

SEMINARIO

Adrián Fidalgo-Díaz

Universidad de Valladolid

Distributed matrix multiplication with straggler mitigation using multivariate polynomials

Abstract: We consider the problem of multiplying matrices using a parallelizable algorithm. These algorithms split the original multiplication in smaller ones, share them with a bunch of worker nodes who simultaneously compute each part and, finally, the original product is recovered gathering all the results. When the amount of worker nodes increases, also does the expected difference between their execution times (more nodes, more chance that one of them becomes anomaly slow). This motivates designing an algorithm that does not need all the results for retrieving the matrix multiplication, only the results from the first responsive nodes.

This problem can be understood from the coding theory perspective as correcting the erasures of a word: each worker node computes a symbol, with the unreceived symbols from slower nodes representing erasures. With this idea, an algorithm using Reed-Solomon codes was proposed in [1] and in our posterior generalization using AG codes [2]. Both methods have different pros and cons, but they agree on some parameter restrictions that made them unapplicable for multiplying matrices defined over some small fields. This motivates the constuction of algorithms filling this gap.

In this work we suggest a different generalization of [1] that uses evaluation codes from polynomials in several variables (similar from going from Reed-Solomon to Reed-Muller codes). We translate the problem of defining such algorithm to a combinatorial context and we use it for constructing the "code" and studying its "minimum distance". We obtain bounds and some constructions attaining them, specially when matrices are defined over \mathbb{F}_2 , a very relevant case since this is the "field of bits" and the existing algorithms do not work with it. This is a joint work with Umberto Martínez-Peñas.

[1] Yu, Qian and Maddah-Ali, Mohammad and Avestimehr, Salman. Polynomial codes: an optimal design for high-dimensional coded matrix multiplication. In: Advances in Neural Information Processing Systems, 30 (2017)

[2] Fidalgo-Díaz, Adrián and Martínez-Peñas, Umberto. Distributed matrix multiplication with straggler tolerance using algebraic function fields. In: arXiv preprint arXiv:2401.13573 (2024).

Seminario del IMUVa, edificio LUCIA
Jueves 27 de Junio de 2024 (13:00)

