# Client Availability in Federated Learning: It Matters!

**Dhruv Garg***, Debopam Sanyal*, Myungjin Lee[◊], Alexey Tumanov, Ada Gavrilovska
Georgia Institute of Technology, [◊]Cisco Research

*5th Workshop on Machine Learning and Systems (EuroMLSys), co-located with EuroSys '25*
*Rotterdam, The Netherlands*

# Federated Learning

- Multiple use-cases
  - Keyboard personalization, virtual assistants

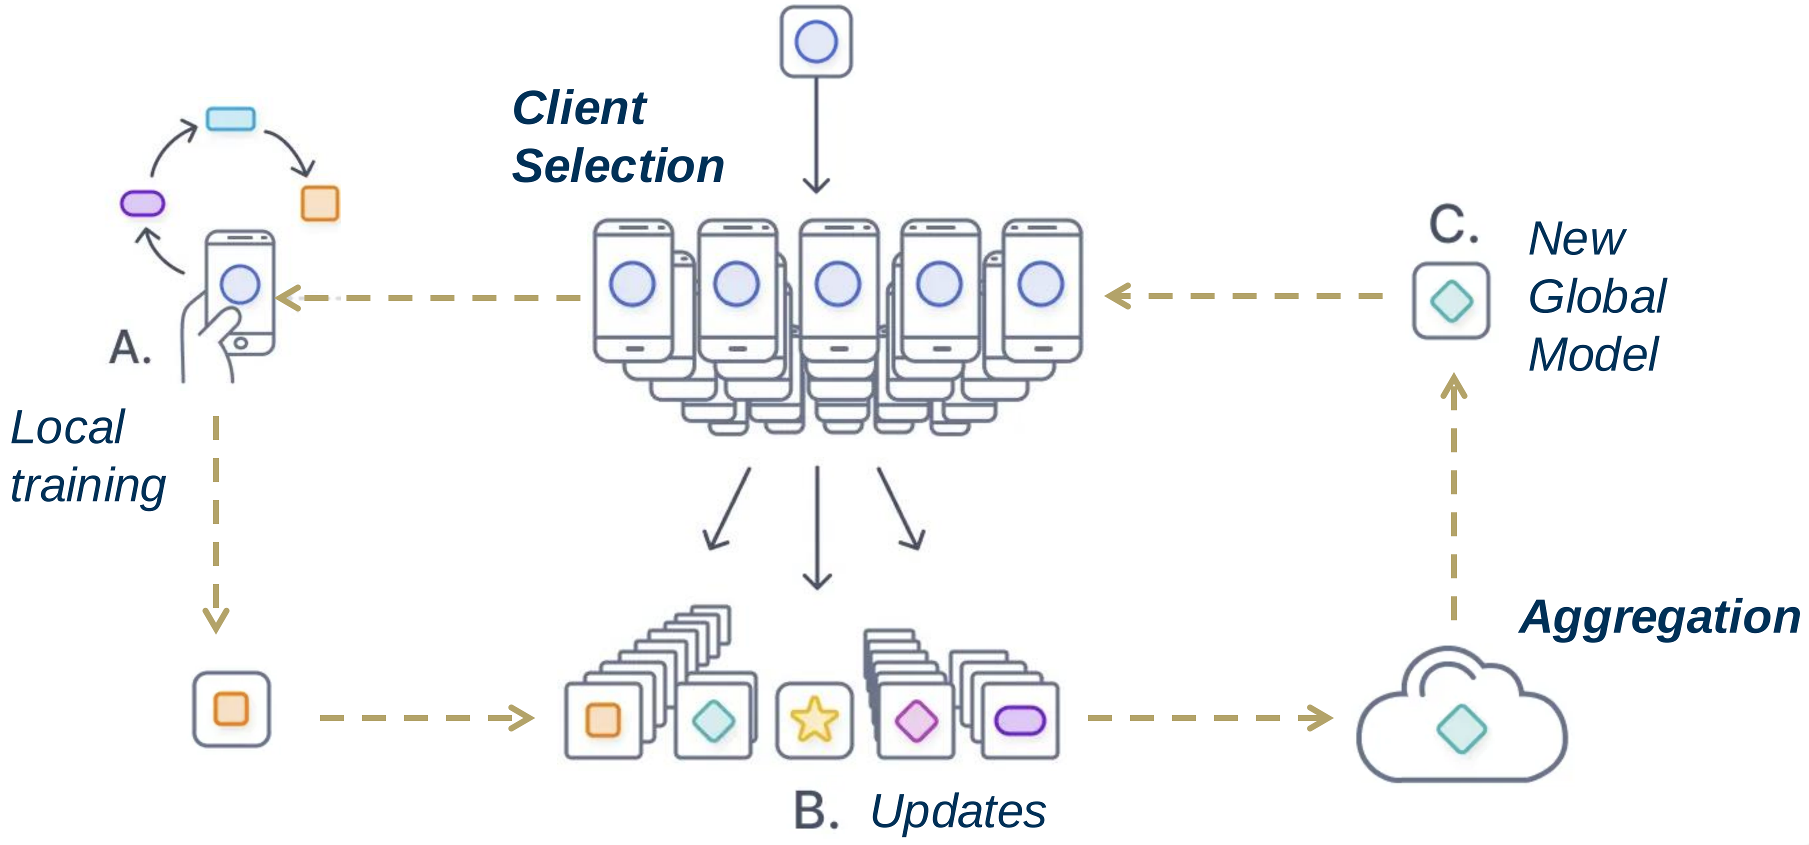




https://research.google/blog/federated-learning-collaborative-machine-learning-without-centralized-training-data/

# Client Dynamics in FL Training

FL jobs run for several ***days***

Clients become unavailable intermittently

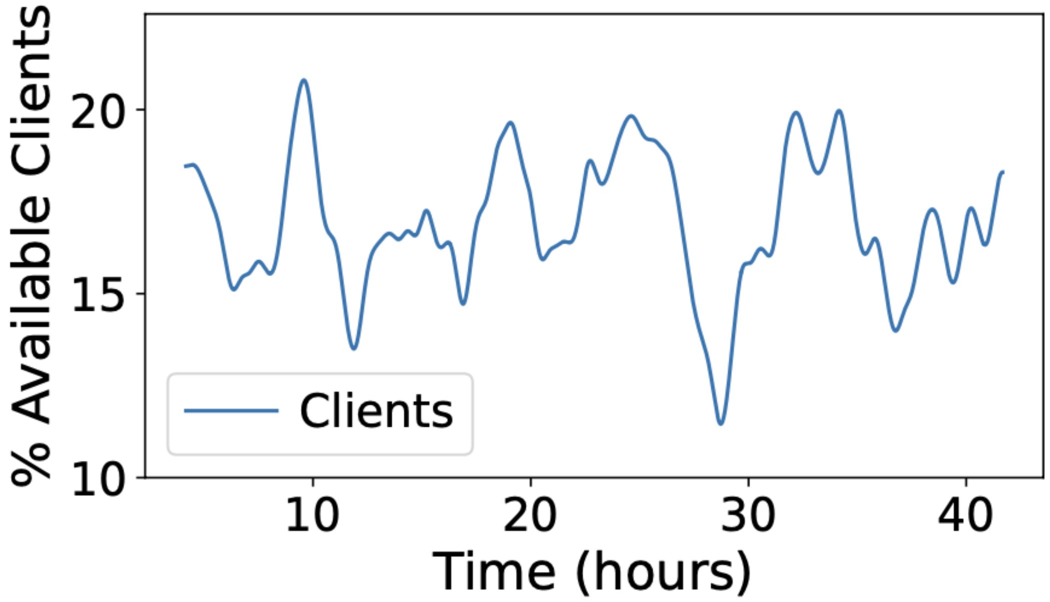


| Trace | FedScale [1] | LinkedIn [2] | Google [3] |
|---|---|---|---|
| Client Availability | 10-20% | 20-80% | 10-60% |


[1] Lai, F., Dai, Y., Singapuram, S., Liu, J., Zhu, X., Madhyastha, H., & Chowdhury, M. (2022, June). Fedscale: Benchmarking model and system performance of federated learning at scale. In International conference on machine learning (pp. 11814-11827). PMLR.
[2] Wang, E., Chen, B., Chowdhury, M., Kannan, A., & Liang, F. (2023). Flint: A platform for federated learning integration. Proceedings of Machine Learning and Systems, 5, 21-34.
[3] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ... & Roselander, J. (2019). Towards federated learning at scale: System design. Proceedings of machine learning and systems, 1, 374-388.

# Is SOTA FL Robust to Realistic Client Unavailability?

Georgia Tech

# Prolonged Straggler = Unavailability?

- ***Straggler*** clients
  - Return updates with a delay

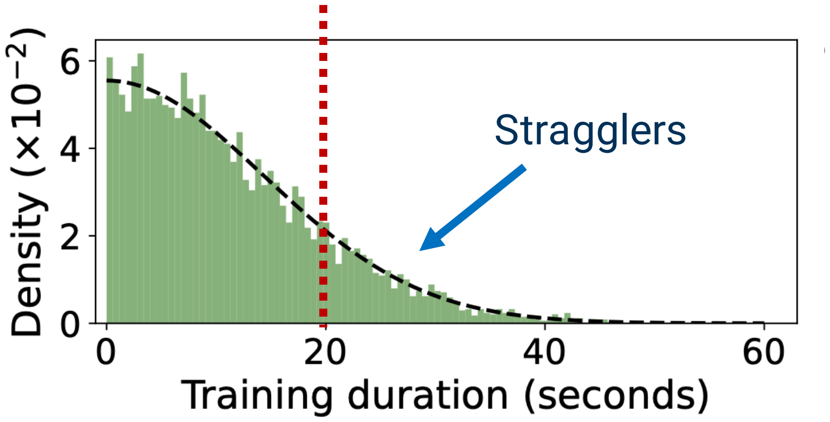


- AsyncFL mitigates stragglers

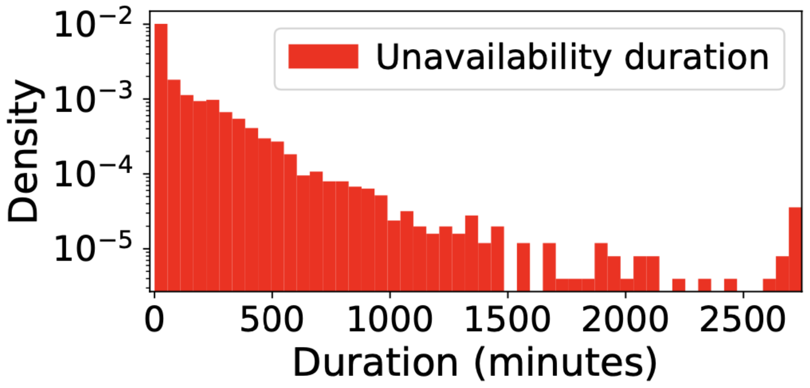


- ***Unavailable*** clients
  - Cannot participate at all

# FL Selection Algorithms Are Availability Unaware

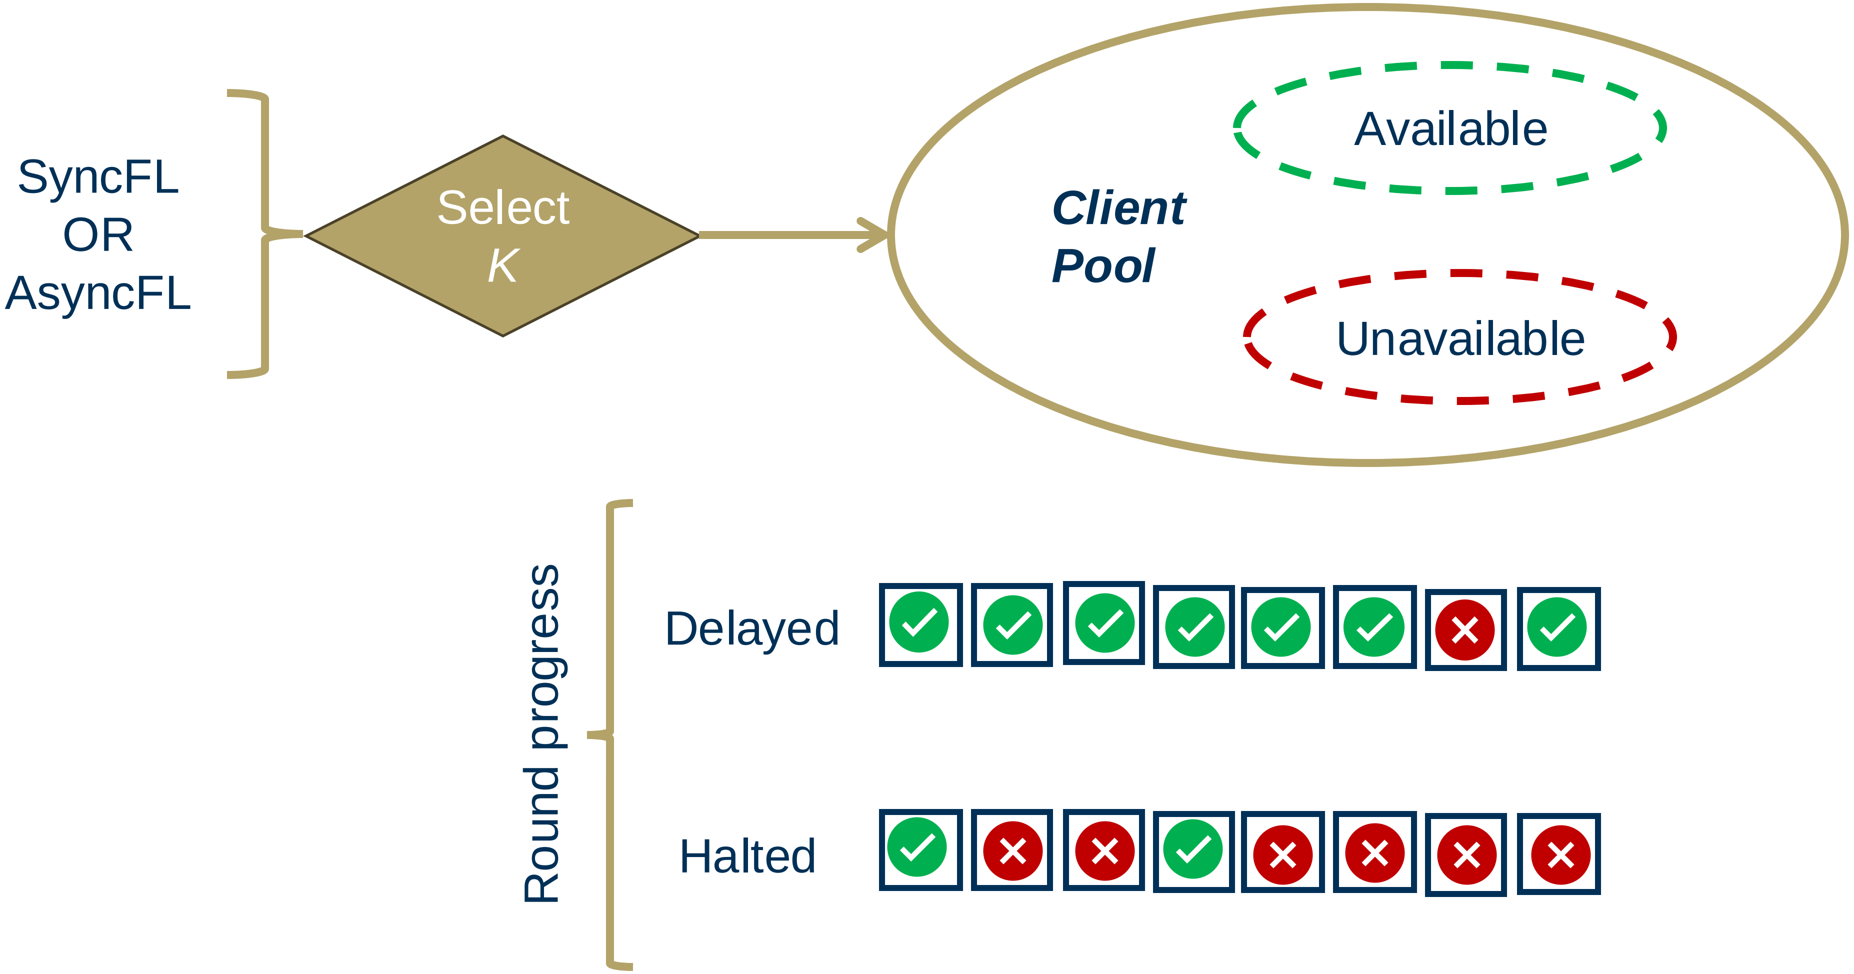

# Systems Mechanisms for Client Availability

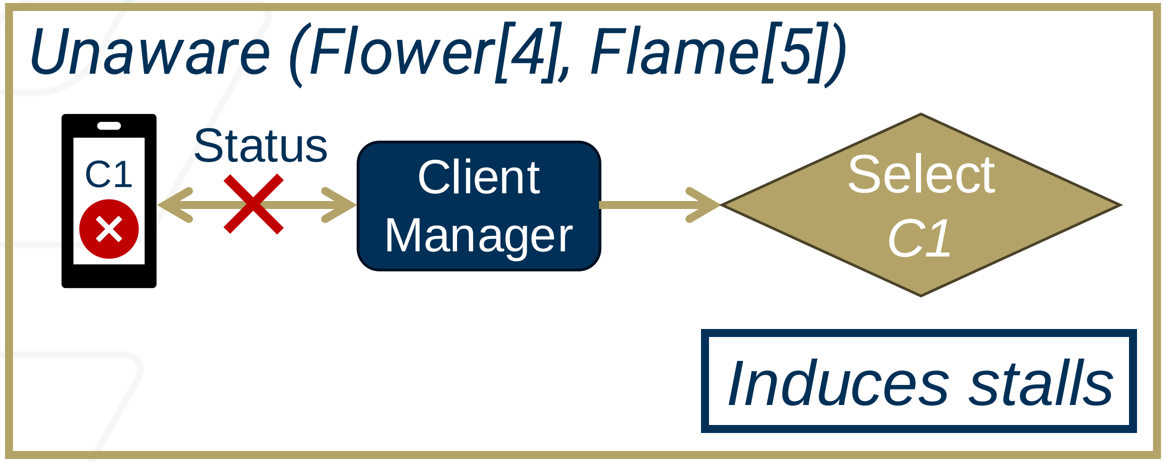


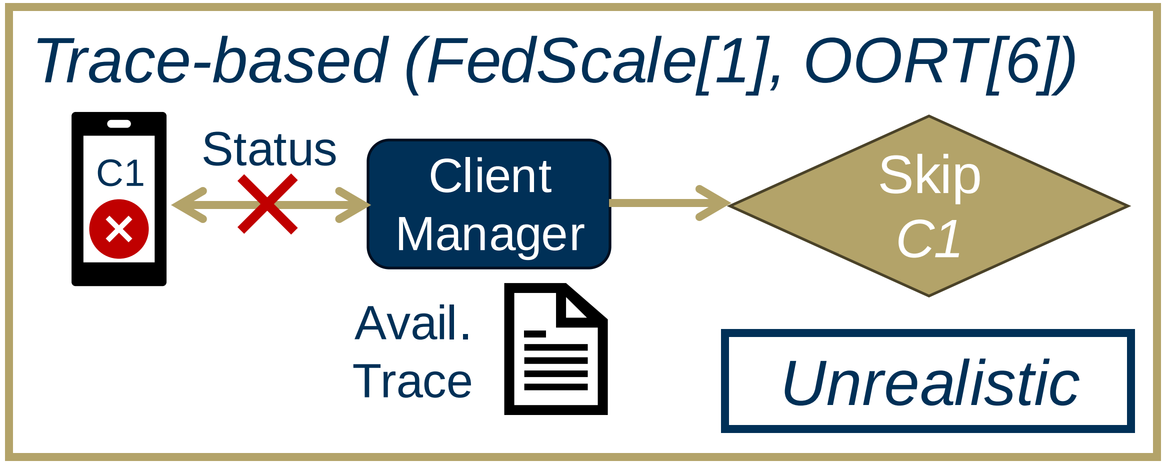


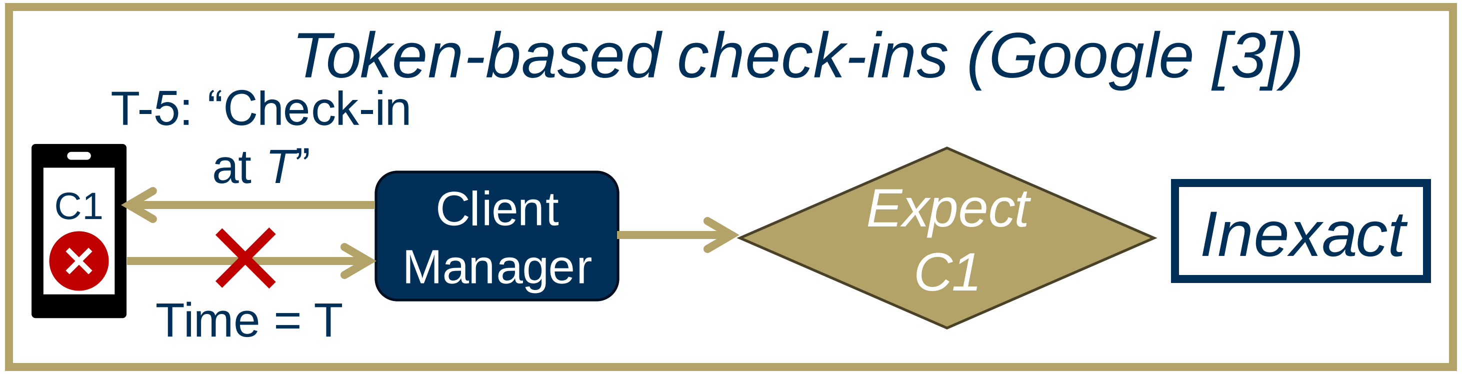



[4] Beutel, D. J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., ... & Lane, N. D. (2020). Flower: A friendly federated learning research framework. arXiv preprint arXiv:2007.14390.

[5] Daga, H., Shin, J., Garg, D., Gavrilovska, A., Lee, M., & Kompella, R. R. (2023, October). Flame: Simplifying topology extension in federated learning. In Proceedings of the 2023 ACM Symposium on Cloud Computing (pp. 341-357).

[6] Lai, F., Zhu, X., Madhyastha, H. V., & Chowdhury, M. (2021). Oort: Efficient federated learning via guided participant selection. In 15th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 21) (pp. 19-35).

# Impact of Client Unavailability

# Experiments: Setup

***Success Metric***

Time-to-accuracy

- Better ML convergence
- Resource efficiency

***Task***

- Image classification on CIFAR-10
- 300 clients

***Execution Strategies***

| Paradigm | Strategy | Availability Awareness (Oracular) |
|---|---|---|
| ***SyncFL*** | OORT[6] | ✗ |
| | OORT✷ | ✓ |
| ***AsyncFL*** | A-OORT | ✗ |
| | A-OORT✷ | ✓ |



Georgia Tech

# Experiments: Workload Characteristics

### *Client Availability Traces*

- Synthetic: Availability ~80%
- Real-world: Availability 10-22%

### *Data heterogeneity*

- Homogenous ($\alpha$=100)
- Heterogeneous ($\alpha$=0.1)



Georgia Tech

# Large Accuracy Fall With Modest Unavailability

~10% accuracy drop in unaware strategies

- Modest 20% drop in availability

Training progresses slower due to stalls

AsyncFL: more resilient than SyncFL

- OORT(20%) loses 11%
- A-OORT(20%) loses 9.5%

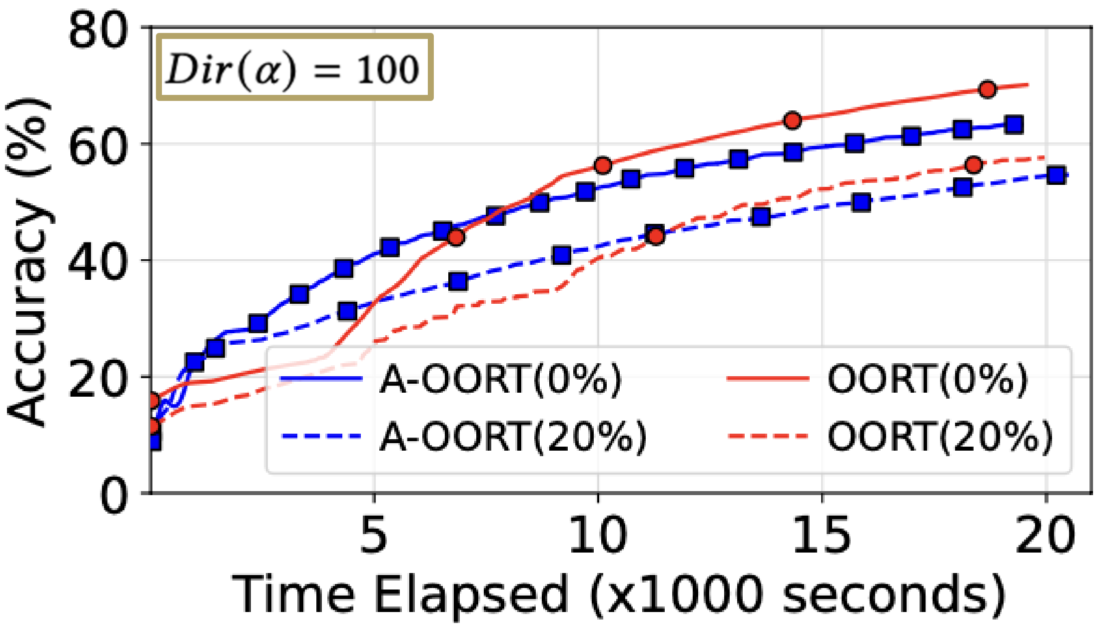


Results on Synthetic Trace

# Strategies Break Down in Real-World Settings

Achieve 46-57% higher accuracy even in 10-22% client availability

- Availability awareness

AsyncFL gains over SyncFL reduce as heterogeneity increases

- Stale updates constrain model training

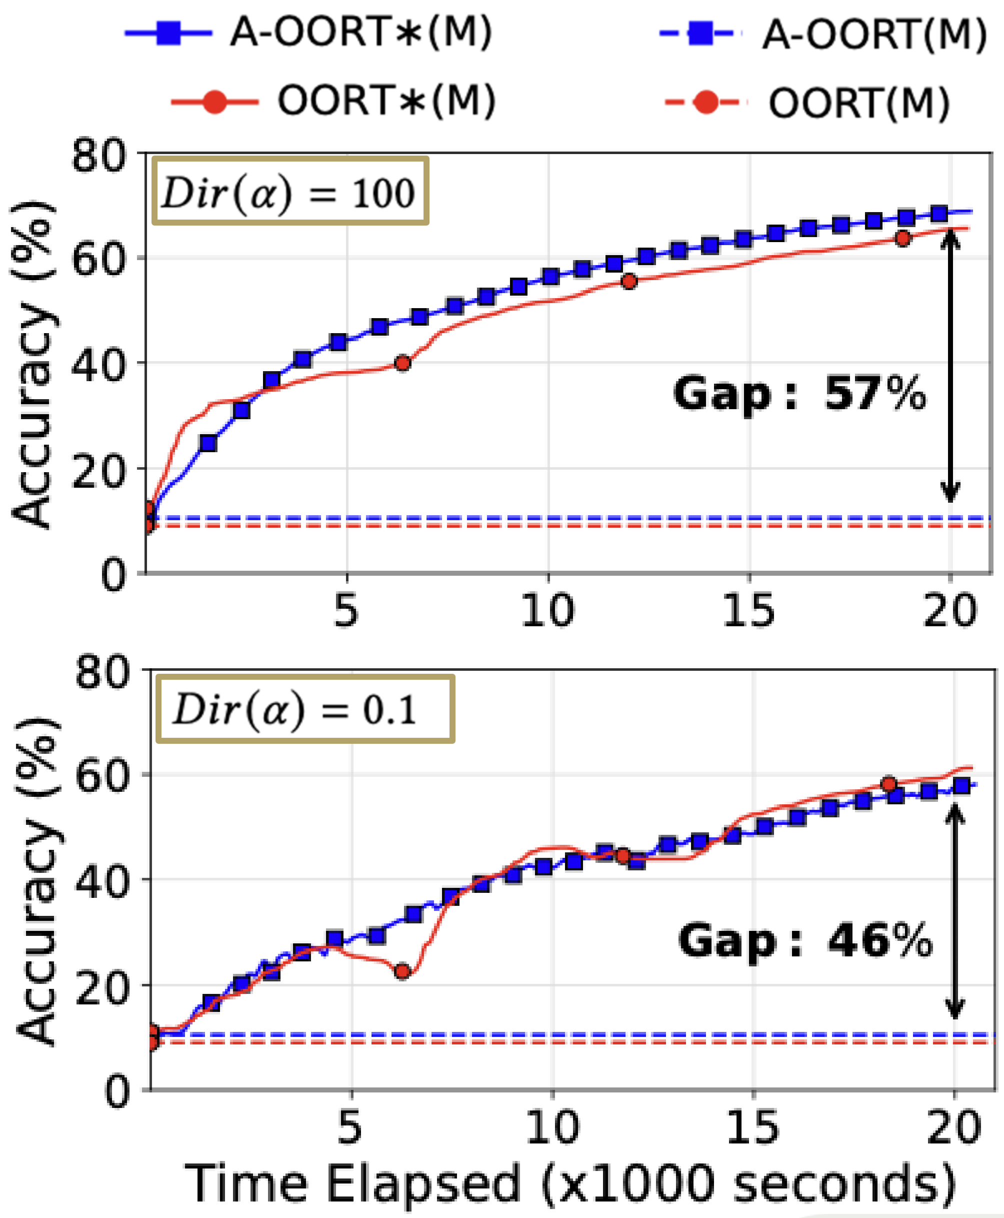

# Opportunities in Resolving Unavailability

# Make FL Systems and Algorithms Robust

- ***Client selection based on holistic tracking***
  - Selector: Fetches accurate, real-time client availability
  - Utilize: Current & historical client capabilities

- ***Efficient aggregation by managing staleness***
  - Not all stale updates are equal
  - Moderately stale updates can contribute to training [7]
  - Mitigate unavailability impact by using received updates



[7] Mitliagkas, I., Zhang, C., Hadjis, S., & Ré, C. (2016, September). Asynchrony begets momentum, with an application to deep learning. In 2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton) (pp. 997-1004). IEEE.

# Summary

- SOTA FL breaks down at high client unavailability
  - Accuracy degrades by up-to 57% in real-world traces
- Data heterogeneity exacerbates training difficulty
- Availability awareness reduces:
  - Aggregation stalls by 94%
  - Staleness of updates by 65%
- Opportunity: Make FL Systems + Algorithms Robust to Unavailability
  - Holistic and tracking-based client selection
  - Efficient aggregation by managing staleness