On the Optimality of Greedy Policies in Dynamic Matching

SÜLEYMAN KERIMOV, Stanford University, USA
ITAI ASHLAGI, Stanford University, USA
ITAI GURVICH, Northwestern University, USA

We study centralized dynamic matching markets with finitely many agent types and heterogeneous match values. Delaying actions to accumulate "inventory" creates a positive externality from forming future matches that generate high value. This delay, however, inevitably compromises short-term value. The goal of this paper is to shed light on this tension within the family of two-way matching networks.

In our model, agents arrive sequentially to the market. The type of an agent is observed upon arrival and independently drawn from a given distribution over finitely many types; we associate each type with a queue that holds waiting agents of that type. A network topology describes which pairs of agent types can match. We assume that agents leave the market when they are matched. A matching policy determines when and which matches to form.

To evaluate a matching policy and the tension between short- and long-term value, we use the notion of regret. The regret at a given time t is measured by the difference between the (expected) total value that can possibly be generated and the (expected) total value generated by the policy until time t. The existence of a policy that achieves a "small" regret *at all times* suggests that the tension between short- and long-term value is essentially moot. We refer to such a policy as *hindsight optimal*.

We first show that the greedy longest-queue policy with a minor variation is hindsight optimal. Importantly, the policy is greedy relative to a residual network, which includes only non-redundant matches with respect to the static optimal matching rates. Moreover, when the residual network is acyclic (e.g., as in two-sided networks), we prescribe a greedy static priority policy that is also hindsight optimal. The priority order of this policy is robust to arrival rate perturbations that do not alter the residual network.

The networks considered in this paper are two-way (each match includes two agents) and satisfy a *general position condition*. General position is a weak (but necessary) condition that holds when the static-planning problem (a linear program that optimizes the first order matching rates) is locally stable and has a non-degenerate optimal solution.

The full version of this paper is available at https://ssrn.com/abstract=3918497.

CCS Concepts: • **Mathematics of computing** → **Queueing theory**; **Markov processes**; • **Theory of computation** → **Design and analysis of algorithms**; **Stochastic control and optimization**.

Additional Key Words and Phrases: dynamic matching markets

ACM Reference Format:
Süleyman Kerimov, Itai Ashlagi, and Itai Gurvich. 2022. On the Optimality of Greedy Policies in Dynamic 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, 1 page. https://doi.org/10.1145/3490486.3538323

ACKNOWLEDGMENTS

This work is supported by the Department of Defense grant STTR A18B-T007 and by NSF grant CMMI-2010940.

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

ACM ISBN 978-1-4503-9150-4/22/07.
https://doi.org/10.1145/3490486.3538323