MSc Projects CWI's Networks and Optimization Group

Budget-feasible Mechanism Design with Predictions

Supervision: Prof. dr. Guido Schafer
Networks and Optimization Group, Centrum Wiskunde & Informatica (CWI)
Institute for Logic, Language and Computation (ILLC), University of Amsterdam

If applicable, an internship ("stagiair") status at CWI can be arranged for the student (to be discussed).

Keywords: Mechanism Design, Algorithmic Game Theory, Algorithms with Predictions, Combinatorial Optimization

Description: We consider the problem of designing *budget-feasible mechanisms* for procurement auctions. In this model, an auctioneer is interested in buying services (or goods) from a set of agents N. Each agent $i \in N$ specifies a cost c_i to be paid by the buyer for using their service; these costs are assumed to be private information. The auctioneer has a budget B and a valuation function $v(\cdot)$, where $v(S)$ specifies the value derived from the services of the agents in $S \subseteq N$.
Given the (reported) costs of the agents, the goal of the auctioneer is to choose a *budget-feasible* subset $S \subseteq N$ of the agents, such that the valuation $v(S)$ is maximized. Budget-feasibility here means that $\Sigma_{i \in S}\, p_i \leq B$, where p_i is the payment issued from the mechanism to agent i.

The agents might try to extract larger payments from the mechanism by misreporting their actual costs, which is undesirable from the auctioneer's perspective. The goal is to design budgetfeasible mechanisms that (i) elicit truthful reporting of the costs by all agents (i.e., strategyproofness) and (ii) achieve a good approximation with respect to the optimal value for the auctioneer (i.e., approximate efficiency). The problem of designing budget-feasible mechanisms was introduced by Singer [2] and has received a lot of attention, both because of its theoretical appeal and of its relevance to several emerging application domains.

In this project, we focus on the most basic setting of the problem, where the valuation function $v(\cdot)$ is additive. The problem then relates to the classic knapsack problem and is generally well understood [1]. However, we add a twist to this setting by assuming that the auctioneer can use some *predictions* of the actual valuations of the agents. Such predictions could be extracted, for example, from the bidding behavior of the agents observed in the past. The question we would like to address in this new model is to which extent these predictions can help to derive improved mechanisms.

Prerequisites: Good knowledge of mechanism design (single-parameter setting, Myerson's lemma) and approximation algorithms (knapsack problem); students having followed the Algorithmic Game Theory course should be well equipped.

Interested? Please get in touch with Guido Schafer (g.schaefer@cwi.nl)

Note: The above is just one specific proposal of a potential research project related to algorithmic game theory. If you have interesting project proposals yourself or want to discuss other possibilities, please feel free to get in touch.

References

[1] Nick Gravin, Yaonan Jin, Pinyan Lu, and Chenhao Zhang. Optimal budget-feasible mechanisms for additive valuations. In ACM Conference on Economics and Computation (EC), pages 887–900. Association for Computing Machinery, 2020. https://dl.acm.org/doi/10.1145/3417746.

[2] Yaron Singer. Budget feasible mechanisms. In IEEE Symposium on Foundations of Computer Science (FOCS), pages 765–774. IEEE Computer Society, 2010. https://ieeexplore.ieee.org/document/5671347.

Inefficiency of Corruption in Single-item Auctions

Supervision: Prof. dr. Guido Schafer
Networks and Optimization Group, Centrum Wiskunde & Informatica (CWI)
Institute for Logic, Language and Computation (ILLC), University of Amsterdam

If applicable, an internship ("stagiair") status at CWI can be arranged for the student (to be discussed).

Keywords: Auction Theory, Algorithmic Game Theory, Equilibrium Analysis, Price of Anarchy

Description: We consider auction settings where a seller wants to sell some items and for this purpose recruits an auctioneer to organize an auction on their behalf. Such settings are widely prevalent in practice as they emerge naturally whenever the seller lacks the expertise (or facilities, time, etc.) to host the auction themselves. The dilemma in such settings is that the incentives of the seller and the auctioneer are rather diverse in general: while the seller is interested in extracting the highest payments for the objects (or getting service at the lowest cost), the agent primarily cares about maximizing their own gains from hosting the auction. This misalignment leads (unavoidably) to fraudulent schemes used by the auctioneer to manipulate the auction to their own benefit.

Corruption in auctions, where an auctioneer engages in bid rigging with one (or several) of the bidders, occurs rather frequently in practice, especially in the public sector (e.g., in construction and procurement auctions). Despite this fact, the negative impact of corrupt auctioneers is still poorly understood theoretically and only a few studies exist. The social welfare loss caused by corrupt auctioneers in fundamental first-price auction settings (i.e., single-item and multiunit auction settings) has recently been studied in [1]. This work focusses on a basic corruption scheme, where the auctioneer colludes with the winning bidders. As it turns out, this setting can be reduced to a *hybrid auction format*, where the items are assigned to the highest bidders and the payments are a convex combination of the first-price and the second-price payments.

In this project, we want to investigate (slightly) more sophisticated corruption schemes in the single-item auction settings. The goal is to study the inefficiency of equilibria of such schemes by deriving bounds on their *(robust) price of anarchy (POA)* through smoothness techniques. Also, an interesting question will be to understand under which circumstances the relation between specific corruption schemes and hybrid auction forms continues to hold.

Prerequisites: Good knowledge of standard auction formats (first-price, second-price single-item auctions) and price of anarchy analysis (robust price of anarchy, smoothness technique); students having followed the Algorithmic Game Theory course should be well equipped.

Interested? Please get in touch with Guido Schafer (g.schaefer@cwi.nl).

Note: The above is just one specific proposal of a potential research project related to algorithmic game theory. If you have interesting project proposals yourself or want to discuss other possibilities, please feel free to get in touch.

References

[1] Andries van Beek, Ruben Brokkelkamp, and Guido Schafer. Capturing corruption with hybrid auctions: Social welfare loss in multi-unit auctions. CoRR, abs/2106.01822, 2021. https://arxiv.org/abs/2106.01822.

Computing real solutions to polynomial equations

Supervisors: Simon Telen and Monique Laurent

Key words: Polynomial functions, Julia implementation

Problem description:

We consider the problem of solving a system of multivariate equations

$$f_1(x_1, \ldots, x_n) = \cdots = f_s(x_1, \ldots, x_n) = 0, \qquad (1.1)$$

where $f_i \in \mathbb{R}[x_1, \ldots, x_n]$ are polynomial functions in the unknowns $x_1, \ldots, x_n$. Geometrically, each of the polynomials f_i defines a hypersurface V_i in $\mathbb{R}^n$, and our solutions form the intersection $V_1 \cap \cdots \cap V_s$ of all these hypersurfaces. In the real plane $\mathbb{R}^2$, we intersect curves.

Example 1.1 ($s = n = 2$). *Let $f_1 = x_1^2 + x_1x_2 + x_2^2 - 1$ and $f_2 = (x_1 - 1)^2 + x_2^2 - 1$. The solutions to $f_1 = f_2 = 0$ are the intersection points of an ellipse and a circle in the real plane. Figure 1 shows that there are two solutions. A more complicated situation is illustrated in the right part of that figure, where the equations each have degree 20.* ◇

Many standard algorithms for solving (1.1) intrinsically work over the complex numbers. That is, they compute all points $(x_1, \ldots, x_n) \in \mathbb{C}^n$ satisfying (1.1), and then select the real solutions. This holds for homotopy continuation methods, as well as for normal form methods. See for instance [3] for an overview. One way to avoid computing all complex solutions is to add equations $f_{s+1}, \ldots, f_t$ which are only satisfied by the real solutions.

Example 1.2. *The equations $f_1 = f_2 = 0$ from Example 1.1 have 4 complex solutions, 2 of which are real. Adding the equation $f_3 = 1 - (1 + \alpha)x_1 - \alpha x_2 = 0$, with $\alpha \approx 0.5652$, only the two real solutions are left. Geometrically, we add the dotted line to the picture.* ◇

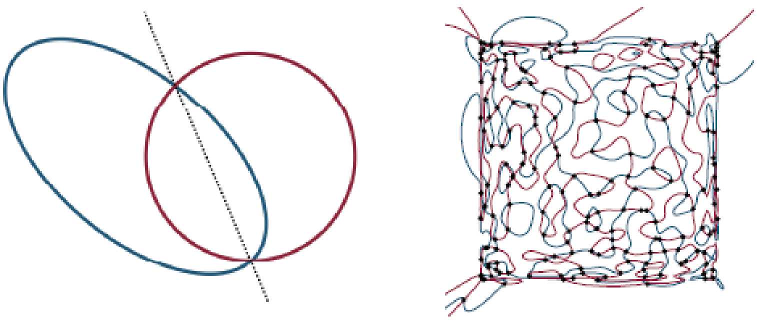

Figure 1: Systems of bivariate equations represent intersecting plane curves.

Project goals: The goal of this project is to compute the equations $f_{s+1}, \ldots, f_t$ needed to eliminate all nonreal solutions of (1.1), using adapted versions of the moment method [2]. We will compare with the recent Julia implementation in [1]. We will use this in a numerical method for finding all real solutions of (1.1), and test our algorithms on real-life problems.

Prerequisits: This project is suitable for students in mathematics or related subjects, with a strong knowledge of basic linear algebra and an interest in algebra and computation.

Contact: For questions, please send an e-mail to Simon.Telen@cwi.nl.

References

[1] L. Baldi and B. Mourrain. Computing real radicals by moment optimization. In Proceedings of the 2021 on International Symposium on Symbolic and Algebraic Computation, pages 43–50, 2021.

[2] M. Laurent and P. Rostalski. The approach of moments for polynomial equations. In Handbook on Semidefinite, Conic and Polynomial Optimization, pages 25–60. Springer, 2012.

[3] S. Telen. Solving Systems of Polynomial Equations. PhD thesis, KU Leuven, Leuven, Belgium, 2020. Available at https://simontelen.webnode.page/publications/.

Geodesic Optimization and the Paulsen Problem in Frame Theory

Supervisors: Daniel Dadush, Akshay Ramachandran, Michael Walter

Key words: Geodesic Optimization, Paulsen Problem in Frame Theory

We consider the following natural linear algebraic question:

Question 0.1. *Let $U = \{u_1, ..., u_n\} \subseteq \mathbb{R}^d$ be a spanning set of vectors satisfying*

$$\forall x \in \mathbb{R}^d : \frac{1-\varepsilon}{d} \leq \sum_{j=1}^{n} |\langle u_j, x\rangle|^2 \leq \frac{1+\varepsilon}{d}, \qquad \forall j \in [n] : \frac{1-\varepsilon}{n} \leq \|u_j\|_2^2 \leq \frac{1+\varepsilon}{n}. \tag{1}$$

What is the minimum dist$(V,U) := \sum_{j=1}^{n} \|v_j - u_j\|_2^2$ over all V exactly satisfying these conditions:

$$\forall x \in \mathbb{R}^d : \sum_{j=1}^{n} |\langle v_j, x\rangle|^2 = \frac{1}{d}, \qquad \forall j \in [n] : \|v_j\|_2^2 = \frac{1}{n}.$$

This is known as the Paulsen problem in frame theory and was listed as a major open problem ([4], [2]), for which little was known despite considerable effort. Frames satisfying the $\varepsilon = 0$ conditions above are known as doubly balanced frames. These give information theoretically optimal constructions for certain recovery tasks in signal processing and coding theory. On the other hand, doubly balanced frames can be difficult to construct explicitly, whereas it is easy to generate random frames satisfy 1 for some small ε with high probability. The Paulsen problem attempts to validate this approach by asking whether the frames satisfying 1 truly do approximate doubly balanced frames, and whether they can be rounded to nearby doubly balanced frames.

A priori, the distance bound may depend on the parameters (d, n, ε). There are known lower bounds showing that in the worst case $\text{dist}(V,U) \gtrsim \varepsilon$, whereas in practice the answer seems to be $\lesssim \varepsilon^2$. In a series of works [6, 5, 7], it was shown that the worst case lower bound ε is tight up to a constant factor.

Prior to these works, there were two partial results on the distance function [3, 1], which showed the distance bound $\text{poly}(d,n)\varepsilon^2$, but only in certain special cases. Note that this bound gives a better exponent for ε than the known worst case examples, and so cannot hold in general.

Project Goal: The procedures and analysis given in [3, 1] are slightly ad-hoc and unrelated to each other. We believe that the framework of group scaling and geodesic optimization (see [6, 5, 7]) can be used to simultaneously generalize and improve these results. This would give a unified and principled approach to the Paulsen problem. Our main goal is to exactly characterize the situations in which it is possible to improve the dependence on ε and prove a beyond worst case bound of the form ε^2. In particular, this would give theoretical justification for the improved performance of frames seen in practice.

References

[1] Bodmann, B. G., and Casazza, P. G. The road to equal-norm parseval frames. Journal of Functional Analysis 258, 2 (2010), 397-420.

[2] Cahill, J., and Casazza, P. The paulsen problem in operator theory. Operators and Matrices (2013).

[3] Casazza, P. G., Fickus, M., and Mixon, D. G. Auto-tuning unit norm frames. Applied and Computational Harmonic Analysis 32, 1 (2012), 1-15

[4] Casazza, P. G., and Kutyniok, G., Eds. Finite Frames: Theory and Applications. Birkhauser Basel, 2013.

[5] Hamilton, L., and Moitra, A. The paulsen problem made simple. In Innovations in Theoretical Computer Science (ITCS) (2019).

[6] Kwok, T., Lau, L., Lee, Y., and Ramachandran, A. The Paulsen problem, continuous operator scaling, and smoothed analysis. In Symposium on Theory of Computing (STOC) (2017), ACM.

[7] Ramachandran, A. Geodesic Convex Analysis of Group Scaling for the Paulsen Problem and Tensor normal model. PhD thesis, University of Waterloo, 2021.