Dynamic Spatial Matching

YASH KANORIA, Columbia University, USA

Motivated by a variety of online matching platforms, we consider demand and supply units which are located i.i.d. in d-dimensional space, and each demand unit needs to be matched with a supply unit. The goal is to minimize the expected average distance between matched pairs (the "cost"). We model dynamic arrivals of one or both of demand and supply with uncertain locations of future arrivals, and characterize the scaling behavior of the achievable cost in terms of system size (number of supply units), as a function of the dimension d. Our achievability results are backed by concrete matching algorithms. Across cases, we find that the platform can achieve cost (nearly) as low as that achievable if the locations of future arrivals had been known beforehand. Furthermore, in all cases except one, cost nearly as low as the expected distance to the nearest neighboring supply unit is achievable, i.e., the matching constraint does not cause an increase in cost either. The aberrant case is where only demand arrivals are dynamic, and space is one-dimensional $d = 1$; excess supply significantly reduces cost in this case.

We now elaborate a little on our settings and results. We study the following three settings: (i) Static matching between N demand units and $N + M$ supply units, where all demand units must be matched. (ii) A *semi-dynamic* model where $N + M$ supply units are present beforehand and N demand units arrive sequentially and must be matched immediately. (iii) A *fully dynamic* model where there are always m supply units present in the system, one supply and one demand unit arrive in each period, and the demand must be matched immediately. We show that one can achieve nearly the same cost under the semi-dynamic model as under the static model, despite uncertainty about the future, and that, under these two models, $d = 1$ is the only case where cost far exceeds the expected distance to the nearest neighboring supply unit (which is $\Theta(1/N^{1/d})$) and where adding excess supply M substantially reduces cost (by smoothing stochastic fluctuations at larger spatial length scales). In the fully dynamic model, we show that, remarkably, for all d we can achieve a cost only slightly more than the optimistic expected distance to the nearest neighbor $\Theta(1/m^{1/d})$. Thus, excess supply m reduces cost in the fully dynamic model for all d by reducing the distance to the nearest neighbor. This is a fundamentally different phenomenon than that seen in the other two models, where – only for $d = 1$ – excess supply reduces cost while leaving the distance to the nearest neighbor unchanged. Our achievability results are based on analysis of a certain "Hierarchical Greedy" algorithm which separately handles stochastic demand-supply imbalances at different length scales.

The full version of the paper can be found at https://arxiv.org/abs/2105.07329.

CCS Concepts: • **Theory of computation** → *Random network models*; *Online algorithms*; • **Mathematics of computing** → *Stochastic processes.*

Additional Key Words and Phrases: dynamic matching, distance, random spatial heterogeneity, length scales

ACM Reference Format:
Yash Kanoria. 2022. Dynamic Spatial Matching. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3490486.3538278

EC '22, July 11–15, 2022, Boulder, CO, USA

ACM ISBN 978-1-4503-9150-4/22/07.
https://doi.org/10.1145/3490486.3538278

ACKNOWLEDGMENTS

The author gratefully acknowledges the support of the National Science Foundation via grant CMMI-1653477.