DECENTRALIZED FEDERATED LEARNING SYSTEM

DECENTRALIZED ORCHESTRATED FEDERATED LEARNING

SOLUTION VALUE

Results demonstrate that TDFLS scales with the number of participants, is robust against intermittent connectivity and dynamic participant departures/ arrivals, requires minimal resources, and guarantees that the accuracy of the trained model quickly converges to that of a centralized FL framework with an accuracy drop of less than 1%.

- Allow the military to adapt current and future ML algorithms in a DFL model
- Allows for rapid development of new algorithms in a DFL model
- Provides system resilience for D-DIL and contested environments, and in the event assets are lost
- Povides DFL across heterogenous network environments and platforms
- Provides communication efficiency

EXECUTIVE SUMMARY

Technica Corporation specializes in delivering tactical edge computing solutions that can operate in Denied-Disrupted, Intermittent, and Limited (D-DIL) environments. This includes Artificial Intelligence (AI) algorithms for austere operations and the capability for distributing and analyzing data via decentralized federated learning.

Technica has developed the Technica Decentralized Federated Learning System (TDFLS) to capitalize on tactical edge fused data.

This federated learning solution requires no central learning node. Instead, it partitions the machine learning (ML) model to minimize bandwidth requirements, allow more significant platform heterogeneity through Knowledge Distillation (KD) and provides a decentralized, redundant, and agnostic file name system to add needed resiliency at the edge.

Technica built TDFLS on scalable, open-source architectures with APIs capable of federating various ML algorithms seamlessly across platforms from the edge to the cloud and supporting bi-directional cross-domain operations.

The platform is based on research initially conducted by Christodoulos Pappas of the University of Thessaly in Greece. Technica worked in collaboration with Pappas to further develop TDFLS.

SOLUTION DETAILS

Federated learning, introduced by Google in 2017, extended the concept of data parallelism to scenarios where data is generated and processed at different physical locations. Importantly, this means that the data is not guaranteed to be identically independently distributed (IID). Federated learning addresses this issue by frequently aggregating models trained at all locations and redistributing the resulting model.

Decentralized Federated Learning (DFL) aggregates all the models at all locations to eliminate a single point of failure. In centralized federated learning, this aggregation occurs at a single location. A naïve approach to DFL requires that each site broadcast its models to all other locations. For moderately sized networks, this can be prohibitive and is simply infeasible at the scale of a military operation and in a D-DIL battlefield environment.

WE LISTEN. **WE APPLY. WE SOLVE.**

Technica's approach using TDFLS addresses the problems associated with traditional DFL. Technically partially based TDFLS on the InterPlanetary File System (IPFS), a protocol and peer-to-peer network for storing and sharing data in a distributed file system. By using IPFS and connecting to the corresponding private network, any party can initiate an ML training process or join an ongoing training process.

ALGORITHM

The Technica Decentralized Federated Learning System balances system resilience with bandwidth efficiency. The algorithm divides the model into multiple partitions, each with consecutive layers, with each partition assigned to at least one compute node.

The nodes send model weights for each partition to the nodes responsible. Each node then aggregates the partitions they are responsible for and sends updated weights to all other nodes. The assigning of each partition to multiple nodes enhances the system's resilience but increases the bandwidth required. The user can turn off this parameter if needed.

Centralized Federated Learning (CFL) requires models from each compute node to be sent to a single server. The results are then distributed to all compute nodes during each round of aggregation. This results in an inefficient system with the number of model copies transmitted for aggregation reaching O(N). Classical DFL models demand that each node transmit its model to all other nodes in the network. This means N(N-1) models are transmitted, making typical DFL bandwidth usage O(N2).

Due to the partitioning of the model, TDFLS only needs each node to transmit 1/N to each of the other nodes and receive 1/N of the model from each of the other nodes. The number of models sent is then kN(1/N)(N-1), where k is the number of nodes each partition is assigned. This results in a bandwidth usage O(kN), a low requirement suited for the tactical edge.

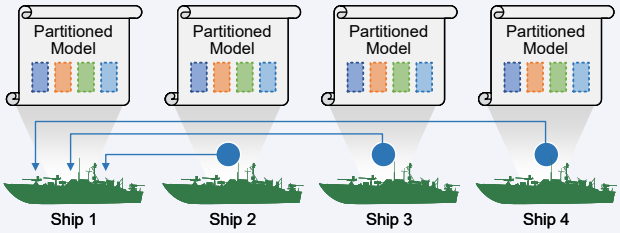

Model Partition Aggregation

Inspired by IPFS, Technica's Decentralized Federated Learning System (TDFLS) framework allows a large number of mobile nodes to collaborate in the training of a model without relying on any central entity.

In TDFLS, each node is in charge of one or more model partitions, and nodes communicate with one another by exchanging partition weights and/or model updates. The nodes in TDFLS, as shown in the figure above, collectively train a model in a peer-to-peer fashion without the use of a server. Any node can start the training process by defining the model, loss function, and algorithm to use. After that, interested nodes can register and take part in the training. In this decentralized FL framework, the model is split into multiple partitions that are replicated on multiple nodes, as opposed to centralized FL, where only the server is responsible for storing, updating, and broadcasting the model to the participating nodes.