

Title:

Solving very large scale discrete optimal transport problems in linear time

Speaker:

Filippo Zanetti
School of Mathematics, University of Edinburgh
F.Zanetti@sms.ed.ac.uk
<https://www.maths.ed.ac.uk/~s1990287/>

Abstract:

Discrete Optimal Transport problems give rise to very large linear programs (LP) with a particular structure of the constraint matrix. In this talk we present a hybrid algorithm that mixes an interior point method (IPM) and column generation, specialized for the LP originating from the Kantorovich Optimal Transport problem. Knowing that optimal solutions of such problems display a high degree of sparsity, we propose a column-generation-like technique to force all intermediate iterates to be as sparse as possible. The algorithm is implemented nearly matrix-free. Indeed, most of the computations avoid forming the huge matrices involved and solve the Newton system using only a much smaller Schur complement of the normal equations. We prove theoretical results about the sparsity pattern of the optimal solution, exploiting the graph structure of the underlying problem. We use these results to mix iterative and direct linear solvers efficiently, in a way that avoids producing preconditioners or factorizations with excessive fill-in and at the same time guaranteeing a low number of conjugate gradient iterations. We compare the proposed method with two state-of-the-art solvers and show that it can compete with the best network optimization tools in terms of computational time and memory usage. We perform experiments with problems reaching more than four billion variables and demonstrate the robustness of the proposed method.

References:

- [1] F. Zanetti and J. Gondzio,
An interior-point-inspired algorithm for linear programs
arising in discrete optimal transport,
(submitted for publication: 22 June 2022, revised 06 December 2022)
<https://arxiv.org/abs/2206.11009>

- [2] F. Zanetti and J. Gondzio,
A new stopping criterion for Krylov solvers applied in interior point
methods,
SIAM J on Sci Computing, (accepted: 20 October 2022).
<https://arxiv.org/abs/2106.16090>