ISSN 2186-7437

NII Shonan Meeting Report

No. 2017-5

Algorithms and Optimization under Uncertainty

Niv Buchbinder
Nikhil Devanur
Debmalya Panigrahi

May 22–25, 2017

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Algorithms and Optimization under Uncertainty

Organizers:
Niv Buchbinder (Tel-Aviv University)
Nikhil Devanur (Microsoft Research)
Debmalya Panigrahi (Duke University)

May 22–25, 2017

In the classical *offline* computational model, an algorithm operates on a specified set of input data to produce a desired output. While this model has propelled much of computer science, modern applications typically do not afford the luxury of complete knowledge and certainty of the input data. Thus, the typical real world algorithm today has to make optimization decisions with incomplete information. This workshop brought together world-class researchers from multiple domains of algorithm design with the goal of discussing and furthering research in this broad area of *algorithms under uncertainty*.

Online algorithms and Competitive analysis. The field of online computation in which the input is presented piece-by-piece to the algorithm has played an important role in theoretical computer science. Recently, there has been a renewed interest in this area due to new applications that have emerged in online systems. For example, the design of data centers for large-scale computing poses new challenges motivated by aspects such as fair and efficient online scheduling of heterogeneous clients in environments such as Hadoop and Map-Reduce. In addition to these applications, recent research has also focused on new generic techniques such as the import of the linear programming toolkit to online algorithms and competitive analysis. Furthermore, new models have been proposed that incorporate delays, limited recourse, and so on to make online problems more practical and amenable to algorithm design. The discussion at the workshop on online algorithms focused on the twin goals of of making fundamental progress on longstanding questions and developing a principled approach for applied problems of practical importance.

Beyond worst-case analysis and inter-connections. While research in online algorithms and competitive analysis has yielded rich dividends in both theory and practice, there is a growing belief among researchers that we also need to explore possibilities beyond worst case (competitive) analysis. This has led to a growing body of work on various stochastic input models and on employing ideas from online learning. Discussions at the workshop focused on the inter-connections between these different subareas of algorithms under uncertainty, and on how ideas from one area can lead to breakthroughs in other areas of algorithms under uncertainty.

Overview of Talks

Discrimination in Machine Decision Making

Krishna Gummadi, MPI

Machine (data-driven learning-based) decision making is increasingly being used to assist or replace human decision making in a variety of domains ranging from banking (rating user credit) and recruiting (ranking applicants) to judiciary (profiling criminals) and journalism (recommending news-stories). Recently concerns have been raised about the potential for discrimination and unfairness in such machine decisions.Against this background, in this talk, I will pose and attempt to answer the following high-level questions:

1. How do machines learn to make discriminatory decision making?
2. How can we quantify discrimination in machine decision making?
3. How can we control machine discrimination? i.e., can we design learning mechanisms that avoid discriminatory decision making?
4. Is there a cost to non-discriminatory decision making?

Dynamic and Online Algorithms

Anupam Gupta, CMU

In online algorithms, the requests arrive one by one, and each request must be serviced without knowledge of the future. And all actions are irrevocable. In this setting, many combinatorial optimization problems become harder than their offline counterparts: e.g., the spanning tree and set cover problems become logarithmically "harder" in the online setting.

But what if we could undo a small number of decisions: could we do much better? We show that for both the above problems, as well as other load-balancing problems, we can get close to the offline guarantees by allowing a small number of changes. Some of these results extend to the setting of dynamic graph algorithms, where we show that the changes can also be implemented with small update times.

This survey-ish talk is based on works with Albert Gu, Guru Guruganesh, Amit Kumar, Ravishankar Krishnaswamy, Debmalya Panigrahi, Cliff Stein, and David Wajc.

Plane Gossip: Approximating rumor spread in planar graphs

R Ravi, CMU

We study the design of schedules for multi-commodity multicast, the problem of efficient point-to-point communication among a collection of source-destination pairs. In this problem, we are given an undirected graph and a collection of source-destination pairs, and the goal is to schedule a minimum-length sequence of matchings that connects every source with its respective destination. The multi-commodity multicast formulation models a classic information dissemination problem in networks where the primary communication constraint is the number of connections that a given node can make, not link bandwidth. Multi-commodity multicast and its special cases, broadcast and multicast, are all NP-complete and the best approximation factors known are $2^{\widetilde{O}(\sqrt{\log n})}$, $O(\log n/\log\log n)$, and $O(\log k/\log\log k)$, respectively, where n is the number of nodes and k is the number of terminals in the multicast instance.

The multi-commodity multicast problem is closely related to the problem of finding a subgraph of optimal poise, where the poise is defined as the sum of the maximum degree of the subgraph and the maximum distance between any source-destination pair in the subgraph. We first show that for any instance of the multicast problem, the minimum poise subgraph can be approximated to within a factor of $O(\log k \log n)$ with respect to value of the natural LP relaxation in an n-node graph with k source-destination pairs. Using this integrality gap upper bound, we obtain an $O(\log^2 k \log^2 n)$-approximation for multi-commodity multicast for planar graphs.

Joint work with Jennifer Iglesias (CMU), Rajmohan Rajaraman (Northeastern) and Ravi Sundaram (Northeastern)

Speeding up MWU for Packing/Covering LPs, and Applications

Chandra Chekuri, UIUC

We consider several implicit fractional packing problems and obtain faster implementations of approximation schemes by a general framework based on multiplicative-weight updates. This leads to new algorithms with near-linear running times for some fundamental problems in optimization. We highlight some concrete applications, in particular the problem of finding a maximum fractional packing of spanning trees in a capacitated graph.

This is joint work with Kent Quanrud.

Matching with delays

Yossi Azar, Tel Aviv University

The problem of min-cost matching with delays is an online problem defined on an underlying metric space as follows. Requests arrive online at points of the metric space, and the algorithm is required to match them, possibly after keeping them waiting for some time. The cost incurred is the sum of the distances between matched pairs of requests (the connection cost), and the sum of the waiting times of the requests (the delay cost). This objective exhibits a natural trade-off between minimizing the distances and the cost of waiting for better matches. The problem comes in two flavors: the homogeneous one (MPMD) and the two sided one (MBPMD). In the homogeneous version, any two requests can be matched, while on the other hand, in the two-sided version, each request is either positive or negative, and two requests may be matched only if they have opposite polarities.

The homogeneous version of the problem was recently introduced by Emek *et al.*, (STOC'16). For an n-point metric space in which the largest distance is

Δ times the smallest, Emek et al. give an algorithm with a competitive ratio $O(\log^2 n + \log \Delta)$. We present an improved $O(\log n)$ competitive algorithm, removing the dependence on Δ. We also prove a lower bound of $\Omega(\log n / \log \log n)$ on the competitive ratio of any randomized algorithm, almost matching the upper bound. For two-sided version, we give an $O(\log n)$-competitive algorithm, and prove a lower bound of $\Omega(\sqrt{\log n / \log \log n})$.

The results for MPMD appeared in SODA'17 (joint work with A. Chiplunkar and H. Kaplan), and the results for MBPMD will appear in APPROX'17 (joint work with I. Ashlagi, M. Charikar, A. Chiplunkar, O. Geri, H. Kaplan, R. Makhijani, Y. Wang, R. Wattenhofer).

Online Boosting Algorithms

Satyen Kale, Google

We initiate the study of boosting in the online setting, where the task is to convert a "weak" online learner into a "strong" online learner. The notions of weak and strong online learners directly generalize the corresponding notions from standard batch boosting. For the classification setting, we develop two online boosting algorithms. The first algorithm is an online version ofboost-by-majority, and we prove that it is essentially optimal in terms of the number of weak learners and the sample complexity needed to achieve a specified accuracy. The second algorithm is adaptive and parameter-free, albeit not optimal.

For the regression setting, we give an online gradient boosting algorithm which converts a weak online learning algorithm for a base class of regressors into a strong online learning algorithm which works for the linear span of the base class. We also give a simpler boosting algorithm for regression that obtains a strong online learning algorithm which works for the convex hull of the base class, and prove its optimality.

When the Optimum is also Blind: a New Perspective on Universal Optimization

Fabrizio Grandoni, IDSIA Lugano

Consider the following variant of the set cover problem. We are given a universe $U = \{1, ..., n\}$ and a collection of subsets $C = \{S_1, ..., S_m\}$ where $S_i \subseteq U$. For every element $u \in U$ we need to find a set $\phi(u) \in C$ such that $u \in \phi(u)$. Once we construct and fix the mapping $\phi : U \to C$ a subset $X \subseteq U$ of the universe is revealed, and we need to cover all elements from X with exactly $\phi(X) := \cup_{u \in X} \phi(u)$. The goal is to find a mapping such that the cover $\phi(X)$ is as cheap as possible.

This is an example of a universal problem where the solution has to be created before the actual instance to deal with is revealed. Such problems appear naturally in some settings when we need to optimize under uncertainty and it may be actually too expensive to begin finding a good solution once the input starts being revealed. A rich body of work was devoted to investigate such problems under the regime of worst case analysis, i.e., when we measure how good the solution is by looking at the worst-case ratio: universal solution for a given instance vs optimum solution for the same instance.

As the universal solution is significantly more constrained, it is typical that such a worst-case ratio is actually quite big. One way to give a viewpoint on the problem that would be less vulnerable to such extreme worst-cases is to assume that the instance, for which we will have to create a solution, will be drawn randomly from some probability distribution. In this case one wants to minimize the expected value of the ratio: universal solution vs optimum solution. Here the bounds obtained are indeed smaller than when we compare to the worst-case ratio.

But even in this case we still compare apples to oranges as no universal solution is able to construct the optimum solution for every possible instance. What if we would compare our approximate universal solution against an optimal universal solution that obeys the same rules as we do? We show that under this viewpoint, but still in the stochastic variant, we can indeed obtain better bounds than in the expected ratio model. For example, for the set cover problem we obtain an $O(\log n)$ approximation which matches the approximation ratio from the classic deterministic offline setup. Moreover, we show this for all possible probability distributions over U that have a polynomially large carrier, while all previous results pertained to a model in which elements were sampled independently. Our result is based on rounding a proper configuration IP that captures the optimal universal solution, and using tools from submodular optimization.

The same basic approach leads to improved approximation algorithms for other related problems, including Vertex Cover, Edge Cover, Directed Steiner Tree, Multicut, and Facility Location.

This is joint work with Marek Adamczyk, Stefano Leonardi and Michal Wlodarczyk.

Online Auctions and Multi-scale Online Learning

Nikhil Devanur, MSR Redmond

We consider revenue maximization in online auctions and pricing. A seller sells an identical item in each period to a new buyer, or a new set of buyers. For the online posted pricing problem, we show regret bounds that scale with the best fixed price, rather than the range of the values. We also show regret bounds that are almost scale free, and match the offline sample complexity, when comparing to a benchmark that requires a lower bound on the market share. These results are obtained by generalizing the classical learning from experts and multi-armed bandit problems to their multi-scale versions. In this version, the reward of each action is in a different range, and the regret w.r.t. a given action scales with its own range, rather than the maximum range.

This is joint work with Sebastien Bubeck, Zhiyi Huang and Rad Niazadeh, and appeared in EC 2017.

Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs

Thomas Kesselheim, MPI

We present a general framework for stochastic online maximization problems with combinatorial feasibility constraints. The framework establishes prophet inequalities by constructing price-based online approximation algorithms, a natural extension of threshold algorithms for settings beyond binary selection. Our analysis takes the form of an extension theorem: we derive sufficient conditions on prices when all weights are known in advance, then prove that the resulting approximation guarantees extend directly to stochastic settings. Our framework unifies and simplifies much of the existing literature on prophet inequalities and posted price mechanisms, and is used to derive new and improved results for combinatorial markets (with and without complements), multi-dimensional matroids, and sparse packing problems. Finally, we highlight a surprising connection between the smoothness framework for bounding the price of anarchy of mechanisms and our framework, and show that many smooth mechanisms can be recast as posted price mechanisms with comparable performance guarantees.

Joint work with Paul Dtting, Michal Feldman, and Brendan Lucier.

Online Service with Delay

Debmalya Panigrahi, Duke University

In the online service with delay problem, points in a metric space service requests over time, and a server serves these requests by traveling to the location of the request. The goal is to minimize the sum of distance traveled by the server and the total delay(or a penalty function thereof) in serving the requests. This problem models the fundamental trade-off between batching requests to improve locality and reducing delay to improve response time, that has many applications in operations management, operating systems, logistics, supply chain management, and scheduling. Our main result is to show a polylogarithmic competitive ratio for the online service with delay problem. This result is obtained by an algorithm that we call the preemptive service algorithm. The salient feature of this algorithm is a process called preemptive service, which uses a novel combination of (recursive) time forwarding and spatial exploration on a metric space. We also generalize our results to $k > 1$ servers, and obtain stronger results for special metrics such as uniform and star metrics that correspond to(weighted) paging problems.

This is joint work with Yossi Azar, Arun Ganesh, and Rong Ge.

Some Open problems

Multicast

R. Ravi, CMU

In the *Multicast* problem, we are given graph $G(V, E)$ which represents a telephone network on V, where there can be a phone call between two nodes if there is an edge between them. We are also given a source vertex r and a set of terminals $R \subseteq V$. The source vertex has a message and it wants to inform all the terminals from the message. To do this, the vertices of the graph can communicate in rounds: In each round, we pick a matching of G and arrange a bidirected phone call between each vertex in the matching and its match. If any of the two vertices knows the message before the phone call, the other one will also know it afterwards. The goal is to deliver the message to all the terminals in the minimum number of rounds.

The best known approximation ratio for finding a minimum multi-commodity multicast scheme is $2^{O\left(\log\log k\sqrt{\log k}\right)}$ where $k = |R|$ (Sending Secrets Swiftly: Approximation Algorithms for Generalized Multicast Problems. Afshin Nikzad and R. Ravi. ICALP 2014). The open problem is to devise the first poly-logarithmic approximation for the problem on general graphs. My talk in the workshop sketched a poly-logarithmic solution for planar graphs.

Metric-TSP

Chandra Chekrui, UIUC

Consider a Metric-TSP induced by the shortest path metric of an edge-weighted undirected graph $G = (V, E)$ with n vertices and m edges. Recently, Chekuri and Quanrud obtained an algorithm that given such an instance, outputs a $(1 + \epsilon)$-approximate solution to the subtour-elimination LP relaxation for such an instance in randomized $O(m\text{polylog}(n)/\epsilon^2)$ time. This raises the question of whether one can round the LP relaxation in near-linear time to obtain a fast approximation algorithm. The simple MST heuristic gives a linear time 2-approximation. Christofides heuristic gives 3/2-approximation but the running time depends on finding a minimum-cost perfect matching.

Is there a near-linear time algorithm that outputs a $2 - \delta$ approximation for Metric-TSP?

It is well-known that the integrality gap of the subtour-elimination LP relaxation is at most 3/2 and it is conjectured that it is at most 4/3. Since the LP relaxation can be solved fast it may be feasible to round it to obtain a near-linear time $3/2 + \epsilon$ approximation.

Online Problems with Delay

Yossi Azar, Tel Aviv University

1. Find deterministic algorithms for matching with delays or show lower bounds (general metric space, line or HST)

2. Find other interesting online problems where the goal function is the sum of costs and delays of serving the requests

Generalizations of Online Matching

Nikhil Devanur, MSR Redmond

Generalizations of the online matching problem, such as the budgeted allocation problem, and capacitated edge weighted matching problem, have been "solved" in stochastic settings such as a random order model [3, 1] and the i.i.d. model [2]. This means one can get a $1 - \epsilon$ competitive algorithm as long as the budgets/capacities are large enough ($\approx \log n/\epsilon^2$), where n is the number of offline nodes). The question is, can we get similar results under non-stochastic assumptions on the input instance. Another approach would be to restrict the optimal offline solution in some meaningful way; this seems stronger. E.g., suppose that we compete only against those offline solutions where the "rate" at which the capacities fill up is not too far away from a linear rate. To be precise, suppose we consider the unweighted B-matching, where each offline node has a capacity of B. Let $\Delta \in (0, 1)$ be some constant. Suppose that there are T online nodes. We restrict the offline solution such that for each offline node, after t online nodes have arrived, the number of matches is at most

$$\left(\frac{t}{T} + \Delta\right) B.$$

We will probably have to consider some kind of bicriteria solution where the restriction on the online solution is weaker. Perhaps, we don't even restrict the online solution in any way. When can we get a $1 - \epsilon$ approximation?

References

[1] Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1405–1424. Society for Industrial and Applied Mathematics, 2015.

[2] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38. ACM, 2011.

[3] Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 303–312. ACM, 2014.

Truthful Edge-Weighted Online Matching

Thomas Kesselheim, MPI

Consider the following variant of bipartite edge-weighted matching. There is a weighted bipartite graph $G = (L \cup R, E)$, $w\colon E \to \mathbb{R}_{\geq 0}$. The vertices in R are known in advance; the vertices in L arrive one at a time in random order. Every time a vertex in L arrives, we get to know its incident edges and their weights. Immediately and irrevocably we have to decide whether to match this respective vertex to a vertex in R or to leave it unmatched. For this problem, a $\frac{1}{e}$-competitive algorithm is known [Kesselheim et al., ESA 2013], which is optimal. However, it is not possible to implement this algorithm *truthfully*. Suppose, for example, the vertices in L correspond to buyers and the vertices in R correspond to items. The selfish buyers will then only report correct weights if this is in their best interest. For the currently known constant-competitive algorithms, there is no way of imposing payments so that this is guaranteed. It is open whether there is a constant-competitive algorithm, let alone a $\frac{1}{e}$-competitive one.

List of Participants

- Yossi Azar, Tel Aviv University.
- Siddhartha Banerjee, Cornell University.
- Nikhil Bansal, TU Eindhoven.
- Sebastien Bubeck, Microsoft Research.
- Niv Buchbinder, Tel Aviv University.
- Chandra Chekuri, UIUC.
- Artur Czumaj, University of Warwick.
- Nikhil Devanur, Microsoft Research.
- Guy Even, Tel Aviv University.
- Fabrizio Grandoni, IDSIA, University of Lugano.
- Krishna Gummadi, Max Planck Institute for Software Systems.
- Anupam Gupta, Carnegie Mellon University.
- Zhiyi Huang, The University of Hong Kong.
- Satyen Kale, Google.
- Ken-ichi Kawarabayashi, National Institute of Informatics.
- Thomas Kesselheim, Max Planck Institute for Informatics.
- Mohammad Mahdian, Google.
- Aranyak Mehta, Google Research.
- Seffi Naor, Technion.
- Debmalya Panigrahi, Duke University.
- R. Ravi, Carnegie Mellon University.
- Devavrat Shah, MIT.
- Balasubramanian Sivan, Google Research.
- Seeun William Umboh, Eindhoven University of Technology.

Meeting Schedule

Check-in Day: May 21 (Sun)

- Welcome Banquet at "Azalea & Camelia" restaurant

Day 1: January May 22 (Mon)

- 7:30 - 9 am: Breakfast in Cafeteria Oak
- 9 - 9:10 am: Introduction movie of Shonan meeting
- 9:10 - 9:15 am: Opening remarks
- 9:15 - 10:30 am: Introduction and open problem session
- 10:30 - 11 am: Coffee break
- 11 - 12 pm: Introduction and open problem session
- 12 - 1:30 pm: Lunch in Cafeteria Oak
- 1:30 - 2 pm: Group photo session
- 2 - 2:45 pm: Anupam Gupta: Dynamic and Online Algorithms: A little change will do you good
- 2:45 - 3:30 pm: Thomas Kesselheim: Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs
- 3:30 - 4 pm: Coffee break
- 4 - 4:30 pm: Chandra Chekuri: Speeding up MWU for Packing/Covering LPs, and Applications
- 4:30 - 5 pm: Nikhil Bansal: The Weighted k-server problem
- 6 - 7:30 pm: Dinner in Cafeteria Oak

Day 2: January May 23 (Tue)

- 7:30 - 9 am: Breakfast in Cafeteria Oak
- 9 - 9:45 am: Seffi Naor: Competitive Algorithms for Online Multi-level Aggregation
- 9:45 - 10:30 am: Sebastien Bubeck: Mirror descent and self-concordant barriers for the linear bandit problem, after Abernathy, Hazan and Rakhlin 2008
- 10:30 - 11 am: Coffee break
- 11 - 11:30 am: Debmalya Panigrahi: Online Service with Delay
- 11:30 am - 12 pm: Yossi Azar: Matching with delays
- 12 - 1:30 pm: Lunch in Cafeteria Oak
- 1:30 - 2 pm: Aranyak Mehta: A Session Auction for In-App Advertising

- 2 - 2:30 pm: Krishna Gummadi: Fairness and Explainability of Algorithmic Decisions
- 2:30 - 3 pm: Balu Sivan: Stability of Service Under Time-of-Use Pricing
- 3 - 3:30 pm: Zhiyi Huang: Online algorithms for covering and packing problems with convex objectives
- 3:30 - 4 pm: Coffee break
- 6 - 7:30 pm: Dinner in Cafeteria Oak

Day 3: January May 24 (Wed)

- 7:30 - 9 am: Breakfast in Cafeteria Oak
- 9 - 9:45 am: Satyen Kale: Online boosting
- 9:45 am - 10:30 am: Siddhartha Banerjee: Stochastic control in two-sided marketplaces
- 10:30 - 11 am: Coffee break
- 11 - 11:30 am: Nikhil Devanur: Online Auctions and Multi-scale Online Learning
- 11:30 am - 12 pm: Mohammad Mahdian: Community Detection on Evolving Graphs
- 12 - 1:30 pm: Lunch in Cafeteria Oak
- 1:30 - 6:30 pm Excursion to Jomyoji Temple with Japanese Tea Ceremony
- 6:30 - 8:30 pm: Conference Banquet

Day 4: January May 25 (Thu)

- 7 - 9 am: Breakfast in Cafeteria Oak
- 9 - 9:45 am: Devavrat Shah: Prediction with Latent Variable Model
- 9:45 - 10:15 am: Fabrizio Grandoni: When the optimum is also blind
- 10:15 - 10:30 am: Coffee break
- 10:30 - 11 am: Seeun William Umboh: LAST but not Least: Online Spanners for Buy-at-Bulk
- 11 - 11:30 am: R. Ravi: Plane Gossip
- 12 - 1:30 pm: Lunch in Cafeteria Oak
- Conference ends