Dynamic Allocation and Matching

Yehua Wei (Duke University)
Azarakhsh Malekian (University of Toronto)
Weina Wang (Carnegie Mellon University)

March 2nd, 2025–March 7th, 2025

1 Overview of the Field

This workshop centers on dynamic resource allocation and matching, a broad and evolving research domain that integrates elements from multiple interdisciplinary fields, including operations research, mathematics, engineering, computer science, and economics. At a high level, the researchers in this domain seek to understand how to allocate and match limited resources, agents, and requests that arrive over time, with the goal of maximizing efficiency, fairness, or other objectives in dynamic, stochastic, or adversarial environments. This domain has proven to be a fertile ground for both theoretical advancements and practical applications. On the practical side, dynamic resource allocation and matching have a broad set of impactful applications across a wide range of settings, including ride-sharing platforms, online marketplaces, e-commerce fulfillment, organ exchange programs, and online advertising. The common thread across these settings is the need to make timely and dynamic decisions under uncertainty, in the presence of strategic agents, capacity constraints, or real-time feedback.

This report highlights recent progress and emerging directions in dynamic resource allocation and matching, as presented during the workshop. We begin by introducing the core research areas under the domain of dynamic resource allocation and matching that are covered at the workshop.

Online Resource Allocation

A foundational area covered in this workshop is the study of online resource allocation. This problem lies at the intersection of combinatorial optimization, stochastic control, and algorithm design, and has been a central topic in computer science, operations research, and applied probability. At its core, the online resource allocation problem involves sequentially assigning requests arriving online to a limited pool of offline resources with the goal of maximizing cumulative reward subject to resource constraints. This framework captures a wide array of applications. A salient example is network revenue management [74], where a service provider, e.g., an airline, must dynamically allocate airfares to a dynamic arrival of itinerary demands, where each demand may consume one or multiple types of airfares, subject to a limited number of airfares. Another example is online advertising faced by a bidder in repeated auctions, where advertisers buy opportunities to display advertisement events, subject to budget constraints. Online resource allocation is also an extension of the stochastic knapsack problem, a well-studied problem in mathematics and operations research that is known for various applications in transportation and scheduling.

The study of online resource allocation algorithms can be divided into two main distinctive directions: the design of algorithms achieving multiplicative approximation and competitive ratios guarantees, and the design of algorithms achieving certain additive regret. Substantial progress has been made in both directions, as surveyed in [52] for approximation algorithms and [17] for regret-based approaches.

Dynamic Matching

Another core area explored in this workshop is dynamic matching. Unlike online resource allocation, where offline resources are matched to online arrivals, dynamic matching typically involves stochastic arrivals on both sides of the market, and agents are matched in real time based on their compatibility and system state. This often leads to queueing dynamics, with agents waiting to be matched with others in the system, possibly subject to abandonment or other time-based constraints. As a result, queueing-related models are common, and key performance metrics include agent queue lengths [7], system stability [58], and abandonment behavior [11, 15]. Another distinctive feature of dynamic matching models is that the underlying interaction network is often non-bipartite, with multi-way matching structures emerging in applications such as carpooling, housing exchanges, and labor platforms. Dynamic matching has been applied to model a wide range of real-world applications, such as kidney exchange programs, housing allocation, ride-sharing platforms, and online labor markets (see, e.g., [69, 2, 21, 60]).

Recent work has led to significant progress in policy design for systems with or without agent abandonment. For systems without abandonment, researchers have proposed batched policies [60, 61], primal-dual policies [68, 80], and sum-of-square policies [45], all of which achieve the optimal regret scaling. For systems with abandonment, researchers have proposed a fluid approximation to analyze systems with high arrival rates [15], and approximate Bellman equations to derive policies with approximation guarantees [7].

Algorithmic Mechanism Design

While the above models often assume truthful behavior and known valuations, many real-world systems involve incentives for the agents that participate in them. The field of algorithmic mechanism design is at the intersection of optimization, computer science, and economics and deals with allocating resources to strategic users. Given the applications in online advertisement (where the resources are the ad slots and the users are advertisers) and online e-commerce platforms (where the resources are goods and the users are buyers), recently, there has been a surge of interest in settings in which users are arriving at the system online, and the designer needs to decide on the mechanism to use for users dynamically.

The field of algorithmic mechanism design is also directly connected with prophet inequalities, a subfield of online resource allocation. To see this, note that the vanilla version of the prophet inequality is as follows: There are n random variables $T_1, T_2, \cdots, T_n$ with known distributions but unknown realizations. These realizations are revealed one by one, in the order from 1 to n. We want to give a strategy (which is a stopping rule) that, upon seeing the value T_i (and all the values before it), decides either to choose i, in which case we get a reward of T_i, and the process stops. Or we can pass, in which case we move on to the next items and are not allowed to come back to i ever again. The goal is to maximize our expected reward. [64], among others, present a strategy with expected reward one half of the best expected offline reward. The main idea is to guarantee that the expected probability of stopping at any T_i is uniformly bounded from below by $1/2$. The prophet inequality, draws an analogy to a simple pricing problem. In this scenario, a seller has a single item to sell to a sequence of potential buyers with private values drawn from known distributions. Buyers and the seller negotiate using various protocols, leading to decisions to sell or not, and at what price. Once the item is sold, future buyers are turned away. The prophet inequality solution establishes a take-it or leave-it selling protocol that guarantees at least half of the optimal welfare in expectation. Moreover, this approach works even if buyers' values are private and has a lot of other desirable properties (e.g., ex-post individual rationality, obviously strategyproof, group strategyproof, etc). The above connection and the direct mapping between the sequential posted pricing setting and the prophet inequality was explored in [29, 4], which demonstrated that any advancement for the prophet inequality problem directly translates into the mechanism version, taking into account the strategic behavior of the users.

Another related area in algorithmic mechanism design, which is also related to dynamic decision-making, involves the acquisition of costly information. In many economic environments with multiple options and uncertain rewards, the challenge is to balance gathering more information to reduce uncertainty about rewards with the associated costs, aiming to achieve a high-quality reward while keeping expenses low. The foundational model of optimal search, called the Pandora's box problem, was introduced by [81]. In this problem, a searcher must choose a prize from n boxes, each containing an unknown-value prize drawn from a known distribution. The searcher can perform actions like opening a box (incurring a cost and revealing the prize value) or selecting a box (yielding the prize value and ending the search). The searcher aims to create an adaptive policy to maximize expected utility, which is the expected selected prize value minus inspection costs, by making decisions based on previous actions and outcomes.

One application of this problem is in mechanism design with costly information acquisition. This includes various settings involving sellers and buyers, where one side of the market incurs costs during the investigation process. In particular, [33] tackles an auction scenario with independent buyers' valuations, reducing the seller's problem to Weitzman's model, where each box corresponds to the Myerson virtual value distribution for each buyer. [63] focus on the sale of items to informed buyers who incur inspection costs. They devise a descending price auction that achieves efficiency similar to a first-price auction with modified value distributions and no inspection costs. [5] extend the revenue maximization setting of [63] to nonobligatory inspection models, offering approximately optimal mechanisms for both single and multiple items, even when buyers arrive in an adversarial order. [13] analyzes buyer-seller interactions in an online marketplace scenario, where buyers determine their value and product prices upon inspection. Armstrong focuses on sellers' strategy decisions and their implications in both monopolistic and competitive settings. [31] examines a pricing game where sellers with substitutable items compete while buyers have partial information about their values. They characterize buyers' optimal behavior and analyze the pricing game among sellers. [30] propose a three-step mechanism for manufacturers outsourcing production to suppliers, with a costly investigation step resembling a variant of the Pandora's box problem. They identify conditions under which Pandora's rule can be optimal for buyer investigations and use its structural properties to design the optimal mechanism. These papers collectively explore diverse scenarios and applications of Pandora's box problem in mechanism design, shedding light on optimizing decision-making processes in various economic settings.

2 Recent Developments and Open Problems

Competitive Analysis for Online Resource Allocation

There have been significant advances in the analysis of approximation and competitive ratios for online resource allocation problems. A major stream of progress stems from the design and analysis of primal-dual algorithms, initiated by the pioneering work of [24, 23, 25] for various online matching problems.

Recent developments can be broadly classified into several categories. In one direction, researchers have developed tighter linear programming (LP) relaxations that improve competitive ratio analyses. These improved relaxations have been applied to settings such as online stochastic matching with i.i.d. arrivals [50, 51], online stochastic matching with correlated arrivals [9], and online matching under high-variance arrival distributions [16]. Moreover, more expressive linear programming relaxations have also been proposed to improve the approximation guarantees relative to dynamic programming benchmarks (the so-called philosopher inequalities), as seen in [70, 75, 22]. Another stream of work extends the primal-dual framework to more general settings. For instance, primal-dual methods have been successfully adapted to vertex-weighted online matching while preserving competitive ratios [34, 53, 50], and further augmented with convex regularization to handle batched arrivals [40] and cancellation of prior commitments [36].

A separate line of development has emerged around creating certificates of optimality that do not rely on LP duality. Originally introduced in [44] for online resource allocation with reusable resources, this framework has since been applied to settings including online matching with stochastic rewards and unknown budgets [76], online assortment optimization [10], inventory pricing [43], and volunteer matching platforms [66].

While most advances to date have focused on online matching and online resource allocation problems, it would be interesting to investigate whether these techniques can be extended to dynamic matching settings, with additional queueing dynamics that arise naturally. In particular, the tighter linear programming relaxations and certificate-based methods developed for online matching could potentially enhance both algorithm design and theoretical guarantees in dynamic matching models.

Regret Analysis for Online Resource Allocation

The study of regret bounds for online resource allocation has also seen major progress in recent years. Early foundational work [56, 54, 55] demonstrated that problems such as network revenue management, a simple re-solving policy that updates the fluid relaxation at every time period can achieve constant regret, independent of the time horizon T, under a non-degeneracy assumption on the initial fluid solution and a fluid scaling. [17] provided a streamlined and unified framework for analyzing such problems, extending the results to a broader class of resource allocation models. Recently, the works of [45] and [46] showed that sum-of-squares and primal-dual-based policies can achieve constant regret guarantees without requiring re-solving, under similar conditions.

Further advances have relaxed the need for non-degeneracy conditions, in the setting where the number of request types is finite. Constant regret policies have been developed for settings such as the multi-secretary problem [12] and network revenue management [26]. For general online resource allocation problems, the *Bayes selector*, also known as “Re-solve and Act Based on Bellman Inequalities”, achieves constant regret without non-degeneracy assumptions [78, 79]. While similar to traditional re-solving policies in recomputing fluid relaxations, the Bayes selector differs in that it employs index-based action selection derived from the primal fluid solution, rather than probabilistic decision rules.

The regret analysis literature has mainly focused on the regime where T and resource capacities are scaled to infinity while keeping the number of request types fixed. An important open direction is to understand how regret bounds behave when additional system parameters, such as the number of request types, also grow large. Progress on this front may require combining tools from both competitive and regret analysis traditions. Moreover, recent work [45, 46] has revealed connections between dynamic resource allocation policies and scheduling policies in stochastic queueing control (e.g., [49, 37]). Further exploration of this connection could inspire new algorithmic frameworks that achieve superior theoretical guarantees and practical performance in large-scale dynamic systems. In addition, recent advances in asymptotic analysis for restless bandit problems under average reward criteria [47, 48] suggest promising new directions for studying online resource allocation problems when resource states or rewards are not fully observable.

Prophet Inequality and Online Contention Resolution Schemes

Significant advancements have been made in prophet inequalities and in particular, a specific setting known as the Online Contention resolution schemes (OCRS). In OCRS, each request would arrive and be “active” with some probability, and one needs to design algorithm that selects the active requests with as high a probability as possible. OCRS and its equivalent prophet inequality setting have been shown to have direct implications for mechanism design and the fairness of resource allocation. To this end, [29] focused on the Bayesian mechanism with a single-item auction, which is later extended to k-items in [4] and was further improved in [57].

The OCRS has also been shown to be fundamentally connected with online resource allocation, as linear programming relaxations typically suggest an ex-ante probability to activate the requests, and combined with an OCRS algorithm with acceptance probability c, it would naturally lead to an online resource allocation algorithm with a competitive ratio of c. The OCRS was considered in [39] to study an online resource allocation problem under general matroid constriants, the OCRS algorthms has since been improved for adversarial vertex arrivals in online matching [38], random-order vertex arrivals online matching for both general and bipartite graph [65], and various edge arrival online matching models [72, 32, 65].

It is worth noting that the analysis of prophet inequalities mainly rely on the competitive ratio for worst-case arrival, and there is an opportunity to identify another benchmark or condition that is similar to the studies in [70, 22]. Additionally, examining scenarios where users arrive in batches or have the potential to arrive multiple times could yield valuable insights and further refine our understanding of the prophet inequality in more complex and dynamic contexts.

Dynamic Matching with Uniform Matching Rewards

Dynamic matching models with uniform rewards have received significant attention in recent years. This line of work is closely connected to the study of two-sided matching queue models [28, 1, 77], which can be viewed as a special case of dynamic matching where all match types yield identical rewards. A central objective in this setting is to minimize agents' queueing times or to maximize the total number of successful matches under the possibility of agent abandonment. Research in this area spans both fixed and random matching network structures.

In fixed matching networks, [67] developed a product-form stationary distribution under the "first-come first-matched" policy, while [27] characterized the optimal matching policies for specific bipartite network structures. Motivated by ride-sharing applications, [59] studied a closed-queueing network and proposed a near-optimal mirror backpressure policy. In addition, stability conditions for general dynamic matching models have been investigated in [58, 18, 73].

In random matching networks, where agents form viable matches randomly upon arrival, [8] analyzed two-way and three-way cycles as well as chains in homogeneous settings, while [14] extended the analysis to heterogeneous agents. Further, [3] introduced models incorporating agent abandonment to study how market thickness affects matching efficiency.

Despite these advances, many important questions remain open. For instance, optimal or near-optimal policies are not yet fully characterized for several general dynamic matching models with either fixed or random networks, especially when incorporating realistic features such as agent impatience, abandonment, and the incentives of the agents in revealing their private information.

Dynamic Matching with Heterogeneous Rewards

Beyond models with uniform rewards, researchers have increasingly focused on dynamic matching problems where match types are associated with heterogeneous rewards. Compared to uniform reward models, heterogeneous rewards introduce a richer set of trade-offs between immediate and delayed matching decisions. In systems without agent abandonment, one line of work aims to bound regret uniformly over time horizons. Batched matching policies have been shown to achieve constant regret at all times [60, 61], with similar guarantees established for sum-of-squares policies [45] and primal-dual-based policies [80].

In dynamic matching models with agent abandonment, one approach is to use large-market approximations. For example, [69] identified asymptotically optimal continuous linear programming-based policies under fluid scaling, while [15] designed fluid-optimal matching policies for heterogeneous agents. Another successful approach is to come up with competitive algorithms against offline benchmarks. To this end, [71] developed a competitive algorithm for the stationary matching problem, achieving a $1 - 1/\sqrt{e}$ ratio, and this is later improved to a strictly better ratio in [6]. For dynamic stochastic models, [62] showed that static threshold-based policies can approximate optimal reward up to a factor of 0.656. Recently, [7] proposed the first Fully Polynomial-Time Approximation Scheme (FPTAS) for a class of bipartite dynamic matching networks with queues on one side only.

Many exciting open questions remain. It would be valuable to design primal-dual-style policies for dynamic matching that achieve strong competitive or approximation guarantees, as primal-dual policies tend to provide naturally interpretable decisions to practitioners. Another important direction is to study settings where system parameters (such as match rewards and arrival rates) are unknown and must be learned online. Like the uniform reward setting, the incorporation of strategic agent behavior, or incorporating incentive compatibilities into dynamic matching models is an emerging area that remains largely unexplored.

Algorithmic Mechanism Design and Sequential Search

Another prominent theme connecting mechanism design and dynamic decision-making is the study of sequential search under costly information acquisition. A classical model in this domain is the Pandora's box problem, where a decision-maker chooses which boxes to inspect at a cost to maximize expected net utility [81]. The Gittins index policy is known to be optimal in this setting.

Applications of sequential search have appeared in mechanism design with inspection costs. [33] modeled auction settings with costly valuation inspection, reducing the problem to Weitzman's framework. [63] proposed a descending price auction for informed buyers that achieves efficiency similar to a first-price auction without inspection costs. However, it has been shown that the problem becomes significantly more difficult when the inspection of boxes is nonobligatory [35], and a recent stream of literature has proposed various approximation algorithms [20, 19, 41]. Furthermore, [42] has provided sample complexity bounds for auction design with costly inspections by connecting the descending price auction to the first price auction in the classic setting. [5] extend the revenue maximization setting of [63] to nonobligatory inspection models, offering approximately optimal mechanisms for both single and multiple items, even when buyers arrive in an adversarial order.

Recent developments have also applied the sequential search with search cost framework to other domains, such as Bayesian optimization [83], product bundling [82], online market dynamics and competition under costly consumer search [13, 31], and supply chain procurement [30]. These lines of research highlight that the Pandora's box framework will continue to inspire new applications, particularly in the space of dynamic information acquisition, market design, and reinforcement learning.

3 Talk Highlights

Online Matching Models

One stream of talks in the workshop presented new online matching models and benchmarks that are motivated by real-world marketplaces. Specifically, they examined models such as philosopher inequalities, imbalanced matching networks, and matching markets with impatient agents with known sojourn times.

Tristan Pollner presented several new breakthroughs in approximation algorithms for online bipartite matching, focusing on the advancements from the classical prophet inequalities to the philosopher inequalities that do not know the future but have infinite computational time to work out the Bellman optimal policy. He demonstrated that novel applications of the pivotal sampling-based algorithms can achieve a 0.67-approximation in a discrete time model and surpass the $1{-}1/e$ threshold in the stationary continuous-time model.

Daniel Freund presented a model with new perspectives on market imbalance in online matching. He introduced a parameterized measure of imbalance and showed that higher imbalance can increase the achievable competitive ratio, both in adversarial and stochastic arrival models. His theoretical results are complemented by empirical analysis using data from a volunteer matching platform, revealing that platform design can benefit from engineering imbalance to improve performance.

Alejandro Toriello presented two widely used heuristic policies, batching and greedy, in a dynamic, non-bipartite stochastic matching model motivated by ride-sharing and freight marketplaces. Using fluid limits and randomization, he showed that both policies are asymptotically optimal in terms of sojourn time and converge exponentially fast to this benchmark. Extensions to settings with impatient agents showed that practical systems can achieve strong performance with minimal policy complexity.

Queueing and Stochastic Control

Another stream of talks in the workshop focused on queueing and stochastic control. Talks in this area highlighted algorithmic design informed by state-dependent features, approximation schemes for intractable control problems, and the interplay between learning and control in dynamic systems.

Alireza Amanihamedani presented adaptive approximation schemes for dynamic matching queues in continuous time. Unlike earlier static policies, his approach exploits queue length information and achieves near-optimal performance in polynomial time. A key technical contribution is a hybrid LP relaxation that combines static and state-dependent approximations. He applied this technique to obtain FPTAS results in both constant-size networks and spatial (Euclidean) network models, with applications to organ donation and ride-hailing systems.

Neil Walton presented scheduling policies for quantum communication networks, focusing on the design of optimal policies for a quantum switch responsible for entanglement distribution. He analyzed the system using fluid limit techniques and showed that the resulting fluid dynamics converge to an optimal solution of a specific average-reward Markov Decision Process, revealing a separation of timescales intrinsic to the system's structure.

Weina Wang's presentation considers a fully heterogeneous weakly-coupled Markov Decision Processes (WCMDPs), a hard class of problems due to the curse of dimensionality. She introduced a novel projection-based Lyapunov function and showed that a generalized Index-based (ID) policy achieves an $O(1/\sqrt{N})$

optimality gap in the average-reward setting, even when the subproblems have distinct dynamics. This is the first asymptotic result for such fully heterogeneous systems.

Weiliang Liu presented a data-driven approach to two-sided matching in networks with heterogeneous and impatient agents. The main challenge is that the optimal policy depends on the full distribution of patience times, not just their means. He showed that fluid-scale optimality can be approached asymptotically, but the rate of convergence depends critically on the learnability of the patience time distribution, highlighting the subtleties of distributional learning in stochastic control.

New Developments in Prophet Inequality and Online Contention Resolution Schemes

Several presentations introduced novel theoretical tools and improved the existing results in prophet inequality and online contention resolution schemes (OCRS). These results, in turn, leads to performance guarantees for several matching models and stochastic submodular optimization.

Andrés Cristi presented a unifying result in the study of prophet inequalities. He showed that for all major variants of the single-selection problem—including prophet-secretary and free-order models—a constant number of samples per distribution is sufficient to achieve approximation ratios close to the known optima. His proof departs from classical constructions and establishes polynomial-time implementability, significantly expanding the known landscape of sample-based prophet inequalities.

Rajan Udwani's presentation studied online submodular welfare maximization with stochastic elements and showed that non-adaptive greedy algorithms can achieve the best possible competitive ratio among all polynomial-time algorithms, including adaptive ones. His results hold even when the submodular structure is weakened to a submodular order property. He introduced a lifting technique that transfers competitive bounds from deterministic to stochastic settings, and vice versa, enabling new insights into the role of adaptivity.

Calum MacRury presented a continuous-time algorithmic framework for random-order contention resolution. Using an exact-time selection mechanism and sampling past execution histories, he designed schemes achieving tight bounds (e.g., a 0.567-selectable RCRS) for bipartite matching under vertex arrivals. The framework generalizes to other applications, such as network revenue management and L-bounded auctions, where it improves on classical benchmarks.

Applications in Market Design and Equity

Another stream of talks focused on real-world applications of online decision-making tools to platform design, workforce coordination, fairness in hiring, and human-in-the-loop machine learning. These talks exemplified how theoretical insights translate into impactful mechanisms and policy interventions.

Scott Rodilitz presented a practical redesign of VolunteerMatch's opportunity ranking algorithm. By balancing fairness (number of opportunities with at least one connection) and efficiency (total number of connections), his proposed algorithm, SmartSort, applies dynamic penalties to connected listings. Field experiments in multiple regions showed an 8% increase in equitable access without hurting overall connection rates.

Chamsi Hssaine presented a new algorithm for online resource allocation problems where allocations must follow time-varying, nonstationary targets. Her "proxy assignment" algorithm anticipates future arrivals using current inputs and achieves optimal $O(\sqrt{T})$ regret bounds. Applications include warehouse fulfillment and truck routing under capacity constraints and time windows.

Yash Kanoria's presentation investigated information deadlocks in matching markets where agents conduct costly inspections before forming matches. He characterized a sharp phase transition between deadlock-free and deadlock-prone regimes. The model, based on techniques from statistical physics and message passing, provides practical design rules for mitigating congestion in platforms with costly compatibility checks.

Shang Liu's presentation examined how to assess and incentivize human annotators in the context of LLM alignment. His principal-agent model with continuous action spaces rationalizes the use of bonus schemes and provides tight performance bounds under various contract structures. The results are validated on real language preference datasets and underscore the importance of jointly designing evaluation and incentive mechanisms.

Jackie Baek presented a dynamic model of statistical discrimination in hiring. Her analysis reveals a feedback loop whereby small initial differences in evaluation ability across groups can amplify into persistent disparities. She highlighted the importance of aggressive short-term interventions to correct such imbalances and offered a formal framework to assess their long-term impact across decision horizons.

Tutorials on Methodologies Related to Dynamic Resource Allocation and Matching

In addition to the research presentations, the workshop featured a rich lineup of tutorials aimed at providing participants with foundational tools and new perspectives relevant to the study of dynamic resource allocation and matching. These tutorials introduced both classical frameworks and modern algorithmic techniques, with applications ranging from graph inference to stochastic control. The sessions were designed to be pedagogical while also offering technical depth and highlighting open research directions.

R. Srikant introduced Stein's method as a versatile tool for obtaining non-asymptotic central limit theorems in settings with dependent data, including martingales and Markov chains. The tutorial focused on bounding the Wasserstein distance between the distribution of scaled sums and a Gaussian, with extensions to vector-valued martingales and functions of Markov chains. Applications discussed included graph sampling, stochastic gradient descent, and temporal-difference learning, demonstrating how these probabilistic tools offer sharp, finite-sample guarantees in machine learning and network analysis.

Lele Wang offered an overview of the graph alignment (or matching) problem, where the goal is to recover vertex correspondences between correlated graphs. The session discussed average-case thresholds, efficient algorithms, and applications ranging from knowledge graphs to brain imaging. It emphasized the interplay between phase transitions and algorithmic feasibility. Tools from probability, optimization, and statistical inference were highlighted as central to current and future advances in this area.

Richard Peng presented in a tutorial session cover the recent results enabling sub-polynomial updates for convex min-cost flow problems, with implications for online matching, optimal transport, and matrix rescaling.

Rad Niazadeh presented a two-part tutorial for the primal-dual method in online decision-making. The tutorials built up the primal-dual formulation for a wide range of online resource allocation models, including those with matching and cancellation costs, replenishment, and reusable resources. The speaker emphasized the flexibility of the primal-dual view in providing insight and provable guarantees, particularly in settings with uncertainty and dynamics. The session connected classical optimization theory with modern applications in online matching and operations.

Finally, Ziv Scully presented the applications of the Gittins index beyond its classical optimality regime, highlighting recent results in Bayesian optimization, tail-latency scheduling, and generalized Pandora's box models.

4 Meeting Activities and Scientific Progress

4.1 Introduction to Topics of the Breakout Sessions

On Monday evening, we organized an "Introduction to Topics of the Breakout Sessions." Prior to the workshop, we had recruited eight participants to serve as the leading discussants for the small group breakout sessions. During this event, each discussant gave an informal presentation, introducing the topic or theme they planned to facilitate during the small group breakout discussions later in the week. The introduction for each small group session lasted about ten to fifteen minutes, during which leaders outlined emerging research ideas and directions, while engaging the audience for some discussion and questions.

Beyond introducing research topics, this session also served as an effective and important ice-breaker. Held in an informal and relaxed evening setting, the event encouraged open conversations among participants, allowing them to engage more freely without the formalities often associated with structured talks. This atmosphere helped junior researchers and newcomers to the community feel more at ease and connected, fostering a greater sense of inclusion early in the workshop. After hearing all of the presentations, each BIRS participant signed up for a breakout discussion group based on their interests, ensuring that everyone had an opportunity to contribute to research conversations throughout the week.

4.2 Breakout Sessions

On Wednesday afternoon, we held the breakout sessions, where participants were divided into small groups to discuss the topics introduced earlier in the week. Although the sessions were initially scheduled for one hour, many discussions continued well into the afternoon, spilling over into the unstructured collaboration time that followed. During the breakout discussions, the discussion groups identified a wide range of intriguing open problems and research directions. These included the queueing dynamics of matching networks, designing matching policies for impatient agents, ensuring fairness in resource allocation, and incorporating strategic behavior into matching models. Other topics that generated strong interest were preference formation in matching markets, the structure and scheduling of quantum switch networks, sampling methods for inference problems, and techniques for quantifying uncertainty in online decision-making. Several collaborative efforts and new research initiatives are already underway as a direct result of these discussions.

4.3 Lightning Talks

On Wednesday evening, we organized a lightning talks session. As the name suggests, the session featured a series of rapid five-minute talks, where any participant could present a snippet of their research, introduce a specific subfield, or highlight an open problem. This session provided an excellent opportunity for participants to share their ideas to everyone at the workshop, particularly those who were not scheduled for full-length talks. The lightning talks also provided another informal opportunity for participants to share and discuss their research, complementing the earlier breakout discussions. Many of the ideas presented during these talks helped spark further conversations and collaborations during the subsequent unstructured time.

4.4 Unstructured time for collaboration

The afternoon sessions on Tuesday, Thursday, and part of Wednesday (immediately after the breakout sessions) were reserved for unstructured collaboration. These sessions provided participants with a valuable chance to explore unresolved issues, embark on new projects, and carry on the conversations that began in the breakout sessions and lightning talks. Many participants commented that this format worked extremely well, and that the opportunity to discuss research while enjoying a walk around the BIRS facilities and the scenic surroundings of BIRS helped inspire new research ideas.

5 Outcome of the Meeting

First, we would like to express our sincere gratitude to BIRS for the opportunity to organize this workshop at Banff, Alberta. The staff at BIRS were tremendously helpful and supportive throughout the event. The workshop enabled us to bring together a group of outstanding researchers to spend a week presenting new results, collaborating on open problems, and contemplating how the field of dynamic allocation and matching will continue to evolve.

The talks at BIRS were well attended. In addition to the in-person participants, we also had multiple virtual participants joining via Zoom or streaming the talks. Two presenters, unable to attend in person due to visa or passport issues, delivered their talks virtually. Thanks to the excellent technical support from BIRS, the virtual sessions ran smoothly, and the audience remained highly engaged, asking numerous questions throughout.

Overall, the participants thoroughly enjoyed the workshop. Those who had previously attended two-day BIRS workshops remarked that the five-day format provided a much better opportunity to interact, collaborate, and devote more focused time to exploring research ideas. Another notable benefit of the workshop was that many junior researchers were introduced to the broader community, fostering new connections and collaborations.

In addition, the meeting helped identify a wide range of open problems and research directions in resource allocation and matching. These include advances in using Stein's method to study limit theorems for the total reward in online resource allocation and matching; applying online resource allocation methodologies to develop asymptotically optimal open-loop policies for inventory control with long lead times; exploring new approaches to queueing aspects of resource allocation; leveraging the Gittins index and mechanism design for dynamic resource allocation; and pursuing innovative applications of resource allocation models, such as in the watermarking of large language models. We anticipate that many of the questions and directions discussed during the meeting will lead to significant scientific progress in the coming months and years.

References

[1] Philipp Afeche, Rene Caldentey, and Varun Gupta. On the optimal design of a bipartite matching queueing system. *Operations Research*, 70(1):363–401, 2022.

[2] Mohammad Akbarpour, Yeganeh Alimohammadi, Shengwu Li, and Amin Saberi. The value of excess supply in spatial matching markets. *arXiv preprint arXiv:2104.03219*, 2021.

[3] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):783–815, 2020.

[4] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[5] Saeed Alaei, Ali Makhdoumi, and Azarakhsh Malekian. Revenue maximization under unknown private values with non-obligatory inspection. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 27–28, 2021.

[6] Alireza AmaniHamedani, Ali Aouad, Tristan Pollner, and Amin Saberi. Improved approximations for stationary bipartite matching: Beyond probabilistic independence. *arXiv preprint arXiv:2411.08218*, 2024.

[7] Alireza AmaniHamedani, Ali Aouad, and Amin Saberi. Adaptive approximation schemes for matching queues. *arXiv preprint arXiv:2501.08775*, 2025.

[8] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.

[9] Ali Aouad and Will Ma. A nonparametric framework for online stochastic matching with correlated arrivals. *arXiv preprint arXiv:2208.02229*, 2022.

[10] Ali Aouad and Daniela Saban. Online assortment optimization for two-sided matching platforms. *Management Science*, 69(4):2069–2087, 2023.

[11] Ali Aouad and Ömer Saritaç. Dynamic stochastic matching under limited time. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 789–790, 2020.

[12] Alessandro Arlotto and Itai Gurvich. Uniformly bounded regret in the multisecretary problem. *Stochastic Systems*, 9(3):231–260, 2019.

[13] Mark Armstrong. Ordered consumer search. *Journal of the European Economic Association*, 15(5):989–1024, 2017.

[14] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019.

[15] Angelos Aveklouris, Levi DeValve, Maximiliano Stock, and Amy Ward. Matching impatient and heterogeneous demand and supply. *Operations Research*, 2024.

[16] Yicheng Bai, Omar El Housni, Billy Jin, Paat Rusmevichientong, Huseyin Topaloglu, and David P Williamson. Fluid approximations for revenue management under high-variance demand. *Management Science*, 69(7):4016–4026, 2023.

[17] Santiago Balseiro, Omar Besbes, and Dana Pizarro. Survey of dynamic resource constrained reward collection problems: Unified model and analysis. *Available at SSRN 3963265*, 2021.

[18] Jocelyn Begeot, Irène Marcovici, and Pascal Moyal. Stability regions of systems with compatibilities and ubiquitous measures on graphs. *Queueing Systems*, 103(3):275–312, 2023.

[19] Hedyeh Beyhaghi and Linda Cai. Pandora's problem with nonobligatory inspection: Optimal structure and a ptas. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 803–816, 2023.

[20] Hedyeh Beyhaghi and Robert Kleinberg. Pandora's problem with nonobligatory inspection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 131–132, 2019.

[21] Jose H Blanchet, Martin I Reiman, Virag Shah, Lawrence M Wein, and Linjia Wu. Asymptotically optimal control of a centralized dynamic matching market with general utilities. *Operations Research*, 2022.

[22] Mark Braverman, Mahsa Derakhshan, Tristan Pollner, Amin Saberi, and David Wajc. New philosopher inequalities for online bayesian matching, via pivotal sampling. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3029–3068. SIAM, 2025.

[23] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *European Symposium on Algorithms*, pages 253–264. Springer, 2007.

[24] Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing problems. In *European Symposium on Algorithms*, pages 689–701. Springer, 2005.

[25] Niv Buchbinder, Joseph Seffi Naor, et al. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science*, 3(2–3):93–263, 2009.

[26] Pornpawee Bumpensanti and He Wang. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science*, 66(7):2993–3009, 2020.

[27] Arnaud Cadas, Josu Doncel, and Ana Bušić. Analysis of an optimal policy in dynamic bipartite matching models. *Performance Evaluation*, 154:102286, 2022.

[28] René A Caldentey and Edward H Kaplan. A heavy traffic approximation for queues with restricted customer-server matchings. 2007.

[29] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 243–251, 2007.

[30] Qi Chen, Damian R Beil, and Izak Duenyas. Procurement mechanisms with post-auction pre-award cost-reduction investigations. *Operations Research*, 70(6):3054–3075, 2022.

[31] Michael Choi, Anovia Yifan Dai, and Kyungmin Kim. Consumer search and price competition. *Econometrica*, 86(4):1257–1281, 2018.

[32] José Correa, Andrés Cristi, Andrés Fielbaum, Tristan Pollner, and S Matthew Weinberg. Optimal item pricing in online combinatorial auctions. *Mathematical Programming*, 206(1):429–460, 2024.

[33] Jacques Crémer, Yossi Spiegel, and Charles Z Zheng. Auctions with costly information acquisition. *Economic Theory*, 38(1):41–72, 2009.

[34] Nikhil R Devanur, Kamal Jain, and Robert D Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 101–107. SIAM, 2013.

[35] Laura Doval. Whether or not to open pandora's box. *Journal of Economic Theory*, 175:127–158, 2018.

[36] Farbod Ekbatani, Yiding Feng, and Rad Niazadeh. Online matching with cancellation costs. *arXiv preprint arXiv:2210.11570*, 2022.

[37] Atilla Eryilmaz and Rayadurgam Srikant. Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Systems*, 72:311–359, 2012.

[38] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet matching with general arrivals. *Mathematics of Operations Research*, 47(2):878–898, 2022.

[39] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.

[40] Yiding Feng and Rad Niazadeh. Batching and optimal multistage bipartite allocations. *Management Science*, 2024.

[41] Hu Fu, Jiawei Li, and Daogao Liu. Pandora box problem with nonobligatory inspection: Hardness and approximation scheme. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 789–802, 2023.

[42] Hu Fu and Tao Lin. Learning utilities and equilibria in non-truthful auctions. *Advances in Neural Information Processing Systems*, 33:14231–14242, 2020.

[43] V Goyal, G Iyengar, and R Udwani. Pricing a finite inventory of substitutable products with show-all constraint. *Available at SSRN*, 2022.

[44] Vineet Goyal, Garud Iyengar, and Rajan Udwani. Asymptotically optimal competitive ratio for online allocation of reusable resources. *Operations Research*, 2025.

[45] Varun Gupta. Greedy algorithm for multiway matching with bounded regret. *arXiv preprint arXiv:2112.04622*, 2021.

[46] Siqi He, Yehua Wei, Jiaming Xu, and Sophie H Yu. Online resource allocation without re-solving: The effectiveness of primal-dual policies. *Available at SSRN*, 2025.

[47] Yige Hong, Qiaomin Xie, Yudong Chen, and Weina Wang. Restless bandits with average reward: Breaking the uniform global attractor assumption. *Advances in Neural Information Processing Systems*, 36:12810–12844, 2023.

[48] Yige Hong, Qiaomin Xie, Yudong Chen, and Weina Wang. Achieving exponential asymptotic optimality in average-reward restless bandits without global attractor assumption. *arXiv preprint arXiv:2405.17882*, 2024.

[49] Longbo Huang and Michael J Neely. Delay reduction via lagrange multipliers in stochastic network optimization. In *2009 7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, pages 1–10. IEEE, 2009.

[50] Zhiyi Huang and Xinkai Shu. Online stochastic matching, poisson arrivals, and the natural linear program. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 682–693, 2021.

[51] Zhiyi Huang, Xinkai Shu, and Shuyi Yan. The power of multiple choices in online stochastic matching. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 91–103, 2022.

[52] Zhiyi Huang, Zhihao Gavin Tang, and David Wajc. Online matching: A brief survey. *ACM SIGecom Exchanges*, 22(1):135–158, 2024.

[53] Zhiyi Huang and Qiankun Zhang. Online primal dual meets online matching with stochastic rewards: configuration lp to the rescue. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1153–1164, 2020.

[54] Stefanus Jasin. Reoptimization and self-adjusting price control for network revenue management. *Operations Research*, 62(5):1168–1178, 2014.

[55] Stefanus Jasin. Performance of an LP-based control for revenue management with unknown demand parameters. *Operations Research*, 63(4):909–915, 2015.

[56] Stefanus Jasin and Sunil Kumar. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345, 2012.

[57] Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1221–1246. SIAM, 2022.

[58] Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, and Nahuel Soprano-Loto. Generalized max-weight policies in stochastic matching. *Stochastic Systems*, 13(1):40–58, 2023.

[59] Yash Kanoria and Pengyu Qian. Blind dynamic resource allocation in closed networks via mirror backpressure. *Management Science*, 70(8):5445–5462, 2024.

[60] Süleyman Kerimov, Itai Ashlagi, and Itai Gurvich. Dynamic matching: Characterizing and achieving constant regret. *Available at SSRN 3824407*, 2021.

[61] Süleyman Kerimov, Itai Ashlagi, and Itai Gurvich. On the optimality of greedy policies in dynamic matching. *Available at SSRN 3918497*, 2021.

[62] Kristen Kessel, Ali Shameli, Amin Saberi, and David Wajc. The stationary prophet inequality problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 243–244, 2022.

[63] Robert Kleinberg, Bo Waggoner, and E Glen Weyl. Descending price optimally coordinates search. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 23–24, 2016.

[64] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

[65] Calum MacRury and Will Ma. Random-order contention resolution via continuous induction: Tightness for bipartite matching under vertex arrivals. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1629–1640, 2024.

[66] Vahideh Manshadi, Scott Rodilitz, Daniela Saban, and Akshaya Suresh. Redesigning volunteermatch's ranking algorithm: Toward more equitable access to volunteers. *Available at SSRN 4497747*, 2023.

[67] Pascal Moyal, Ana Bušić, and Jean Mairesse. A product form for the general stochastic matching model. *Journal of Applied Probability*, 58(2):449–468, 2021.

[68] Mohammadreza Nazari and Alexander L Stolyar. Reward maximization in general dynamic matching systems. *Queueing Systems*, 91:143–170, 2019.

[69] Erhun Özkan and Amy R Ward. Dynamic matching for real-time ride sharing. *Stochastic Systems*, 10(1):29–70, 2020.

[70] Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 763–764, 2021.

[71] Neel Patel and David Wajc. Combinatorial stationary prophet inequalities. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4605–4630. SIAM, 2024.

[72] Tristan Pollner, Mohammad Roghani, Amin Saberi, and David Wajc. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 321–322, 2022.

[73] Nahuel Soprano-Loto, Matthieu Jonckheere, and Pascal Moyal. Online matching for the multiclass stochastic block model. *arXiv preprint arXiv:2303.15374*, 2023.

[74] Kalyan T Talluri, Garrett Van Ryzin, and Garrett Van Ryzin. *The theory and practice of revenue management*, volume 1. Springer, 2004.

[75] Alfredo Torrico and Alejandro Toriello. Dynamic relaxations for online bipartite matching. *INFORMS Journal on Computing*, 34(4):1871–1884, 2022.

[76] Rajan Udwani. Adwords with unknown budgets and beyond. *Management Science*, 71(2):1009–1026, 2025.

[77] Sushil Mahavir Varma, Pornpawee Bumpensanti, Siva Theja Maguluri, and He Wang. Dynamic pricing and matching for two-sided queues. *Operations Research*, 71(1):83–100, 2023.

[78] Alberto Vera and Siddhartha Banerjee. The bayesian prophet: A low-regret framework for online decision making. *Management Science*, 67(3):1368–1391, 2021.

[79] Alberto Vera, Siddhartha Banerjee, and Itai Gurvich. Online allocation and pricing: Constant regret via bellman inequalities. *Operations Research*, 69(3):821–840, 2021.

[80] Yehua Wei, Jiaming Xu, and Sophie H Yu. Constant regret primal-dual policy for multi-way dynamic matching. *Available at SSRN 4357216*, 2023.

[81] Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

[82] Chenguang Wu, Chen Jin, and Ying-Ju Chen. Managing customer search via bundling. *Manufacturing & Service Operations Management*, 24(4):1906–1925, 2022.

[83] Qian Xie, Raul Astudillo, Peter Frazier, Ziv Scully, and Alexander Terenin. Cost-aware bayesian optimization via the pandora's box gittins index. *Advances in Neural Information Processing Systems*, 37:115523–115562, 2024.