satellite connectivity: 6G wireless and distributed intelligence

Petar Popovski

Connectivity Section

Department of Electronic Systems

petarp@es.aau.dk

Visiting Excellence Chair
University of Bremen, Germany

AALBORG UNIVERSITY
DENMARK

Universität Bremen

thanks

AAU

- Israel Leyva-Mayorga
- Beatriz Soret

U Bremen

- Bho Matthiesen
- Nasrin Razmi
- Armin Dekorsy

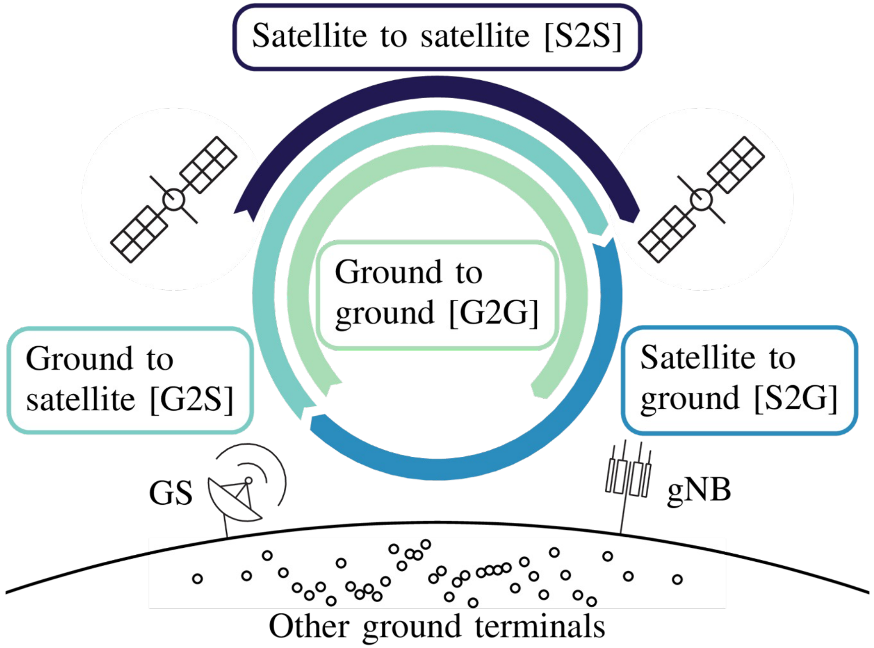

outline

- towards 6G
- revival of satellite connectivity
- satellite constellations
- federated learning in space
- satellite edge computing

outline

towards 6G

revival of satellite connectivity

satellite constellations

federated learning in space

satellite edge computing

6G: more than communications

the 5G triangle was about communications

this triangle is about to be augmented in 6G

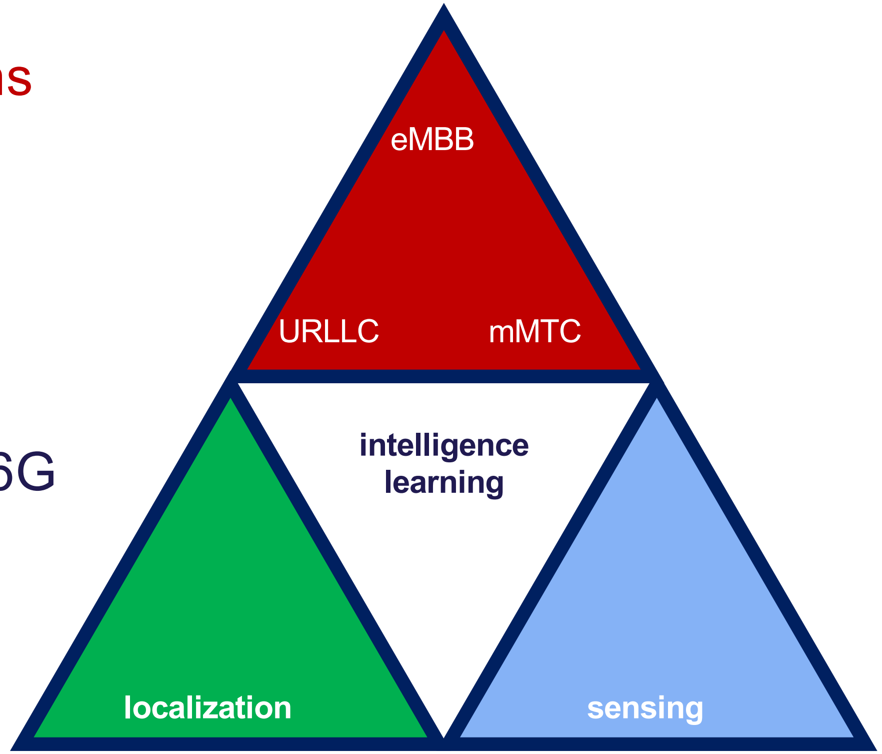

some initial steps towards semantics

6G Summit 2019

start making sense:
semantic plane filtering and control for post-5G connectivity

Petar Popovski
AALBORG UNIVERSITY

Osvaldo Simeone
KING'S College LONDON

European Research Council

6G Wireless Summit @ Levi, Finland, March 24-26, 2019

6G's search for meaning

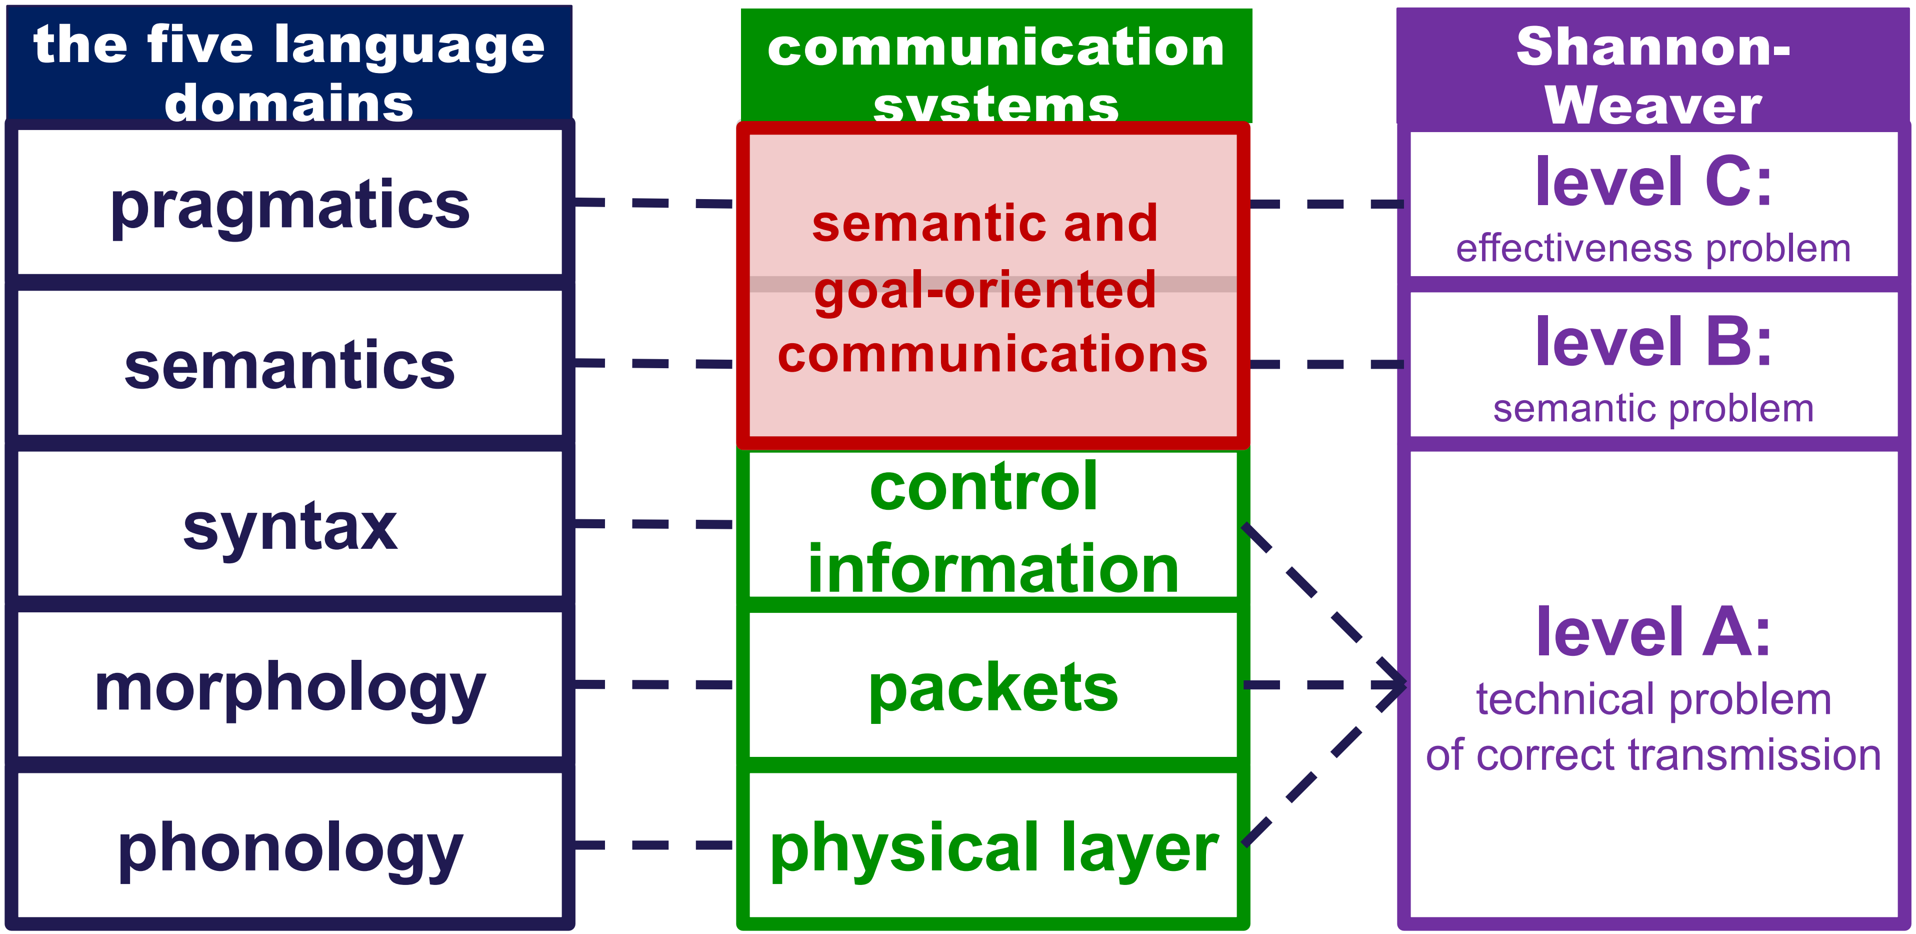

D. Gündüz, F. Chiariotti, K. Huang, A. E. Kalør, S. Kobus and P. Popovski, "Timely and Massive Communication in 6G: Pragmatics, Learning, and Inference," in IEEE BITS the Information Theory Magazine, doi: 10.1109/MBITS.2023.3322667, 2023.

IoT integral part of cellular since 5G

IoT as a micro-tunnel between
the physical and digital world

- physical information → digital data
- data+algorithms → physical actions

data used in three principal ways

- learning and training of AI models
- inference and command actuation
- value storage and exchange

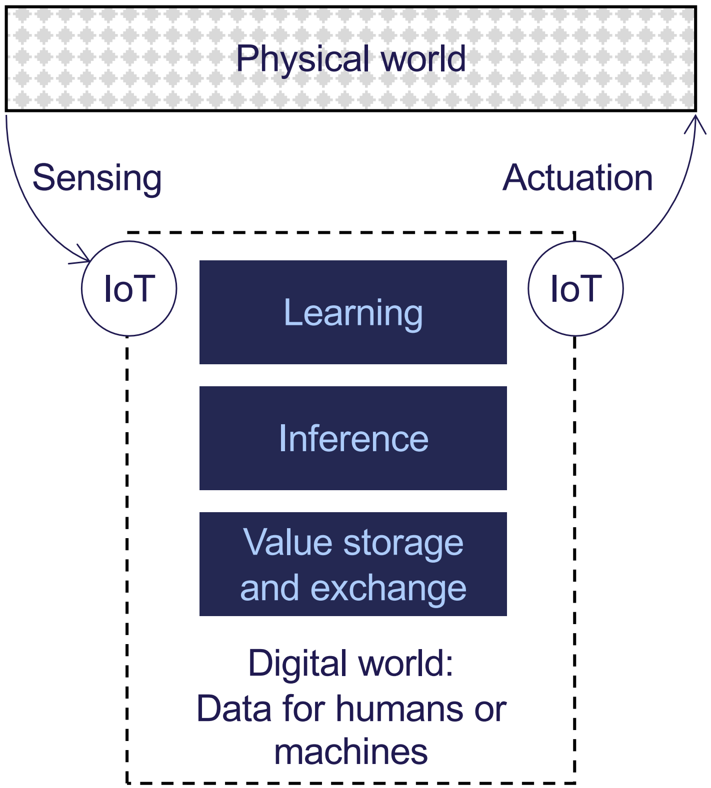

P. Popovski, F. Chiariotti, V. Croisfelt, A. E. Kalør, I. Leyva-Mayorga, L. Marchegiani, S. R. Pandey, B. Soret, “Internet of Things (IoT) Connectivity in 6G: An Interplay of Time, Space, Intelligence, and Value”, available on ArXiv, https://arxiv.org/pdf/2111.05811.pdf, 2021.

intelligence in 6G

how are the communication protocols affected
by the growing intelligence in the nodes?

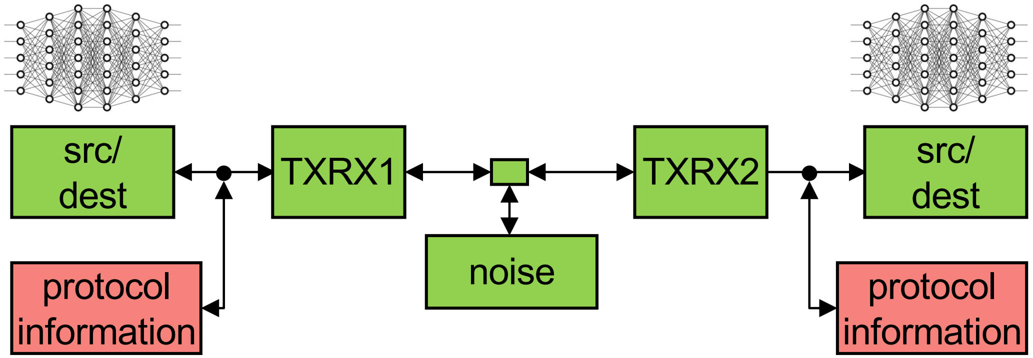

- data representation and compression to be conveyed within a specific context of knowledge or side information

[*] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, and O. Sahin, "Semantic-Effectiveness Filtering and Control for Post-5G Wireless Connectivity", Journal of the Indian Institute of Science, invited paper, 2020.
[**] Q. Lan, D. Wen, Z. Zhang, Q. Zeng, X. Chen, P. Popovski, and K. Huang, "What is Semantic Communication? A View on Conveying Meaning in the Era of Machine Intelligence", Journal of Communications and Information Networks (JCIN), invited paper, accepted, 2021.

time in 6G: beyond latency

- perception of time by humans and machines
    - Tactile Internet or Internet of Senses
- wireless connectivity **augments** the natural **time-space** context
- digital time gets intertwined with physical time
    - revisiting simultaneity, presence, causality
- increased interest in various timing measures
    - latency, Age of Information and its derivatives

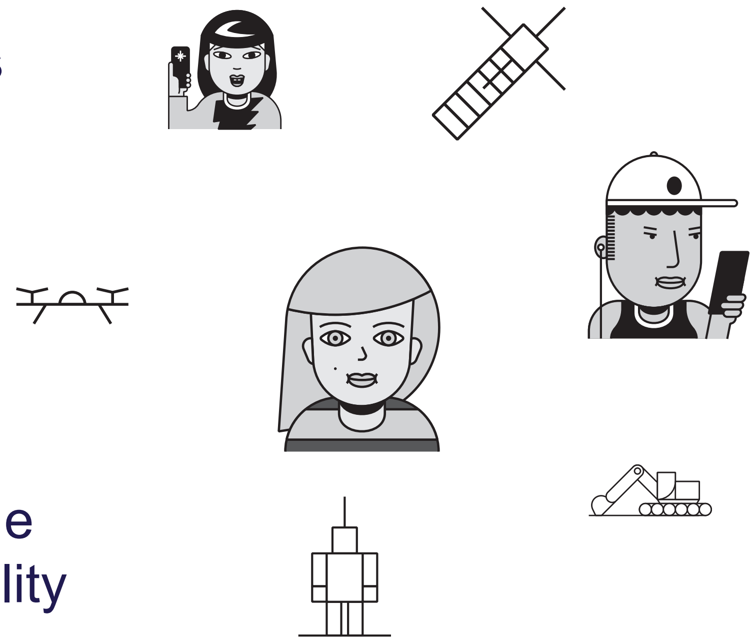

P. Popovski, F. Chiariotti, K. Huang, A. E. Kalør, M. Kountouris, N. Pappas, and B. Soret, “A Perspective on Time towards Wireless 6G,” arXiv:2106.04314, June 2021.

time in 6G: beyond latency

5G was (is) much about latency and ultra-high reliability

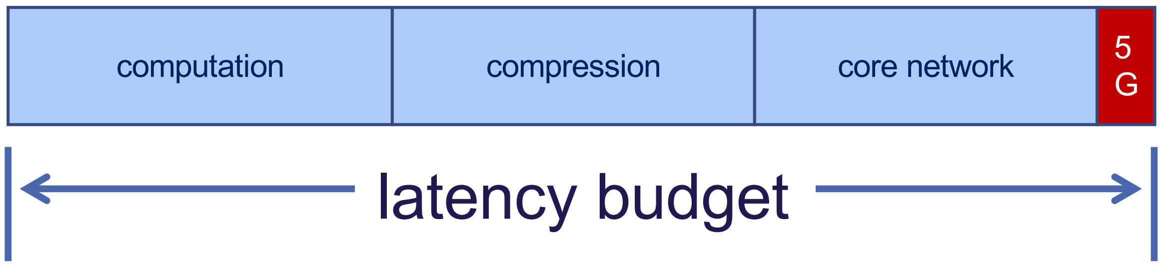

the idea with low values (~1 ms) is to
cut a low, predictable part of the latency buget

- invest latency budget into other operations

to mitigate communication failures

- paradoxically, communication should work very well! (ultra-reliable)

latency vs. age

- latency performance historically characterized with packet delays
- tracking applications and sense-compute-actuate cycles
are not sensitive to packet delay,
but to the freshness of the information at the receiver

other timing measures:

- freshness
- value of information
- interplay with intelligence and prediction

value in 6G

- enormous data amounts used in various inference and learning tasks
- privacy vs. economic value of data
- future IoT devices may become **autonomous** sellers and buyers of data

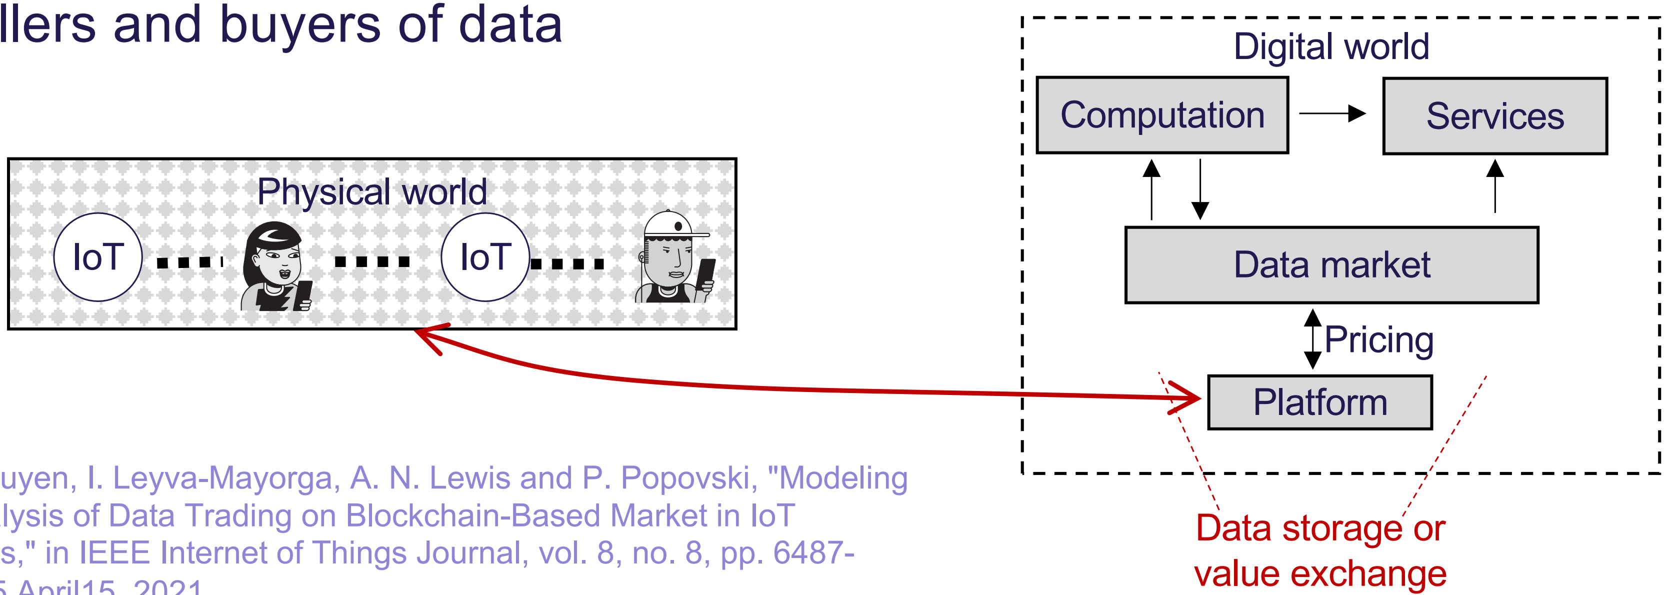

L. D. Nguyen, I. Leyva-Mayorga, A. N. Lewis and P. Popovski, "Modeling and Analysis of Data Trading on Blockchain-Based Market in IoT Networks," in IEEE Internet of Things Journal, vol. 8, no. 8, pp. 6487-6497, 15 April15, 2021

space in 6G: two aspects

controlling the propagation space

- reconfigurable intelligent surfaces (RIS), metasurfaces

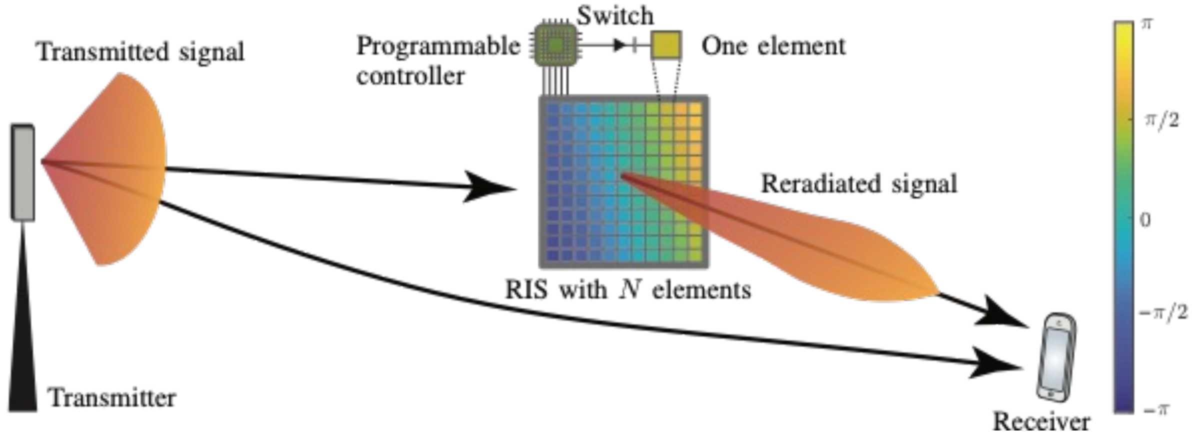

general objective: cause constructive interference where desirable

E. Björnson, H. Wymeersch, B. Matthiesen, P. Popovski, L. Sanguinetti and E. de Carvalho, "Reconfigurable Intelligent Surfaces: A signal processing perspective with wireless applications," in IEEE Signal Processing Magazine, vol. 39, no. 2, pp. 135-158, March 2022

space in 6G: two aspects

bringing 6G into space

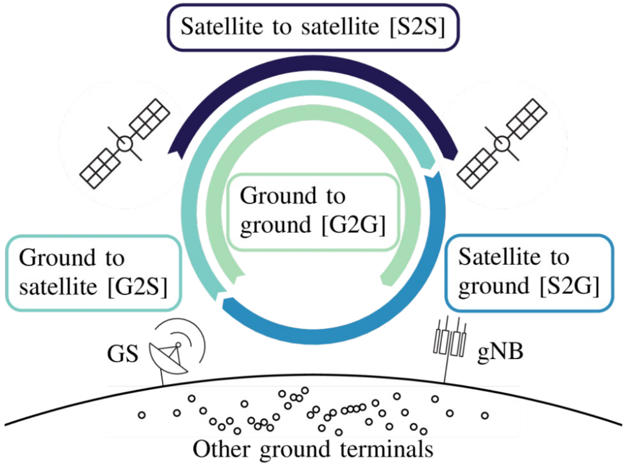

I. Leyva-Mayorga, B. Soret, M. Röpper, D. Wübben, A. Dekorsy, and P. Popovski, "LEO Small-Satellite Constellations for 5G and Beyond-5G Communications," in IEEE Access, vol. 8, pp. 184955-184964, 2020.

or even a combination between the two

- RIS tracks predictable satellite

B. Matthiesen, E. Björnson, E. De Carvalho and P. Popovski, "Intelligent Reflecting Surface Operation under Predictable Receiver Mobility: A Continuous Time Propagation Model," in IEEE Wireless Communications Letters, 2020.

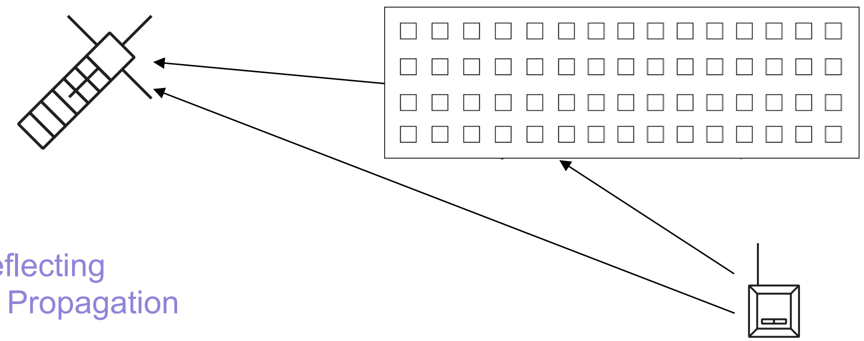

outline

- towards 6G
- **revival of satellite connectivity**
- satellite constellations
- federated learning in space
- satellite edge computing

the era of new space

Old Space

- expensive rockets, expensive satellites, long deployment times
- national agencies and states
- Inmarsat launch mass: 6100 kg

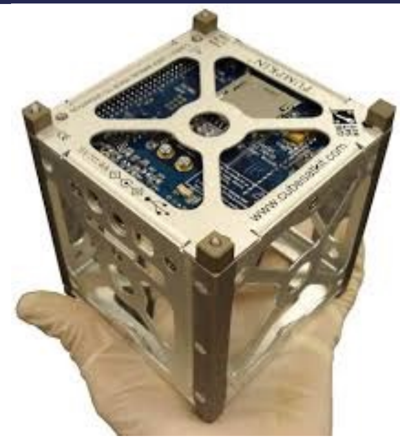

New Space

- space miniaturization
- space privatization
- novel services based on space data
- Starlink launch mass: 260 kg

SmallSat

they are **small**
few kg

they are **cheap**
commercial off-the-shelf components
launched as secondary payloads

development times are **short**

non-terrestrial networks and 3GPP

- in 2018 3GPP jumped on the bandwagon of NewSpace
  - Non-Terrestrial Networks (NTN) for the integration of satellite and terrestrial networks
  - spaceborne (i.e., GEO, MEO, LEO) or airborne (i.e., UAS and HAPS) vehicles

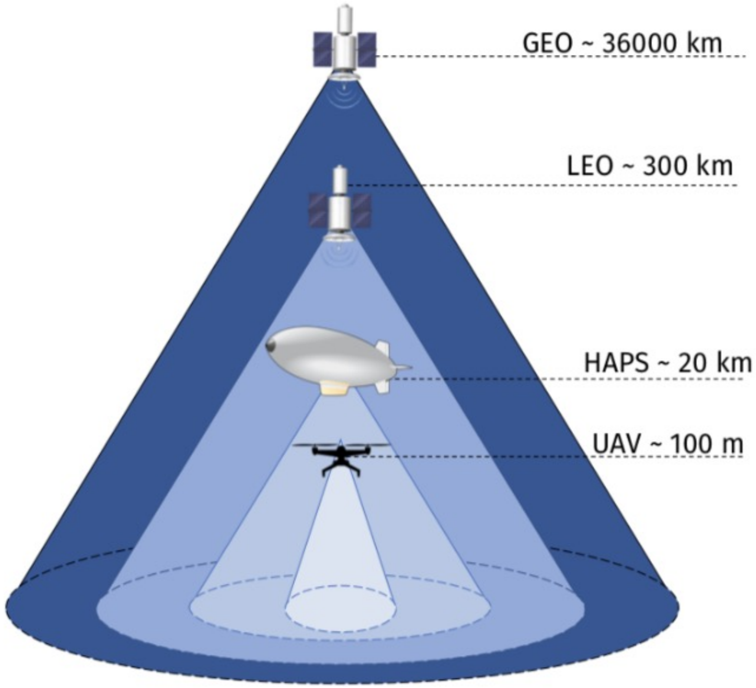

M. Giordani and M. Zorzi, "Non-Terrestrial Networks in the 6G Era: Challenges and Opportunities," in IEEE Network, vol. 35, no. 2, pp. 244-251, March/April 2021

application scenarios

- 3D orbital-aerial-terrestrial networks
  - offloading, backhauling, resilience
  - counteracts densification
- global connectivity
  - worldwide connectivity (direct access)
  - backhaul remote base stations
- Internet of Things
  - collect data
  - provide intelligence as a service
- Earth observation
  - distributed sensors
  - low latency propagation of results

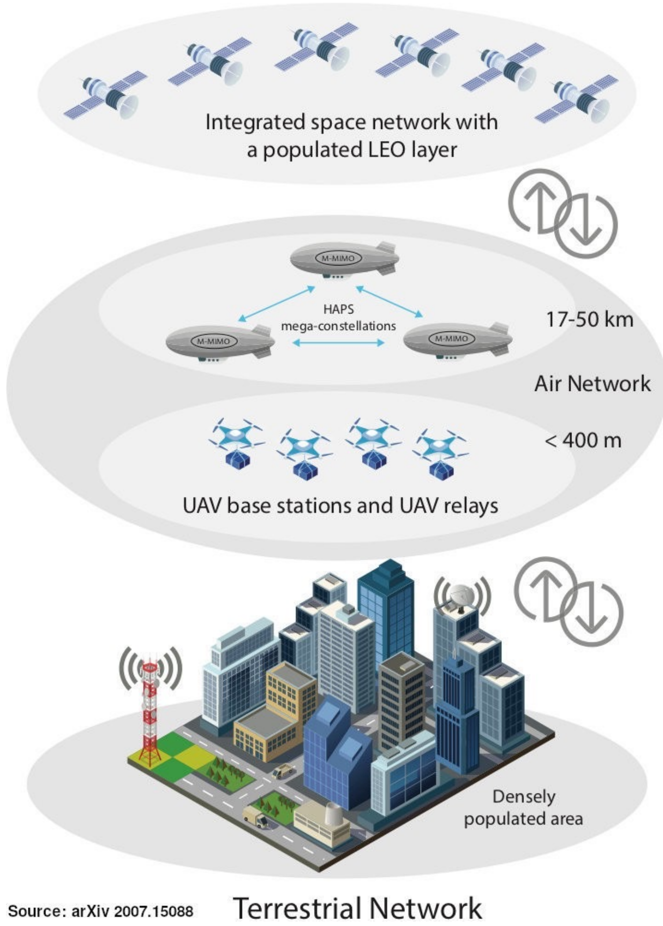

Source: arXiv 2007.15088

global, resilient, low-cost internet access

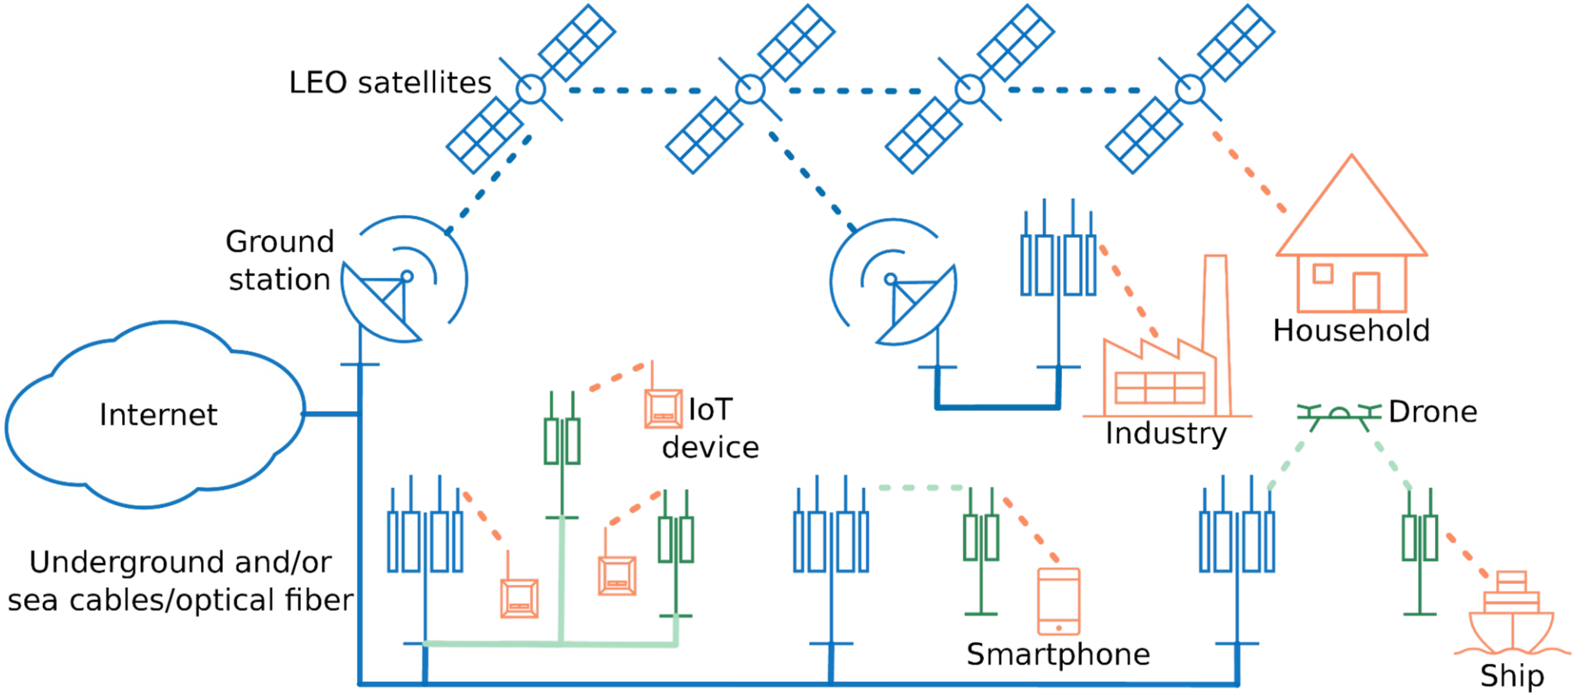

M.S. Abildgaard, C. Ren, I. Leyva-Mayorga, Č. Stefanović, B. Soret, and P. Popovski, “Arctic connectivity: A frugal approach to infrastructural development,” Arctic Journal, 2022.

mobile device and satellites

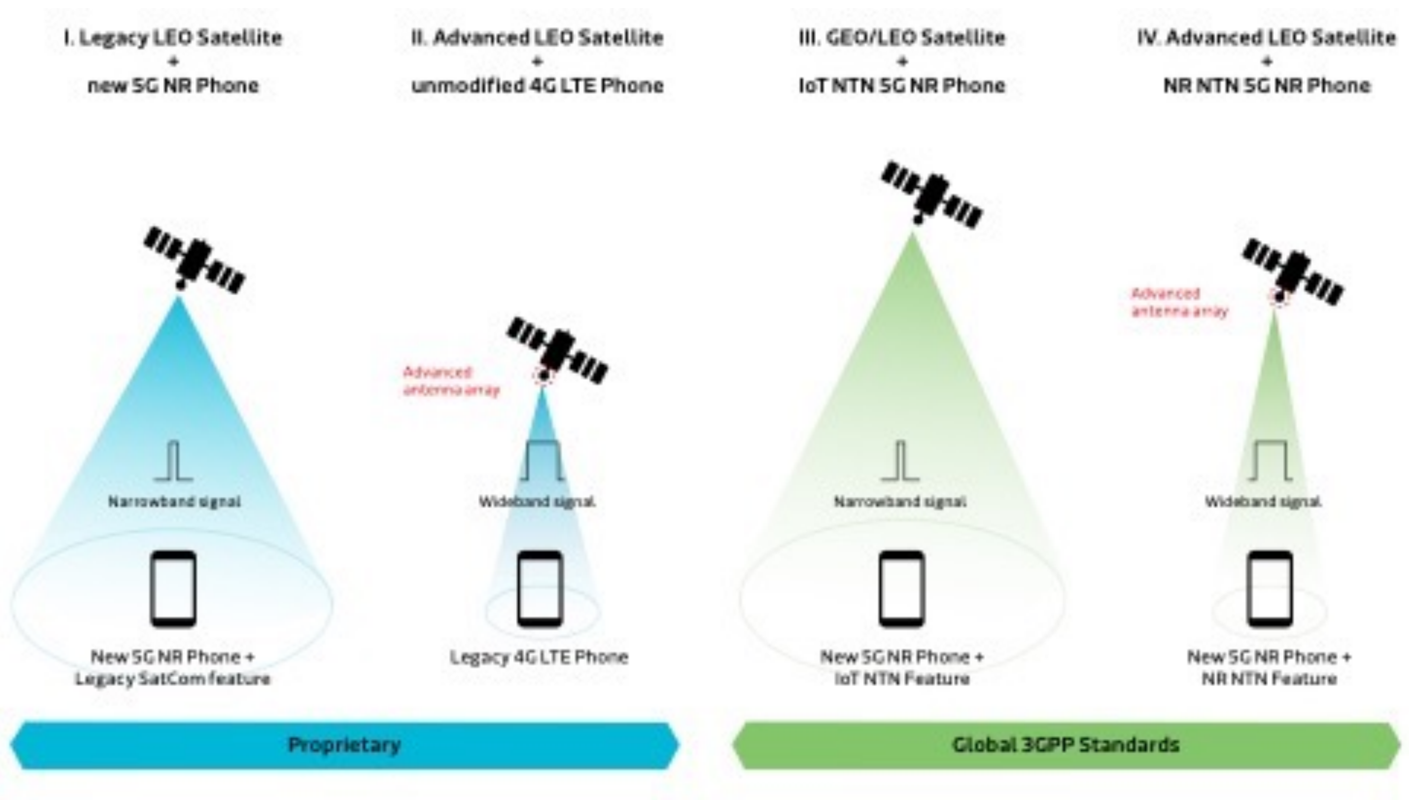

MediaTek 6G Technology White Paper, "Satellite and Terrestrial Network Convergence," April 2023

5G satellites for IoT

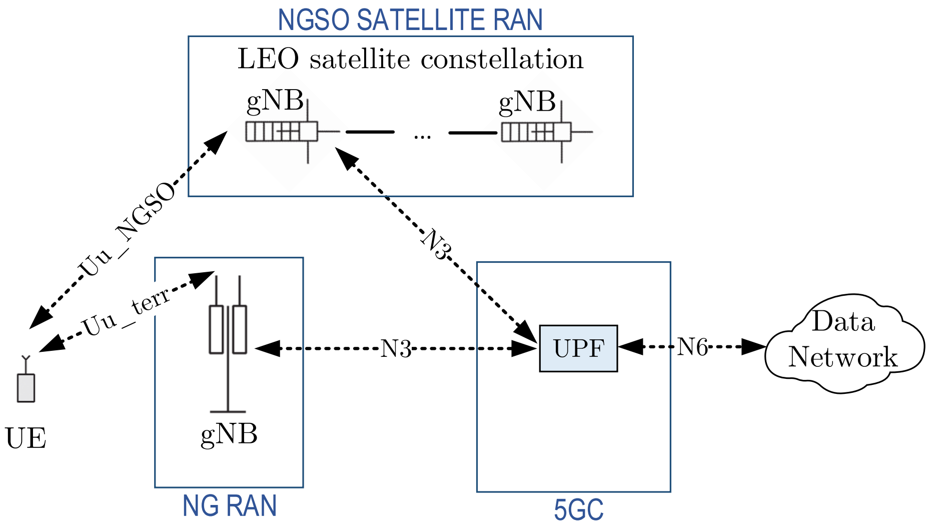

offloading

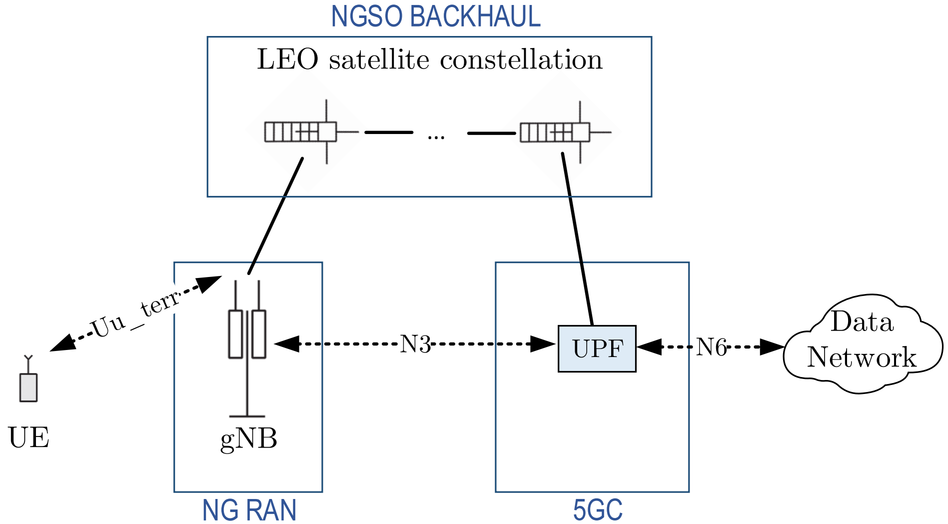

backhauling

B. Soret, I. Leyva-Mayorga, S. Cioni, and P. Popovski, "5G Satellite Networks for IoT: Offloading and Backhauling", International Journal of Satellite Communications and Networking, vol. 39, no. 4, pp. 431-444, Jul/Aug 2021.

NTN IoT connectivity in Europe

Gatehouse Satcom

3GPP compliant and proven implementations of:

- 5G NB-IoT NTN UE SW
- 5G NB-IoT NTN NodeB SW

Pre-launch Feasibility and Validation Support

- 5G NB-IoT NTN Emulator
- Feasibility Studies / Performance Validations

https://gatehousesatcom.com

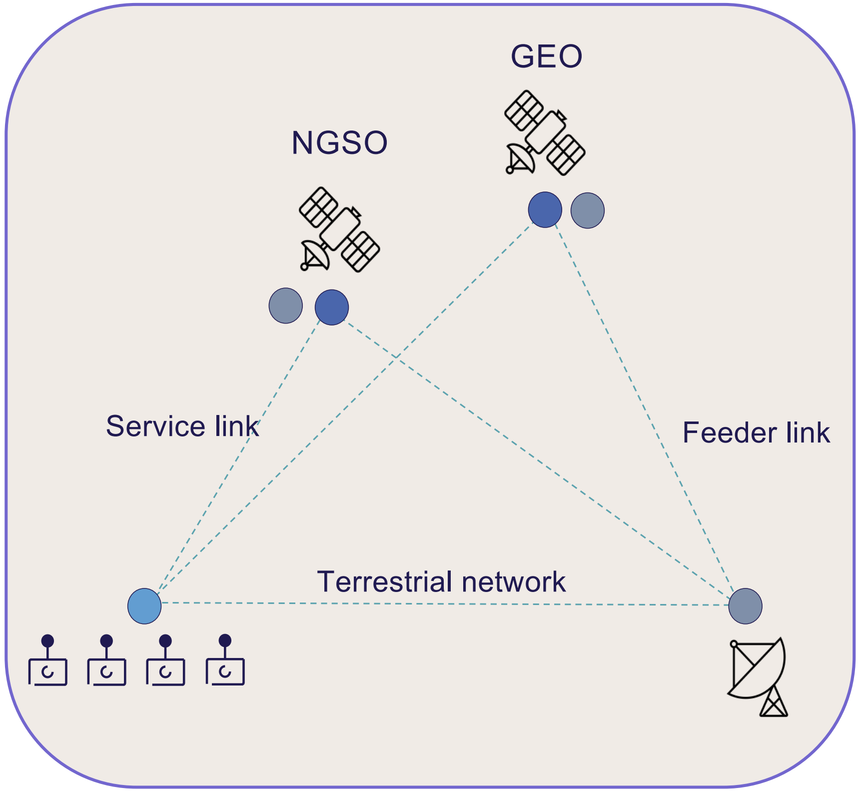

outline

- towards 6G
- revival of satellite connectivity
- **satellite constellations**
- federated learning in space
- satellite edge computing

satellite orbits

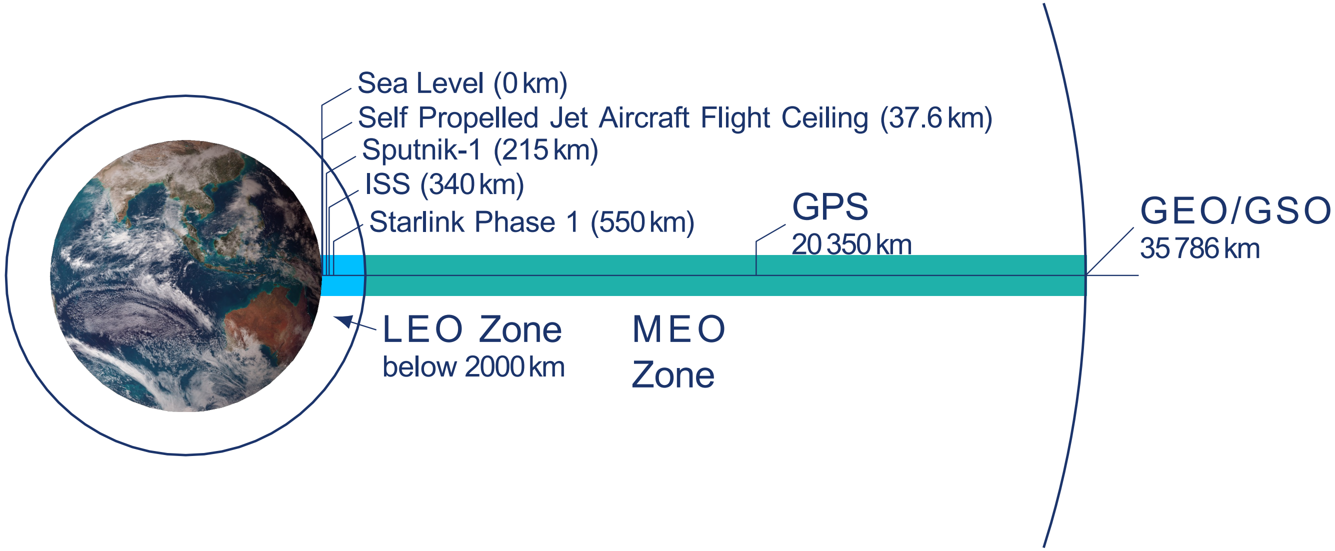

Low Earth Orbit (LEO): Orbital period ≤ 128 min
Medium Earth Orbit (MEO): Between LEO and GSO
Geosynchronous Orbit (GSO): Orbital period 23 h 56 min 4 s (1 sidereal day)
Geostationary Orbit (GEO): circular GSO above Equator
High Earth Orbit (HEO): Beyond GSO

LEO small satellite constellations

- propagation latency of several ms
- Doppler spread can be very significant

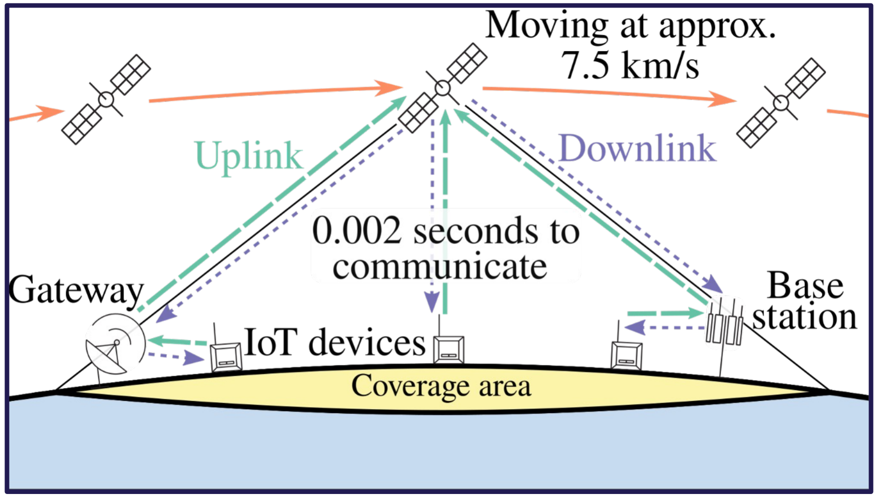

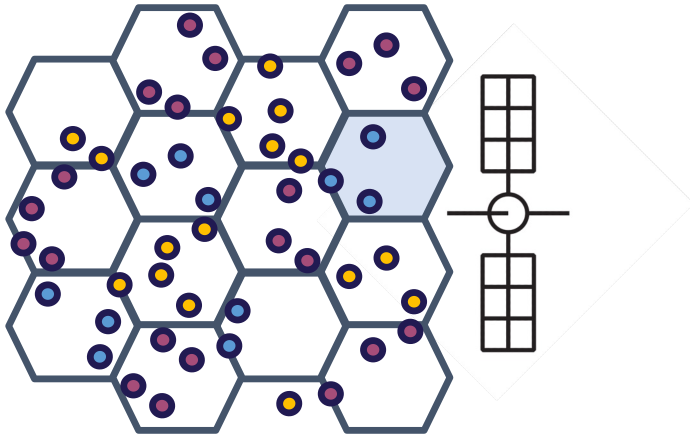

cell types

- **Earth-moving cells**
  follow the satellite as it orbits the Earth
- **quasi-Earth fixed cells**
  fixed on the ground, tracked by satellite beams

I. Leyva-Mayorga, B. Soret, M. Röpper, D. Wübben, A. Dekorsy, and P. Popovski, "LEO Small-Satellite Constellations for 5G and Beyond-5G Communications," in IEEE Access, vol. 8, pp. 184955-184964, 2020.

satellite constellations

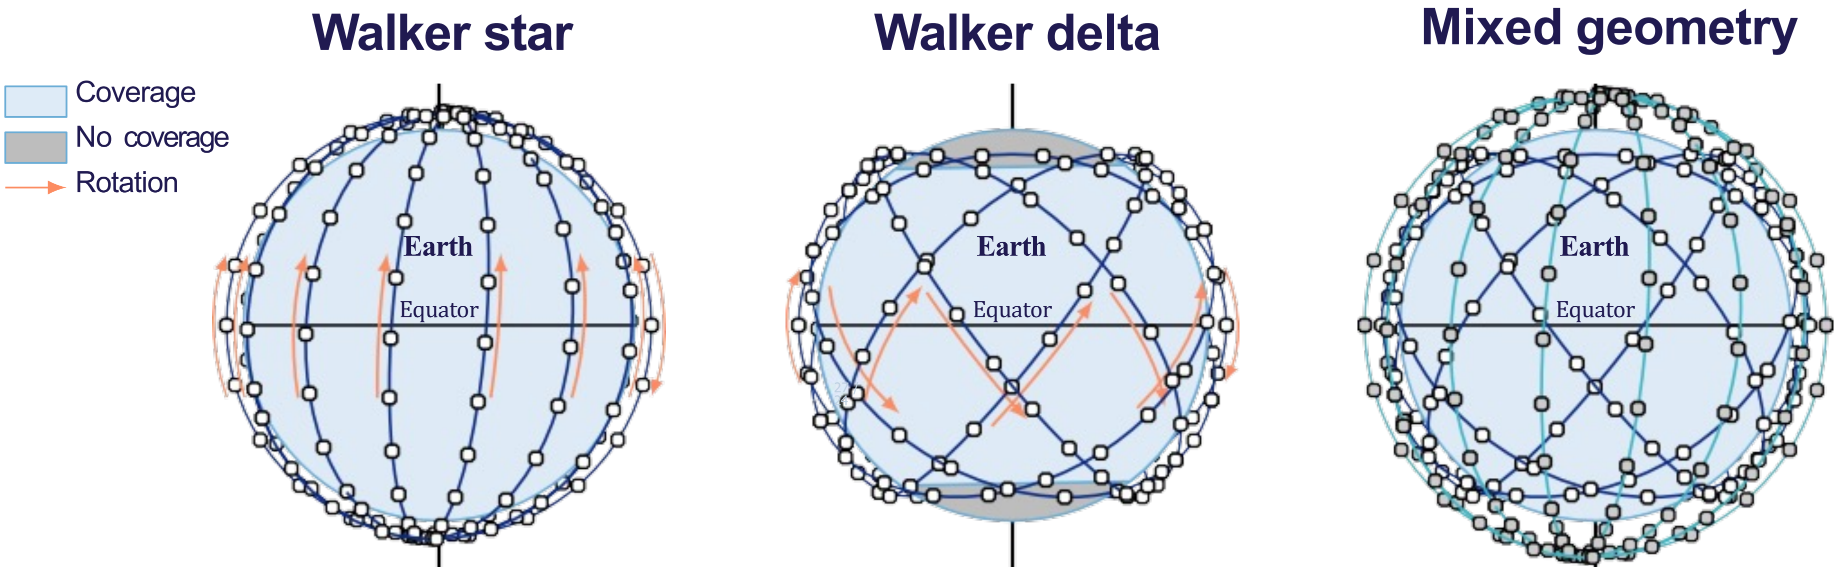

Orbital height	Velocity	Period
200 km	7.78 km/s	88.5 min
550 km	7.59 km/s	95.7 min
2000 km	6.90 km/s	127.2 min

Leyva-Mayorga, Soret, Matthiesen, Röper, Wübben, Dekorsy, Popovski, “NGSO constellation design for global connectivity”, in Non-Geostationary Satellite Communications Systems, Lagunas, Chatzinotas, An, Beidas, Eds., IET, Jul. 2022, to appear.

inter-satellite networking

- link types:
  - Intra-Plane: same orbital plane
  - Inter-Plane: different orbital planes, same orbital shell
  - Inter-Orbit: different orbital altitudes
- Free Space Optical (FSO) and RF interfaces
- Intra-Plane: stable relative position → FSO
- Inter-Plane / Inter-Orbit:
  - high relative velocity
  - short contact times

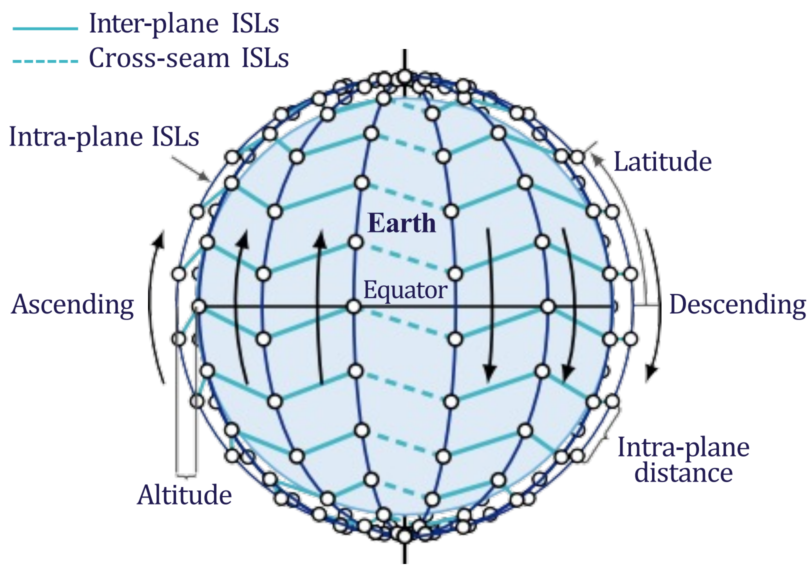

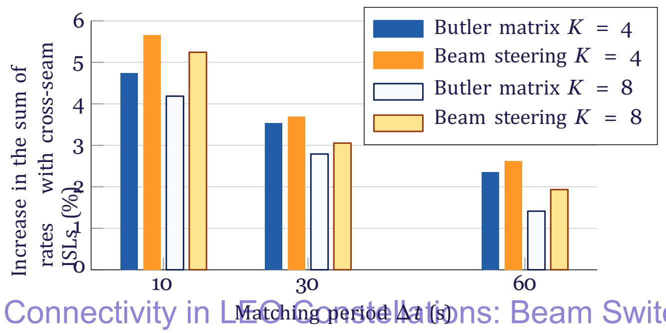

Leyva-Mayorga, Röper, Matthiesen, Dekorsy, Popovski, Soret, “Inter-Plane Inter-Satellite Connectivity in LEO Constellations: Beam Switching vs. Beam Steering,” Globecom 2021.

satellite-to-ground communication

- LOS channel, low rank
- distributed beamforming creates virtual array
- problem: propagation delay
  - Sat-to-Ground RTT: ≈ 4 ms
  - Sat-to-Sat: 50 km intersatellite distance ≙ 0.17 ms
- exploit position knowledge → beamspace MIMO
- AoA & AoD based precoding & linear equalization
- perfect position knowledge:
  99.8 % of optimal beamforming

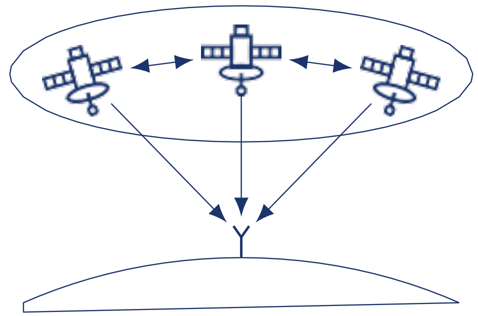

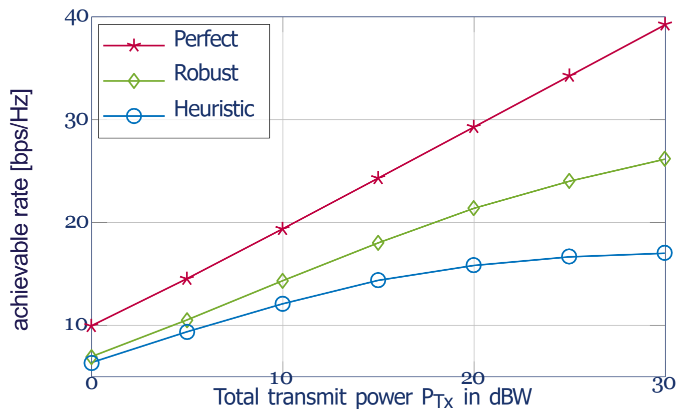

Röper, Matthiesen, Wübben, Popovski, Dekorsy, “Beamspace MIMO for Satellite Swarms,” WCNC 2022

satellite-to-ground communication

- Reconfigurable Intelligent Surface (RIS) as second path
- exploit predictable position of satellite
- LEO Satellite: Doppler Shift -> multipath-> Doppler Spread
- continuous time propagation model

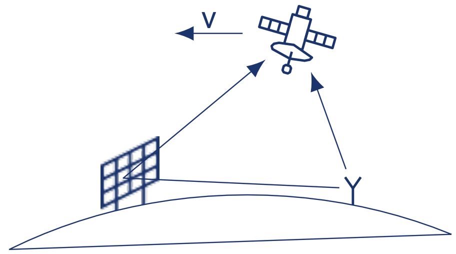

- optimal configuration: Power, Doppler Spread, Delay Spread
- Pareto optimal lexicographic solution: $\phi_{m,n}(t) = 2\pi \bmod(f_c(\tau_0(t) - \tau_{m,n}(t), 1)$
  - maximizes received power
  - no Doppler spread
  - small delay spread

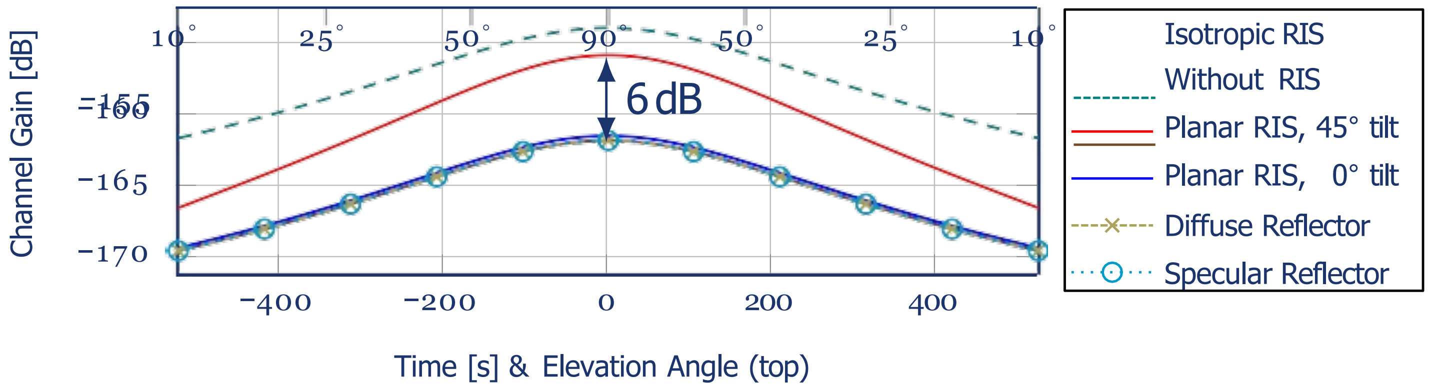

Height Sat: 1500 km
Dist. Tx–RIS: ≈1 km
$f_c = 2$ GHz

Matthiesen, Björnson, De Carvalho, Popovski, "Intelligent Reflecting Surface Operation under Predictable Receiver Mobility: A Continuous Time Propagation Model", IEEE WCL 2021.

outline

- towards 6G
- revival of satellite connectivity
- satellite constellations
- **federated learning in space**
- satellite edge computing

federated learning

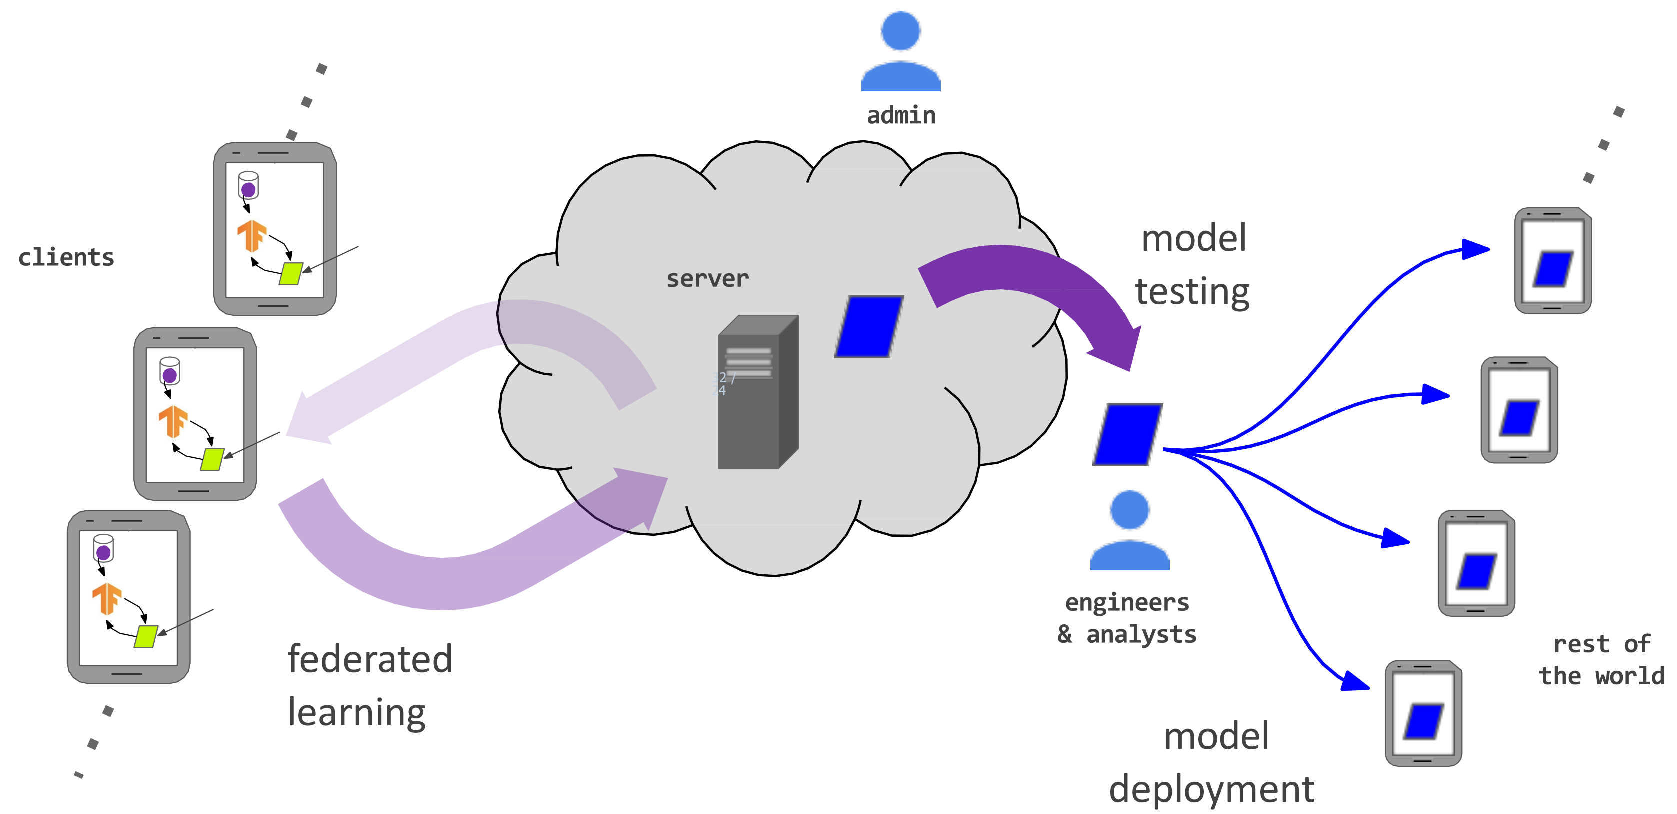

Source: Kairouz, et. al., “Advances and Open Problems in Federated Learning,” NOW Publishers, arXiv:1912.04977.

federated learning

- **privacy:** raw training data remains local
- **non-IID:** local dataset not representative of population distribution
- **unbalanced:** varying amounts of local training data
- **massively distributed:** #devices > # local data points
- **limited communication:** random device participation

total population size	10^6–10^{10} devices
devices selected for one round of training	50 – 5000
total devices that participate in training one model	10^5–10^7
number of rounds for model convergence	500 – 10000
wall-clock training time	1 – 10 days

table 1: order-of-magnitude sizes for typical cross-device federated learning applications.

Source: Kairouz, et. al., “Advances and Open Problems in Federated Learning,” NOW Publishers, arXiv:1912.04977.

federated optimization

- centralized machine learning: solve $\min_{\boldsymbol{w}\in\mathbb{R}^d} \frac{1}{n}\sum_{i=1}^{n} f_i(\boldsymbol{w})$
  - n data points
  - data set $\{\boldsymbol{x}_i, y_i\}_{i=1}^{n}$
  - f_i cost function of ith point, e.g. quadratic

- distributed ML / optimization: data center
  - K clients
  - partition data set and distribute to clients
  - distributed solution: heavy on communications

- federated optimization
  - natural partition of data set: D_k
  - $n_k = |D_k|$
  - K large, n_k unbalanced

$$\min_{\boldsymbol{w}\in\mathbb{R}^d} \sum_{k=1}^{K} \frac{n_k}{n} \cdot \frac{1}{n_k} \sum_{i\in\mathcal{D}_k} f_i(\boldsymbol{w}) = \min_{\boldsymbol{w}\in\mathbb{R}^d} \sum_{k=1}^{K} \frac{n_k}{n} F_k(\boldsymbol{w})$$

Konečny, McMahan, Ramage, “Federated Optimization: Distributed Optimization Beyond the Datacenter,” arXiv:1511.03575, 2015.

synchronous and asynchronous algorithms

synchronous model

- clients work in the same model
- update after
  all clients have delivered
- waiting and latency

McMahan, Moore, Ramage, Hampson, Aguera y Arcas, “Communication-efficient learning of deep networks from decentralized data,” AISTATS, 2017.

asynchronous model

- ClientUpdate:
  - wait for task
  - run local SGD
  - return result and timestamp
- clients work on
  different model versions
- updates whenever results arrive

Xie, Koyejo, Gupta, “Asynchronous Federated Optimization,” OPT2020, arXiv:1903.03934, 2020.

setup for federated learning with satellites

- **privacy:** raw training data remains local,
  but privacy is not the motivation
- **non-IID:** sometimes.
- **unbalanced:** sometimes
- **massively distributed:** orders of magnitude less devices
- **limited communication:**
  - deterministic device participation
  - long delay, high transmission costs, limited energy
  - no control over device availability
- **control:** devices owned by operator

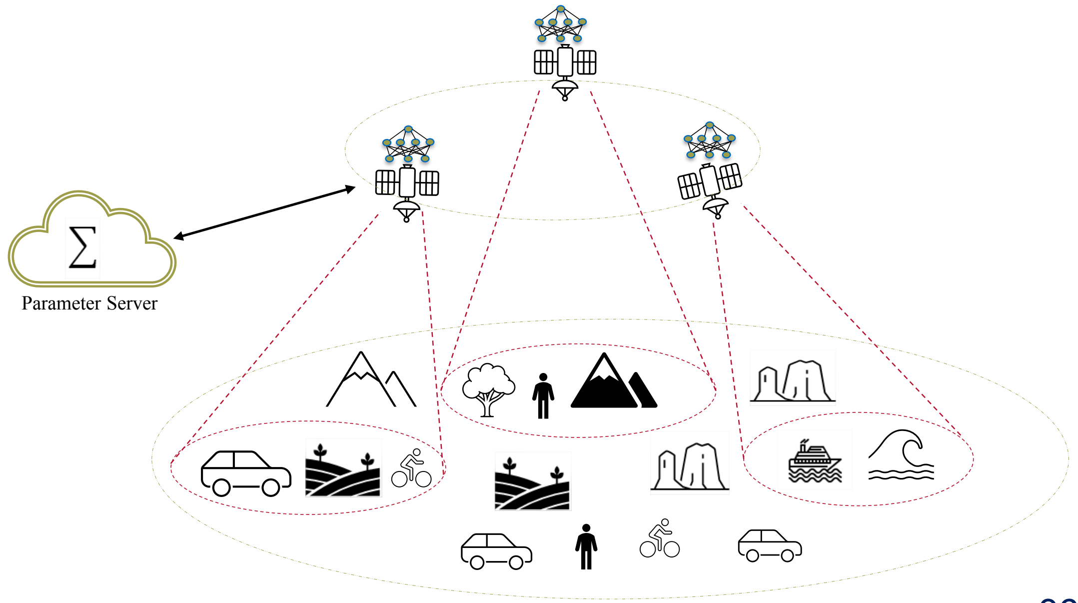

two generic options

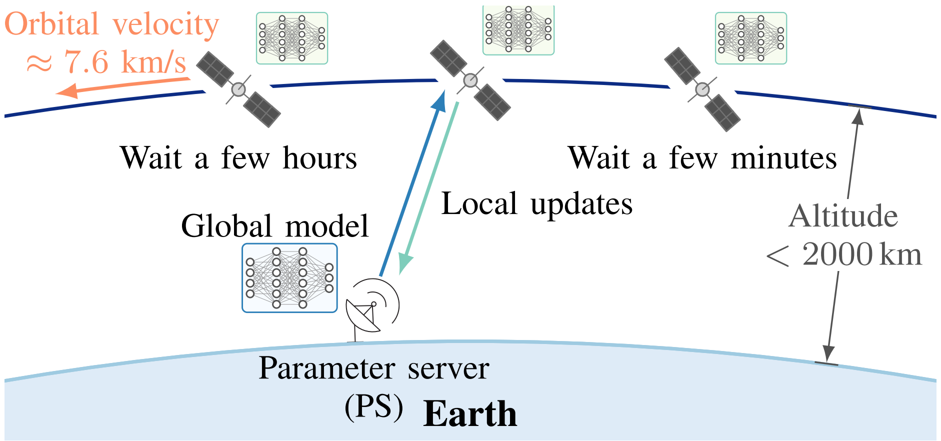

without inter-satellite links

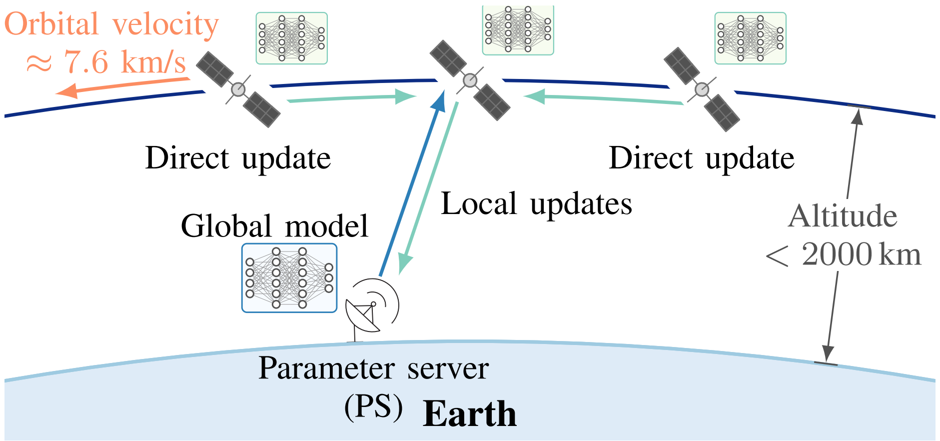

with inter-satellite links

ground-assisted federated learning

- single ground-station as server
- no inter-satellite communication
- data exchange during pass
- run local SGD during offline time
- time between contacts
  - orbital period: ≈ 90 min to 128 min
  - behind horizon: ≲ 12 h

- distinctive features
  - "not learning" not an option
  - full client participation
- synchronous learning:
  - 1 - 2 orbital periods per global epoch
  - asynchronous learning

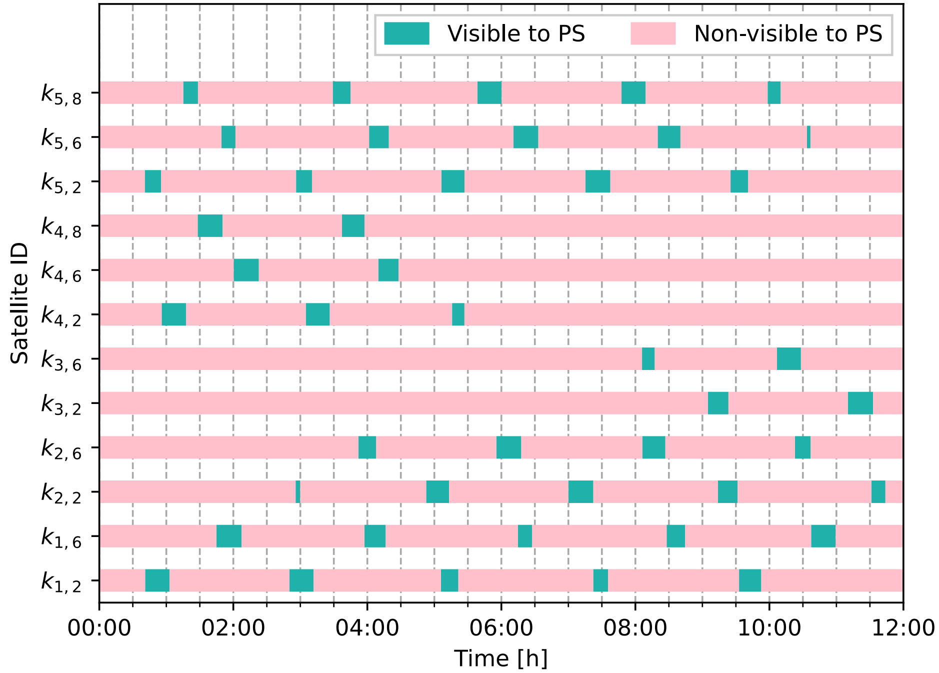

Razmi, Matthiesen, Dekorsy, Popovski, "Ground-Assisted Federated Learning in LEO Satellite Constellations," IEEE WCL, 2022.

FedSat: asynchronous FedAvg

idea:

- GS at North pole, single orbital shell → symmetric
- cyclic contact sequence: 1→2→3→⋯→K→1→...
- FedAvg update rule: $w_{t+1} = \sum_{i=1}^{K} \frac{n_k}{n_{t+1}} w_{t+1}^{i}$
- "Unroll" FedAvg: Incremental update rule
  - Satellite k visits at $t_{i_1}, t_{i_2}, \ldots$
  - At t_{i_2}: $w_{i_2+1} = w_{i_2} - \frac{n_k}{n}(w_{i_1}^k - w_{i_2}^k)$
  - After K iterations: Same as w_{t+1} of FedAvg

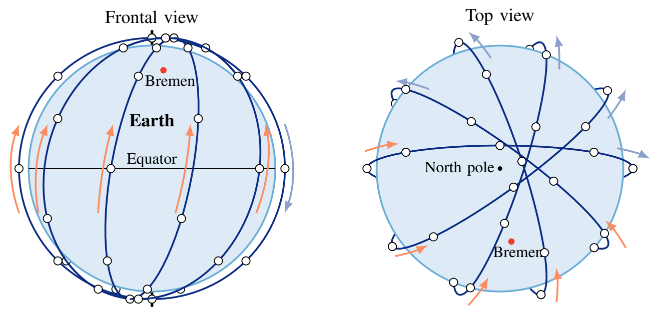

algorithm:

- satellite k transmits weight update $\Delta w^k = n_k(w_{i_1}^k - w_{i_2}^k)$
- GS updates global model $w_{i+1} \leftarrow w_i - \frac{1}{n}\Delta w^k$
- GS sends w_{i+1} to satellite k

convergence:

• established (Nedić et. al. 2001): single orbital shell, arbitrary GS location

• open: multiple orbital shells

FedSat: Numerical Results

Top-1 accuracy for a GS in Bremen with Non-IID CIFAR data

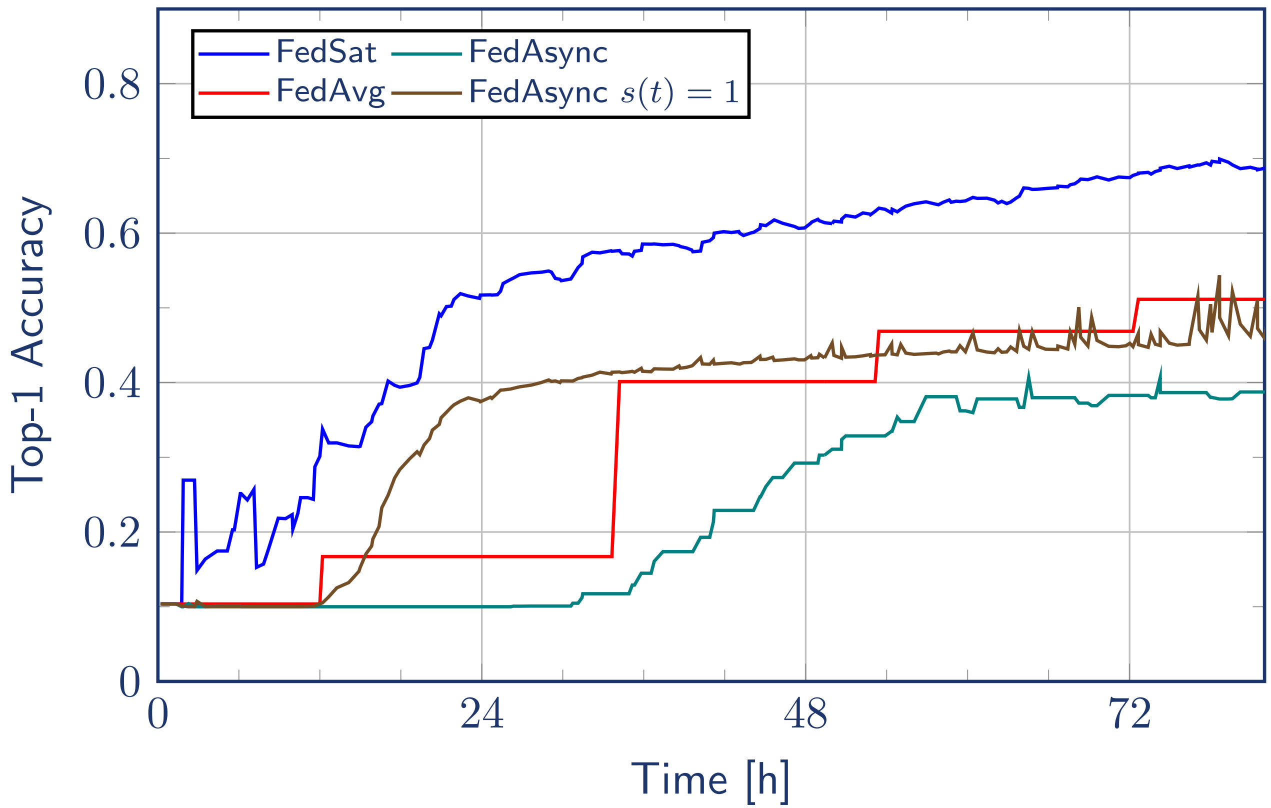

scheduling for FedSat

previous assumptions:

- data exchange during pass
- run local SGD during offline time
- → is it possible during pass?

algorithm:

- after delivering model update:
- next pass long enough for computation?
- **Yes:** receive model at next visit, work online
- **No:** receive model now, work offline

result:

- reduces model staleness
- improves convergence

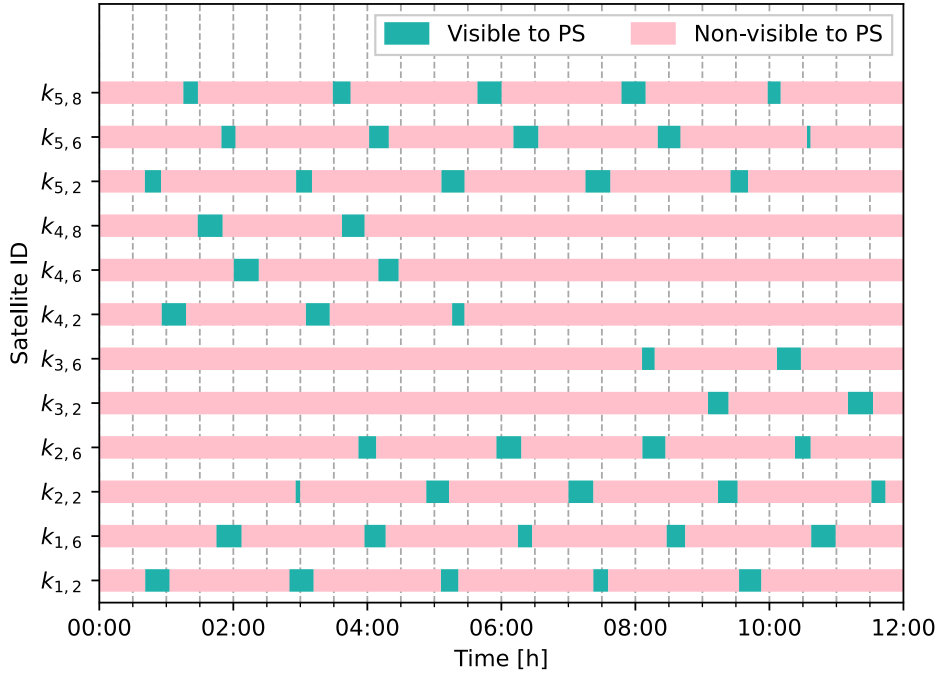

Razmi, Matthiesen, Dekorsy, Popovski, "Scheduling for Ground-Assisted Federated Learning in LEO Satellite Constellations," EUSIPCO, 2022.

scheduling for FedSat: results

Top-1 accuracy for a GS in Bremen with Non-IID CIFAR data

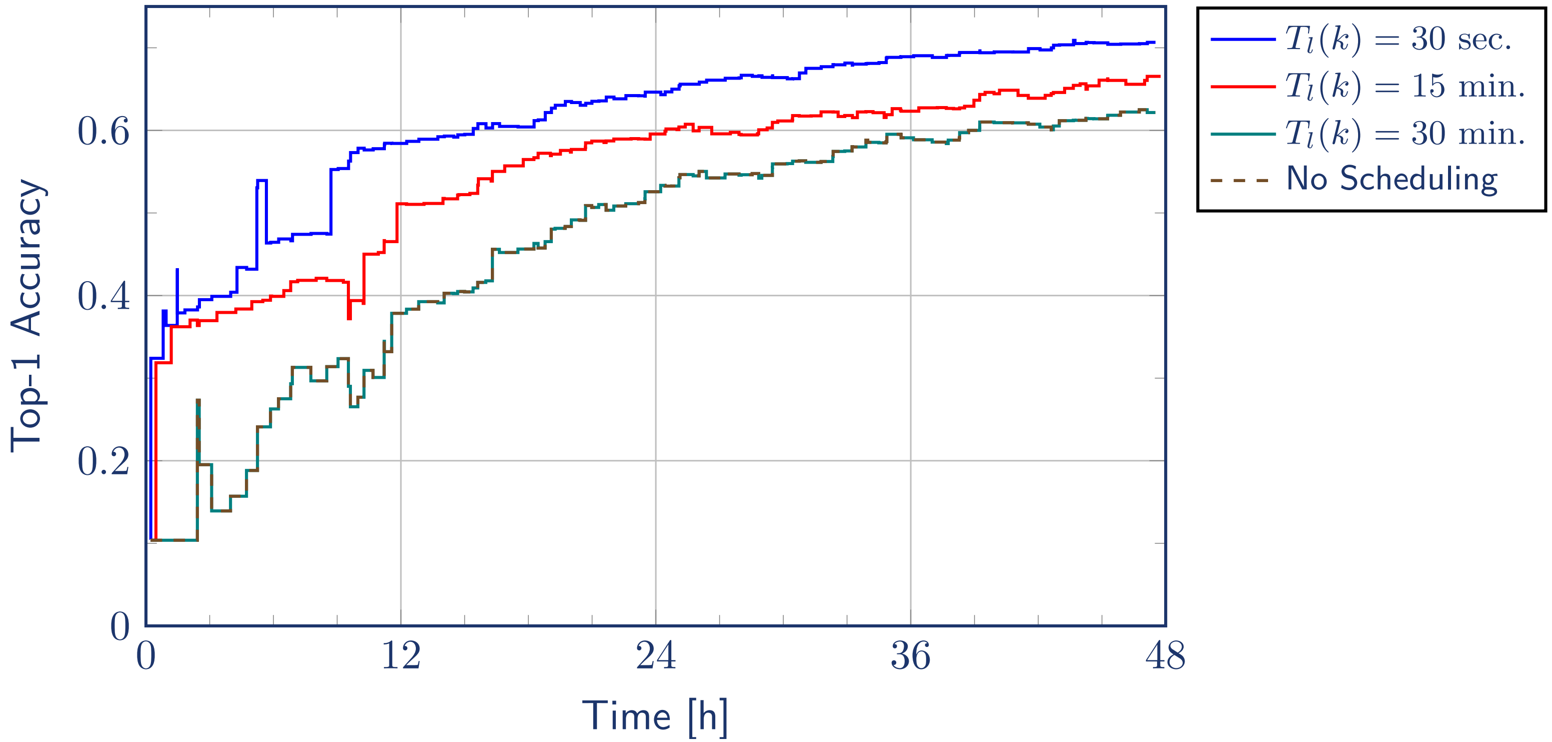

federated learning with inter-satellite links

Intra-Plane Inter-Satellite Links (ISLs):

- connects adjacent satellites within orbital plane
- stable relative position

idea:

- one satellite per orbit connects to PS (GS, MEO, GEO)
- multi-hop intra-orbit routing
- predictive routing determines sink satellite (per orbit)

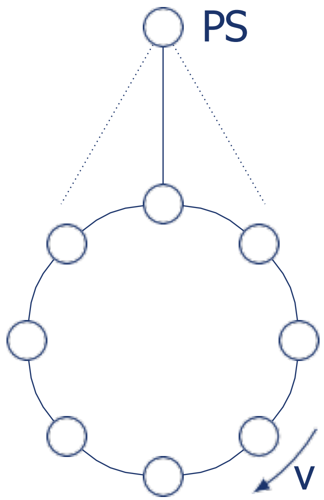

Razmi, Matthiesen, Dekorsy, Popovski, "On-Board Federated Learning for Dense LEO Constellations," IEEE ICC, 2022.

types of connectivity with inter-satellite

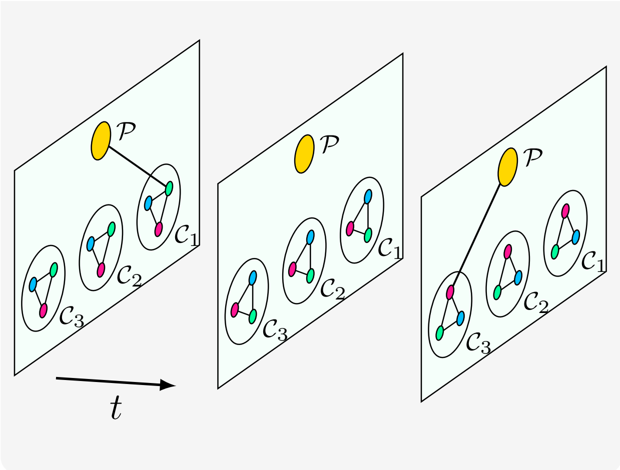

sporadic direct connection to PS

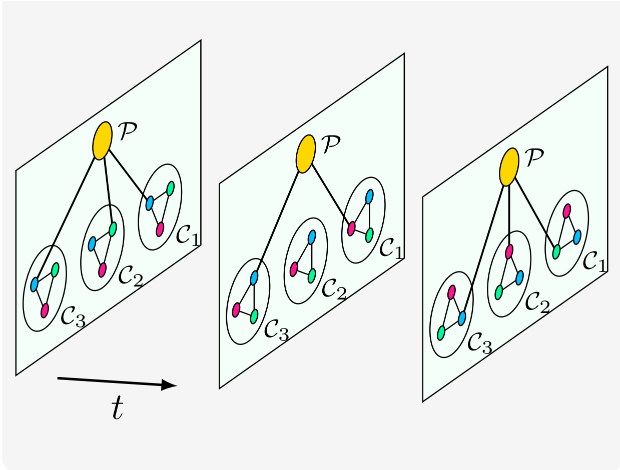

Near-persistent direct connection to PS

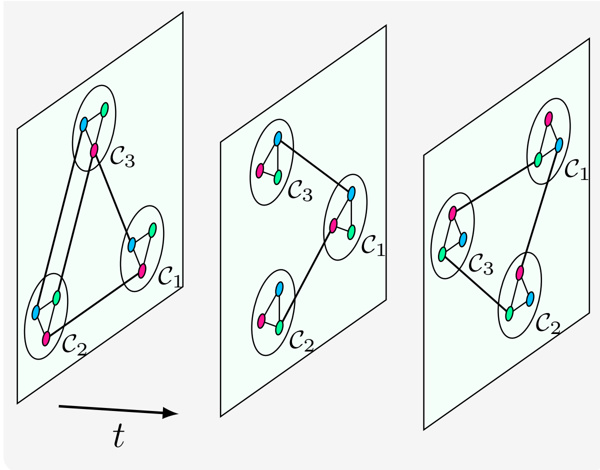

Multi-hop connection to PS via inter-cluster connectivity

B. Matthiesen, N. Razmi, I. Leyva-Mayorga, A. Dekorsy, and P. Popovski, "Federated Learning in Satellite Constellations", in IEEE Network Magazine, accepted, 2023.

federated learning with inter-satellite links

algorithm:

- parameter distribution:

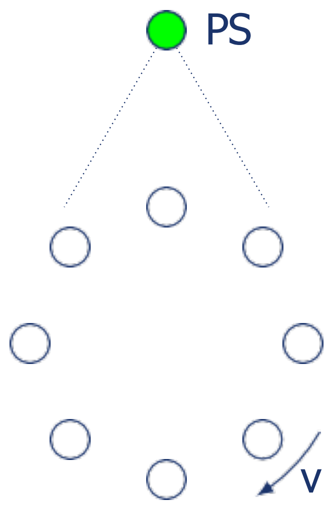

time: t_{now}

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS

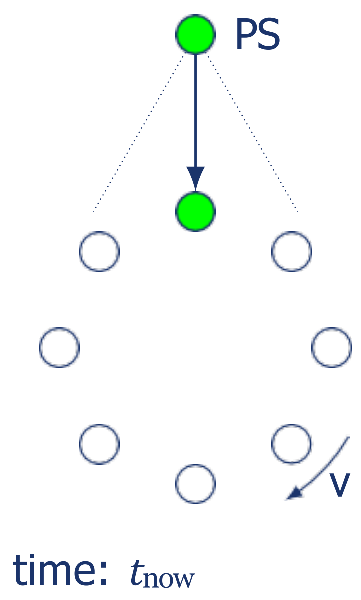

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning

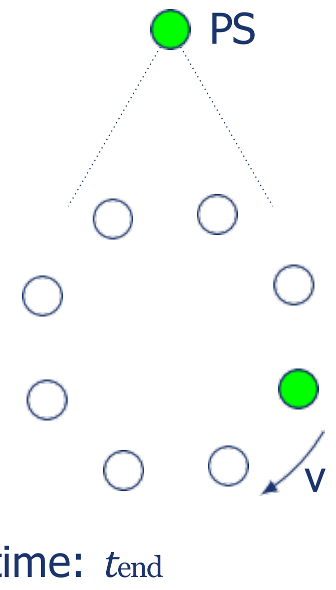

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}

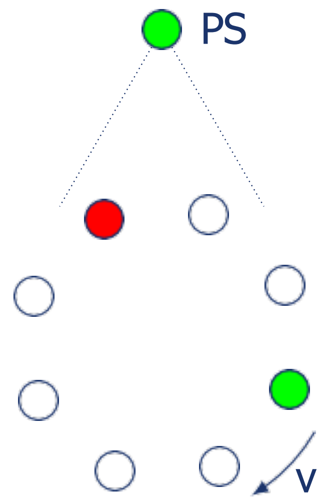

time: t_{end}

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL

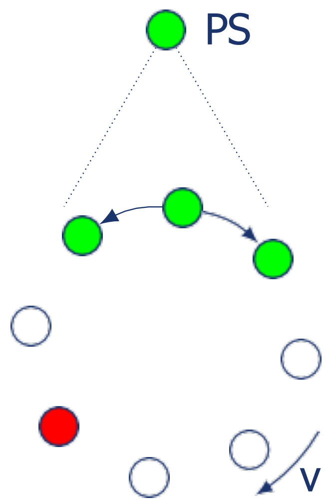

time: $t_{\text{now}} + T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL

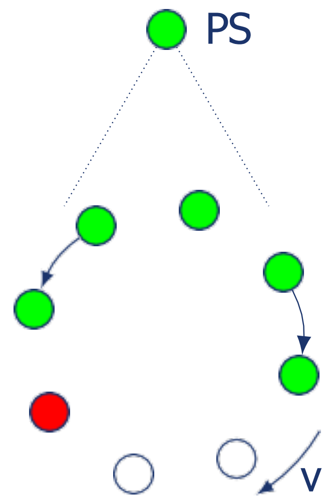

time: $t_{\text{now}} + 2T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL

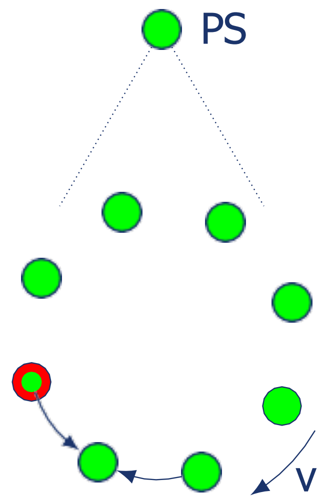

time: $t_{\text{now}} + 4T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights

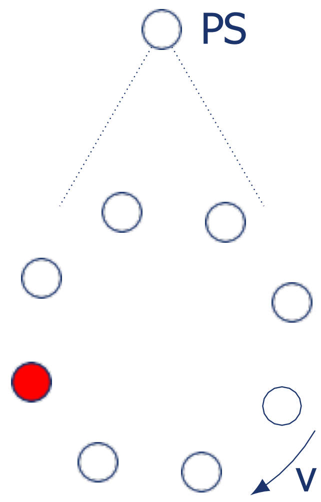

time: $t_{\text{now}} + 4T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path

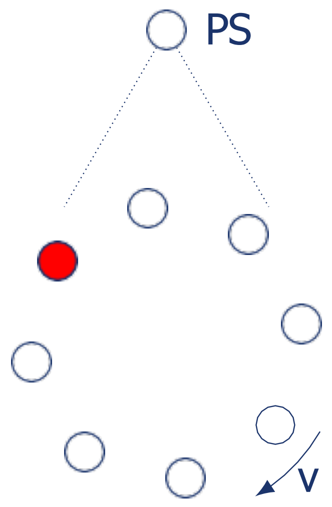

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path

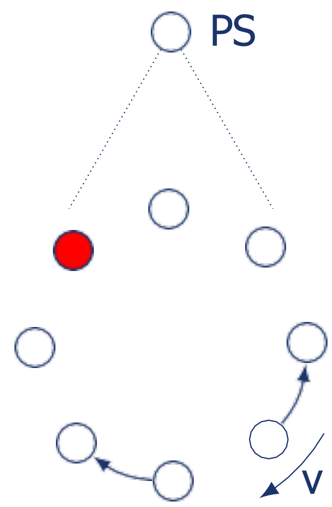

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path
  - sink satellite
    - wait for all the results

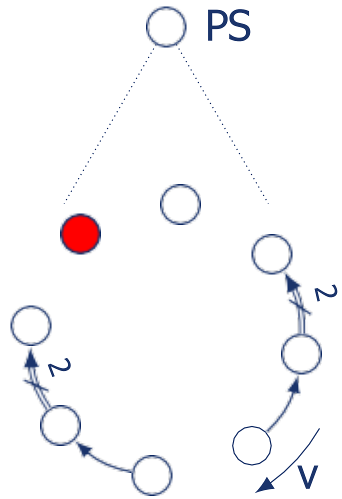

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + 2T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path
  - sink satellite
    - wait for all the results

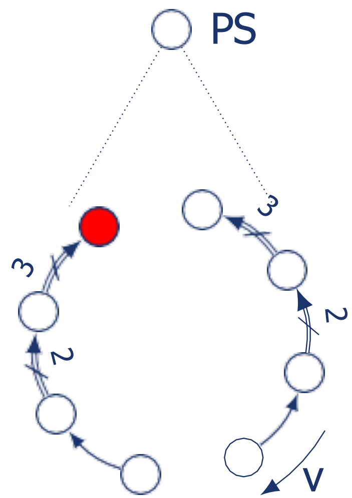

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + 3T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path
  - sink satellite
    - wait for all the results

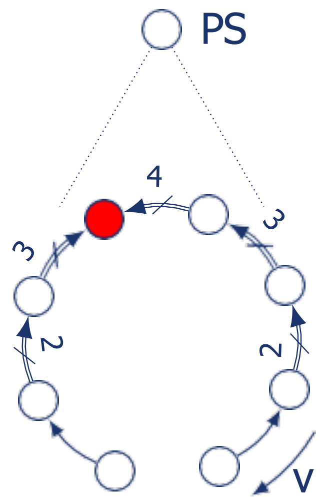

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + 4T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path
  - sink satellite
    - wait for all the results
    - forward to the PS on the ground

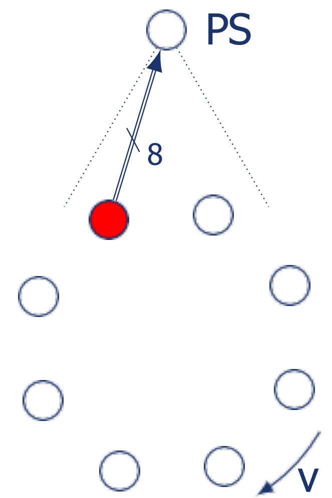

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + 4T_{c,p}$

federated learning with inter-satellite links

algorithm:

- parameter distribution:
  - 1 satellite per orbit: get weight from PS
  - estimate constellation state at time t_{end}, based on expected times for communication, processing, and learning
  - select the satellite with the best connection at t_{end}
  - distribute parameters and selected sink satellite through ISL
- computation: every satellite updates weights
- aggregation
  - after computation, send update to sink satellite over shortest path
  - sink satellite
    - wait for all results
    - forward to the PS on the ground

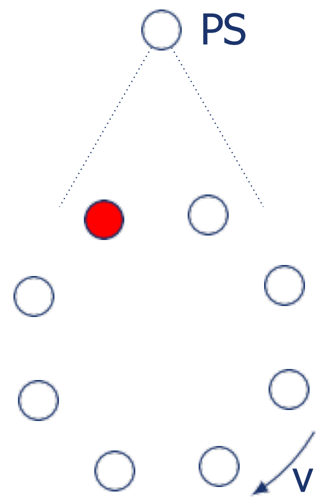

time: $t_{\text{now}} + 4T_{c,p} + T_{L,p} + 4T_{c,p} + T_{PS\text{-}ground}$

result

- short (or no) offline period per orbit
- synchronous federated learning (FedAvg)

incremental aggregation

central aggregation:

- PS receives all weights w_{t+1}^{i}
- computes new global weights $\boldsymbol{w}_{t+1} = \sum_i \frac{n_k}{n} \boldsymbol{w}_{t+1}^{\boldsymbol{i}}$
- commmunication effort scales as $O\left(\left\lfloor \frac{K^2}{2} \right\rfloor\right)$ per orbit

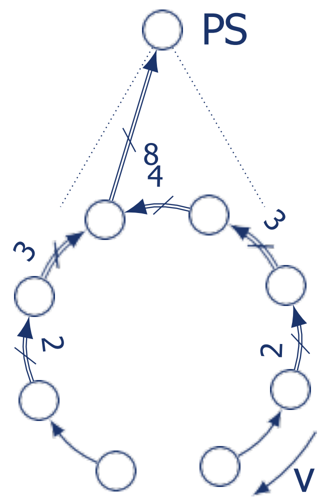

incremental aggregation

central aggregation:

- PS receives all weights w_{t+1}^{i}
- computes new global weights $\boldsymbol{w}_{t+1} = \sum_i \frac{n_k}{n} \boldsymbol{w}_{t+1}^{\boldsymbol{i}}$
- commmunication effort scales as $O\left(\left\lfloor \frac{K^2}{2} \right\rfloor\right)$ per orbit

incremental aggregation:

- satellite i collects incoming weights in $\mathcal{I}_{t+1}^{i}$
- transmits $\boldsymbol{w}_{t+1}^{i,out} = n_k \boldsymbol{w}_{t+1}^{\boldsymbol{i}} + \sum \boldsymbol{w}$
- PS receives one weight per orbit
- PS computes a single update based on all received weights
- exactly K transmissions per orbit

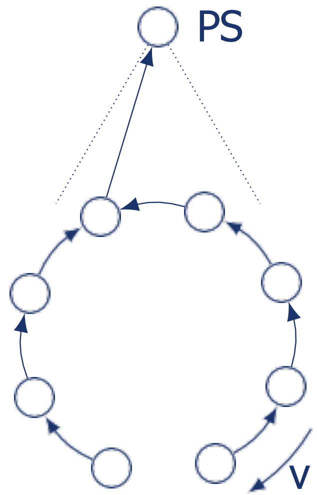

numerical results

ResNet-18 with non-IID CIFAR-10 data.

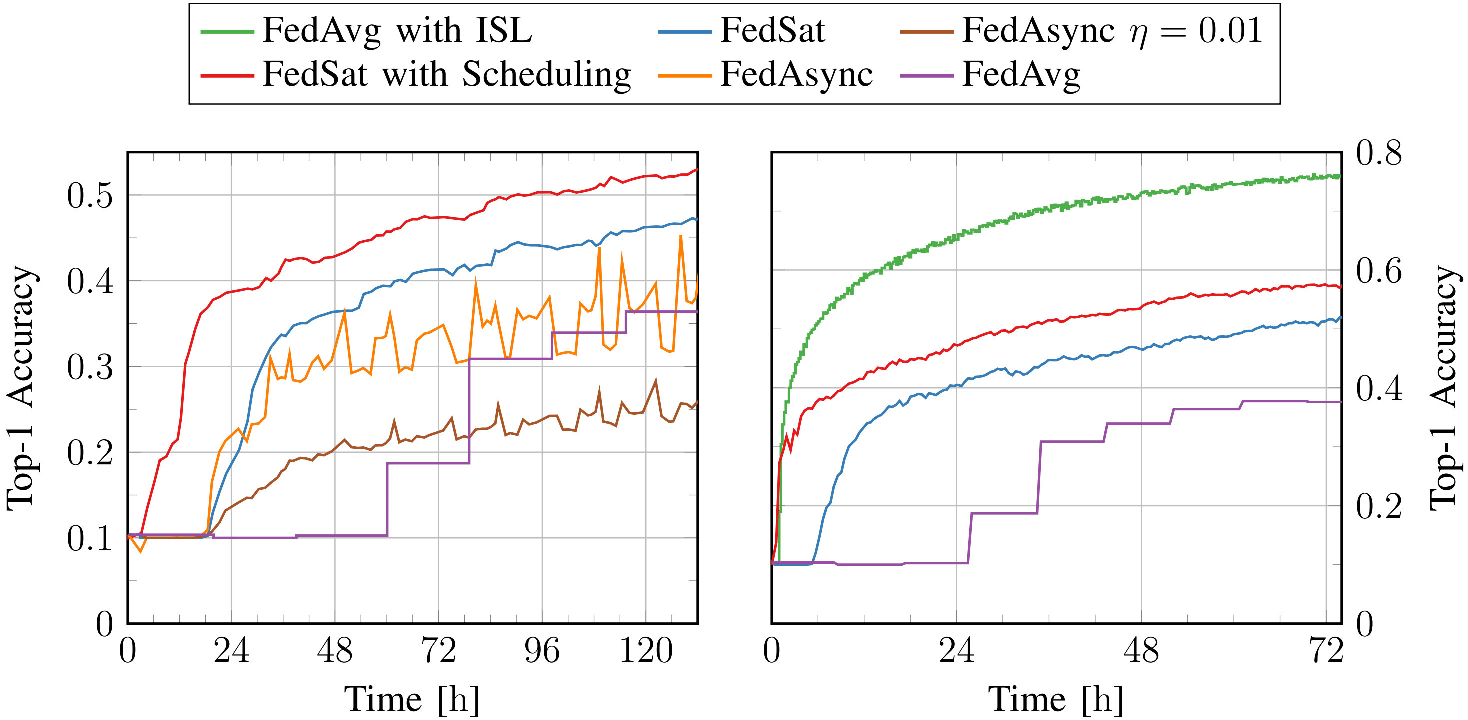

PS located in a GS in Bremen, Germany. No inter-satellite connectivity.

PS located in a satellite in equatorial orbit at 2000 km altitude.

outline

- towards 6G
- revival of satellite connectivity
- satellite constellations
- federated learning in space
- **satellite edge computing**

global intelligence and edge computing

data source
an IoT device on ground, a satellite, or gateway

edge infrastructure
gateways and/or satellites

computing resources
mobile terminals, gateways, satellites, or the cloud
assumed to have the algorithms/program

destination
mobile terminal, satellite, or cloud

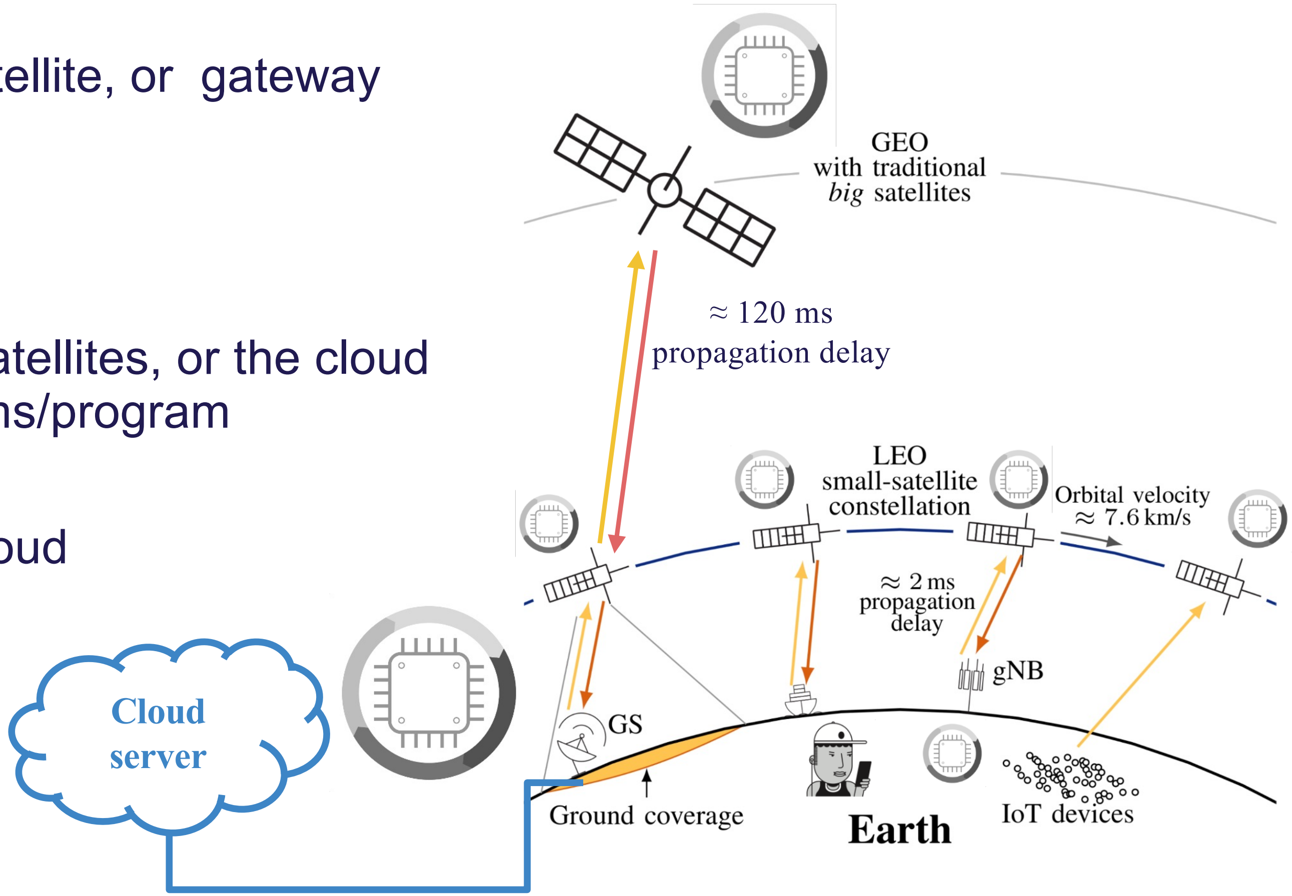

traditional-style mobile edge computing (MEC)

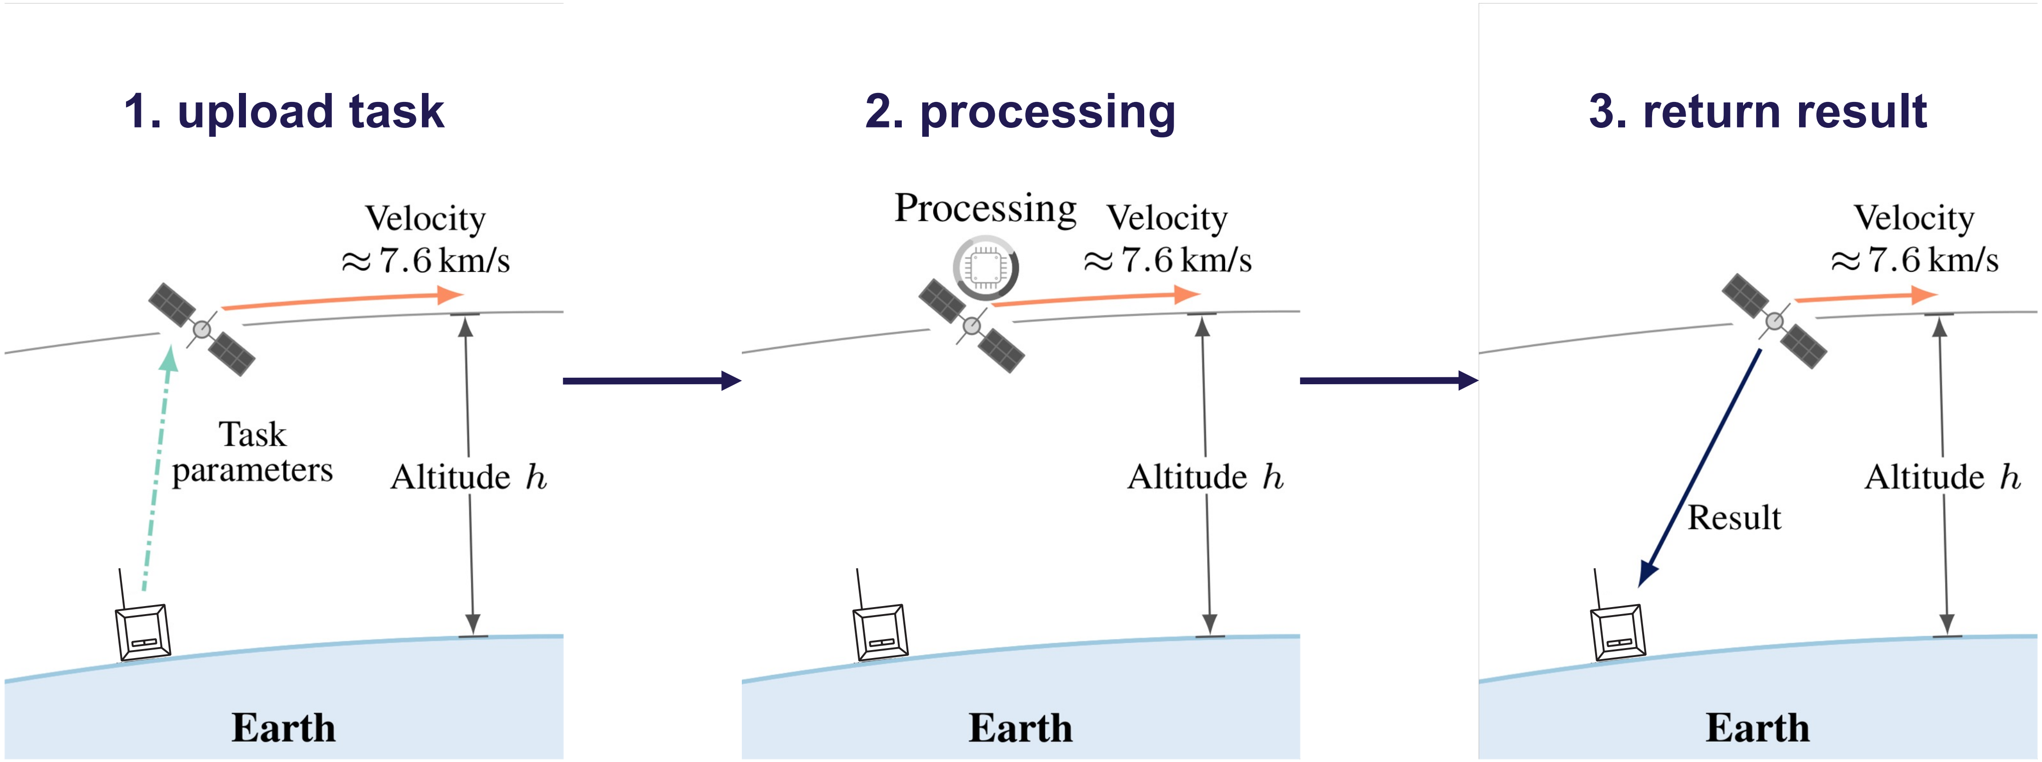

distributed MEC

1. segmentation: partition the data
2. allocation: segment-to-satellite allocation
3. scatter: transmission of the segments
4. processing: each satellite in parallel
5. gather: send the result to the destination

scatter and gather: limited by the connectivity

orbital planes
Ring topology with stable links

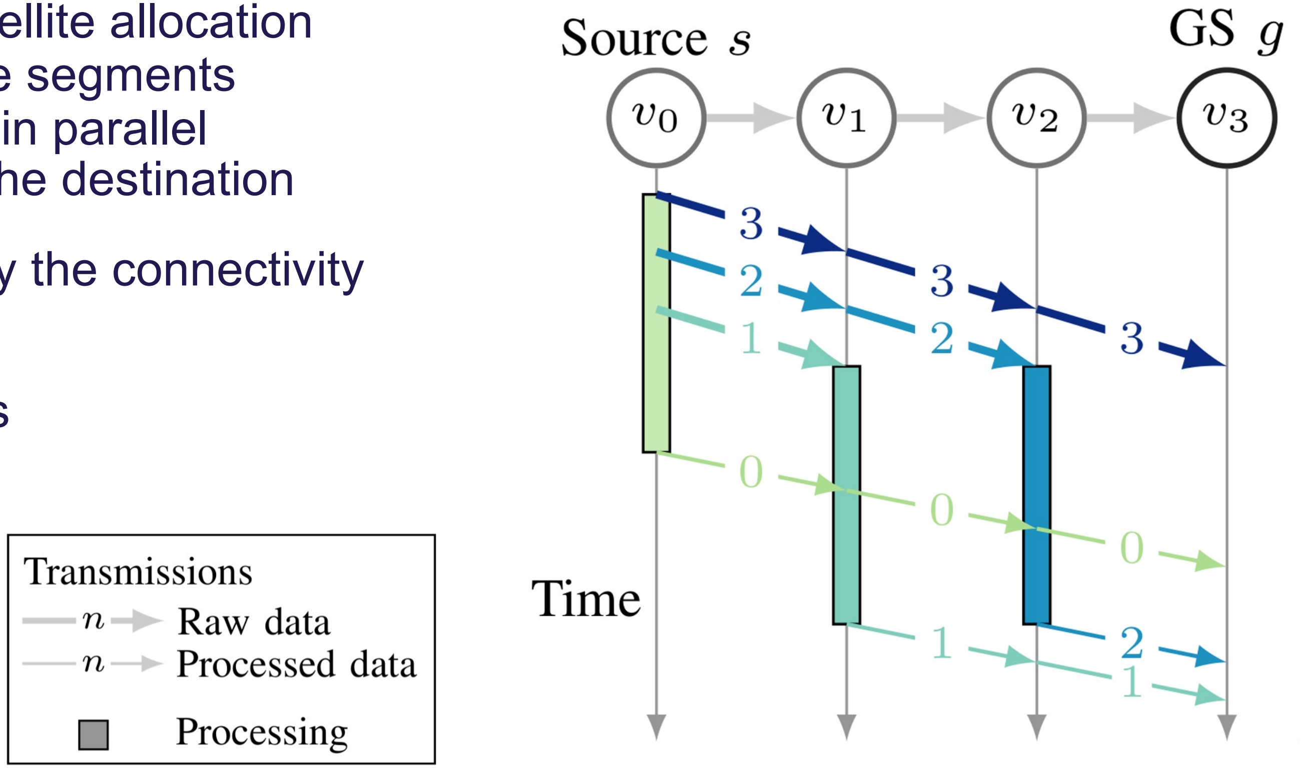

optimization objectives

aimed for feasibility, efficiency, or stability
minimize latency
or
Minimize energy consumption
subject to real-time constraints

feasibility: real-time constraints per task

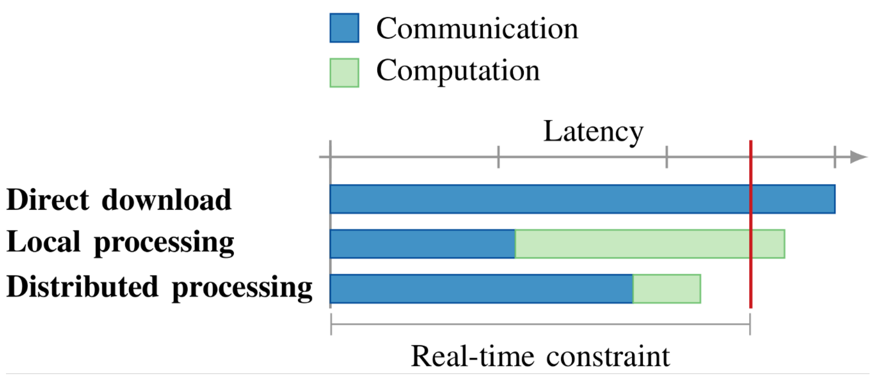

stability: real-time constraints per resource

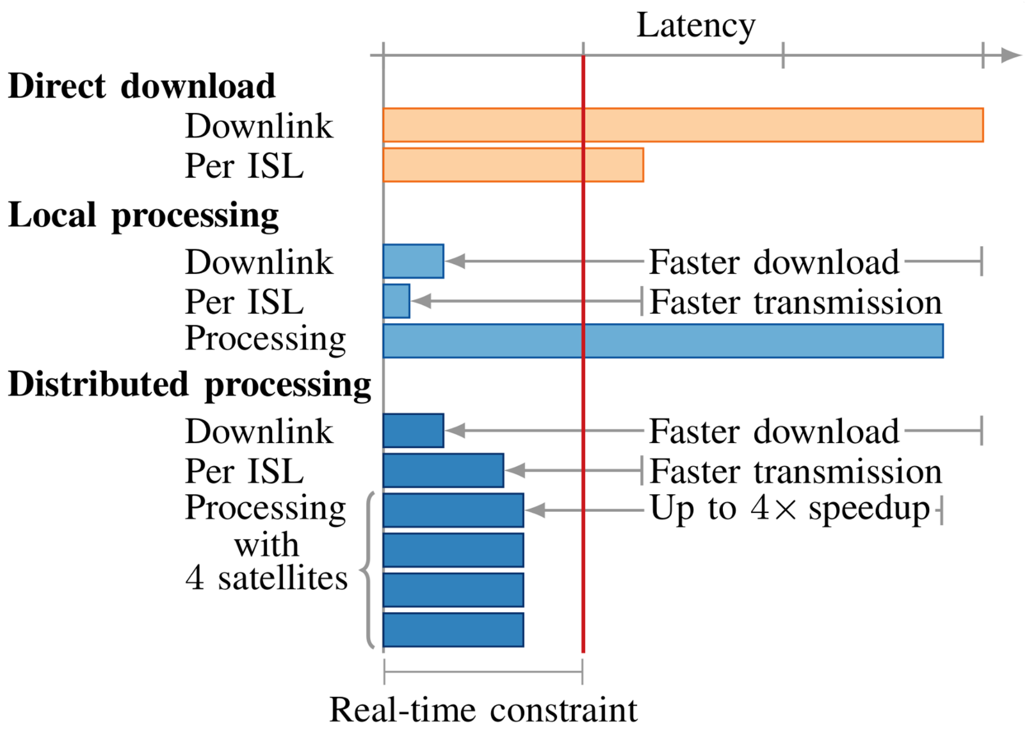

very high-definition Earth observation

data source:
one satellite

computing resources:
at each satellite in the orbit

algorithm:
lossy compression

destination:
gateway

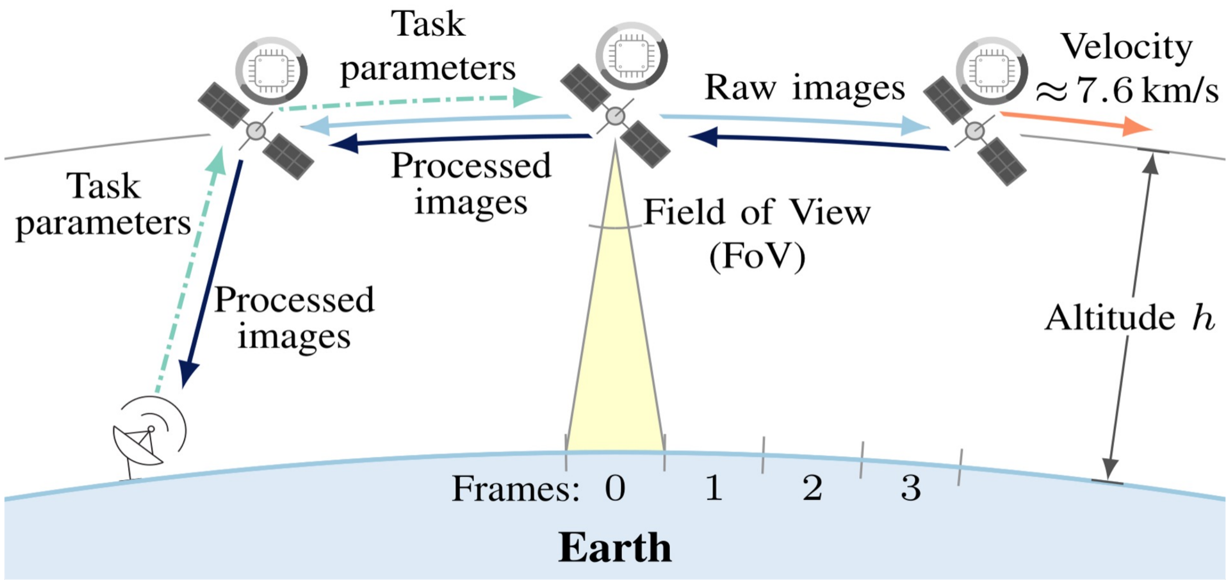

I. Leyva-Mayorga, M.M. Gost, M. Moretti, A. Pérez-Neira, M.Á. Vázquez, P. Popovski, and B. Soret,
“Satellite edge computing for real-time and very-high resolution Earth observation,” IEEE Trans. Commun., 2023.

Earth observation: scanning over K frames

Ground track frame period (GTFP)
Period between consecutive frames
to avoid pixel overlap and coverage holes

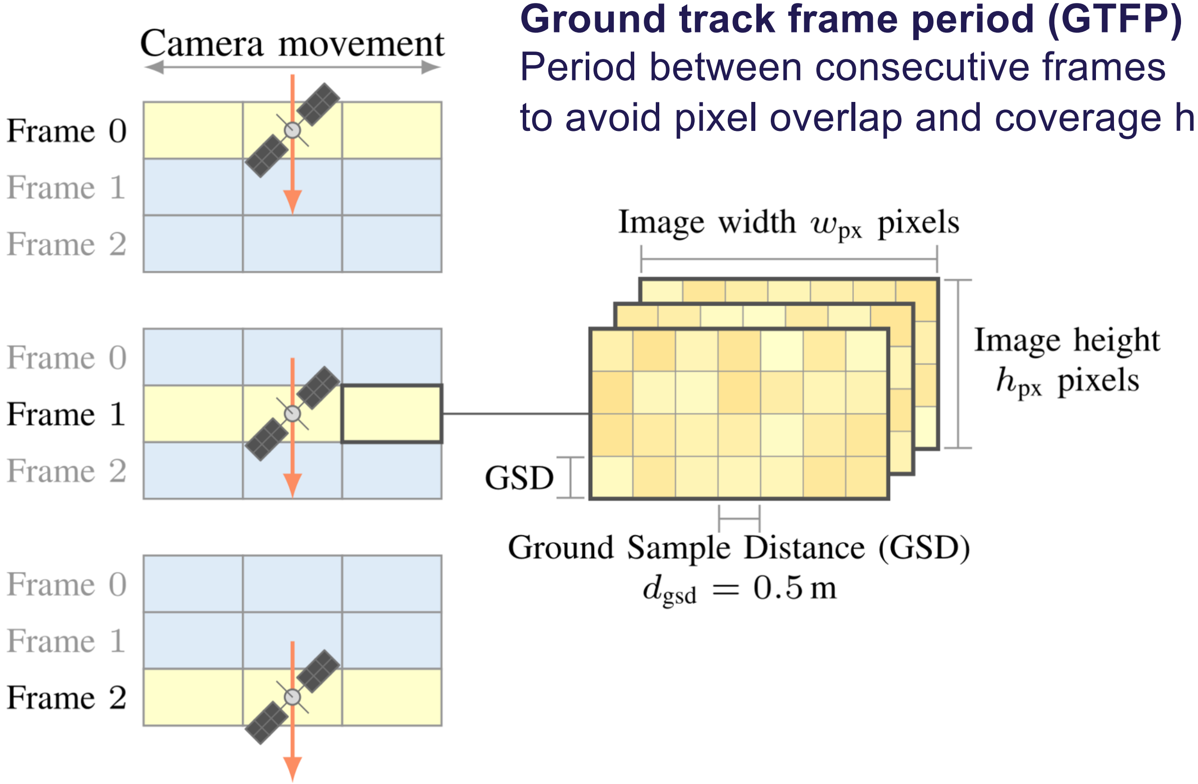

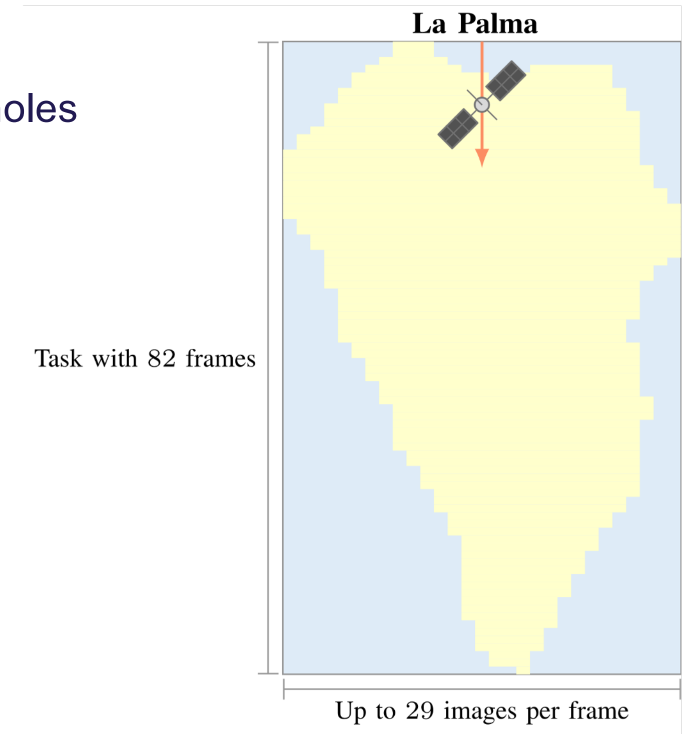

advantages of distributed MEC

Improved system capacity with global minimum energy consumption

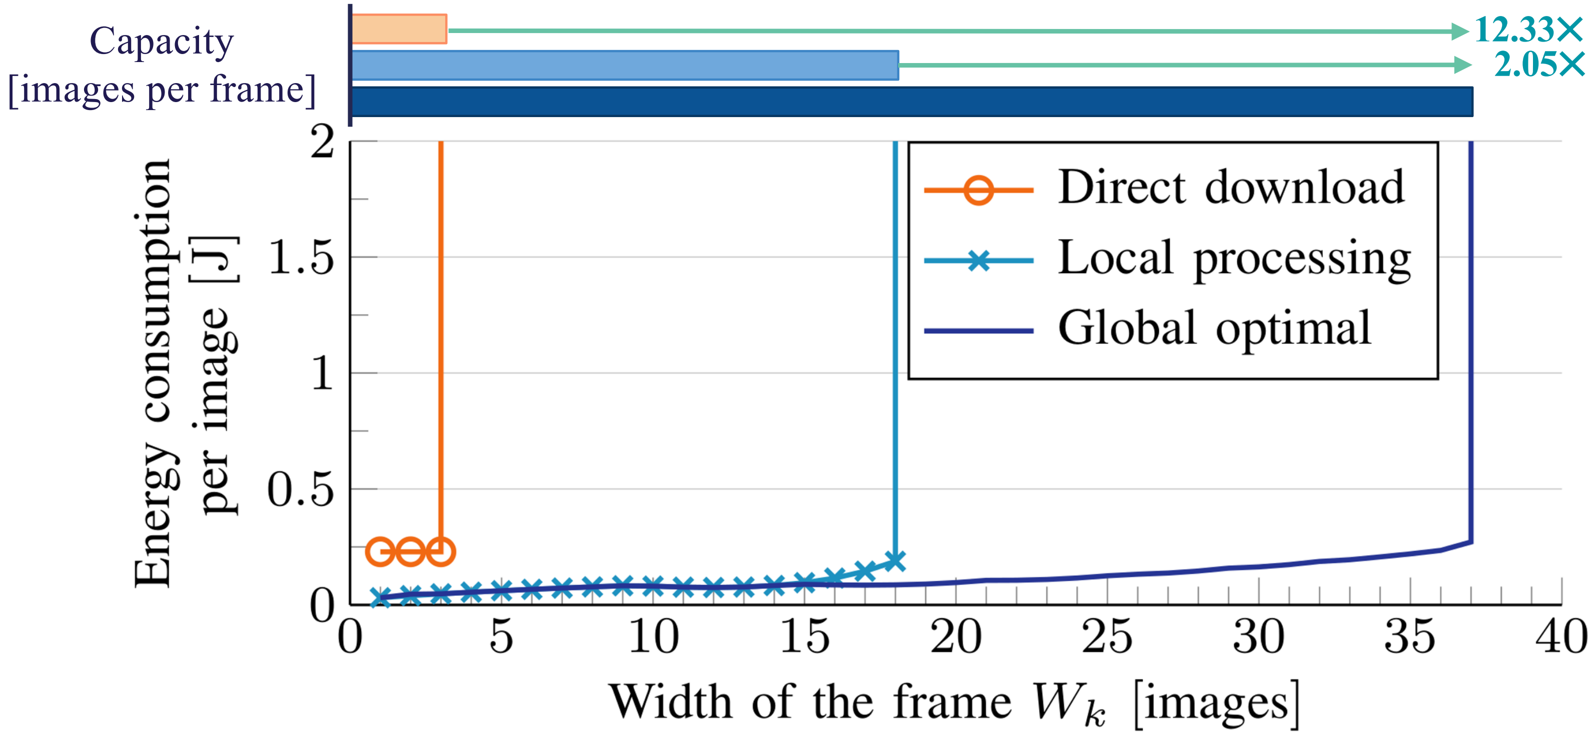

conclusion and outlook

- rekindled interest in satellite connectivity
  - diversified players and equipment
- predictable satellite connectivity requires rethinking of distributed algorithms
- we have built the case for federated learning that operates under predictable satellite connectivity
- plethora of new research problems
  - distributed algorithms, satellite IoT, edge computing with satellites

recent books