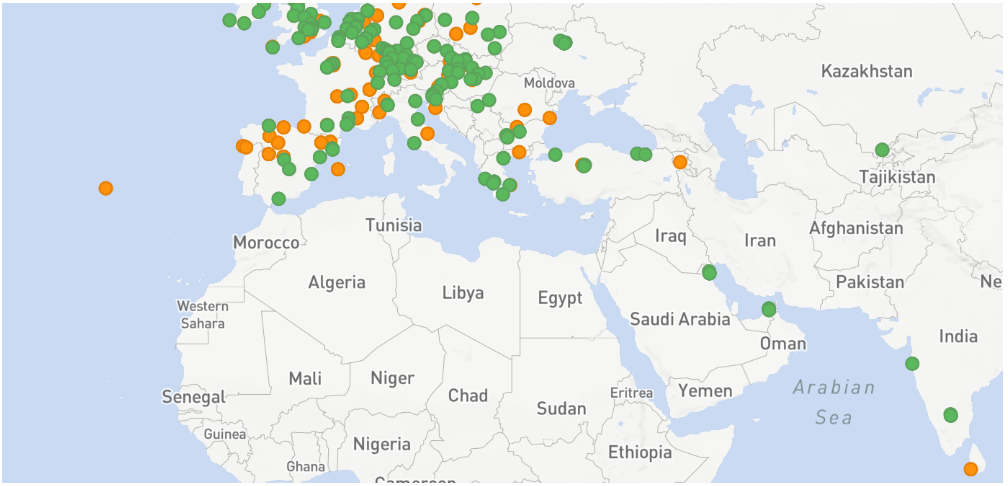

Dataset & Models

- 12 real *SkySat* satellites
- 30 real GSs from *SatNOGS*
- ITU-R P.838 rain attenuation model

Baseline

- M-Umbra/ISLCO/Hurry
- GreedyISL/RAN (random)
- BFS (brute-force oracle)

Metrics

- Accumulated downlink data
- Throughput per phase
- Evenness of cached data

Simulation results

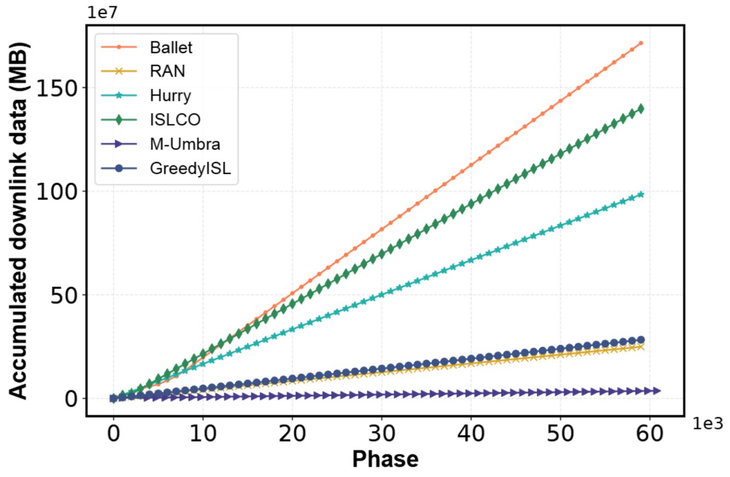

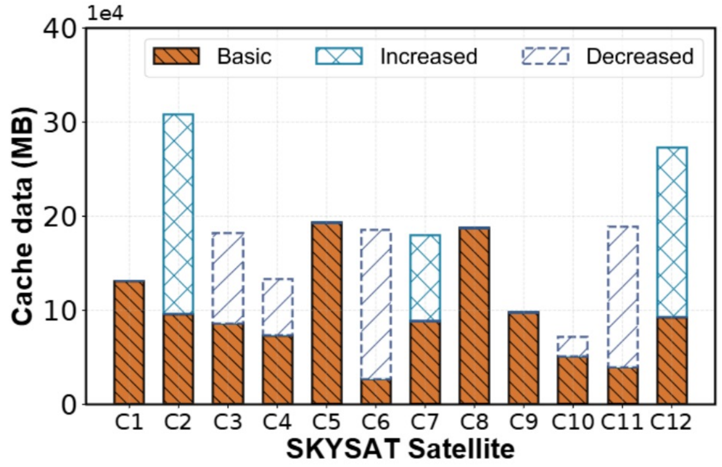

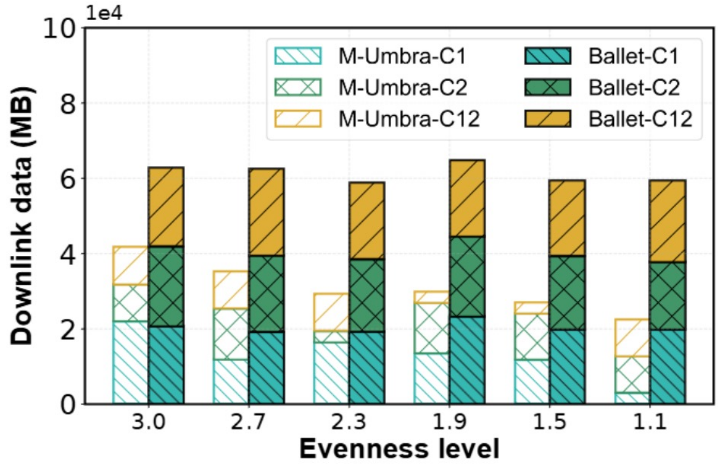

- Ballet achieves the **highest** accumulated throughput among all methods; At least **29.8%** average improvement over the best baseline; The gap widens over time as rate estimates become more accurate.
- **Before-and-after cached data** distribution: Overloaded satellites (C3, C4, C6, C10, C11) offload data to underloaded ones with better downlink opportunities (C2, C7, C12)
- **Throughput under uneven data** distribution: ***Ballet*** + ***FlowDance*** achieves **1.98×** throughput compared to M-Umbra

Simulation results

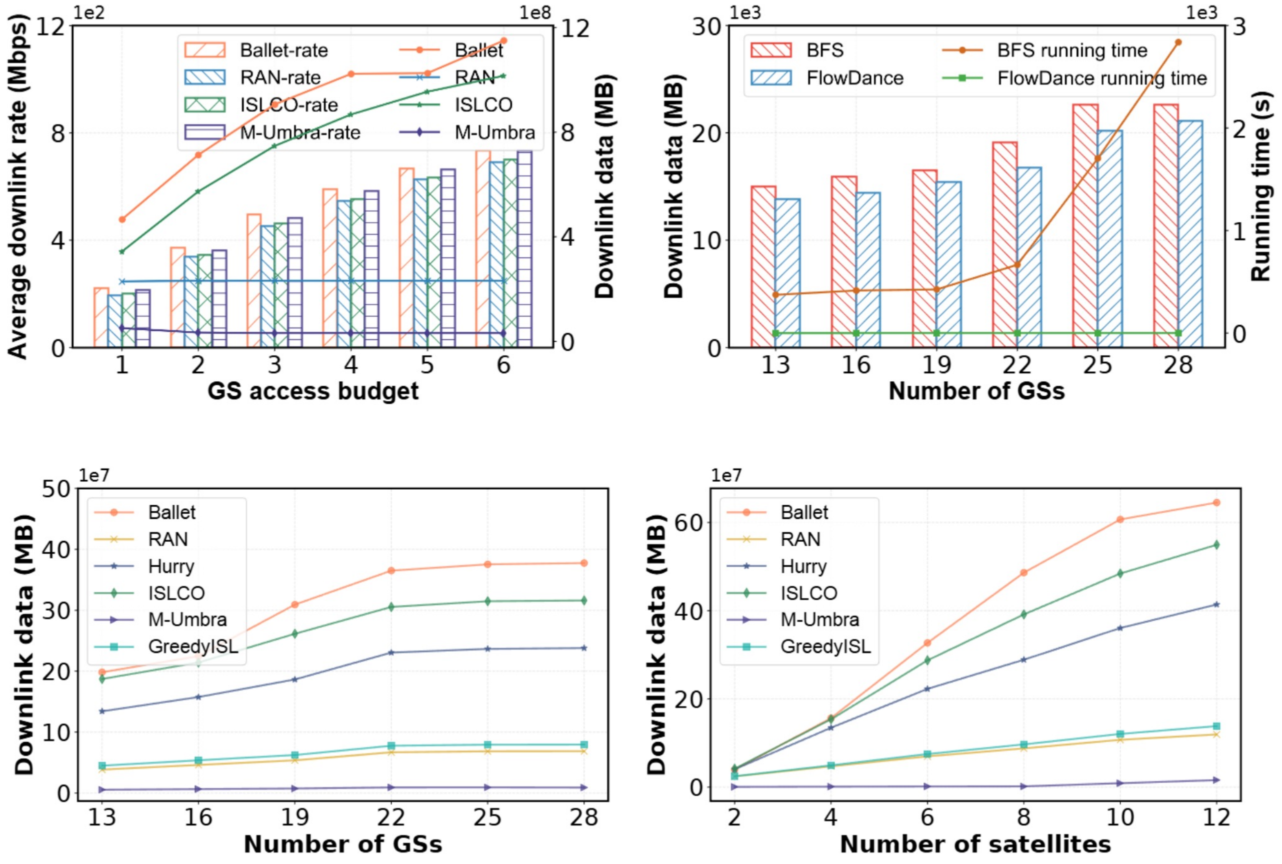

- ***Ballet*** outperforms ISLCO by **2.1×** and RAN by **3.4×** in aggregated SGL rate.
- ***FlowDance*** achieves **91.1%** of BFS throughput, far exceeding its theoretical worst-case bound of **73.5%**.
- As number of GSs and satellites increase, ***Ballet*** outperforms all baselines with an average improvement ≥**12.8%**.

Outline

- ✓ **Key Observation**: Satellite data downlink suffers from an imbalance between uneven data distribution and unbalanced downlink opportunities, exacerbated by uncertain SGL quality.
- ✓ **Our Solution**
  - ***FlowDance***: A provably good offline approximation algorithm for the oracle case, jointly optimizing ISL balancing and GS selection.
  - ***Ballet***: A partition-batch exploration policy that learns SGL rates over time, enabling adaptive, foresighted scheduling.
- ✓ **Key Contributions**
  - Joint optimization of ISL data balancing and GS group selection under SGL uncertainty.
  - Design FlowDance for oracle case and Ballet for online learning with a sublinear regret.
  - Extensive validation showing ≥29.8% average throughput improvement over baselines.
- ✓ **Future Work**
  - Incorporate predictive models (e.g., weather forecasts) for more accurate SGL estimation.
  - Extend to dynamic ISL topologies and heterogeneous satellite capabilities.
  - Address energy consumption and mission priority constraints.

Thanks for your listening!

Q & A