UNIVERSITEIT VAN AMSTERDAM
Complex Cyber Infrastructure
FABRIC
Co-funded by the European Union

Towards Efficient Vertical Federated Learning on Distributed Scientific Infrastructure

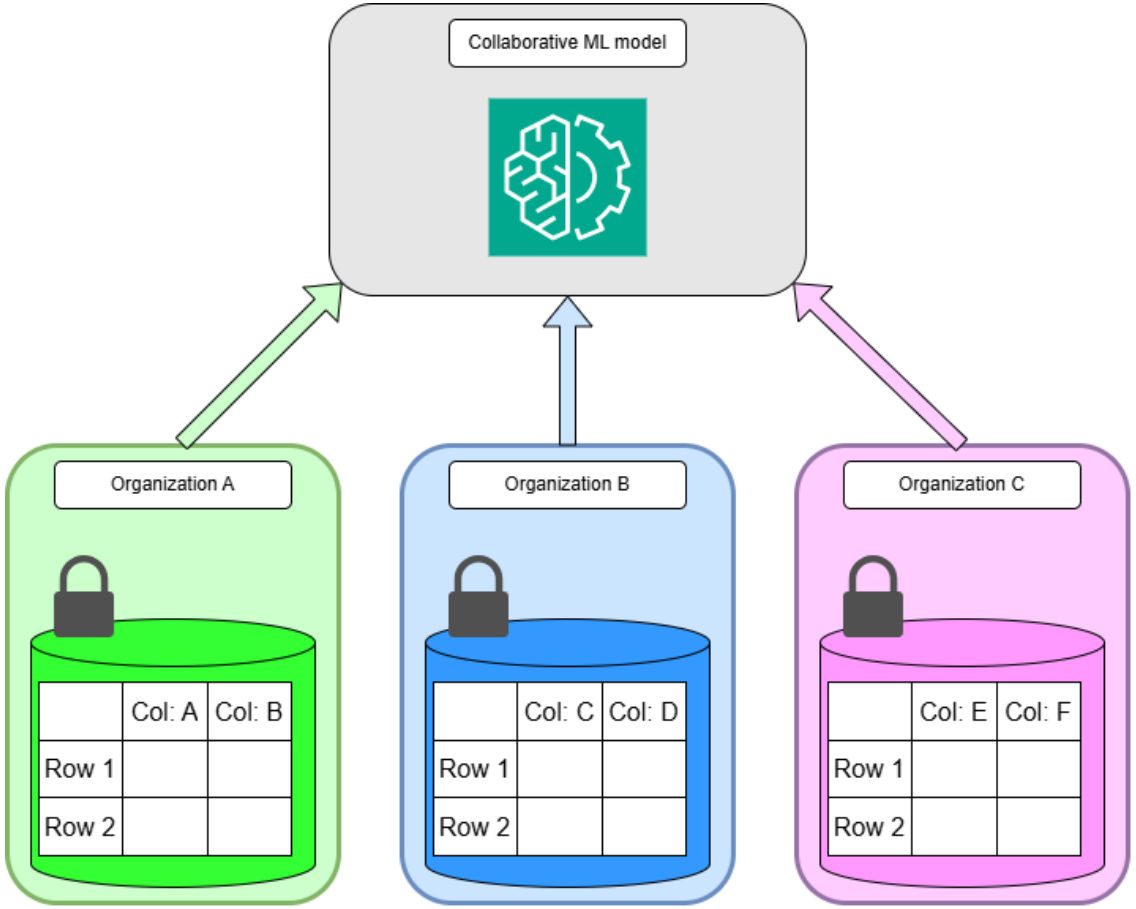
Evangelos Pipilikas
Alexandros Koufakis
Ana Oprescu

Complex Cyber-Infrastructure (CCI)
Universiteit van Amsterdam

prets embedded lending
ciena
OIVA OPEN INNOVATION FOR VERIFIED AUTOMATION

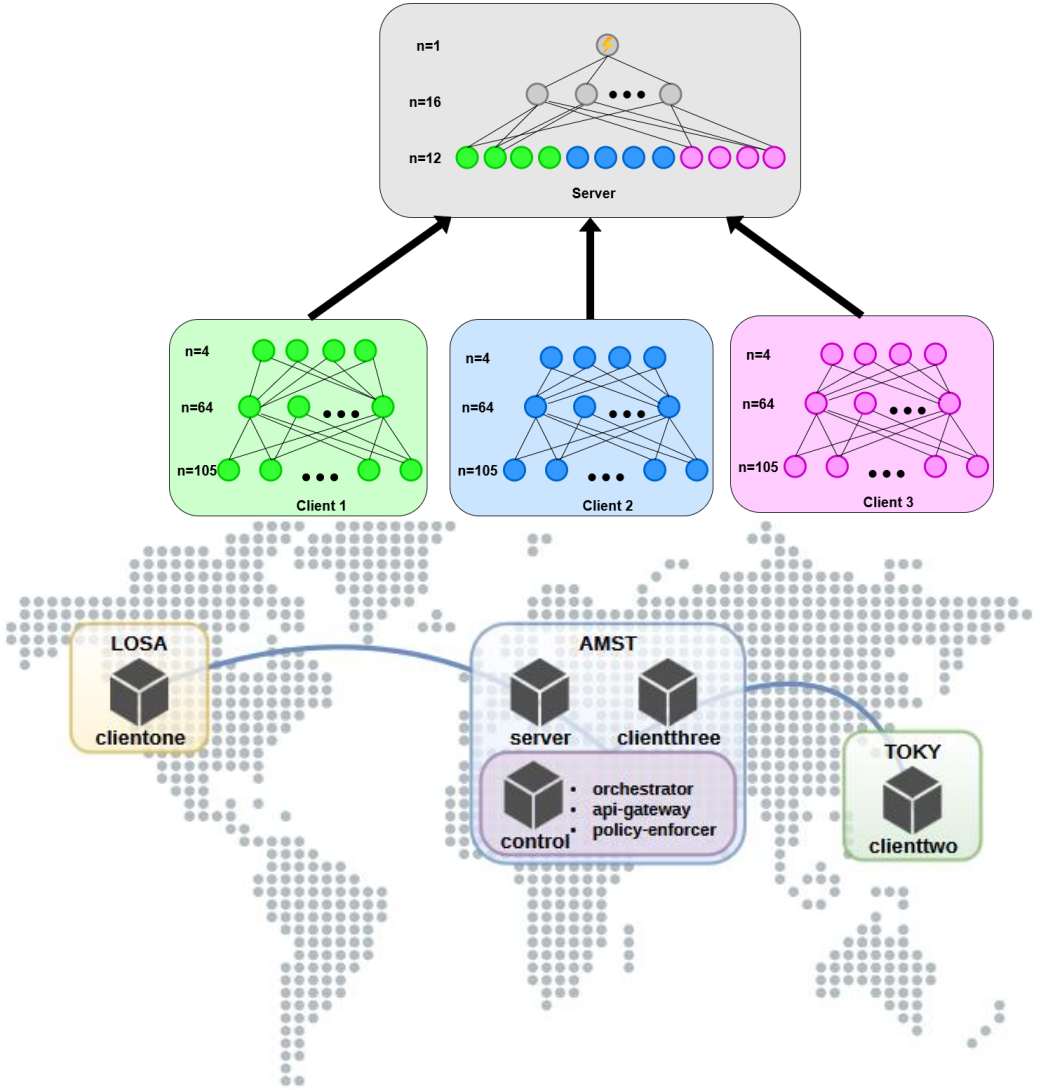
Vertical Federated Learning

- Different organizations hold different features about the same entities
- They train a shared model using those complementary views
- This is useful when data is distributed across organizations or cannot easily be centralized
- Raw datasets remain local, but training still requires coordination and information exchange



Communication and synchronization challenge

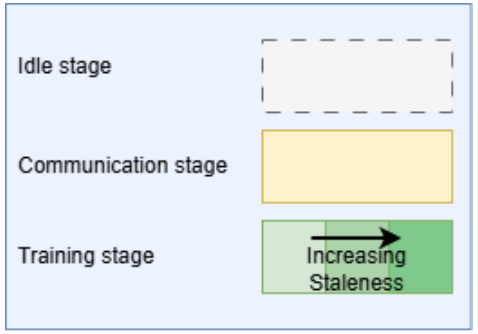
- Across wide-area networks, coordination can become a VFL bottleneck
- Machine learning model is split across all participants
- Case study configuration:
  - Control: AMST
  - Client 1: LOSA
  - Client 2: TOKY
  - Client 3: AMST
  - Server: AMST

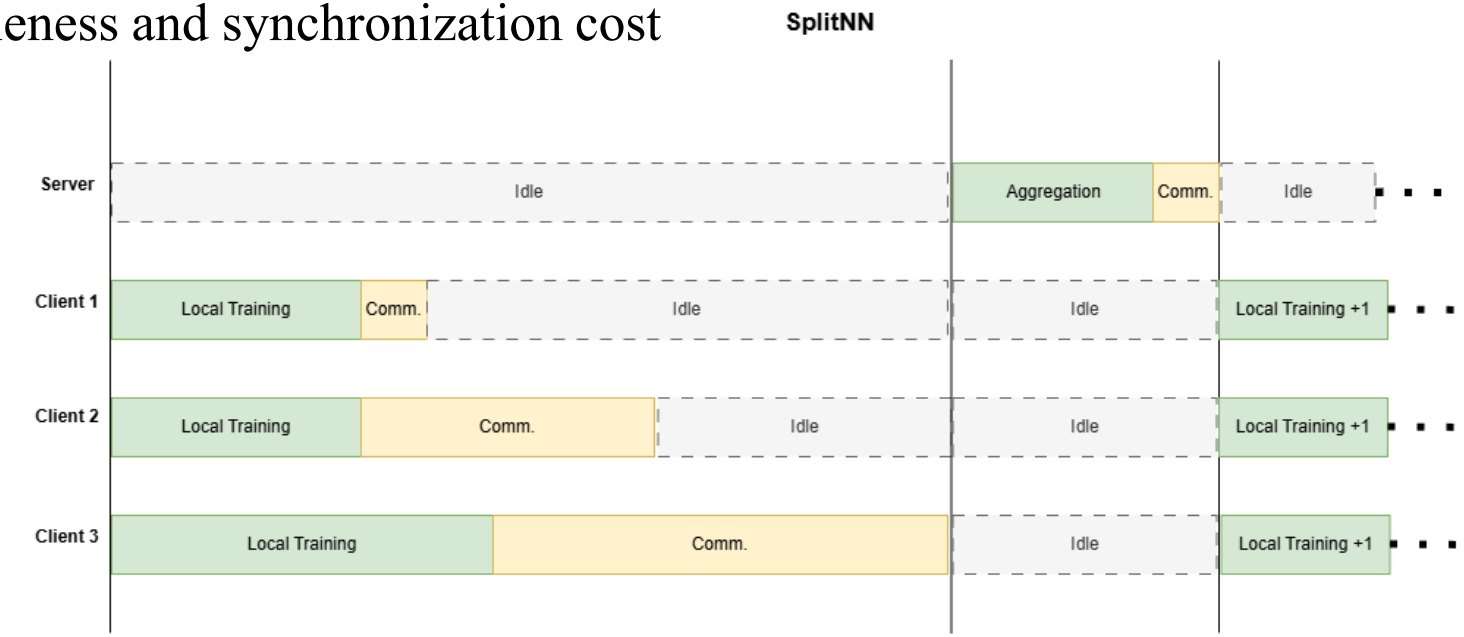


VFL Execution Strategies

Baseline (SplitNN) – synchronization at every training step

- Each client computes a local representation from its private features and sends it to the server
- The server combines the representations, computes the loss using its labels, then sends gradients back
- Clients use those gradients to update their local models.
- This exchange repeats at every training step—potentially costly across distant sites
- Analysis on earlier work [1] showed big idleness and synchronization cost

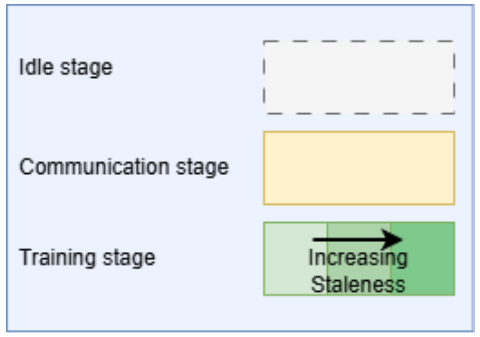


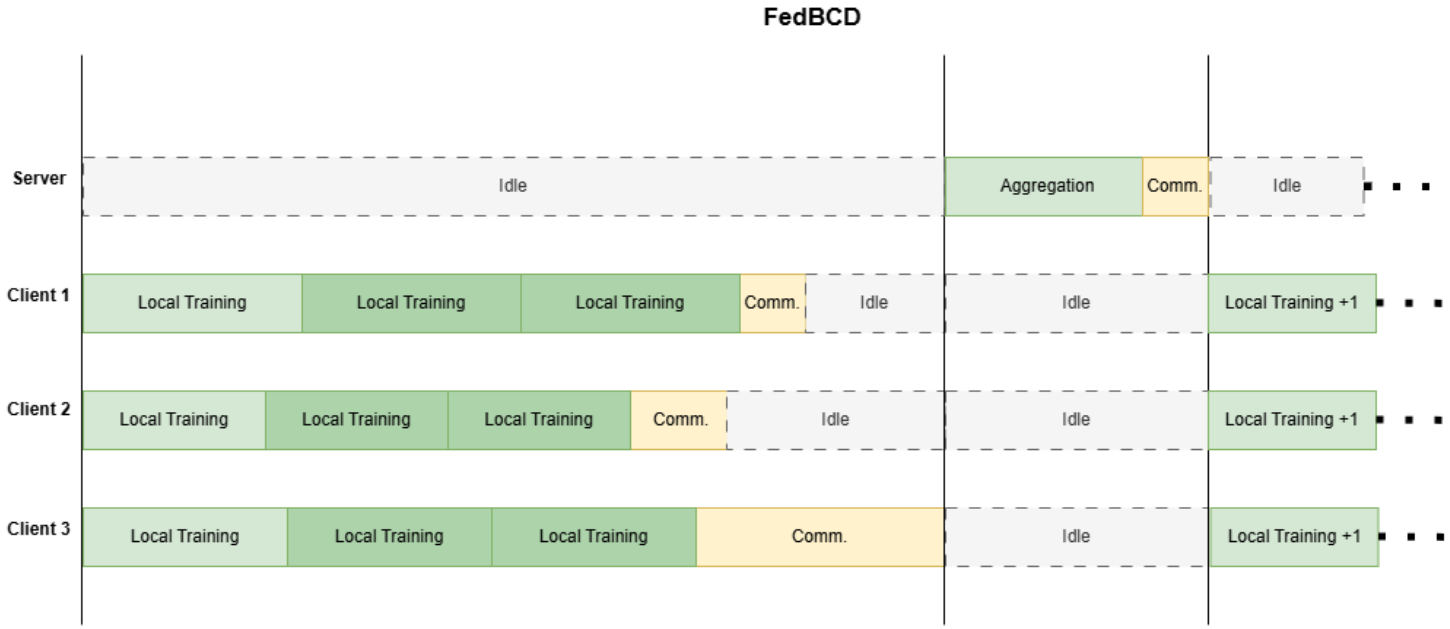


[1] Jongejans, J., Koufakis, A., Oprescu, A.: Vertical federated learning on scattered directive: Enforcing policie on vfl workflows. In: 2025 IEEE International Conference on Big Data (BigData). pp. 267–274 (2025). https://doi.org/10.1109/BigData66926.2025.11401076

FedBCD – Synchronization after multiple local updates

- FedBCD [2] reduces communication frequency
- Multiple local training steps per communication step
- Tradeoff: fewer communication rounds, but updates rely partly on **stale** information

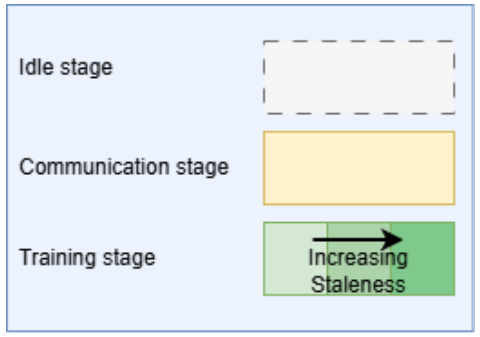


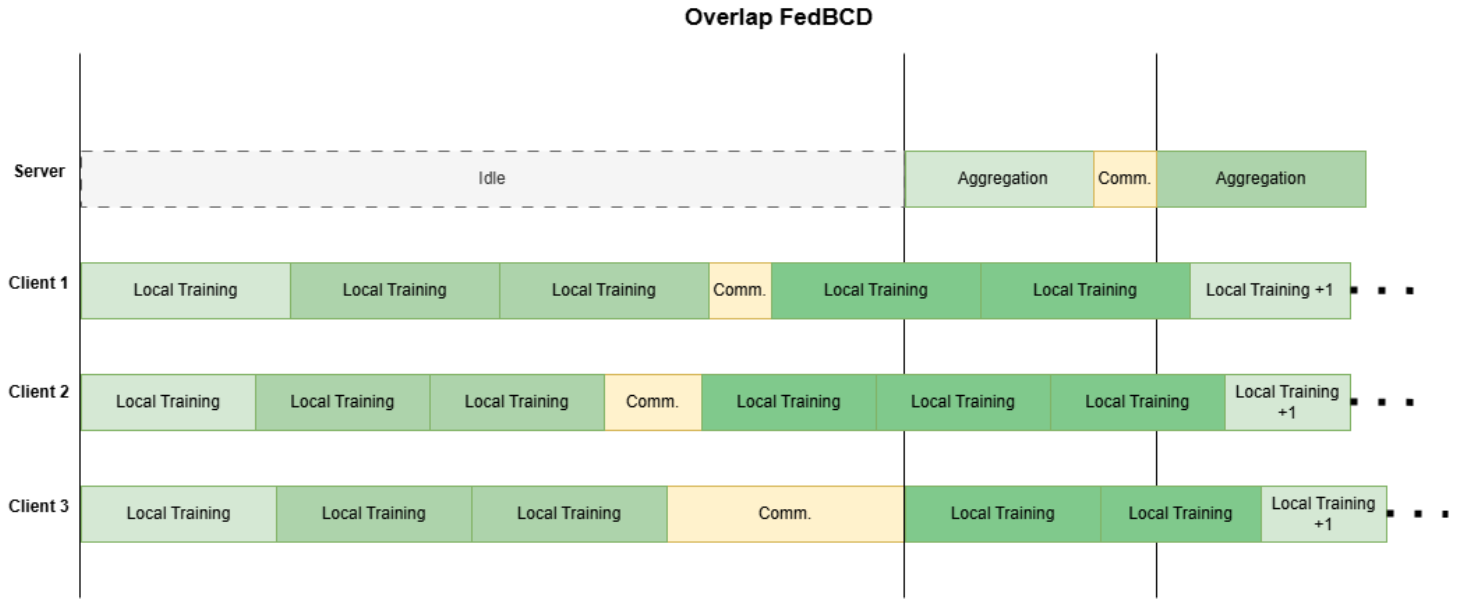


[2] Liu, Y., Zhang, X., Kang, Y., Li, L., Chen, T., Hong, M., Yang, Q.: Fedbcd: A communication-efficient collaborative learning framework for distributed features. IEEE Transactions on Signal Processing 70, 4277– 4290 (2022). https://doi.org/10.1109/TSP.2022.3198176

Novel VFL execution strategy: *Overlap FedBCD*

- Asynchronous VFL with Periodic Averaging (Overlap-FedBCD)
  - Novel adaptation from HFL [3] to VFL that reduces idleness
  - Challenges in moving from HFL to VFL: Different synchronization needs
- Adds more gradient staleness
  - Prioritazation of new gradients by clients



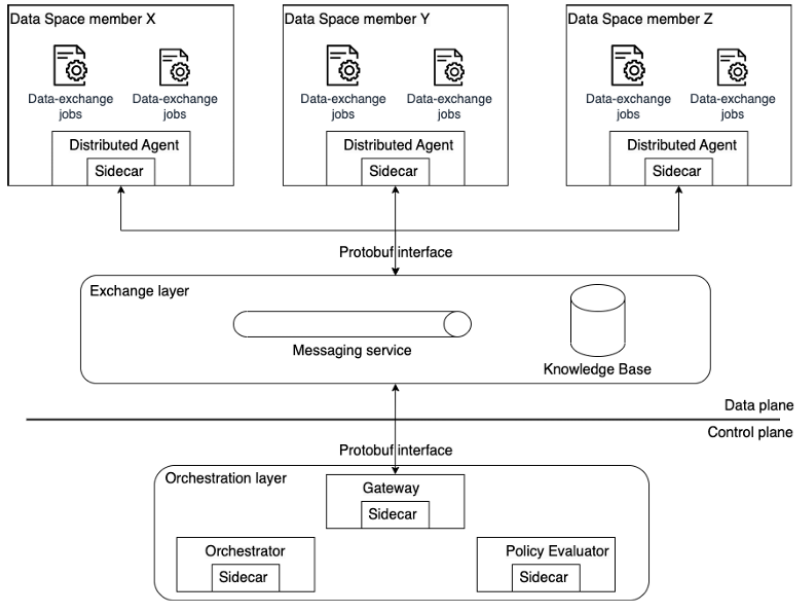


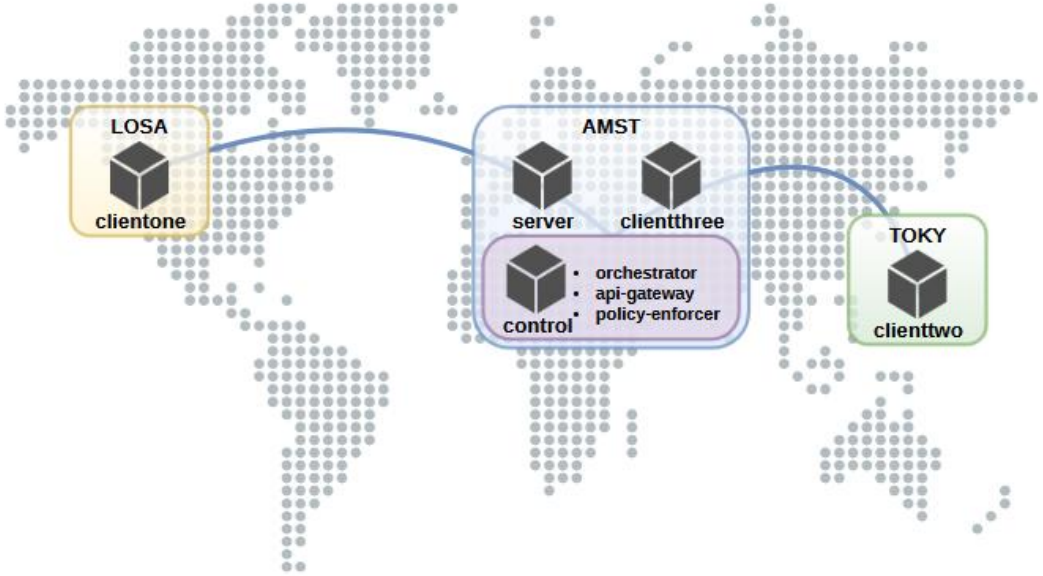
[3] Zhou, Y., Ye, Q., Lv, J.: Communication-efficient federated learning with compensated overlap-fedavg. IEEE Transactions on Parallel and Distributed Systems 33(1), 192–205 (2022). https://doi.org/10.1109/TPDS.2021.3090331

Evaluation

Experimental setup

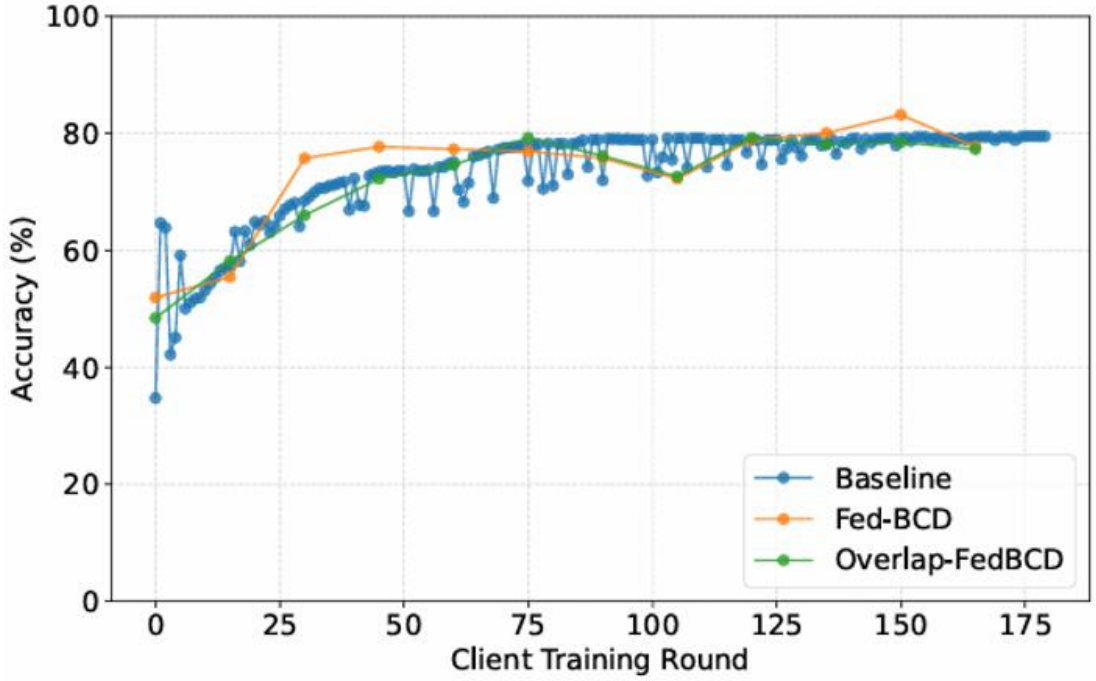
- Deployment on DYNAMOS
  - Middleware for data exchanges
  - www.github.com/DYNAMOS-UVA
- *"Titanic"* dataset
  - https://www.kaggle.com/datasets/yasserh/titanic-dataset
- Three Clients & one Server (Aggregator)
- Local and FABRIC experiments
  - https://portal.fabric-testbed.net/





Predictive Performance Evaluation

- Accuracy reached good levels for all methods
- Staleness did not seem to damage accuracy
- Test set accuracy (hidden Kaggle test set)
  - Consistent to literature



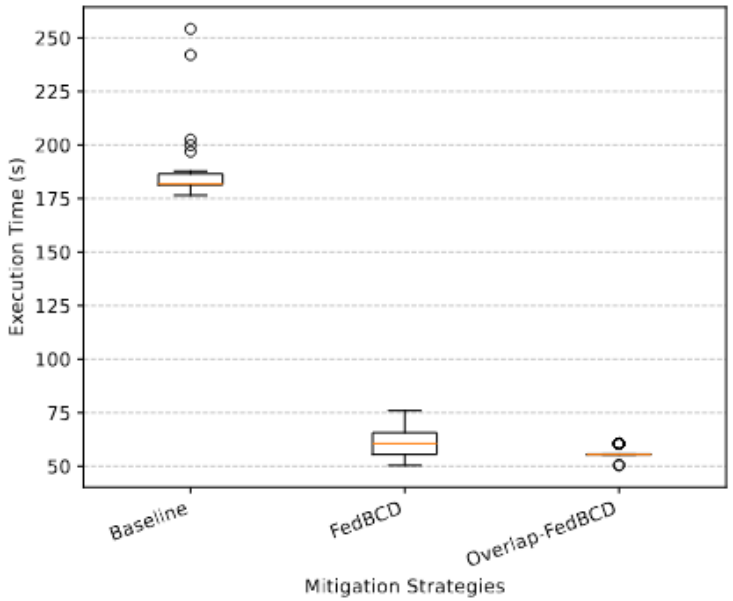
VFL Strategy	Validation Set	Test Set
Baseline	77.66	76.31
FedBCD	79.39	76.31
Overlap-FedBCD	80.15	76.31

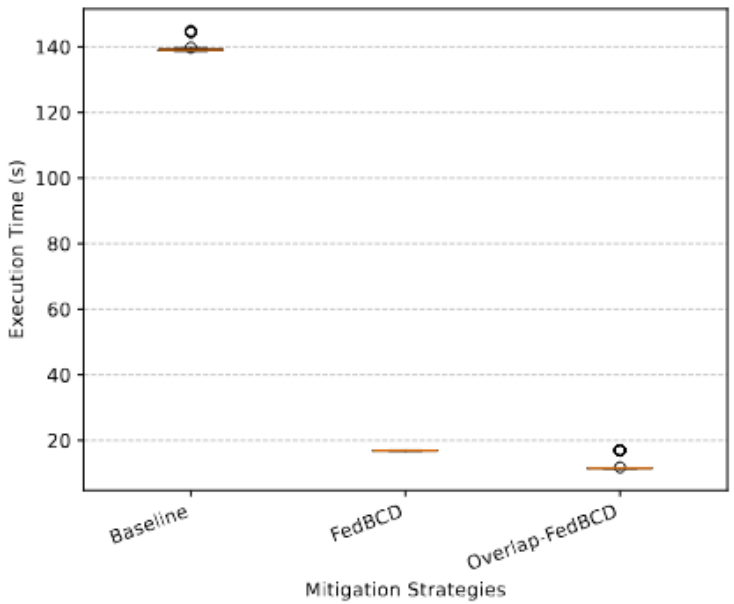
Execution Time Evaluation

- The proposed methods decreased execution time significantly both on local and the distributed FABRIC setup
- On distributed setup the improvement was even more pronounced

Method	Local (s)		FABRIC (s)		Δ
	Mean	Std.	Mean	Std.	(%)
Baseline	187.55	14.70	139.71	1.73	-25.51†
FedBCD	62.05*	6.38	16.91*	0.06	-45.14†
Overlap-FedBCD	56.39*,‡	2.58	12.71*,‡	2.28	-77.45†

Method	Local	FABRIC
FedBCD	66.91%*	87.89%*
Overlap-FedBCD	69.93%*	90.89%*

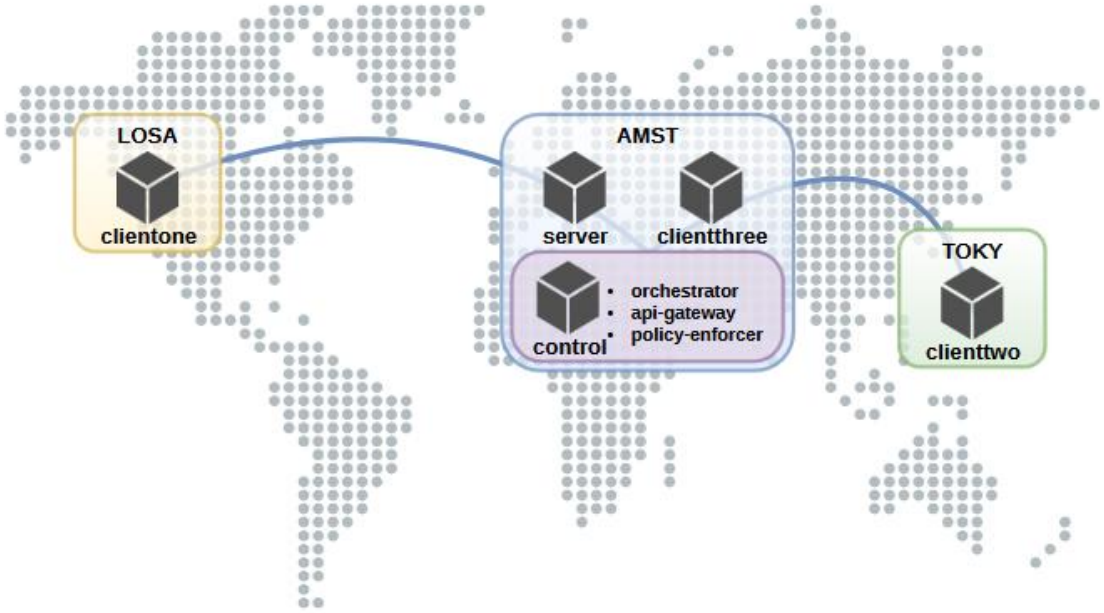


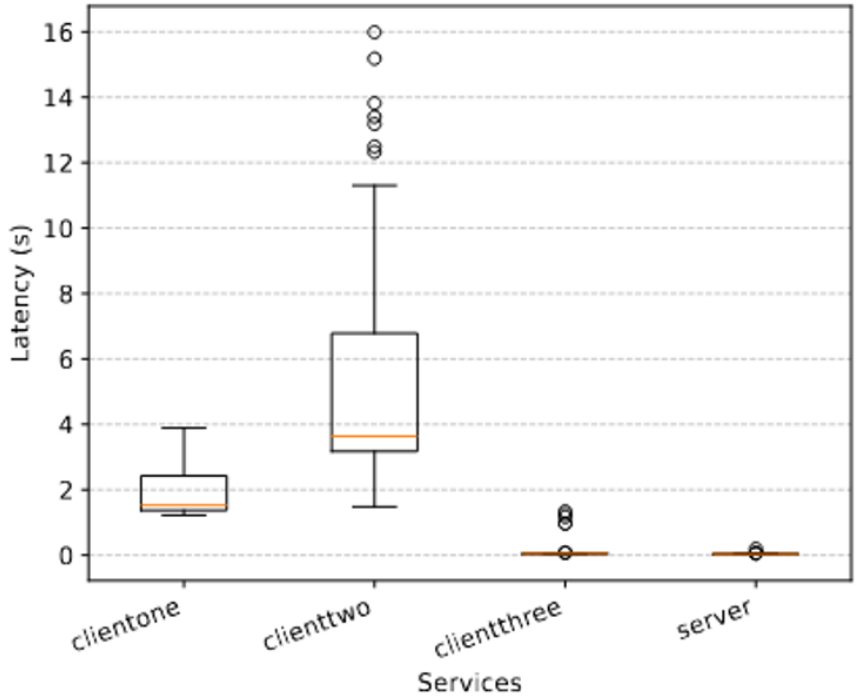


Heterogeneous network delays

- Geographically distributed clients with heterogeneous delays
- Remote clients exhibited much higher delay
- Increased benefit compared to baselines

Method	Local	FABRIC
FedBCD	66.91%*	87.89%*
Overlap-FedBCD	69.93%*	90.89%*





Acknowledgements

- This work was made possible by **FABRIC** and specifically FABRIC Across Borders.
- This work was partially supported by the project: "*BNVw: Budget-neutral sustainable home improvements (**MOOI**-224049)*".
- This work has been partially carried out as part of the "*Open Innovation for Verified Automation (**OIVA**)*" project, co-funded by the European Union under Grant Agreement No 101298830. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.
- Thanks to **CIENA** for hosting us at their booth on the exhibition floor.

FABRIC

prets embedded lending

OIVA
OPEN INNOVATION FOR VERIFIED AUTOMATION

Co-funded by the European Union

ciena

Thank you!

- FedBCD and OverlapFedBCD resulted in statistical improvement of execution time compared to baseline in both local and FABRIC environment
- Future work:
  - Policy awareness across different VFL strategies – decision boundaries
  - Explore efficient Hybrid (Horizontal and Vertical) Federated Learning
  - Exploration of Energy efficiency vs sovereignty tradeoffs

Method	Local (s)		FABRIC (s)		Δ
	Mean	Std.	Mean	Std.	(%)
Baseline	187.55	14.70	139.71	1.73	-25.51†
FedBCD	62.05*	6.38	16.91*	0.06	-45.14†
Overlap-FedBCD	56.39*,‡	2.58	12.71*,‡	2.28	-77.45†

Evangelos Pipilikas
Alexandros Koufakis
Ana Oprescu

Complex Cyber-Infrastructure (CCI)

a.m.koufakis@uva.nl

- Codebase: github.com/Apipilikas/scattered-directive-energy-monitoring
- More resources: github.com/DYNAMOS-UVA

Co-funded by the European Union

Bonus slide: *"Titanic"* Dataset

- *"Titanic"* dataset
- Target variable:
  - "*Survived*"
- Rows:
  - 891
- Features:
  - Pclass, Name, Sex, Age, SibSp, Parch, Ticket, Fare, Cabin, Embarked
- Source:
  - https://www.kaggle.com/datasets/yasserh/titanic-dataset



Source: https://www.kaggle.com/datasets/yasserh/titanic-dataset

Bonus slide: Architectural Diagram of FedBCD

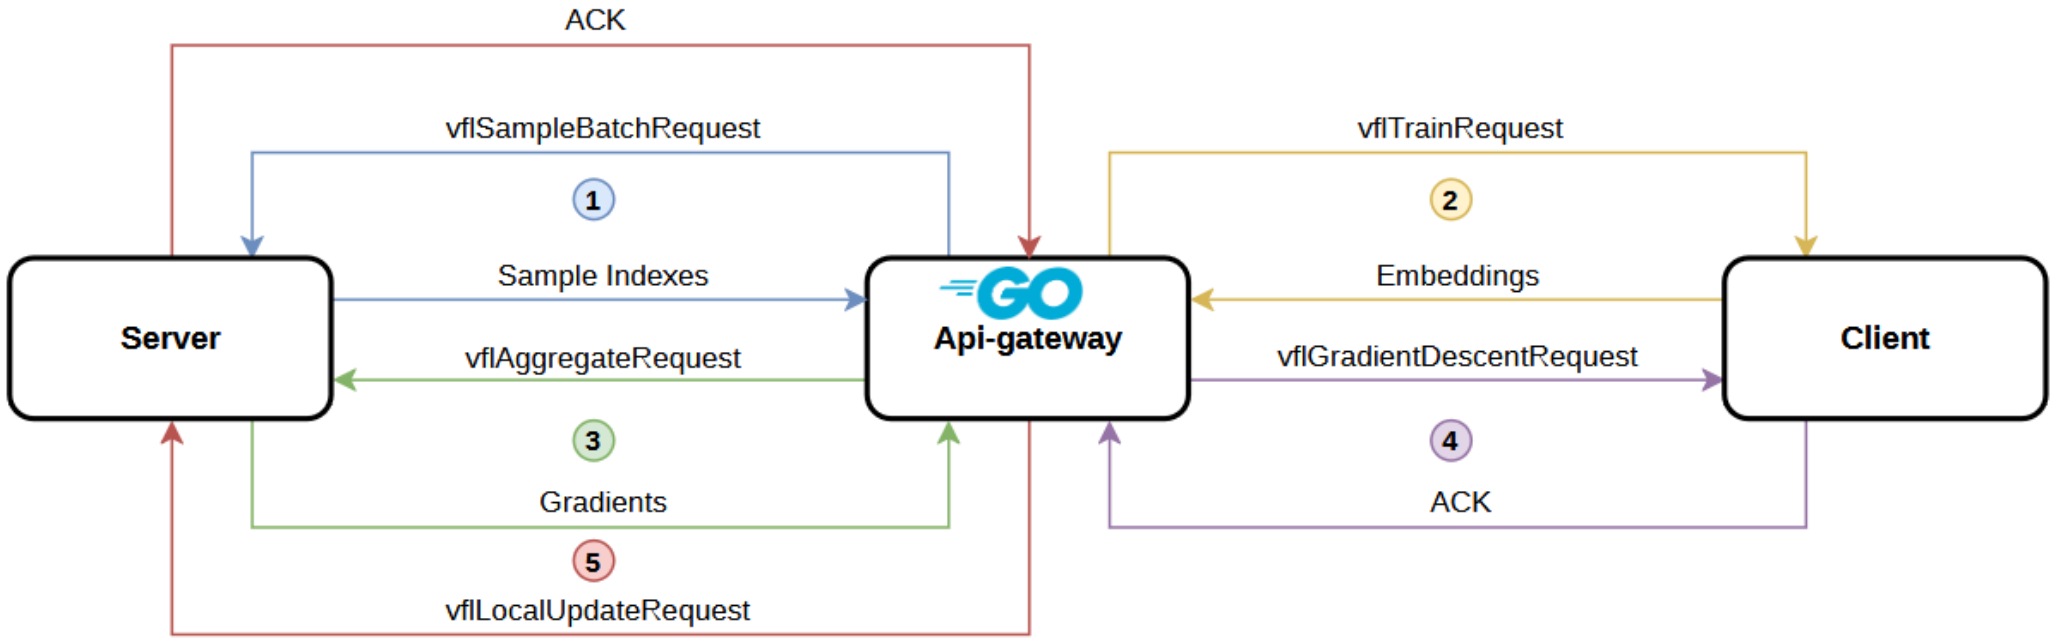


Fig. 3: Demonstration of the FedBCD execution pipeline.

Bonus slide: Coefficient of Variation

- Runtime consistency depended on the deployment: FedBCD had the lowest relative variability on FABRIC, whereas Overlap-FedBCD had the lowest locally but the highest on FABRIC.

$$CV = \frac{\sigma}{\mu}$$

Method	Local (s)		FABRIC (s)		Δ
	Mean	Std.	Mean	Std.	(%)
Baseline	187.55	14.70	139.71	1.73	-25.51†
FedBCD	62.05*	6.38	16.91*	0.06	-45.14†
Overlap-FedBCD	56.39*,‡	2.58	12.71*,‡	2.28	-77.45†

	Local	FABRIC
Baseline	0.078	0.012
FedBCD	0.103	0.004
Overlap-FedBCD	0.046	0.179