

Online Stochastic Matching

Report of Contributions

Contribution ID: 10 Contribution code: 5

Type: **not specified**

Characterizing the Typewise Top-Trading-Cycles Mechanism for Multiple-Type Housing Markets

Tuesday, September 24, 2024 3:30 PM (45 minutes)

We consider the generalization of the classical Shapley and Scarf housing market model (Shapley and Scarf, 1974) to so-called multiple-type housing markets (Moulin, 1995). Throughout the paper, we focus on strict preferences. When preferences are separable, the prominent solution for these markets is the typewise top-trading-cycles (tTTC) mechanism.

We first show that for lexicographic preferences, a mechanism is unanimous (or onto), individually rational, strategy-proof, and non-bossy if and only if it is the tTTC mechanism. Second, we obtain a corresponding characterization for separable preferences. We obtain additional characterizations when replacing [strategy-proofness and non-bossiness] with self-enforcing group (or pairwise) strategy-proofness. Finally, we show that for strict preferences, there is no mechanism satisfying unanimity, individual rationality, and strategy-proofness. We obtain further impossibility results for strict preferences based on weakening unanimity to ontonecess and on extending the tTTC solution.

Our characterizations of the tTTC mechanism constitute the first characterizations of an extension of the prominent top-trading-cycles (TTC) mechanism to multiple-type housing markets.

Co-authors: KLAUS, Bettina; FENG, Di (Dongbei University of Finance and Economics)

Presenter: KLIJN, Flip

Session Classification: Afternoon Session

Contribution ID: 11 Contribution code: 6

Type: **not specified**

Stable Solutions for Kidney Exchanges

Tuesday, September 24, 2024 4:15 PM (45 minutes)

In this talk we will survey the main results of three recent papers on stable solutions for kidney exchange programmes (KEPs). These programmes have been established in most of the developed countries in the world during the last two decades to facilitate the exchange of living kidney donors for those recipients who have willing but immunologically incompatible donors. The UK has the largest KEP in Europe, where optimal solutions with 2-cycles, 3-cycles and altruistic chains are computed for about 300-400 registered pairs in every three months. The concept of stability (or core in the corresponding cooperative game) is an alternative approach, that can provide group fairness and good incentives for the recipients to bring more valuable donors or multiple registered donors to the pool. In this talk, besides some new theorems on the so-called respecting improvement property, we will present novel IP techniques for computing stable (core) solutions for bounded and unbounded exchanges. Furthermore, we will show simulation results on the price of stability, and on the question to what extent the respecting improvement property holds for various solution concepts regarding the stability and optimality requirements.

Presenter: BIRO, Peter (HUN-REN KRTK)

Session Classification: Afternoon Session

Contribution ID: 12 Contribution code: 24

Type: **not specified**

Greedy Algorithm for Multiway Matching with Bounded Regret

Friday, September 27, 2024 4:15 PM (45 minutes)

We consider a finite horizon online resource allocation/matching problem where the goal of the decision maker is to combine resources (from a finite set of resource types) into feasible configurations. Each configuration is specified by the number of resources consumed of each type and a reward. The resources are further subdivided into three types - offline, online-queueable (which arrive online and can be stored in a buffer), and online-nonqueueable (which arrive online and must be matched on arrival or lost). We prove the efficacy of a simple greedy algorithm when the corresponding static planning linear program (SPP) exhibits a non-degeneracy condition called the general position gap (GPG). In particular we prove that, (i) our greedy algorithm gets bounded any-time regret when no configuration contains both an online-queueable and an online-nonqueueable resource, and (ii) $O(\log t)$ expected any-time regret otherwise (we also prove a matching lower bound). By considering the three types of resources, our matching framework encompasses several well-studied problems such as dynamic multi-sided matching, network revenue management, online stochastic packing, and multiclass queueing systems.

Presenter: GUPTA, Varun (Northwestern University)

Session Classification: Afternoon Session

Contribution ID: 13 Contribution code: 11

Type: **not specified**

Optimal Queueing Auctions

Wednesday, September 25, 2024 3:30 PM (45 minutes)

This paper studies theoretically the optimal design of a queue system for providing goods or services to buyers that arrive stochastically over time. In a standard auction problem, a fixed number of goods are allocated to a fixed number of potential buyers. However, many services, such as online or phone customer service, repair, food delivery, and restaurant service, take a stochastic amount of time to complete, and many goods, such as public housing, become available at stochastic times. Buyers for these services and goods also arrive stochastically over time. Specifically, we consider an M/M/1 model in which a potential buyer arrives at a Poisson rate, and the arrival of a good takes exponential-distributed time. Each buyer is privately informed of his valuation of the good, which is distributed according to some atomless distribution, and incurs a waiting cost per unit of time.

The seller commits to a Markovian mechanism that charges a monetary fee as a function of the type a buyer reports. Each feasible mechanism induces a Markov chain on the number and types of buyers in the queue. We look for a mechanism that maximizes the expected revenue at the steady state—the stationary distribution of the Markov chain. The extant literature from Operations Research considers screening buyers with a static policy not contingent on the state. We solve for a dynamically optimal screening mechanism that involves general dynamic allocations, and, in the process, we pin down the associated stationary distribution of the (infinite-dimensional) state. The optimal mechanism uses a reserve price (or minimum bid) that increases with the number of buyers in a queue and an auction to select the buyer with the highest valuations among those in the queue. The former feature means that the arrival of a new buyer either triggers the queue length to increase if all buyers have sufficiently high valuations or triggers an exit of a buyer with the lowest valuation. As the cost of waiting goes to zero, the optimal mechanism converges to the efficient auction rule with a reserve price set at the standard monopoly price.

Presenter: CHOI, Andy (Bocconi University)

Session Classification: Afternoon Session

Contribution ID: 14 Contribution code: 2

Type: **not specified**

Simulation is All You Need

Tuesday, September 24, 2024 10:30 AM (45 minutes)

Motivated by online matching markets and network revenue management (NRM) problems with many types (e.g., fulfillment optimization), we study dynamic spatial matching (DSM) in which supply and demand live in d dimensional space and need to be matched with each other dynamically. If demand and supply have the same spatial distribution, greedy matching suffices, and achieves average match distance of the same order as the distance to the nearest neighbor, *except* for the case of $d=1$ and both supply and demand arriving dynamically over time. If demand and supply have different spatial distributions, the matching constraint has bite and greedy matching fails. We introduce a unifying and practical algorithmic principle for NRM and DSM dubbed SOAR: Simulate, Optimize, Assign, Repeat, which repeatedly simulates the future to enable good matching decisions. Simulating one sample path at each stage already enables SOAR to produce near optimal regret for the majority of NRM models in the literature, and for DSM with non-atomic demand and supply distributions. For more challenging NRM and DSM models, SOAR with multiple simulated sample paths at each stage achieves near optimal regret.

Co-authors: BESBES, Omar (Columbia Business School); CHEN, Yilun (The Chinese University of Hong Kong); KUMAR, Akshit (Columbia Business School); ZHANG, Wenxin (Columbia Business School)

Presenter: KANORIA, Yash (Columbia Business School)

Session Classification: Morning Session

Contribution ID: 17 Contribution code: 20

Type: **not specified**

Algorithms for Autobidding in Online Advertising

Friday, September 27, 2024 10:30 AM (45 minutes)

Online advertising has recently grown into a highly competitive and complex multi-billion-dollar industry, with advertisers bidding for ad slots at large scales and high frequencies. This has resulted in a growing need for efficient “auto-bidding” algorithms that determine the bids on behalf of advertisers for incoming queries to maximize advertisers’ objectives subject to their specified constraints. We consider the problem from the perspective of a single advertiser, and we focus on two of the most prevalent constraints in practice: Budget and Return-on-Spend (RoS). Technically, we take the approach of the primal-dual framework and first-order low-regret dynamics from online learning, and show how the problem-specific structures in “auto-bidding” can lead to simple online algorithms with provable performance guarantees.

Presenter: WANG, Di (Google Research)**Session Classification:** Morning Session

Contribution ID: **18** Contribution code: **19**Type: **not specified**

Online Allocation in Online Advertising: Matching, Auctions, and Autobidding

Friday, September 27, 2024 9:00 AM (1 hour)

Internet ad auctions are a fascinating innovation with a huge impact. Advertising auctions match buyers and sellers under various constraints, at scale, enabling highly valuable services for their users. The talk will give an overview of this area, from online budgeted matching to the interplay between matching, autobidding, and auctions. We will present the theory and practical aspects of these settings, including optimal algorithms and equilibrium properties.

Presenter: MEHTA, Aranyak (Google Research)

Session Classification: Morning Session

Contribution ID: **19** Contribution code: **21**Type: **not specified**

Stochastic Matching: A Polytope Perspective

Friday, September 27, 2024 11:15 AM (45 minutes)

Stochastic dynamic matching has numerous applications, ranging from supply-chain management to kidney exchange programs. We consider a generic model where items of different classes arrive randomly. Two items can be matched if their classes are compatible. Unmatched items are queued.

We study the stability of such systems and the matching rates between classes. Our main results, which rely on the conservation equation between arrivals and departures, are:

- We give simple relationships between departure and arrival rates, the shape of the compatibility graph, and the existence of a stable matching policy.
- We describe the polytope of non-negative solutions of the conservation in the degenerated and non-degenerated cases. Any stable policy necessarily yields one of these solutions.
- Reciprocally, we show that any part of the polytope can be reached (possibly asymptotically) by using non-greedy policies. The same result does not hold for greedy policies.

This is a joint work with Céline Comte, Sushil Varma, and Ana Bušić.

Co-authors: COMTE, Céline (CNRS and LAAS); VARMA, Sushil (Georgia Tech); BUŠIĆ, Ana (Inria)

Presenter: MATHIEU, Fabien (Swapcard)

Session Classification: Morning Session

Contribution ID: 20 Contribution code: 4

Type: **not specified**

Happy 50th birthday TTC! & Characterizing the top trading cycles rule for housing markets with lexicographic preferences when externalities are limited

Tuesday, September 24, 2024 2:00 PM (1 hour)

In my talk, I would like to commemorate the 50th birthday of the famous top trading cycles algorithm/rule and will mix a survey with a recent paper. The abstract for the paper is as follows. We consider a housing market model with limited externalities where agents care both about their own consumption via demand preferences and about the agent who receives their endowment via supply preferences (we extend the associated lexicographic preference domains introduced in Klaus and Meo, 2023). If preferences are demand lexicographic, then our model extends the classical Shapley-Scarf housing market (Shapley and Scarf, 1974) with strict preferences model. Our main result is a characterization of the corresponding top trading cycles (TTC) rule by individual rationality, pair efficiency, and strategy-proofness (Theorem 1), which extends that of Ekici (2023) from classical Shapley-Scarf housing markets with strict preferences to our model. Two further characterizations are immediately obtained by strengthening pair efficiency to either Pareto efficiency or pairwise stability (Corollaries 1 and 2). Finally, we show that as soon as we extend the preference domain to include demand lexicographic as well as supply lexicographic preferences (e.g., when preferences are separable), no rule satisfying individual rationality, pair efficiency, and strategy-proofness exists (Theorem 2).

Presenter: KLAUS, Bettina (University of Lausanne)

Session Classification: Afternoon Session

Contribution ID: 21 Contribution code: 14

Type: **not specified**

Online Matching in Geometric Random Graphs

Thursday, September 26, 2024 10:30 AM (45 minutes)

We investigate online maximum cardinality matching, a central problem in ad allocation. In this problem, users are revealed sequentially, and each new user can be paired with any previously unmatched campaign that it is compatible with. Despite the limited theoretical guarantees, the greedy algorithm, which matches incoming users with any available campaign, exhibits outstanding performance in practice. Some theoretical support for this practical success was established in specific classes of graphs, where the connections between different vertices lack strong correlations - an assumption not always valid. To bridge this gap, we focus on the following model: both users and campaigns are represented as points uniformly distributed in the interval $[0,1]$, and a user is eligible to be paired with a campaign if they are similar enough, i.e. the distance between their respective points is less than c/N , with $c>0$ a model parameter. As a benchmark, we determine the size of the optimal offline matching in these bipartite random geometric graphs. In the online setting and investigate the number of matches made by the online algorithm closest, which greedily pairs incoming points with their nearest available neighbors. We demonstrate that the algorithm's performance can be compared to its fluid limit, which is characterized as the solution to a specific partial differential equation (PDE). From this PDE solution, we can compute the competitive ratio of closest, and our computations reveal that it remains significantly better than its worst-case guarantee. This model turns out to be related to the online minimum cost matching problem, and we can extend the results to refine certain findings in that area of research. Specifically, we determine the exact asymptotic cost of closest in the ϵ -excess regime, providing a more accurate estimate than the previously known loose upper bound.

Presenter: SENTENAC, Flore (HEC Paris)

Session Classification: Morning Session

Contribution ID: 22 Contribution code: 1

Type: **not specified**

Challenges in Dynamic Matching

Tuesday, September 24, 2024 9:00 AM (1 hour)

Presenter: ASHLAGI, Itai (Stanford)

Session Classification: Morning Session

Contribution ID: 23 Contribution code: 3

Type: **not specified**

Designing Efficient Interview Graphs for Interim Stable Matchings: A Signaling Approach

Tuesday, September 24, 2024 11:15 AM (45 minutes)

In many two-sided matching markets agents initially engage in costly interviews in order to refine their interim preferences. We study how signaling mechanisms can reduce the number of interviews in random markets. We are interested in interim stability, which expands the notion of stability to ensure that no two agents regret not interviewing with each other. The final match is almost interim stable if it is interim stable after removing a vanishing small fraction of agents.

For single-tiered markets, we study the impact of short-side (resp. long-side, both-side) signaling, where agents from short-side (resp. long-side, both-side) of the market signal their top d preferred candidates. The interview graphs are formed by including all pairs where at least one party has signaled to the other. When $d = \omega(1)$, we show that short-side signaling leads to almost interim stable matchings for every stable matching. Long-side signaling is effective in weakly imbalanced markets but fails in strongly imbalanced markets. We also demonstrate the failure of both-side signaling when the impact of interviews is negligible. When $d \geq \Omega(\log^2 n)$, short-side signaling achieves interim stability, while long-side signaling fails in imbalanced markets. We extend our results to multi-tiered markets where a proposed signaling mechanism achieves interim stability across tiers, and identify conditions for truthful signaling.

Our analysis for sparse interview markets develops a message-passing algorithm that efficiently determines interim stability and match outcomes by leveraging their local neighborhood structure.

Co-authors: ALLMAN, Maxwell (Stanford); ASHLAGI, Itai (Stanford); SABERI, Amin (Stanford)

Presenter: YU, Sophie (University of Pennsylvania)

Session Classification: Morning Session

Contribution ID: 24 Contribution code: 7

Type: **not specified**

Dynamic Matching with Post-Allocation Service and its Application to Refugee Resettlement

Wednesday, September 25, 2024 9:00 AM (1 hour)

Motivated by our collaboration with a major refugee resettlement agency in the U.S., we study a dynamic matching problem where each new arrival (a refugee case) must be matched immediately and irrevocably to one of the static resources (a location with a fixed annual quota). In addition to consuming the static resource, each case requires post-allocation services from a server, such as a translator. Given the uncertainty in service time, a server may not be available at a given time, thus we refer to it as a dynamic resource. Upon matching, the case will wait to avail service in a first-come-first-serve manner. Bursty matching to a location may result in undesirable congestion at its corresponding server. Consequently, the central planner (the agency) faces a dynamic matching problem with an objective that combines the matching reward (captured by pair-specific employment outcomes) with the cost for congestion for dynamic resources and over-allocation for the static ones. Motivated by the observed fluctuations in the composition of refugee pools across the years, we aim to design algorithms that do not rely on distributional knowledge. We develop learning-based algorithms that are asymptotically optimal in certain regimes, easy to interpret, and computationally fast. Our design is based on learning the dual variables of the underlying optimization problem; however, the main challenge lies in the time-varying nature of the dual variables associated with dynamic resources. Our theoretical development brings together techniques from Lyapunov analysis, adversarial online learning, and stochastic optimization. On the application side, when tested on real data from our partner agency, our method outperforms existing ones making it a viable candidate for replacing the current practice upon experimentation.

Presenter: MANSHADI, Vahideh (Yale School of Management)

Session Classification: Morning Session

Contribution ID: 25 Contribution code: 8

Type: **not specified**

Adaptive Policies and Approximation Schemes for Dynamic Matching

Wednesday, September 25, 2024 10:30 AM (45 minutes)

In the dynamic matching problem, customers arrive according to a Poisson process and must be matched to available servers within a network of queues. These servers independently replenish and renege from the market. Previous literature on the dynamic matching problem often concentrates on static policies—where the matching decisions are not adapted to dynamic information on queue lengths. In this work, we formulate tighter, yet efficient linear programming relaxations that optimize an adaptive matching process, where the queue lengths inform the matching decisions in real-time.

We develop two applications of this linear programming (LP) framework that provide new formal approximation algorithms for the dynamic matching problem. First, we devise a bi-criteria polynomial-time approximation scheme for a cost-minimization version of dynamic matching on Euclidian networks of fixed dimensions. The goal is to find a matching policy that approaches any feasible target for matching the cost rate and throughput rate arbitrarily closely. The approximation scheme combines our new LP framework with an instance decomposition that allocates the cost-throughput rate targets across localized subproblems. Second, we devise a $(1-1/\epsilon)$ -balanced contention resolution scheme for our linear program. This result matches the best-known approximation for the reward maximization problem, but the new matching policy is guided by queue length information and is simpler to analyze than in previous literature.

Co-authors: AMANIHAMEDANI, Alireza (London Business School); SABERI, Amin (Stanford)

Presenter: AOUAD, Ali (London Business School)

Session Classification: Morning Session

Contribution ID: 26 Contribution code: 9

Type: **not specified**

Teacher Redistribution in a Public School System

Wednesday, September 25, 2024 11:15 AM (45 minutes)

The Right to Free and Compulsory Education Act (2009) (RTE) of the Government of India prescribes teacher-student ratios for state-run schools. One method advocated by the Act to achieve its goals is the redeployment of teachers from surplus to deficit schools. I'll discuss a model in which teachers can either remain in their initially assigned schools or be transferred to a deficit school in their acceptable set. Transfers cannot turn a surplus school into a deficit school, and a deficit school cannot be made a surplus school. We show the existence of a transfer policy whose (post-transfer) deficit vector Lorenz dominates all achievable (post-transfer) deficit vectors. We also show that the Lorenz-dominant deficit vector can be achieved as the outcome of a strategy-proof mechanism.

Co-authors: MISHRA, Debasis; SARKAR, Soumendu; SEN, Arunava; YADAV, Sonal

Presenter: SETHURAMAN, Jay (Columbia University)

Session Classification: Morning Session

Contribution ID: 27 Contribution code: 10

Type: **not specified**

Optimal Queue Design

Wednesday, September 25, 2024 2:00 PM (1 hour)

We study the optimal method for rationing scarce resources through a queue system. The designer controls agents' entry into a queue and their exit, their service priority—or queueing discipline—as well as their information about queue priorities, while providing them with the incentive to join the queue and, importantly, to stay in the queue, when recommended by the designer. Under a mild condition, the optimal mechanism induces agents to enter up to a certain queue length and never removes any agents from the queue; serves them according to a first-come-first-served (FCFS) rule; and provides them with no information throughout the process beyond the recommendations they receive. FCFS is also necessary for optimality in a rich domain. We identify a novel role for queueing disciplines in regulating agents' beliefs and their dynamic incentives, and uncover a hitherto unrecognized virtue of FCFS in this regard.

Presenter: TERCIEUX, Olivier (CNRS & Paris School of Economics)

Session Classification: Afternoon Session

Contribution ID: 28 Contribution code: 12

Type: **not specified**

Optimal Stopping with Interdependent Values

Wednesday, September 25, 2024 4:15 PM (45 minutes)

Consider a single-item auction, say for a piece of art, where buyers arrive online. The goal is to sell the item to the agent with the highest value, while making the decisions about whether to select each buyer immediately on their arrival. Additionally, buyers' values can depend on one another: a buyer interested in decorating their living room might be influenced by the impression of those arriving before him, and spontaneously attribute a higher value to the item if it is very popular. A buyer that sees the item as a pure investment, on the other hand, will be interested in its resale value alone, which is fully determined only after the arrival of the very last buyer. We analyze settings of the above type, formally, of online selection processes with interdependent values. This means, we combine concepts from the areas of online selection/optimal stopping with those from the theory on interdependent values, both of which have raised a lot of recent interest due to their central and important applications in economics.

Presenter: MAURAS, Simon (Inria Saclay)

Session Classification: Afternoon Session

Contribution ID: 29 Contribution code: 16

Type: **not specified**

Match Now or Later?

Thursday, September 26, 2024 2:00 PM (1 hour)

The talk will have two parts, both focusing on centralized markets, the tension between short-term and longer-term objectives, and the value of waiting to thicken the market and increase match opportunities.

In the first part I'll talk about hindsight optimality in networks without departures. A matching policy is hindsight optimal if the policy can (nearly) maximize the total collected value, simultaneously at all times. I will present (non-greedy) policies that achieve hindsight optimality in multiway networks and greedy policies that achieve the same goal in two-way networks.

In the second part I'll augment the models with departures/abandonments and address two questions: (1) how does the finite-patience of agents impact the collected match value (this is the cost of impatience), and (2) how does impatience impact the optimality of greedy matching policies.

Co-authors: KERIMOV, Suleyman; KOHLENBERG, Angela; ASHLAGI, Itai

Presenter: GURVICH, Itai (Northwestern University)

Session Classification: Afternoon Session

Contribution ID: **30** Contribution code: **18**Type: **not specified**

Approximating Optimum Online for Online Capacitated Allocation

Thursday, September 26, 2024 4:15 PM (45 minutes)

We study online capacitated resource allocation, a natural generalization of online stochastic max-weight bipartite matching. This problem is motivated by ride-sharing and Internet advertising applications, where online arrivals may have the capacity to serve multiple offline users. Our main result is a polynomial-time online algorithm which is $(0.5+c)$ -approximate to the optimal online algorithm for $c=0.0115$. This can be contrasted to the (tight) 0.5-competitive algorithms to the optimum offline benchmark from the prophet inequality literature. Optimum online is a recently popular benchmark for online Bayesian problems which can use unbounded computation, but not “prophetic” knowledge of future inputs. A key technical challenge in deriving our guarantee is bounding the positive correlation among users introduced when rounding a LP relaxation online. Unlike in the case of unit capacities, this positive correlation is unavoidable for guarantees beyond 0.5. Conceptually, our results show that the study of optimum online as a benchmark can reveal problem-specific insights that are irrelevant to competitive analysis.

Presenter: POLLNER, Tristan (Stanford)

Session Classification: Afternoon Session

Contribution ID: 31 Contribution code: 17

Type: **not specified**

Batching and Optimal Multi-stage Bipartite Allocations

Thursday, September 26, 2024 3:30 PM (45 minutes)

In several applications of real-time matching of demand to supply in online marketplaces — for example, matching delivery requests to dispatching centers in Amazon or allocating video ads to users on YouTube — the platform allows for some latency (or there is an inherent allowance for latency) in order to batch the demand and improve the efficiency of the resulting matching in such non-stationary environments. Motivated by these scenarios, I investigate the optimal trade-off between batching and inefficiency in the context of designing online allocations under non-stationary arrivals in this talk. In particular, I consider a variant of the classic adversarial online bipartite allocation family of problems where demand requests arrive stage-wise and in K batches—in contrast to online arrival. Our main result for this family is an optimal $(1 - (1 - 1/K)^K)$ -competitive (fractional) matching algorithm, improving the classic $(1 - 1/e)$ competitive ratio bound known for the online variants of these problems (Mehta et al., 2007; Aggarwal et al., 2011; Golrezaei et al. 2014). Time permits, I briefly explain a generalization of the base result to multistage configuration allocations and its applications to video display advertising and AdX.

Our main technique at a high level is developing algorithmic tools to vary the trade-off between “greediness” and “hedging” of the matching algorithm across stages. We rely on a particular family of convex-programming-based matchings that distribute the demand in a specifically balanced way among supply in different stages, while carefully adjusting the balancedness of the resulting matching across stages. More precisely, we identify a (unique) sequence of polynomials with decreasing degrees to be used as strictly concave regularizers of the basic linear program for the greedy matching at each stage. At each stage, our multi-stage algorithm returns the corresponding regularized optimal solution as the matching of this stage (by solving the convex program). By providing structural decomposition (tied to these convex programs) of the underlying subgraph at each stage and recursively connecting the regularizers together, we develop a new multi-stage primal-dual framework to analyze the competitive ratio of this algorithm. For the extension to multi-stage configuration allocation, we develop a new idea of regularizing separately at different price levels, and we analyze the resulting regularized price-level convex programming-based multistage algorithm using a new primal-dual argument.

Presenter: NIAZADEH, Rad (University of Chicago)

Session Classification: Afternoon Session

Contribution ID: 32 Contribution code: 15

Type: **not specified**

A Framework for Weighted Matchings In Dynamic Graphs

Thursday, September 26, 2024 11:15 AM (45 minutes)

In dynamic graph algorithms, the goal is to maintain a key property of the graph while an adversary makes changes to the graph in the form of edge deletions or insertions. The goal is to minimize the update time of the algorithm, which is the time taken by the algorithm to adapt to a single adversarial edge insertion or deletion and output accordingly. I will briefly talk about some algorithms that maintain a good approximation to the maximum matching in a dynamic unweighted graph. Then, I will describe a framework that extends these results to the more general case of weighted graphs. Moreover, this reduction is applicable to a number of different models such as distributed, massively parallel and streaming models.

Presenter: DUDEJA, Aditi (University of Salzburg)

Session Classification: Morning Session

Contribution ID: **33** Contribution code: **13**

Type: **not specified**

Online Matching on Large Graphs: Theoretical and Practical Challenges

Thursday, September 26, 2024 9:00 AM (1 hour)

We discuss both some theoretical advances for online matching on random graphs and how to reach practical efficient schemes for critical applications like organ donations based on reinforcement learning and orchestration of expert policies.

Presenter: JONCKHEERE, Matthieu (CNRS & LAAS)

Session Classification: Morning Session

Contribution ID: **34** Contribution code: **22**

Type: **not specified**

Stochastic Matching Models

Friday, September 27, 2024 2:00 PM (1 hour)

Presenter: MAIRESSE, Jean (CNRS)

Session Classification: Afternoon Session

Contribution ID: 35 Contribution code: 23

Type: **not specified**

Online Stochastic Matching Under the First-Come-First-Matched Policy

Friday, September 27, 2024 3:30 PM (45 minutes)

We consider the following online stochastic matching problem: items from different classes arrive according to independent Poisson processes; matching compatibilities between items are described by a simple graph over the classes, so that two items can be matched if and only if their classes are neighbors in the graph; this graph is assumed to be connected and non-bipartite; lastly, each incoming item is matched immediately with the oldest compatible unmatched item, if any. The evolution of the sequence of unmatched item classes, ordered by arrival times, defines an irreducible Markov chain. Previous work has derived (i) a necessary and sufficient condition for this Markov chain to be positive recurrent, which we assume is satisfied, and (ii) a closed-form expression for its stationary distribution. In this presentation, we first provide more direct method to derive the stationary distribution, based on quasi-reversibility, a notion introduced by Kelly (1979). We then derive new closed-form expressions for several performance metrics, such as the waiting probability and the mean waiting time, which can be efficiently implemented using dynamic programming. Both these formulas and the numerical results that they allow us to derive are used to better understand the impact of parameters on performance. In particular, we characterize performance in a so-called heavy-traffic regime, in which the number of items of a subset of the classes goes to infinity while items of other classes become scarce.

Presenter: COMTE, Céline (CNRS and LAAS)

Session Classification: Afternoon Session

Contribution ID: **37** Contribution code: **0**

Type: **not specified**

Opening