14th International Conference on Wireless Communications and Signal Processing (WCSP)

FedHAP: Fast Federated Learning for LEO Constellations Using Collaborative HAPs

Nov. 1-3, 2022, Nanjing, China

Mohamed Elmahallawy, Tony (Tie) Luo

Missouri University of Science and Technology

❑ Low-Earth Orbit (LEO) constellations

- Each satellite collects a **massive amount** of **Earth imagery** up to millions/billion of images every day

- This empowers to use of machine learning (ML) to address global challenges, such as **detecting disasters** in real-time, **weather forecast**, or **climate change**.

- However, downloading these **high-resolution** satellite **images** to a **parameter server** (PS) for training an ML model is impractical due to the following reasons:
  - Satellite-GS **connectivity** is **intermittent** and **irregular**
  - Network **bandwidth is limited**
  - **Constraints on** images **privacy** and **resolutions**

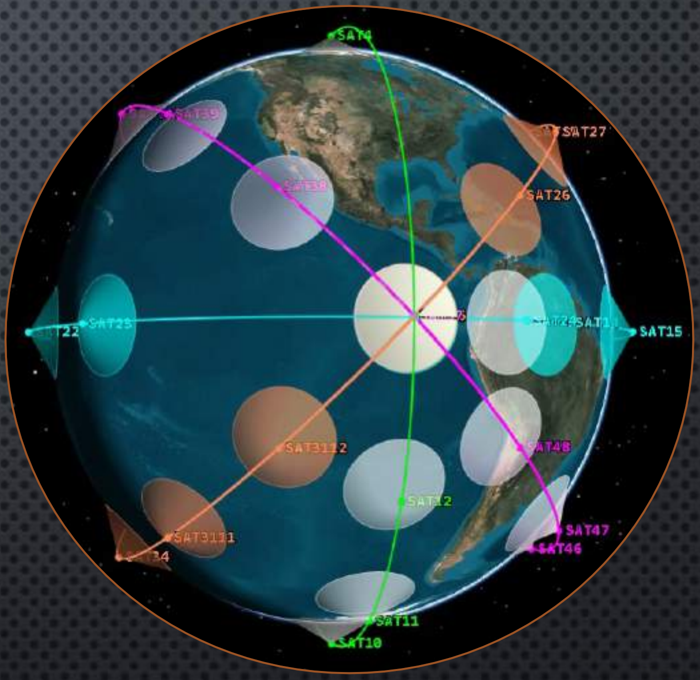

❑ Federated Learning (FL)

- **FL is collaborative learning** among satellites in various **orbits**
- **Not requiring** satellites to **download** their images to a PS (which saves **bandwidth** and **images privacy**)
- Instead, satellites download **only their ML models**

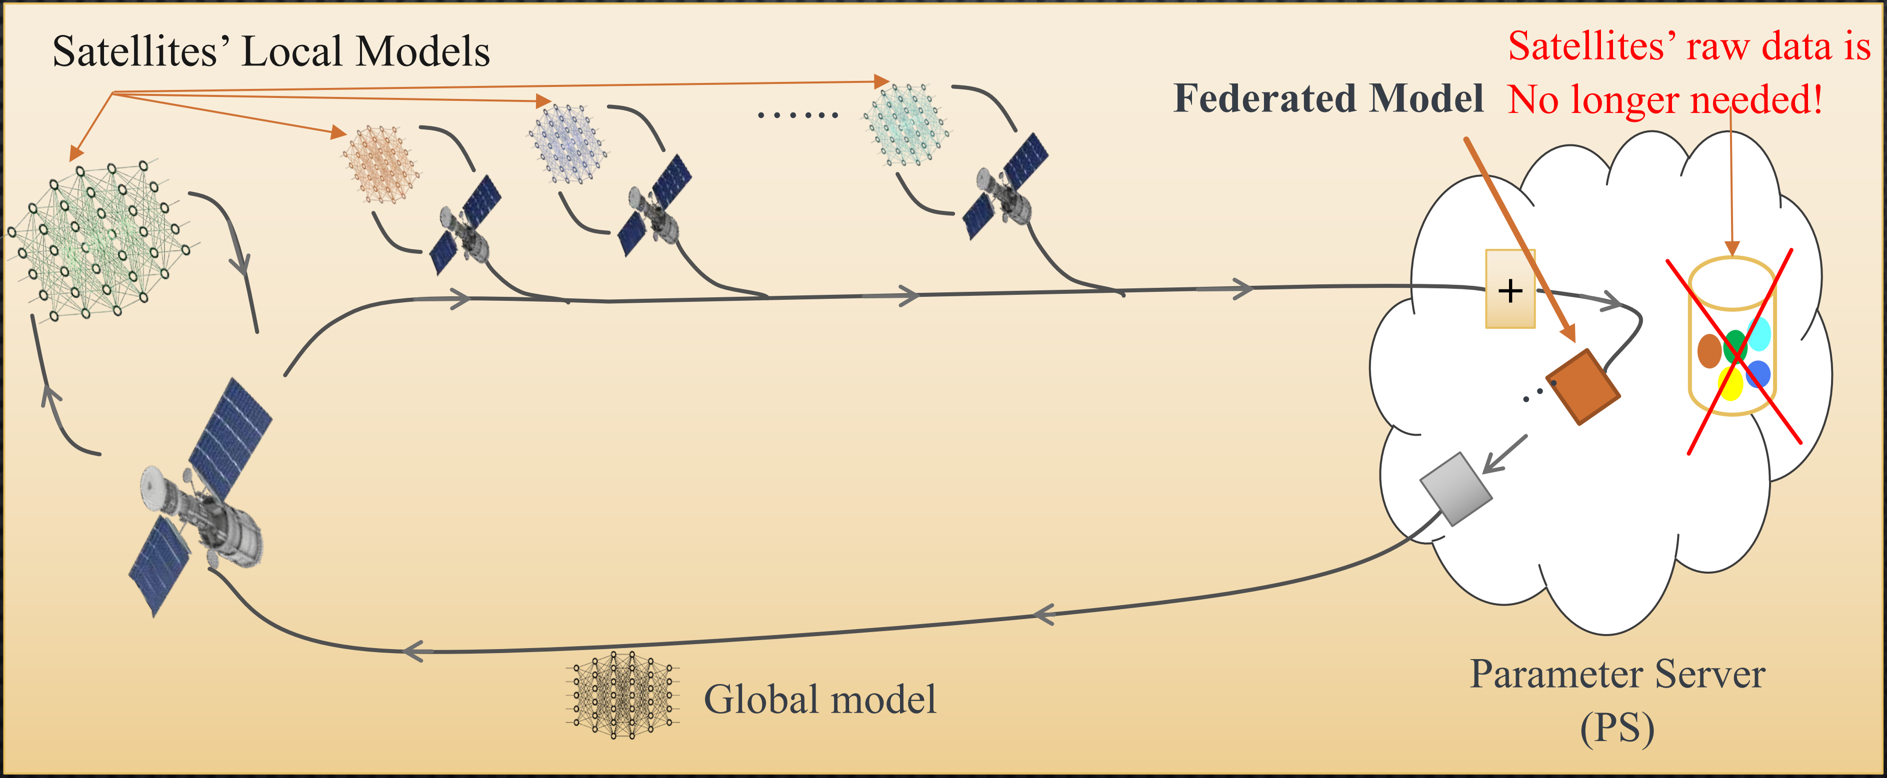

❑ Challenges

Although **FL** seems to be a **promising solution**, applying it to satellite communication (**Satcom**) poses many **challenges**.

- LEO satellites can be **visible** to the PS (e.g., ground station) **only** several times a day due to their very high speed (~**7.8 Km/s**).
- The **communication link** between an LEO satellite and the PS maybe **blocked** for several days due to **the distinct trajectories between them**.
- Communication among satellites located in different orbits, **inter-plane inter-satellite link (ISL),** will be affected by the **Doppler shift**.
- Convergence of the global model will be impacted by **uncertainty in the wireless channel** such as wind turbulence, propagation delay, and transmission delay.

❑ Existing Approaches

- **Synchronous FL** requires all LEO satellites located in different orbits to participate in the generation of the global model (causes **idleness**)
  - **FedISL** leverages intra-plane inter-satellite-link (ISL) among satellites to **reduce idleness**. However, it performs only well where the **PS** is a medium Earth orbit (MEO) satellite above the Equator. In addition, they overlooked the **Doppler shift** between them [1]
- **Asynchronous FL** allows the PS to proceed to the next training rounds without waiting for the model updates from all satellites
  - **FedSat** assumes that the GS is located at the north pole (NP) so that every satellite visits the GS at **periodically** regular intervals [2]
  - **FedSpace** offers a trade-off between **idleness** in synchronous FL and **staleness** in asynchronous FL, but it needs **satellites** to download a fraction of their data to the GS to **schedule model aggregation**, which **conflicts** with **FL principle** [3]

[1] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "On-board federated learning for dense leo constellations," in IEEE International Conference on Communications (ICC), Seoul, South korea, May 2022.
[2] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "Ground-assisted federated learning in leo satellite constellations," IEEE Wireless Communications Letters, April 2022.
[3] J. So, K. Hsieh, B. Arzani, S. Noghabi, S. Avestimehr, and R. Chandra, "Fedspace: An efficient federated learning framework at satellites and ground stations," arXiv preprint arXiv:2202.01267, Feb 2022.

❑ Our contributions

We introduce **HAPs** (High altitude platform) **in lieu of GSs** to act as **PSs** in FL to train ML models collaboratively with satellites, in a **multi-orbit LEO constellation** to achieve **fast convergence**.

- HAP is **quasi-stationary** in space
- A few HAPs can **cover a country**
- HAP can be **easily relocated**

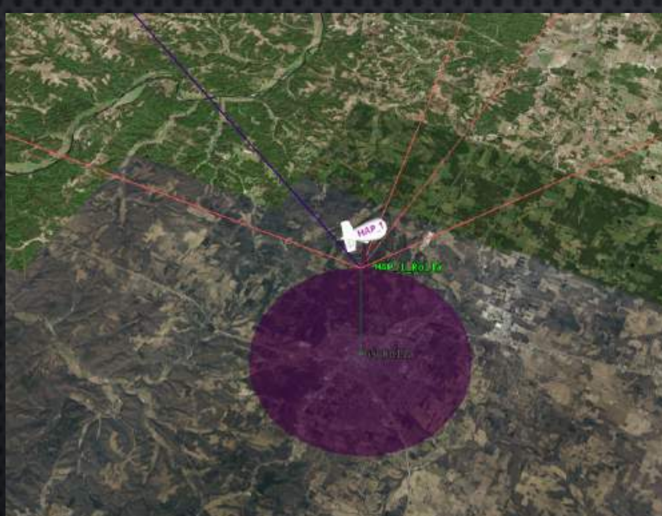

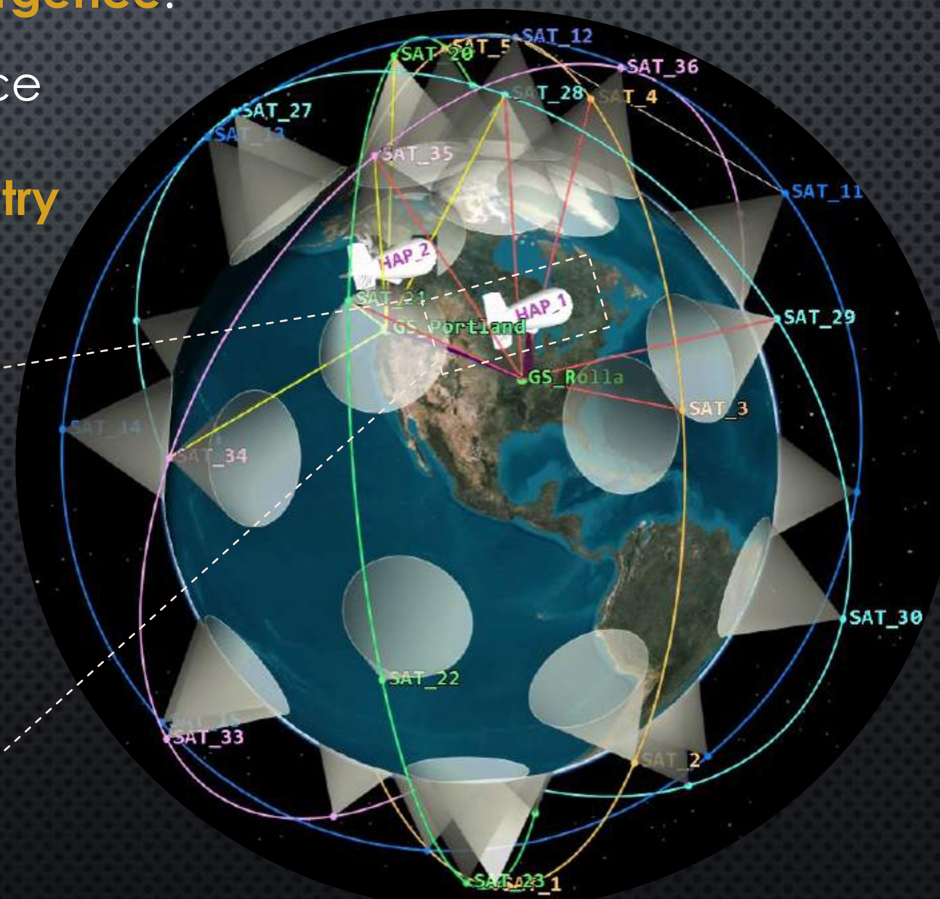

Our contributions

- We propose a framework, **FedHAP**, that consists of **three components**:
  - **Hierarchical** non-star communication architecture
  - **Model dissemination algorithm** that overcomes the challenge of **sporadic satellite-HAP** visits
  - **Partial model aggregation** that **accelerates** global model convergence
- We **evaluate** the performance of **FedHAP** in a wide range of settings (IID vs. non-IID, CNN vs. MLP, GS vs. HAP, single HAP vs. multi-HAP) with multiple state-of-the-art **FL-Satcom** approaches. The results show that **FedHAP** significantly **outperforms them** on both **convergence speed** and **model accuracy.**

❑ FedHAP Communication Architecture

- We consider a hierarchical communication architecture consisting of two tiers (worker tier and server tier).

- In lieu of implementing a traditional star topology in FL, which causes significant delays in the Satcom/LEO constellation due to highly sporadic connectivity, we propose to use multiple point-to-point architectures.

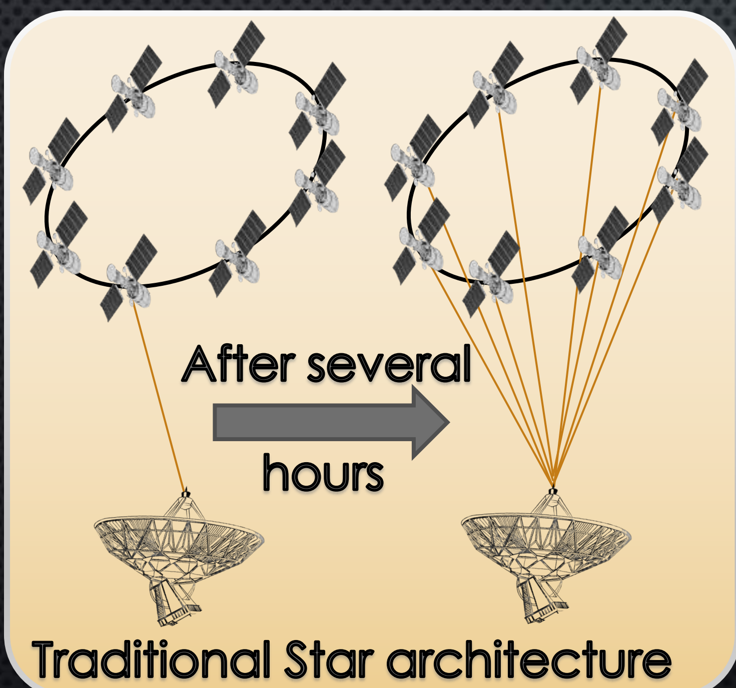

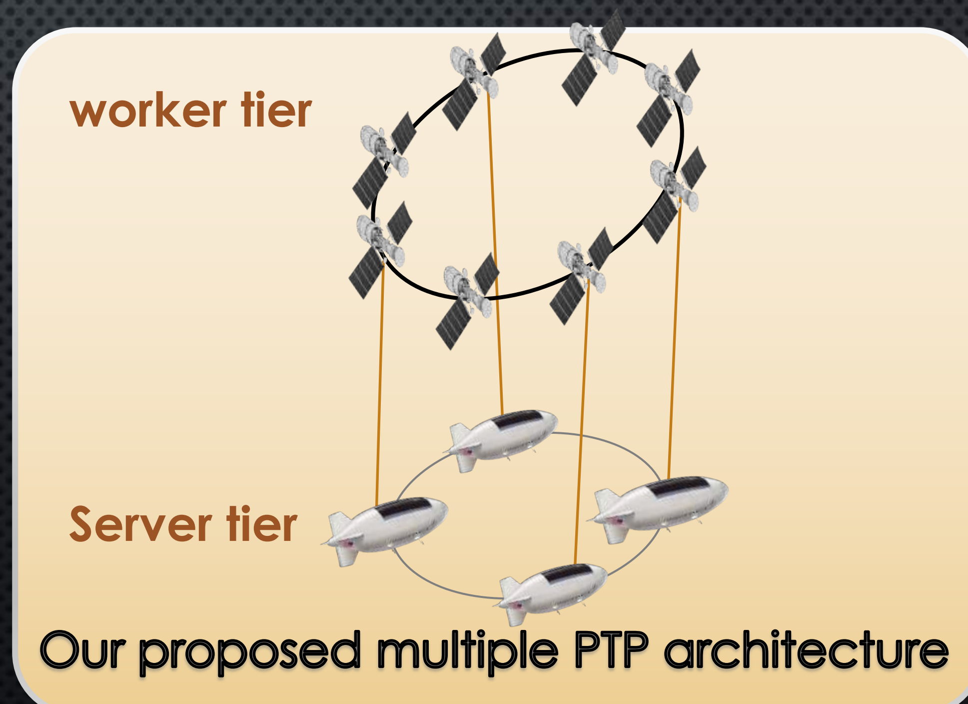

❑ Dissemination Algorithm

Step 1: **Initialization** — **Server tier**

Source HAP generates an initial global model (w^{β}, β=0) and disseminates it to its neighboring HAPs using inter-HAP link (IHL). In parallel, the source HAP transmits w^{0} to all satellites in its visible zone

Step 2: **Disseminating** w^{β} — **Server tier**

Each HAP will disseminate w^{0} to its neighbor and its visible satellites until reaches the sink HAP

Step 3: **Disseminating** w^{β} — **Worker tier**

Each visible satellite k performs two tasks:

1. Retrain w^{β} on its own data to generate an updated local model w_k^{β}.
2. Send both w^{β} & w_k^{β} to its next-hop satellite, which might be invisible, using ISL.

Server tier

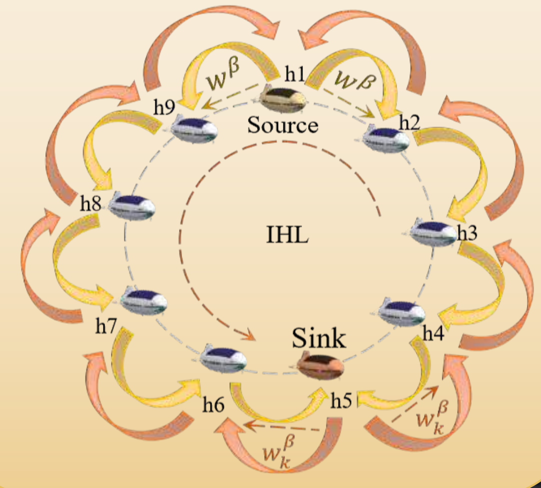

Worker tier

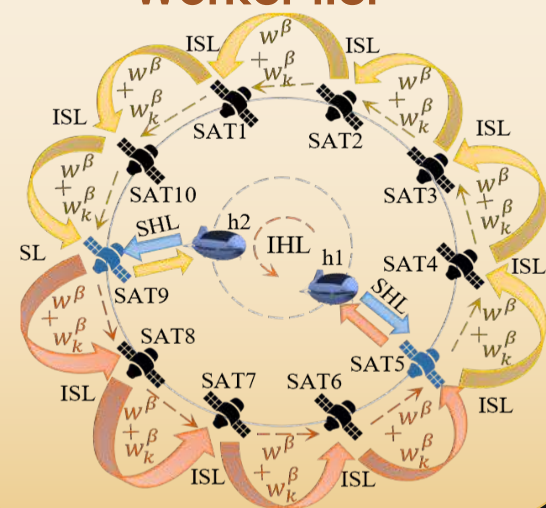

❑ Dissemination Algorithm

Step 4: Generation of *Partial global model* — Worker tier

Each of the invisible satellite k' will perform three tasks:

1. Retrain w^{β} to generate an updated local model $w_{k'}^{\beta}$.
2. Generate a partial global model as:

$$w^{\beta} = (1-\gamma_{k'})w_{k}^{\beta} + \gamma_{k'}\, w_{k'}^{\beta}\,, \qquad \gamma_{k'} = \frac{m_{k'}}{m}$$

$$m = \textit{total size of satellites' data}$$

3. Send both w^{β} & w_{k}^{β} to its next-hop satellite using ISL.

Step 5: Uploading w_{k}^{β} — Worker tier

Once a visible satellite receives w_{k}^{β} , it will transmit it for its perspective HAP

Step 6: Disseminating w_{k}^{β} — Server tier

When all HAPs have received their expected partial models, they will disseminate them to the *source HAP* for aggregation

Server tier

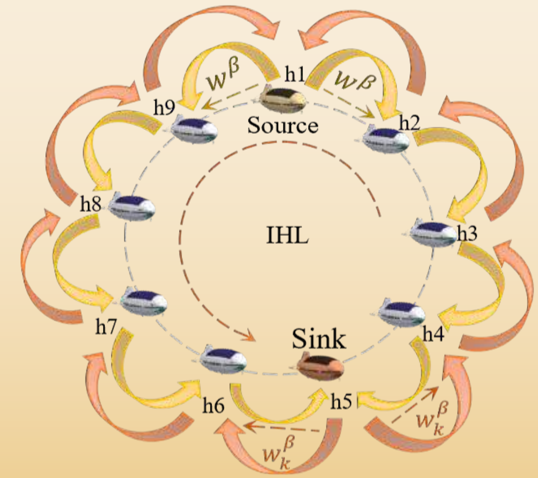

Worker tier

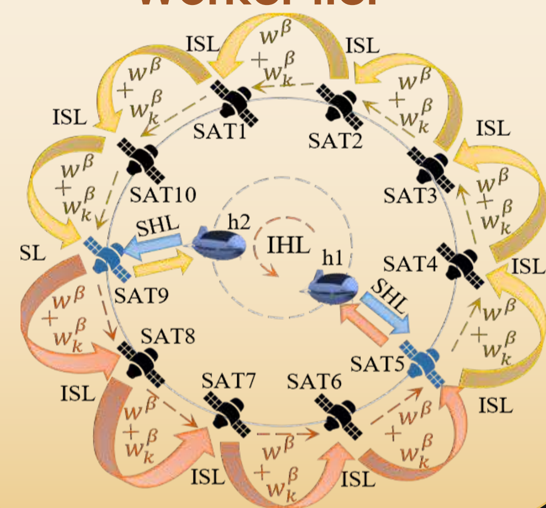

Model aggregation algorithm

Step 1 : **Organizing the partial models** Server tier

Once the source HAP has collected all partial models from their neighing HAPs, it organizes them as follows:

$$\mathcal{U} = \{\underbrace{\{\mathcal{U}_{h1}, \mathcal{U}_{h2}, \ldots \mathcal{U}_{H}\}_{l_1}}_{S_{l_1}}, \ldots \ldots, \underbrace{\{\mathcal{U}_{1}, \mathcal{U}_{2}, \ldots \mathcal{U}_{H}\}_{L}}_{S_L}\}$$

Step 2: **Filtering the redundant models** Server tier

Since satellites may be available to various HAPs, FedHAP will filter out redundant partial models for each S_l using satellite IDs received by HAPs as metadata, which results in:

$$\mathcal{U}' = \{S'_{l_1}, S'_{l_2}, \ldots, S'_{L}\}$$

Step 3: **Checking** Server tier

FedHAP checks if there is any satellite ID being left out of $\mathcal{U}'$. If yes, FedHAP will wait until $\mathcal{U}'$ receives the models of those satellites.

Step 4: **Generation of the global model** Server tier

Once $\mathcal{U}'$ receives all models from all orbits, FedHAP aggregates them as follows:

$$w^{\beta+1} = \sum_{s'_l \in S'_L} \sum_{u'_h=1}^{u'_H} \frac{m_{u'_h}}{m_l} w_k^{\beta}$$

These processes will be repeated for β=1,2,… until the FL converges.

❑ Algorithm:

Algorithm 1: FedHAP Model Dissemination & Aggregation

Initialize: Global epoch β=0, global model w^β, $\mathcal{U}_h|_{h=1}^{H} = \phi$

while *Stopping criteria is not met* **do**

- **foreach** *h from source to sink HAP* **do** ▷ `Inter-HAP dissemination of the global model`
  - Transmit w^β to all visible satellites of h
  - **foreach** *$k \in \mathcal{K}$ that is visible to h* **do** ▷ `Inter-Sat dissemination of local and partial-global models`
    - Retrain w^β on k's own data to obtain w_k^β
    - **foreach** *invisible k' between k and $k+1$* **do** ▷ `Aggregate partial models`
      - Retrain w^β on k''s local data to obtain $w_{k'}^\beta$
      - Aggregate w_k^β and $w_{k'}^\beta$ using (14)
      - Propagate both w^β and w_k^β to next k'
    - $k+1$ transmits w_k^β to its visible HAP
    - Update $\mathcal{U}_h \leftarrow \mathcal{U}_h \cup \{w_k^\beta\}$ and record all the disseminating Sat IDs
- **foreach** *h from sink to source HAP* **do** ▷ `Inter-HAP Dissemination of partial models`
  - Transmit $\mathcal{U}_h$ to the next neighboring HAP
- **if** *Source HAP receives all partial models* **then**
  - Filter out redundant models from $\mathcal{U}$ (15) based on sat ID
  - Aggregate $w^{\beta+1}$ using (16)
- **else**
  - Reschedule model aggregation to the next epoch
- $\beta \leftarrow \beta + 1$

❑ Performance Evaluation:

- We consider a **Walker-Star constellation** consisting of **5 orbits**, each with 8 satellites at an altitude of **2000 km,** with an **inclination of** angle 80°
- For a single **PS**, we consider a **HAP** or a **GS** located in **Rolla, Missouri, USA** with an elevation angle of 10°
- For Multiple PSs, **Two-HAPs.** scenario, we consider one HAP **floats** above **Rolla**, and the second one **floats** above Dallas, Texas, USA (anyplace on the Earth)
- We use the **MNIST dataset** (as used by most of FL-Satcom) that consists of handwritten numbers **from 0-9**
- For training satellites, we use two ML models : convolutional neural network **(CNN**) and multi-layer perceptron (**MLP**)
- For **non-IID** data setting, **two** orbits are trained with **4 classes (0-3)**, and **three** orbits are trained with **6 classes (4-9)**

❑ Comparison with the SOTA:

FL model	Accuracy (%)	Convergence time (h)	Remark
FedISL (ideal) [1]	82.87	3.5	GS at NP or MEO above the equator
FedISL [1]	63.74	72	GS at arbitrary location
FedSat (ideal) [2]	88.83	12	Satellites visit GS periodically
FedSpace [3]	46.1	72	GS need satellites raw data
FedHAP-GS	**83.94**	**40**	GS at arbitrary location
FedHAP-oneHAP	**87.286**	**30**	HAP at arbitrary location
FedHAP-twoHAP	**80.45 (89.83)**	**5 (30)**	HAPs at arbitrary location

[1] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "On-board federated learning for dense leo constellations," in IEEE International Conference on Communications (ICC), Seoul, South korea, May 2022.
[2] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "Ground-assisted federated learning in leo satellite constellations," IEEE Wireless Communications Letters, April 2022.
[3] J. So, K. Hsieh, B. Arzani, S. Noghabi, S. Avestimehr, and R. Chandra, "Fedspace: An efficient federated learning framework at satellites and ground stations," arXiv preprint arXiv:2202.01267, Feb 2022.

❑ Results

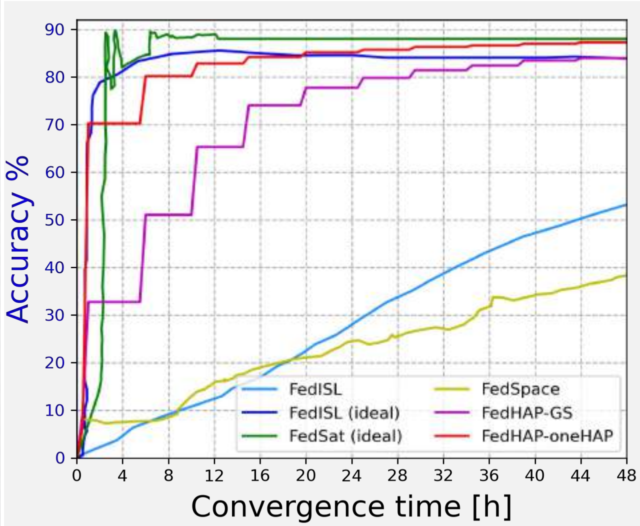

(a) Comparison with baselines.

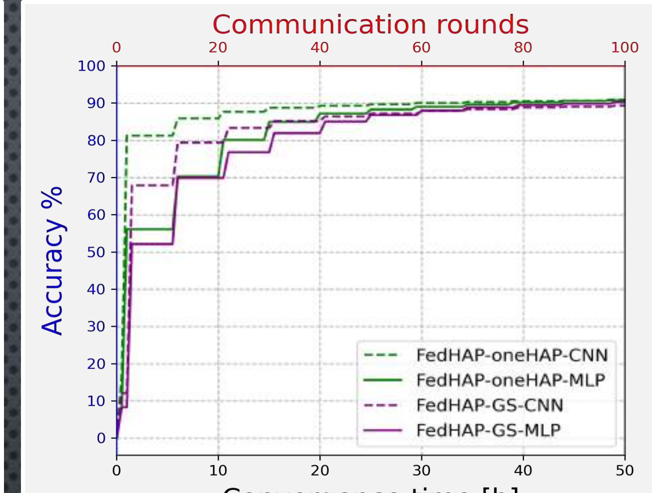

(b) IID data.

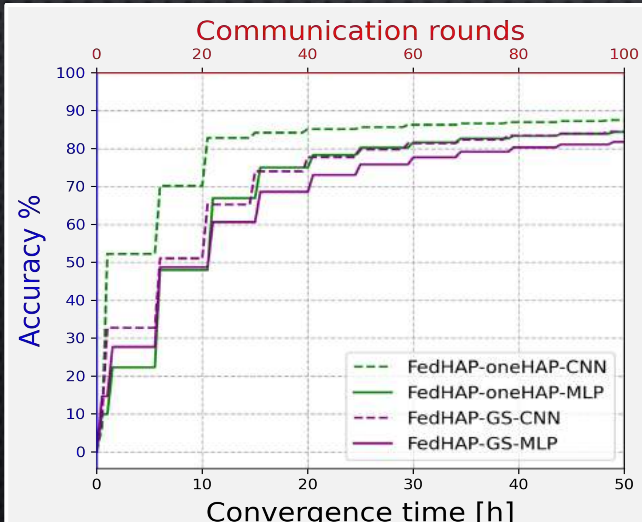

(c) Non-IID data.

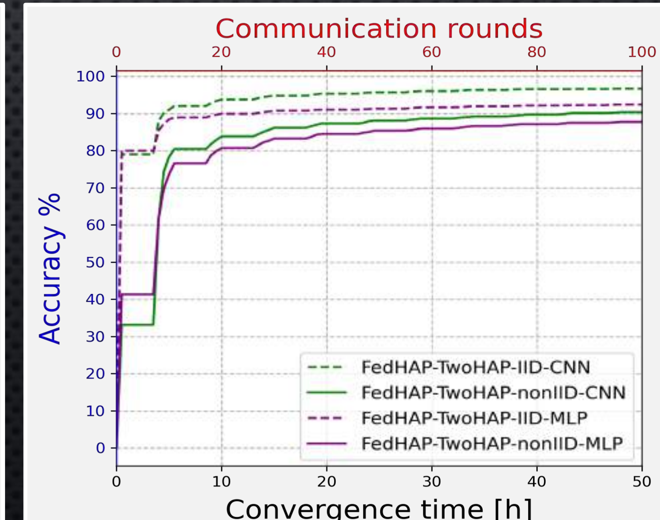

(d) Two HAPs (Rolla and Dallas) Non-IID.

❑ Conclusion

- This paper introduces **HAPs into FL-Satcom** to orchestrate the iterative learning process and proposes a **novel synchronous FL** framework called **FedHAP**

- **FedHAP** leverages inter-satellite/HAP collaborations to accelerate FL convergence and improve model accuracy.

- In addition, **FedHAP** tackles the challenge of highly sporadic and irregular satellite-GS connectivity in LEO constellations using a **hierarchical communication architecture**, **model dissemination scheme**, and **model aggregation algorithm.**

- Our **simulation results** demonstrate promising results of FedHAP as compared to the state of the art (**5 times faster** with an **accuracy** as high **as 97\%**), as well as its **robustness** to **non-IID data** as is typical in FL-Satcom settings.

Questions

14 WCSP 2022

Thank you!