Towards Resilient, Secure, and Private Distributed Learning: A Coding-Theoretic Approach

Abolfazl Changizi
changizi@kth.se

Presentation Outline

Introduction

Big Data and Huge Learning Models

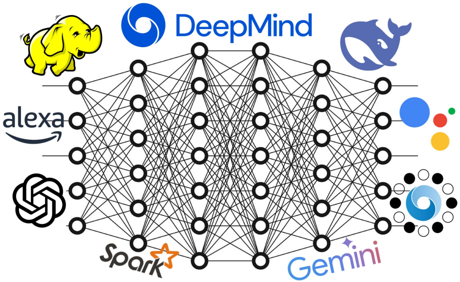

Huge learning models

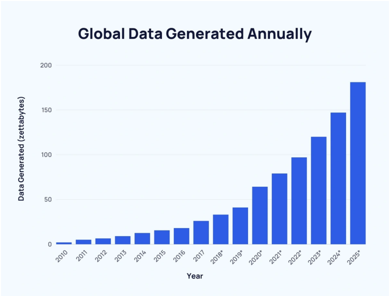

(Duarte, 2025): 4×10^{20} bytes per day!

Introduction

Big Data and Huge Learning Models

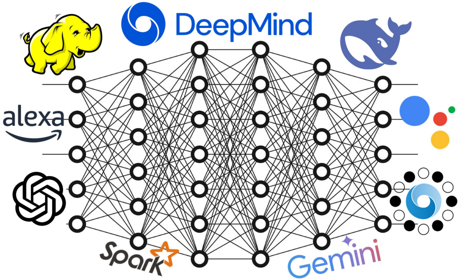

Huge learning models

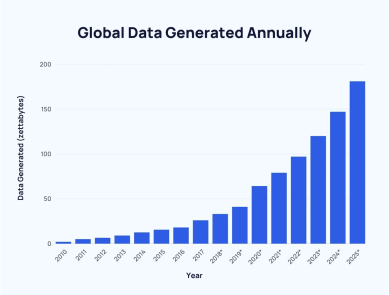

(Duarte, 2025): 4×10^{20} bytes per day!

There is a need for large-scale computations

Introduction

Big Data and Huge Learning Models

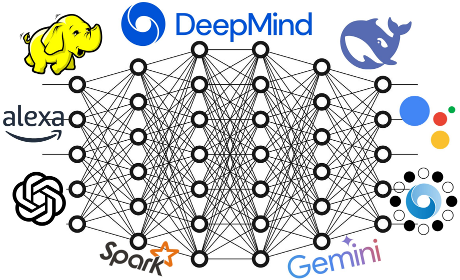

Huge learning models

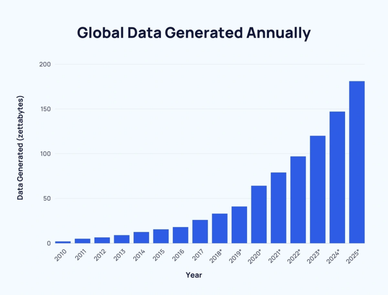

(Duarte, 2025): 4×10^{20} bytes per day!

There is a need for large-scale computations

Computations cannot be done in a centralized manner

Introduction

Computation Offloading Framework: Opportunities and Bottlenecks

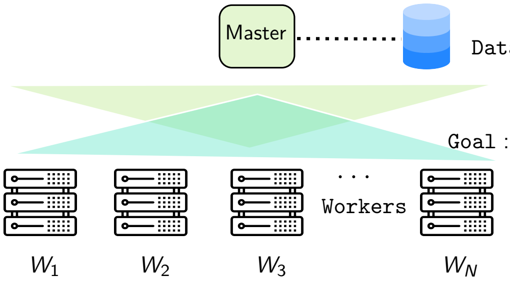

Data $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_K$

Goal : $\mathrm{f}(\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_K)$

— Loss Function

— Matrix Algebra

— ...

Introduction

Computation Offloading Framework: Opportunities and Bottlenecks

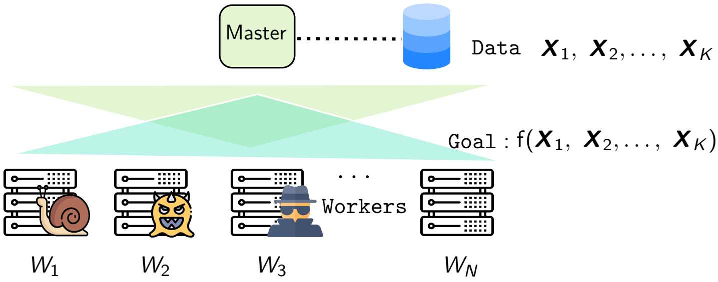

Challenges?

- Stragglers
- Adversaries
- Privacy Concerns
- Comm Efficiency
- ...

Introduction

Computation Offloading Framework: Opportunities and Bottlenecks

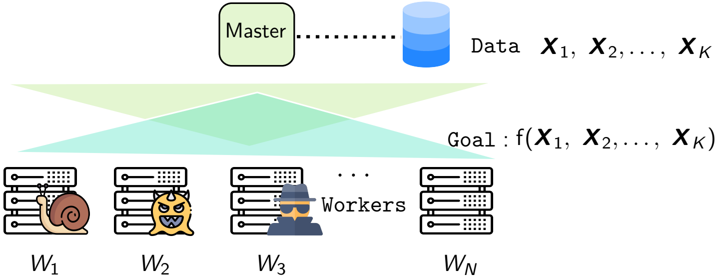

Challenges?

- Stragglers
- Adversaries
- Privacy Concerns
- Comm Efficiency
- ...

How to address these challenges?

Introduction

Computation Offloading Framework: Adding Redundancy

Redundancy → Coded version of data (~~raw data~~)

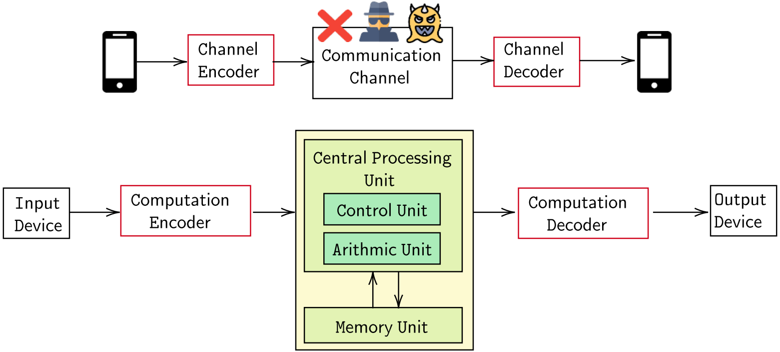

Adapted from [S. Avestimehr, ICML 2019, SlidesLive]

Introduction

Computation Offloading Framework: Adding Redundancy

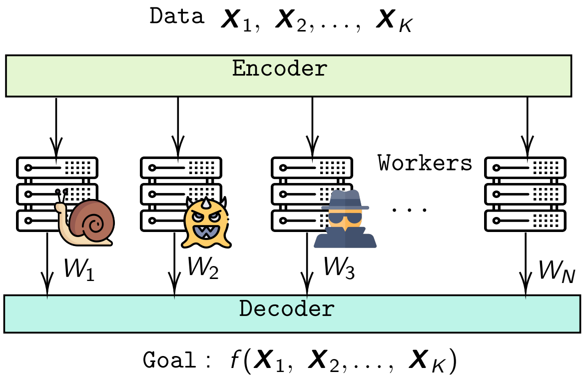

How can data be encoded s.t. computations performed on them remain meaningful?

Introduction

Computation Offloading Framework: Adding Redundancy

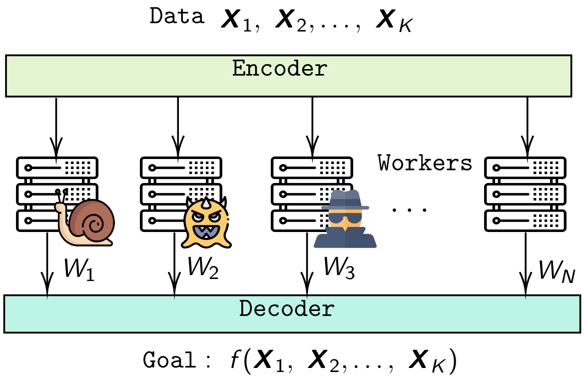

Codes over Finite Field

- Short-Dot (Dutta et al., 2016)
- Polynomial Codes (Yu et al., 2017)
- **LCC (Yu et al., 2019)**
- CSA Codes (Jia and Jafar, 2021)
- ...

Coded Computation over Finite Field

Lagrange Coded Computing (LCC)

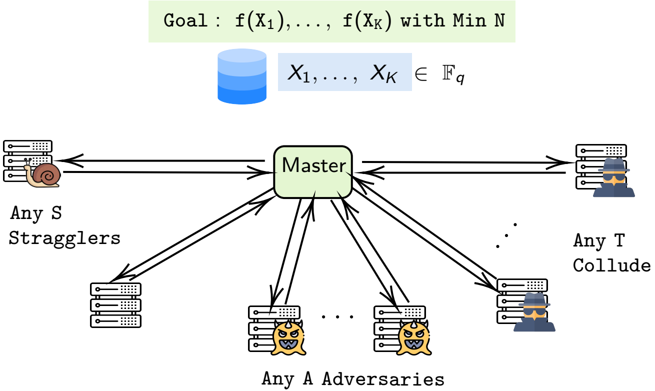

Q. Yu, S. Li, N. Raviv, S. M. M. Kalan, M. Soltanolkotabi, and S. A. Avestimehr, "Lagrange coded computing: Optimal design for resiliency, security, and privacy," in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1215–1225

Coded Computation over Finite Field

Lagrange Coded Computing (LCC)

Theorem - LCC (Yu et al., 2019)

Given N workers and a dataset $\boldsymbol{X} = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_K)$, LCC framework provides an S-resilient, A-secure, and T-private scheme for computing $\{f(\boldsymbol{X}_i)\}_{i=1}^{K}$ for any polynomial f functoion, as long as

$$(K + T - 1)\deg f + S + 2A + 1 \leq N.$$

Coded Computation over Finite Field

Lagrange Coded Computing (LCC)

- Embedding Data with some randomness into a polynomial

$$u(z) \triangleq \sum_{j\in[K]} \boldsymbol{X}_j \cdot \prod_{k\in[K+T]\setminus\{j\}} \frac{z-\beta_k}{\beta_j-\beta_k} + \sum_{j=K+1}^{K+T} \boldsymbol{U}_j \cdot \prod_{k\in[K+T]\setminus\{j\}} \frac{z-\beta_k}{\beta_j-\beta_k}$$

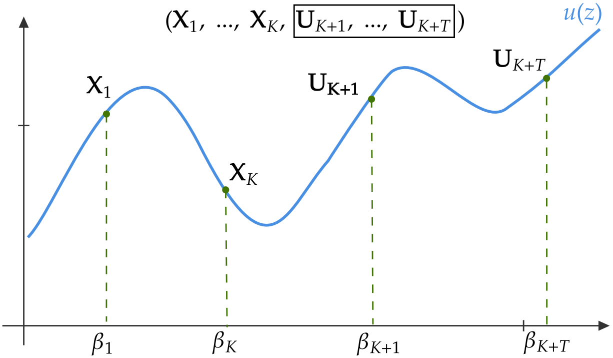

Coded Computation over Finite Field

Lagrange Coded Computing (LCC)

- Select N distinct points on the polynomial $u(z)$.
- Assign each worker a coded data point and have them compute on it: $f(\tilde{\boldsymbol{X}}_1), \ldots, f(\tilde{\boldsymbol{X}}_N)$

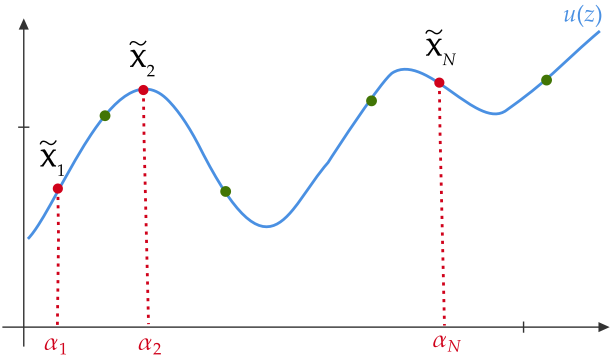

Coded Computation over Finite Field

Lagrange Coded Computing (LCC)

- $\{f(\tilde{\boldsymbol{X}}_i)\}_{i=1}^{N}$, as well as $\{f(\boldsymbol{X}_i)\}_{i=1}^{K}$, lie on $f(u(z))$
- It is enough to interpolate $f(u(z))$
- $N \geq$ #points needed for $f(u(z))$ interpolation

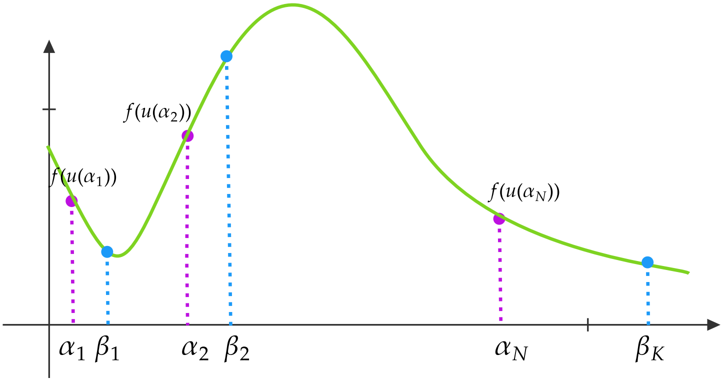

Coded Computation over Finite Field

Challenges and Limitations (Moradi et al., 2024)

- Threshold-dependent: If the number of workers drops below a threshold, recovery fails
- Only works for specific types of computations
- All solutions apply to finite fields: Quantization can cause significant accuracy loss

`Real Field` → `Finite Field` → `Real Field`

- “These methods are unsuitable for approximate computing, where exact computation is neither possible nor necessary”

Coded Computation over Finite Field

Challenges and Limitations (Moradi et al., 2024)

- Threshold-dependent: If the number of workers drops below a threshold, recovery fails
- Only works for specific types of computations
- All solutions apply to finite fields: Quantization can cause significant accuracy loss

Real Field → Finite Field → Real Field

- "These methods are unsuitable for approximate computing, where exact computation is neither possible nor necessary"

> Idea: Compute in the real field instead!

Coded Computation over Real Field

- Can we naively perform these operations in the real field?
  - Lagrange interpolation solves a linear system of equations with a Vandermonde matrix.
  - Condition number grows exponentially with matrix size

$$V = \begin{bmatrix} 1 & a_1 & a_1^2 & \cdots & a_1^{n-1} \\ 1 & a_2 & a_2^2 & \cdots & a_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & a_n & a_n^2 & \cdots & a_n^{n-1} \end{bmatrix}$$

Severe Numerical Instability!

V. Y. Pan, "How bad are vandermonde matrices?" *SIAM Journal on Matrix Analysis and Applications*, vol. 37, no. 2, pp. 676–694, 2016

Coded Computation over Real Field

Works in the literature (Moradi et al., 2024)

- Modifying coding mechanisms to improve numerical stability
  - LCC over real field (Soleymani et al., 2021)
  - Chebyshev polynomials instead of monomial basis (Fahim and Cadambe, 2021)
  - Structured matrices for evaluation points (Ramamoorthy and Tang, 2022)
- Sacrificing exactness and using approximation techniques
  - Embedding the data into a smooth rational function (Jahani-Nezhad and Maddah-Ali, 2023)
  - Embedding the data into a bigger class of functions, i.e., second order Sobolev space (Moradi et al., 2024)
- Approximating non-polynomial functions using polynomials (So et al., 2021)

Our Ongoing Research

- Not all challenges associated with finite-field computation have been fully resolved.
- In the analog domain, existing works address either straggler mitigation, Byzantine robustness, both issues, or privacy. A unified framework that tackles all aspects is still missing (Ulukus et al., 2022).
- Except for a few recent works, existing schemes are coding-theoretic rather than learning-theoretic.
- Most of the existing works in the literature have studied perfect privacy.

Our Ongoing Research

Meanwhile, we are also exploring another research direction in SweWIN's area 4, Resilience and Security, focusing on studying the fundamental limits of designing *fair representations* under different notions of fairness:

A. Zamani, A. Changizi and M. Skoglund, "*On information theoretic fairness: From perfect to bounded demographic parity*," IEEE Transactions on Information Theory. Submitted August 2025.

A. Zamani, A. Changizi, R. Thobaben, and M. Skoglund, "*Information-theoretic fairness with a bounded statistical parity constraint*," in Proc. IEEE WiOpt 2025.

References I

F. Duarte. (2025) Amount of data created daily. [Online]. Available: https://explodingtopics.com/blog/data-generated-per-day

S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.

Q. Yu, M. Maddah-Ali, and S. Avestimehr, "Polynomial codes: an optimal design for high-dimensional coded matrix multiplication," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

Q. Yu, S. Li, N. Raviv, S. M. M. Kalan, M. Soltanolkotabi, and S. A. Avestimehr, "Lagrange coded computing: Optimal design for resiliency, security, and privacy," in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1215–1225.

Z. Jia and S. A. Jafar, "Cross subspace alignment codes for coded distributed batch computation," *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 2821–2846, 2021.

P. Moradi, B. Tahmasebi, and M. A. Maddah-Ali, "Coded computing for resilient distributed computing: A learning-theoretic framework," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 111 923–111 964.

V. Y. Pan, "How bad are vandermonde matrices?" *SIAM Journal on Matrix Analysis and Applications*, vol. 37, no. 2, pp. 676–694, 2016.

References II

M. Soleymani, H. Mahdavifar, and A. S. Avestimehr, "Analog lagrange coded computing," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 283–295, 2021.

M. Fahim and V. R. Cadambe, "Numerically stable polynomially coded computing," *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 2758–2785, 2021.

A. Ramamoorthy and L. Tang, "Numerically stable coded matrix computations via circulant and rotation matrix embeddings," *IEEE Transactions on Information Theory*, vol. 68, no. 4, pp. 2684–2703, 2022.

T. Jahani-Nezhad and M. A. Maddah-Ali, "Berrut approximated coded computing: Straggler resistance beyond polynomial computing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 111–122, 2023.

J. So, B. Güler, and A. S. Avestimehr, "CodedPrivateML: A fast and privacy-preserving framework for distributed machine learning," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 441–451, 2021.

S. Ulukus, S. Avestimehr, M. Gastpar, S. A. Jafar, R. Tandon, and C. Tian, "Private retrieval, computing, and learning: Recent progress and future challenges," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 3, pp. 729–748, 2022.

Thank you!

Questions or comments?