APPROXIMATION IN ALGORITHMIC GAME THEORY: ROBUST APPROXIMATION BOUNDS FOR EQUILIBRIA AND AUCTIONS

TIM ROUGHGARDEN
STANFORD UNIVERSITY
SCRIBED BY QIQI YAN

1. Introduction

1.1. **Motivation.** Many modern computer science applications involve autonomous, self-interested agents. It is therefore important for us to consider agents' strategic behavior in modelling the problems, where non-cooperative game theory can be very helpful. Unfortunately, as one can expect, strategic behavior of the agents often make full optimality difficult or impossible for various reasons. Three common reasons are the following:

(1) Equilibria (e.g., Nash) of noncooperative games are typically suboptimal.
(2) Auctions lose revenue from strategic behavior.
(3) Incentive constraints can make polynomial time approximation of NP-hard problems even harder.

As a result, it is essential for approximation, or approximate optimality, to come into play. In particular, approximation is relevant in several topics of study in algorithmic game theory, notably:

(1) In the area of price of anarchy, we study worst-case approximation guarantees for equilibria.
(2) In revenue maximization, we discuss approximately-optimal revenue guarantees for auctions in non-Bayesian settings.
(3) In algorithmic mechanism design, we study approximation algorithms robust to selfish behavior.
(4) In computing approximate equilibria, we study the question of whether there is a PTAS for computing an approximate Nash equilibrium.

In this scribed talk, we focus on the first two topics. The third topic was covered in a FOCS 2010 tutorial given by the speaker and available online.

2. Price of Anarchy and Smoothness

2.1. **Price of Anarchy.** The notion of price of anarchy was introduced by Koutsoupias and Papadimitriou in [26]. It quantifies the inefficiency of equilibria w.r.t. an objective function. E.g., a typical equilibrium concept is Nash equilibrium, where at equilibrium, no player is better off by switching her strategy unilaterally.

These are scribed notes for a talk given by Tim Roughgarden at an approximation algorithms workshop held at Princeton University in June 2011.

Definition 1. The price of anarchy (POA) of a game (w.r.t. some nonnegative minimization objective function) is the supremum of:

$$\frac{\text{objective function value at equilibrium}}{\text{optimal objective function value}}.$$

Here the supremum is over all equilibria.

The POA for maximizing problems can be similarly defined. Either way, a ratio closer to 1 is more desirable.

Example 2. Consider the following network[1], where there are two players who want to go from s to t. Each player chooses exactly one path. Each edge in the network has a congestion function, which is the cost for each player using that edge, as a function of the total number of players using that edge. In Nash equilibrium, both players choose the same path, incurring a cost of $2 \times 2 + 5 \times 2$ each, summing up to 28 in total. On the other hand, to minimize total cost, one player should take the "top" path, while the other player should take the "bottom" path, with total cost of 24. Therefore, the price of anarchy in this case is $\frac{28}{24} = \frac{7}{6}$.

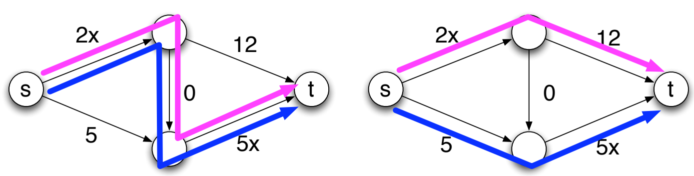

FIGURE 2.1. Nash Equilibrium and Optimal

A remark here is that there might exist multiple equilibria. When this is the case, we look at the worst one w.r.t. our objective.

2.2. **Call for Robustness.** A price of anarchy bound means that if the game is at an equilibrium, then the outcome is near-optimal. However, the reality is often much more complicated, and we may not reach equilibrium for the following potential reasons:

- There may not exist a (pure) equilibrium at all.
- An equilibrium can be computationally hard to find. [19, 15, 11]
- The equilibrium might be hard to learn in a distributed way.

These raise the worry: are our price of anarchy bounds "meaningless"?

We do not think that these concerns render price of anarchy bounds meaningless. However, we do need to look for more robust guarantees than simple POA bounds to address these concerns. Our high-level goal is the following:

> we want worst-case bounds that apply even to non-equilibrium outcomes.

Such outcomes could be:

- best-response dynamics, pre-convergence outcomes [28, 20, 2]
- correlated equilibria [12]

[1] I thank Kshipra Bhawalkar for providing the nicely drawn figures.

- coarse correlated equilibria a.k.a. "price of total anarchy" a.k.a. "no-regret players" [7, 8]

A smoothness condition will be the key to our goal.

2.3. **Smoothness and Price of Anarchy.** To define our smoothness framework, we need to introduce a somewhat abstract setup. In the setup, let there be n players, each picking a strategy s_i. Player i incurs a cost of $C_i(\mathbf{s})$, where $\mathbf{s} = (s_i)_i$ is the strategy vector. We assume that the objective function is the total cost $cost(\mathbf{s}) = \sum_i C_i(\mathbf{s})$. We remark that it is important that the cost is additive.

Our smoothness condition is the following:

Definition 3. A game is (λ, μ)-smooth ($\lambda > 0, \mu < 1$) if for every pair $\mathbf{s}, \mathbf{s}^*$ of outcomes, the following condition holds:

$$\sum_i C_i(s_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot cost(\mathbf{s}^*) + \mu \cdot cost(\mathbf{s})$$

This says that if we start with a strategy profile $\mathbf{s}$, and consider for each agent i the cost of i if she unilaterally deviates to her strategy in profile $\mathbf{s}^*$, then the sum of all such costs is upper-bounded jointly using costs of $\mathbf{s}$ and $\mathbf{s}^*$.

To make sense out of this condition, we first see how the smoothness condition implies price of anarchy bounds, in a canonical way. Such an argument is called a smoothness argument.

Theorem 4. *If a game is (λ, μ)-smooth, then the price of anarchy is at most $\frac{\lambda}{1-\mu}$.*

Proof. We assume that a game is (λ, μ)-smooth. Let $\mathbf{s}$ be a Nash equilibrium, and let $\mathbf{s}^*$ be an optimal solution. Then the following holds:

$$\begin{aligned} cost(\mathbf{s}) &= \sum_i C_i(\mathbf{s}) \\ &\leq \sum_i C_i(s_i^*, \mathbf{s}_{-i}) \\ &\leq \lambda \cdot cost(\mathbf{s}^*) + \mu \cdot cost(\mathbf{s}) \end{aligned}$$

Here the first line is by the additive definition of total cost, the second line is by the fact that $\mathbf{s}$ is a Nash equilibrium, and the third line is by the smoothness condition.

Now by rearranging, it follows that $cost(\mathbf{s}) \leq \frac{\lambda}{1-\mu} \cdot cost(\mathbf{s}^*)$, i.e., the price of anarchy of pure equilibria is at most $\frac{\lambda}{1-\mu}$. □

Remark 5. Note that smoothness condition is a stronger condition than is necessary to derive the above bound on the price of anarchy. To derive the price of anarchy bound, one only needs the smoothness condition to hold in the special case where $\mathbf{s}$ is a Nash equilibrium, and $\mathbf{s}^*$ is optimal. On the other hand, the smoothness condition requires the inequality to hold even when the profile $\mathbf{s}$ is not a pure Nash equilibrium.

For the settings listed below, price of anarchy bounds can be obtained via smoothness argumenta in this same canonical way.

- atomic (unweighted) selfish routing [3, 12, 1, 31]
- nonatomic selfish routing [32, 30, 14]

- weighted congestion games [1, 6]
- submodular maximization games [33, 27]
- coordination mechanisms [13]

2.4. **Beyond Nash Equilibria.** Smoothness arguments can be used to prove worst-case approximation guarantees for outcomes that are not Nash equilibria. We illustrate the idea here using the example of no-regret outcomes.

Definition 6. A sequence $s^1, s^2, \ldots, s^T$ of outcomes is no-regret if for each player i, and each fixed action q_i of the player i, the average cost player i incurs over the sequence is (in the limit as T goes to infinity) no worse than playing q_i every time.

It is well-known for example that if every player uses the multiplicative weights algorithm, then $o(1)$ regret is achieved in polynomial time.

Theorem 7 (An Out-Of-Equilibrium Bound)**.** [31] *In a (λ, μ)-smooth game, the average cost of every no-regret sequence is at most $\frac{\lambda}{1-\mu}$ times the cost of an optimal outcome, in the limit as T goes to infinity. (I.e., the same upper bound proved for pure Nash equilibria.)*

Proof. (sketch) Let $s^1, s^2, \ldots, s^T$ denote a no-regret outcome sequence, and let $\mathbf{s}^*$ denote the optimal outcome.

Assuming (λ, μ)-smoothness, we have:

$$\begin{aligned} \sum_t cost(\mathbf{s}^t) &= \sum_t \sum_i C_i(\mathbf{s}^t) \\ &= \sum_t \sum_i [C_i(s_i^*, \mathbf{s}_{-i}^t) + \Delta_{i,t}] \\ &\leq \sum_t [\lambda \cdot cost(\mathbf{s}^*) + \mu \cdot cost(\mathbf{s}^t)] + \sum_i \sum_t \Delta_{i,t} \end{aligned}$$

Here the first equality is by linear separability of cost, whereas in the second equality, the regret term$\Delta_{i,t}$ is defined as $C_i(\mathbf{s}^t) - C_i(s_i^*, \mathbf{s}_{-i}^*)$, and the inequality follows from the smoothness condition.

Now by the no-regret assumption, $\sum_t \Delta_{i,t} \leq o(T)$ for each i, and the proof finishes by rearranging, and dividing using T. □

Furthermore, the connection of the price of anarchy to smoothness is "intrinsic":

Theorem 8. [31] *For every set C of cost functions, unweighted congestion games with cost functions restricted to C are "tight" in the following sense: the maximum of the pure POA over congestion games with cost functions in C is equal to the minimum of $\frac{\lambda}{1-\mu}$ over (λ, μ) such that all such games are (λ, μ)-smooth.*

One can show that weighted congestion games [6] are also tight in this sense.

2.5. **What's Next.** There are several future research directions along these lines.

One direction is to beat the worst-case POA bounds by always reaching a non-worst-case equilibrium. Non-worst-case equilibria might be reached due to learning dynamics (e.g. [10, 25]) from modest intervention (e.g. [4]), etc.

Another direction is to study POA bounds for auctions. Practical auctions often lack "dominant strategies" (like in sponsored search, combinatorial auctions, etc.), and therefore it would be nice if we can have bounds on their (Bayes-Nash) equilibria [5, 24].

2.6. **Summary of Key Points.** To summarize this section, smoothness is a "canonical way" to bound the price of anarchy (for pure equilibria). It gives robust POA bounds in the sense that smoothness bounds extend automatically beyond Nash equilibria. Furthermore, smoothness bounds provably give optimal POA bounds in fundamental cases, and can be extended to approximate equilibria; best-response dynamics; local smoothness for correlated equilibria; and Bayes-Nash equilibria.

3. Auctions

In this section we study revenue-maximizing auctions. We start by pointing out the failure of traditional competitive analysis in this context.

3.1. **Failure of Traditional Competitive Analysis in Auctions.** Let us consider a very simple problem. We have two bidders and one item for sale. For such problems, which auction (e.g., whether to pick an opening bid of 1 cent or 10 dollars) is best heavily depends on the (unknown) input bids from the bidders.

One could try to use competitive analysis: compare your revenue to that obtained by an omniscient opponent. However, it is easy to prove that no auction can guarantee a meaningful fraction of this benchmark, and all auctions are equally terrible in terms of competitive ratio. Therefore, a novel analysis framework is needed.

3.2. **A New Analysis Framework.** We consider the prior-independent analysis framework of [21, 22], where we compare revenue to that of an opponent with statistical information about input. The goal is then to design a distribution-independent auction that is always near-optimal for the underlying distribution (no matter what the distribution is). Note that here distribution over inputs is not used in the design of the auction, but only used in its analysis.

3.2.1. *Bulow-Klemperer Theorem.* The classical Bulow-Klemperer theorem can be seen as one such prior-independent result, and in the rest of this section we focus on the Bulow-Klemperer theorem. (Technically, the Bulow-Klemperer theorem is not exactly the same as a prior-independent approximation guarantee as it needs the additional bidder. But this need can be easily eliminated [17] at a constant loss in the approximation factor.)

Let us consider the auction problem of selling a single item to one of n bidders. Each bidder i has a valuation v_i for winning the item, where the v_i's are drawn i.i.d. from a distribution F. Myerson characterized the optimal auction for this context, where the optimal auction is dependent on the distribution F.

We make the standard regularity condition on the distribution F, as defined in Myerson [29]. The classical Bulow-Klemperer theorem [9] says the following:

Theorem 9. [9] *For every n, and regular F, the expected revenue of Vickrey's auction with $n+1$ i.i.d. bidders from F is at least the expected revenue of the optimal auction with n i.i.d. bidders from F.*

One interpretation of the theorem is that a small increase in competition is more important than running the optimal auction. This also in some sense gives a "bicriteria bound". Note that the Bulow-Klemperer theorem is prior-independent, in that the Vickrey auction does not need knowledge about the underlying distribution.

3.2.2. *Bayesian Profit Maximization.* To prove the Bulow-Klemperer theorem, we first need to understand the optimal auction given distributional information.

First, consider one bidder and one item. The value of the bidder for the item v is drawn from a known distribution F, and we want to choose an optimal reserve price p to maximize revenue. Suppose we offer price p, the expected revenue of p is $p(1 - F(p))$. We can solve for the optimal p^*, e.g. $p^* = \frac{1}{2}$ for $v \sim uniform[0, 1]$.

But what about cases beyond one bidder and one item (and i.i.d. draws v_i)? Myerson characterized the optimal auction as follows:

Theorem 10. [29] *The auction that gives the maximum expected revenue is the second-price auction (or Vickrey auction) with the above-defined reserve p^*. (Note: p^* is independent of n.)*

Recall that the Bulow-Klemperer theorem says that for every n, Vickrey's expected revenue over $n+1$ i.i.d. bidders is as good as the maximum-possible expected revenue from n i.i.d. bidders. In the following, we prove this theorem for the case of $n = 1$, which is in some sense the most interesting and non-trivial case.

Let $n = 1$ for now. Note that in a Vickrey auction with $n + 1 = 2$ i.i.d. bidders, from each bidder's perspective, she faces a take-it-or-leave-it offer which equals to the bid of the other bidder, which is in turn is a random draw from distribution F. Therefore, the Bulow-Klemperer theorem for the case of $n = 1$ can be reformulated as follows:

Claim: twice the expected revenue from one bidder with a random reserve (drawn from F) is as good as the expected revenue from one bidder with an optimal reserve r^*.

One advantage of this reformulation is that the problem is now purely about selling one item to one bidder.

We define the (expected) revenue function $R(q) = q \cdot F^{-1}(1 - q)$ as a function of sale probability. I.e., if we sell at a price so that the probability of sale is q, the price has to be $F^{-1}(1-q)$, and the expected revenue is $R(q)$. It turns out that the revenue function $R(q)$ is concave in q if and only if the distribution F is regular.

Now let q^* maximize $R(q)$. Then the optimal revenue is $R(q^*)$, or the "height" of the curve. On the other hand, the revenue of using a random reserve drawn from F equals to the expected value of $R(q)$ with q uniformly drawn from $[0, 1]$, which equals to the area under the curve and above the q axis. By concavity of $R(q)$, the area under the curve is at least half of the height of the curve, completing the proof.

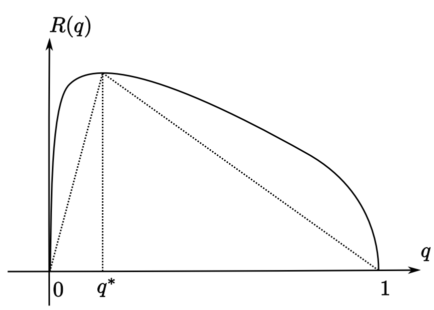

FIGURE 3.1. Revenue Curve

3.3. **Related Work.** There has been a lot of recent progress on prior-independent approximations. More general forms of such "bicriteria bounds" were obtained in [22, 18], and more prior-independent approximations have been achieved in [16, 17, 23].

3.4. **Summary.** Standard competitive analysis is useless for worst-case revenue maximization. But we can get a simultaneous competitive guarantee with all Bayesian-optimal auctions.

Toward future directions, can we thoroughly understand problems in the so-called “single-parameter” domain, include "non-downward-closed" ones, and what about non-i.i.d. settings? Or what about multi-parameter settings, like combinatorial auctions?

Epilogue

The higher-level moral of this survey talk is the following. Worst-case approximation guarantees can be a powerful "intellectual export" from theoretical computer science to other fields (e.g., to economics and game theory). There are many reasons why we should consider approximation (not just computational tractability), and studying approximation in various contexts gives not only robust guarantees, but also new insights into the problem.

References

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *Proceedings of the 23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3884 of *Lecture Notes in Computer Science*, pages 218–229, 2006.

[2] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrkoni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pages 264–273, 2008.

[3] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

[4] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Circumventing the price of anarchy: Leading dynamics to good behavior. In Andrew Chi-Chih Yao, editor, *ICS*, pages 200–213. Tsinghua University Press, 2010.

[5] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In Dana Randall, editor, *SODA*, pages 700–709. SIAM, 2011.

[6] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In Mark de Berg and Ulrich Meyer, editors, *ESA (2)*, volume 6347 of *Lecture Notes in Computer Science*, pages 17–28. Springer, 2010.

[7] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: On convergence to nash equilibria of regret-minimizing algorithms in routing games. *Theory of Computing*, 6(1):179–199, 2010.

[8] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In Cynthia Dwork, editor, *STOC*, pages 373–382. ACM, 2008.

[9] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

[10] M. Charikar, H. Karloff, C. Mathieu, J. Naor, and M. Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th ACM Symposium on Parallelism in Algorithms and Architextures (SPAA)*, pages 70–76, 2008.

[11] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3), 2009.

[12] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

[13] R. Cole, J. R. Correa, V. Gkatzelis, V. S. Mirrokni, and N. Olver. Inner product spaces for minsum coordination mechanisms. In *STOC*, pages 539–548, 2011.

[14] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. On the inefficiency of equilibria in congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.

[15] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM J. Comput*, 39(1):195–259, 2009.

[16] N. Devanur and J. D. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 41–50, 2009.

[17] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*, pages 129–138, 2010.

[18] S. Dughmi, T. Roughgarden, and M. Sundararajan. Revenue submodularity. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 243–252, New York, NY, USA, 2009. ACM.

[19] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

[20] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[21] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84, 2008.

[22] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 225–234, 2009.

[23] J. D. Hartline and Q. Yan. Envy, truth, and profit. In *ACM Conference on Electronic Commerce*, pages 243–252, 2011.

[24] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *ACM Conference on Electronic Commerce*, pages 295–296, 2011.

[25] R. D. Kleinberg, G. Piliouras, and É. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 533–542, 2009.

[26] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[27] J. R. Marden and T. Roughgarden. Generalized efficiency bounds in distributed resource allocation. In *Proceedings of the Conference on Decision and Control (CDC)*, pages 2233–2238. IEEE, 2010.

[28] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 183–194, 2004.

[29] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[30] G. Perakis. The price of anarchy when costs are non-separable and asymmetric. In *Proceedings of the 10th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3064 of *Lecture Notes in Computer Science*, pages 46–58, 2004.

[31] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, pages 513–522, 2009.

[32] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[33] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.