ATA: Adaptive Task Allocation for Efficient Resource Management in Distributed Machine Learning

Artavazd Maranjyan, El Mehdi Saad,
Peter Richtárik, Francesco Orabona

Motivation

Imagine you're running minibatch SGD on a heterogeneous cluster of 1010 workers (GPUs), using a batch size (B) of 10. The fastest way to collect the batch is to do it asynchronously across all workers. However, this approach wastes the computations of at least 1000 workers.

Problem setup

Let $\alpha \geq 0$. For all $i \in [n]$, X_i is a positive random variable satisfying $\|X_i\|_{\psi_1} \leq \alpha$, where $\|X\|_{\psi_1}$ denotes the Orlicz norm of a real-valued random variable X, defined as

$$\|X\|_{\psi_1} := \inf\{C > 0 : \mathbb{E}[\exp(|X|/C)] \leq 2\}.$$

The main objective of this work is to develop an online allocation strategy with small expected total computation time, defined as $\mathcal{C}_K := \sum_{k=1}^{K} \mathbb{E}[C(\boldsymbol{a}_k)]$, where

$$C(\boldsymbol{a}_k) := \max_{i \in \text{supp}(\boldsymbol{a}_k)} \sum_{u=1}^{a_{i,k}} X_{i,k}^{(u)} .$$

If the distributions of the arms were known in advance, the optimal allocation $\boldsymbol{a}^* \in \mathcal{A}$ would be selected to minimize the expected computation time per round, $\mathbb{E}[C(\cdot)]$, and this allocation would be used consistently over K rounds, leading to the optimal total computation time $\mathcal{C}_K^* = K\mathbb{E}[C(\boldsymbol{a}^*)]$.

Based on empirical means, we define the lower confidence bounds $s_{i,k}$ as

$$s_{i,k} = (\hat{\mu}_{i,k} - \text{conf}(i,k))_+ ,$$

where $(x)_+ = \max\{x, 0\}$ and $\text{conf}(\cdot,\cdot)$ is defined as

$$\text{conf}(i,k) = \begin{cases} 4e\alpha\left(\sqrt{\frac{\ln(2k^2)}{K_{i,k}}} + \frac{\ln(2k^2)}{K_{i,k}}\right), & K_{i,k} \geq 1, \\ +\infty, & K_{i,k} = 0 . \end{cases}$$

Theorem 0.1. *Let $\eta := \max_{i \in [n]} \frac{\sigma_i}{\mu_i}$. Then, the total expected computation time after K rounds, using the allocation prescribed by ATA with inputs (B, α) satisfies*

$$\mathcal{C}_K \leq \left(1 + \eta\sqrt{\ln B}\right)\mathcal{C}_K^* + \mathcal{O}(\ln K) .$$

Achieving *optimal training time* with *minimal computation* in *heterogeneous clusters* through *Adaptive Task Allocation* and *online worker speed estimation*.

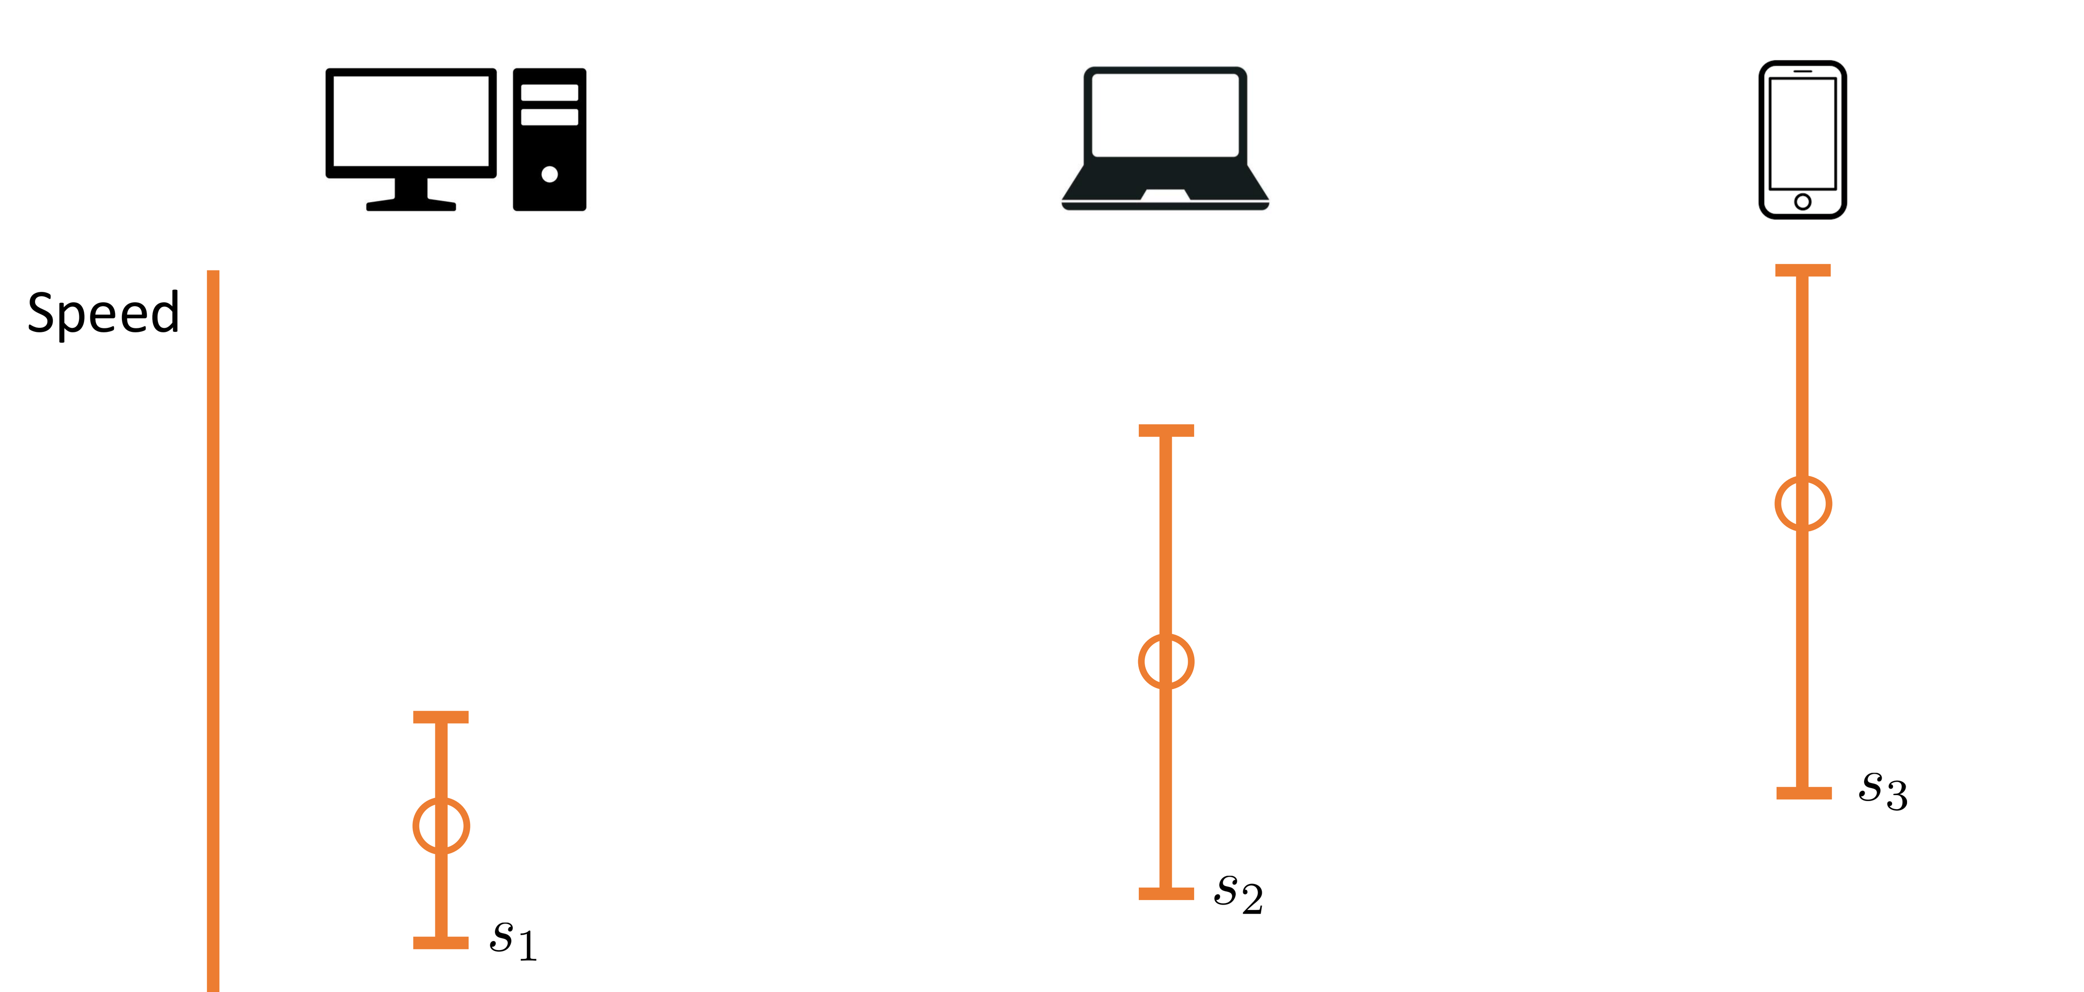

$$\text{Allocation} = \text{argmin}_{\sum a_i = B} \left\{ \max_j a_j s_j \right\}$$

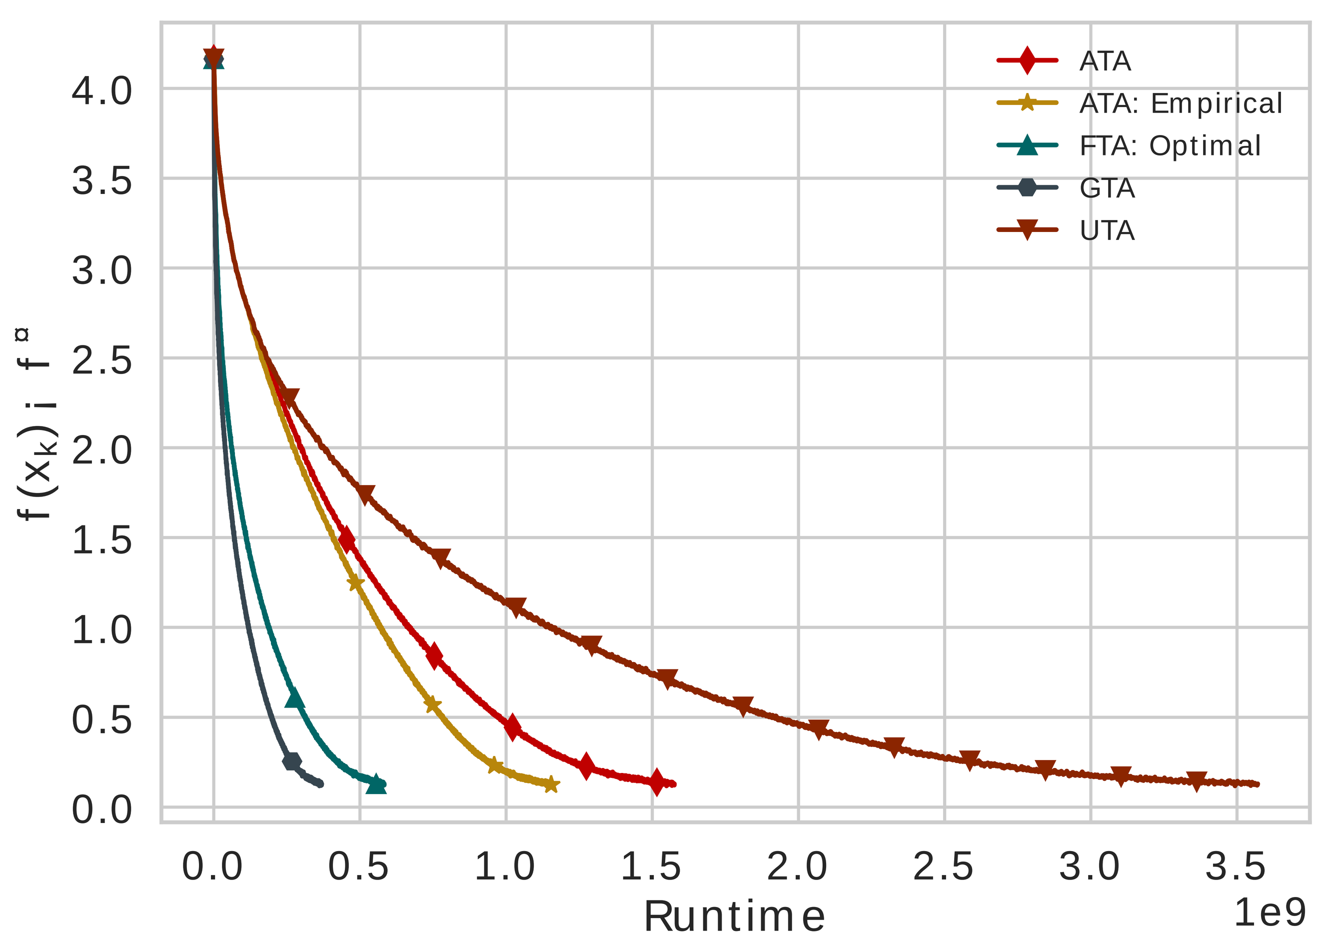

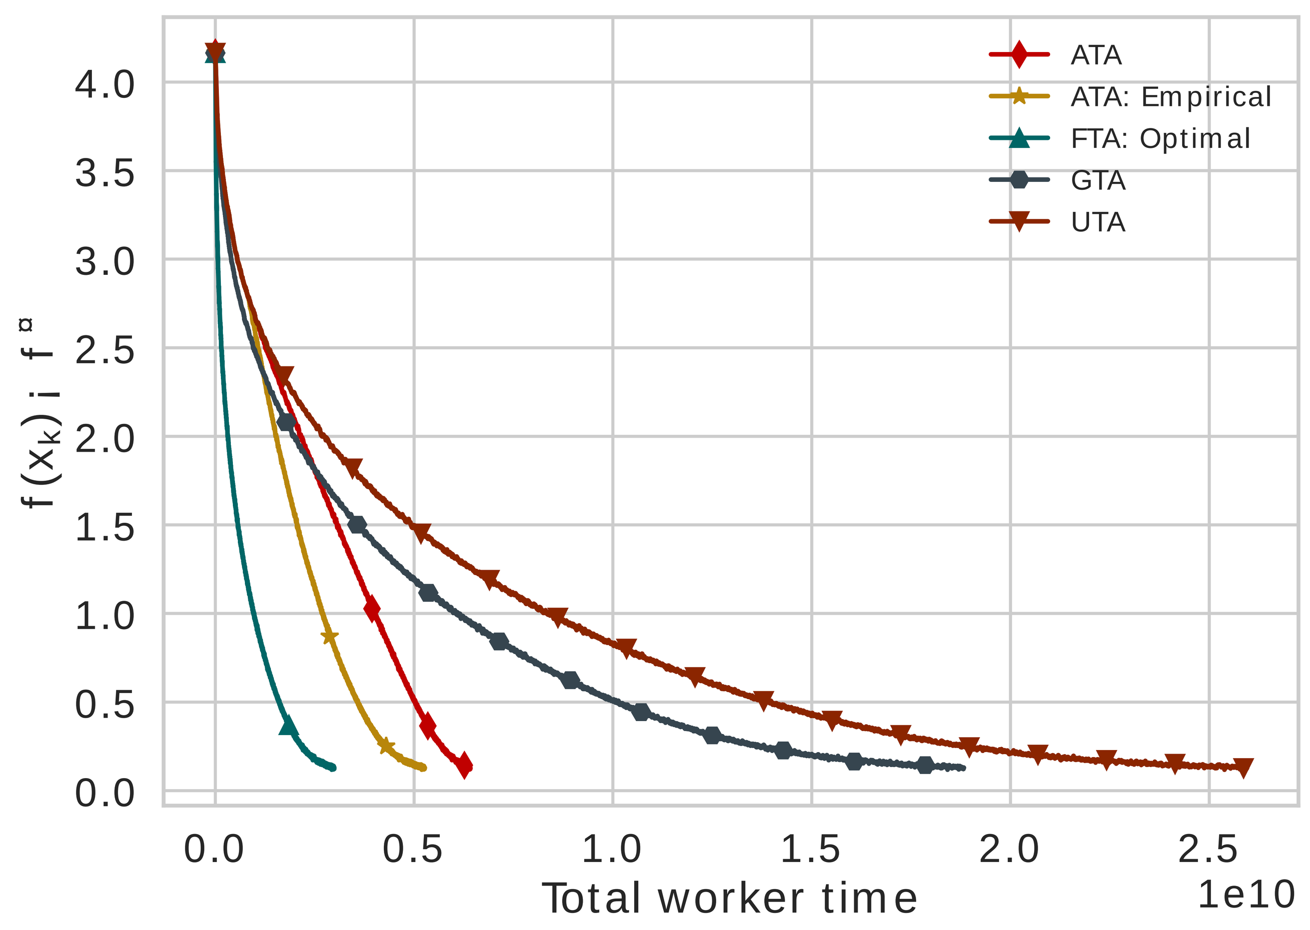

We use the CIFAR-100 dataset and train a CNN with three convolutional layers followed by two fully connected layers, totaling 160k parameters. The Adam optimizer is used with a constant step size of $8 \cdot 10^{-5}$. Each worker's computation time is sampled from an exponential distribution, with the mean increasing linearly across workers. The batch size is fixed at $B = 23$, and the number of workers is $n = 51$.

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology