Parallel approximation of min-max problems with applications to classical and quantum zero-sum games*

Gus Gutoski[†] Xiaodi Wu[‡]

Abstract

This paper presents an efficient parallel algorithm for a new class of min-max problems based on the matrix multiplicative weight (MMW) update method. Our algorithm can be used to find near-optimal strategies for competitive two-player classical or quantum games in which a referee exchanges any number of messages with one player followed by any number of additional messages with the other. This algorithm considerably extends the class of games which admit parallel solutions and demonstrates for the first time the existence of a parallel algorithm for *any* game (classical or quantum) in which one player reacts adaptively to the other.

As a direct consequence, we prove that several competing-provers complexity classes collapse to PSPACE such as $\mathrm{QRG}(2)$, SQG and two new classes called DIP and DQIP. A special case of our result is a parallel approximation scheme for a new class of semidefinite programs whose feasible region consists of n-tuples of semidefinite matrices that satisfy a "transcript-like" consistency condition. Applied to this special case, our algorithm yields a direct polynomial-space simulation of multi-message quantum interactive proofs resulting in a first-principles proof of $\mathrm{QIP} = \mathrm{PSPACE}$. It is noteworthy that our algorithm establishes a new way, called the *min-max* approach, to solve SDPs in contrast to the *primal-dual* approach to SDPs used in the original proof of $\mathrm{QIP} = \mathrm{PSPACE}$.

Competitive multi-turn two-player (say, Alice and Bob) games are often studied in the classical game theory either from the aspect of computing the game values or from the aspect of the complexity classes induced by those game models. For succinct games, exponential-time algorithm exists for finding the exact value [KM92, KMvS94] and it is also EXP-hard to approximate the game value [FIKU08, FKS95]. The situation is much different for shorter games, where succinct two-turn games admit polynomial-space approximation scheme and are also PSPACE-hard to approximate [FK97]. Approximating one-turn games is known to be $\mathrm{S}_2^{\mathrm{P}}$-complete [FIKU08].

For each competitive game model, one can analogously define the corresponding competing provers interactive proofs (also called *refereed games*), where players become competing provers who are trying to convince some randomized polynomial-time verifier to either accept or reject on some input x. Let $\mathrm{RG(k)}$ denote the complexity class of problems that admit classical refereed games of k-turns and RG be short for $\mathrm{RG(poly)}$. Thus the above algorithmic results imply $\mathrm{RG} = \mathrm{EXP}$ and $\mathrm{RG}(2) = \mathrm{PSPACE}$.

Those game settings naturally extend to quantum case where provers and referees are allowed to exchange and process quantum information. It is known that the class of problems that admit quantum refereed games, denoted by QRG, coincide with its classical counterpart RG and henceforth EXP [GW07]. Also there exists a polynomial-space approximation scheme for quantum one-turn refereed games [JW09]. However, much more remains unknown about quantum refereed games of small number of turns.

In this paper, we consider the following class of competitive two-player refereed games, either classical or quantum, that subsumes all the quantum refereed games of small number of turns studied so far [Gut05, GW05, GW07].

(i) The referee exchanges several messages only with Alice.

(ii) After processing this interaction with Alice, the referee exchanges several additional messages only with Bob. After further processing, the referee declares a winner.

[*] Full version available at arXiv:1011.2787 [quant-ph]. Please refer to the most recent version as a new version will be posted very shortly.

[†] Institute for Quantum Computing and School of Computer Science University of Waterloo, Waterloo, Ontario, Canada

[‡] Department of Electrical Engineering and Computer Science University of Michigan, Ann Arbor, Michigan, USA

Due to the similarity with the oft-studied interactive proof model of computation, we denote games of this form by *double interactive proofs*: the referee in such a game executes a standard interactive proof with Alice followed by a second interactive proof with Bob. Ordinary interactive proofs are thus special cases when referee completely ignores Bob. One can also define the corresponding complexity classes associated with classical and quantum double interactive proofs, which are denoted by DIP and DQIP respectively.

Our Results

The main contribution of this paper is an efficient parallel algorithm for a new class of min-max problems associated with the classical and quantum double interactive proofs. If the referee is specified succinctly by circuits then our parallel algorithm can be used to find near-optimal strategies in polynomial space (via the relation $\mathrm{NC(poly)} = \mathrm{PSPACE}$ [Bor77]). This algorithm is optimal in that it is PSPACE-hard even to distinguish games that Alice can win with near certainty from games that Bob can win with near certainty, even in the special case of *two-turn games* [FK97] where the referee exchanges only two messages *synchronously* with each player.

Prior to the present work polynomial-space algorithms were known only for two-turn classical games [FK97], one-turn quantum refereed game [JW09], and for quantum interactive proofs [JJUW10, Wu10a]. Our result unifies and subsumes both of these algorithms. It also demonstrates for the first time the existence of a parallel algorithm for two-turn quantum games and for any game (classical or quantum) in which one player reacts adaptively to the other.

When applied to complexity theory, our result implies the collapse to PSPACE of the newly defined double interactive proofs $\mathrm{DQIP}, \mathrm{DIP}$. A special case of our result yields the equality

$$\mathrm{SQG} = \mathrm{QRG}(2) = \mathrm{PSPACE},$$

thus solving the open problems of Ref. [GW05, JJUW10].

Our result also illustrates a difference in the role of public randomness between *single*-prover interactive proofs and *competing*-prover interactive proofs. Any classical single prover interactive proof can be simulated by another *public coin* interactive proof (known as *Arthur-Merlin games*) where the verifier only sends uniformly random bits to the prover and [GS89]. Extending the notion of public coin interaction to refereed games, it is easy to see that an arbitrary multi-turn public-coin refereed game can be simulated by a double interactive proof. Therefore one has the public-coin version of RG is a subset of DIP, which equals PSPACE. Thus, by contrast to the single-prover case where we have $\mathrm{public\text{-}coin\text{-}IP} = \mathrm{IP}$, in the competing-prover case we have $\mathrm{public\text{-}coin\text{-}RG} \neq \mathrm{RG}$ unless $\mathrm{PSPACE} = \mathrm{EXP}$.

As a special case our algorithm yields a direct polynomial-space simulation of multi-message quantum interactive proofs, resulting in a first-principles proof of $\mathrm{QIP} = \mathrm{PSPACE}$. By contrast, all other known proofs [JJUW10, Wu10a] were based on the simplified yet equivalent model of quantum interactive proofs [KW00, MW05].

Our main result is achieved through the following technical steps. Consider the feasible region $\mathbf{A}$ defined below.

$$\mathbf{A} = \{(X_1, X_2, \cdots, X_n) : \forall i, X_i \succeq 0 \text{ and } \mathrm{Tr}_{\mathcal{C}_1}(X_1) = Q, \mathrm{Tr}_{\mathcal{C}_i}(X_i) = \mathrm{Tr}_{\mathcal{C}_i}(V_{i-1}X_{i-1}V_{i-1}^*) \text{ for } i = 2, \cdots, n\}$$

where $\mathrm{Tr}_{\mathcal{C}_1}, \ldots, \mathrm{Tr}_{\mathcal{C}_n}$ are partial trace maps and $V_1, \cdots, V_{n-1}$ are unitary operators. Note that such set $\mathbf{A}$ corresponds to Kitaev's *transcript* representation of quantum interactions. Our first step can be stated in full generality as follows

Theorem 1 (Informal). *Let $\mathbf{P}$ denote some convex compact set. For any appropriately bounded Q, $\mathbf{P}$, there exists an efficient parallel oracle-algorithm for finding approximate solutions to the min-max problem*

$$\min_{(X_1,\ldots,X_n)\in\mathbf{A}} \max_{P\in\mathbf{P}} \mathrm{Tr}(X_n P) \tag{1}$$

with an oracle for optimization over the set $\mathbf{P}$.

For the purpose of approximating quantum double interactive proofs, our second step demonstrates parallel implementations of the oracle for the corresponding $\mathbf{P}$. This step is established via another use of the algorithms appearing in Theorem 1 in special cases. Namely our main algorithm calls special instances of itself as subroutines.

As a special case of the min-max problem where $\mathbf{P} = \{P\}$ is a singleton set, our algorithm yields a parallel approximation scheme for the following semidefinite programs (SDPs).

$$\min \mathrm{Tr}(X_n P) \text{ s.t. } (X_1, X_2, \cdots, X_n) \in \mathbf{A}$$

It has long since been known [Ser91, Meg92] that the problem of approximating the optimal value of an arbitrary SDP is logspace-hard for P, so there cannot be a parallel approximation scheme for *all* SDPs unless $\mathrm{NC} = \mathrm{P}$. However, the precise extent to which SDPs admit parallel solutions is not known. Our result adds considerably to the set of such SDPs.

Techniques

Our algorithm is an example of the *matrix multiplicative weights update method (MMW)* [AHK05, Kal07, WK06]. We also draw upon the valuable experience of recent applications of this method in quantum computation [JW09, JUW09, JJUW10, Wu10a]. However, our application of the MMW method is somewhat different from all previous ones in the sense that our algorithm is applied *twice* in a two-level recursive fashion. At the top level, our algorithm makes use of the MMW method to solve a min-max problem. At the bottom level, a special case of our algorithm is used to solve a SDP problem as the implementation of the oracle for any min-max problem required by the MMW method.

Previously the MMW was applied to SDPs in the *primal-dual* way, where MMW method is utilized to either find a feasible solution to the primal problem with small objective function value or generate an approximately feasible solution to the dual problem that is used later to bound the optimum value from below. By contrast, we do not take such primal-dual approach—our SDP solution arises as a *special case* of a more general min-max problem. More detailed comparisons between the two methods can be found in the full version paper or the reference [Wu10b].

Competing quantum games with multi-turns admit a natural representation by *quantum strategy* [GW07], which may be viewed as a special type of channel specified by its Choi-Jamiolkowski matrix [Wat08, Lecture 5]. However optimizing over such representation is a task fraught with difficulty [JUW09]. Fortunately, double quantum interactive proofs admit another representation, namely Kitaev's *transcript* representation [Kit02], for strategies that is more suitable for our purpose. Intuitively the actions of a player are represented by a list $\rho_1, \ldots, \rho_n$ of density matrices, corresponding to "snapshots" of the state of the referee's qubits at various times , that satisfy a special consistency condition.

The key property of double quantum interactive proofs that we exploit is the ability to draw a "temporal line" in the interaction just after Alice's last action. Given a transcript $\rho_1, \ldots, \rho_n$ for Alice, the actions of Bob can then be represented by another transcript $\xi_1, \ldots, \xi_m$. By optimizing over all such transcripts one obtains an oracle for "best responses" for Bob to a given strategy of Alice as required by the MMW. Whereas the MMW in its unaltered form can be used to solve min-max problems over the domain of density operators, we introduce a new extension to this method for min-max problems over the domain of transcripts—a domain consisting of lists of *multiple* operators, each drawn from a *strict subset* of the density operators. The high-level approach of our method is as follows:

1. **Extend the domain from a single density matrix to a list of n density matrices.**
   This step is straightforward: the MMW can be applied directly to all n density matrices at the same time.

2. **Restrict the domain to a strict subset of density matrices.**
   This step is more difficult. It is accomplished by relaxing the game so as to allow *all* density matrices, with an additional *penalty term* to remove incentive for the players to use inconsistent transcripts.

3. **Round strategies in the relaxed game to strategies in the original game.**
   For this step one must prove a "rounding" theorem , which establishes that near-optimal, fully admissible strategies can be obtained from near-optimal strategies in the unrestricted domain with penalty term.

Finally, it is noteworthy that the proof of our rounding theorem contains an interesting and nontrivial application of the Bures metric, which is a distance measure for quantum states that is defined in terms of the more familiar fidelity function. Properties of the trace norm, which captures the physical distinguishability of quantum states, are often sufficient for most needs in quantum information. When some property of the fidelity is also required one uses the Fuchs-van de Graaf inequalities to convert between the trace norm and fidelity [FvdG99].

However, every such conversion incurs a quadratic slackening of relevant accuracy parameters. Our study calls for repeated conversions, which would incur an unacceptable exponential slackening if done naively via Fuchs-van de Graaf. Instead, we make only a *single* conversion between the trace norm and the Bures metric and then repeatedly exploit the simultaneous properties of (i) the triangle inequality, (ii) contractivity under quantum channels, and (iii) preservation of subsystem fidelity.

References

[AHK05] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta algorithm and applications. 2005.

[AK07] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC 2007)*, pages 227–236, 2007.

[Bor77] Allan Borodin. On relating time and space to size and depth. *SIAM Journal on Computing*, 6(4):733–744, 1977.

[Fan53] K. Fan. Minimax theorems. *Proceedings of the National Academy of Sciences*, 39:42–47, 1953.

[FIKU08] Lance Fortnow, Russell Impagliazzo, Valentine Kabanets, and Christopher Umans. On the complexity of succinct zero-sum games. *Computational Complexity*, 17(3):353–376, 2008.

[FK97] Uriel Feige and Joe Kilian. Making games short. In *Proceedings of the 29th ACM Symposium on Theory of Computing (STOC 1997)*, pages 506–516, 1997.

[FKS95] Joan Feigenbaum, Daphne Koller, and Peter Shor. A game-theoretic classification of interactive complexity classes. In *Proceedings of the 10th Conference on Structure in Complexity Theory*, pages 227–237, 1995.

[FvdG99] Christopher Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999. arXiv:quant-ph/9712042v2.

[GS89] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In Silvio Micali, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 73–90. JAI Press, 1989.

[Gut05] Gus Gutoski. Upper bounds for quantum interactive proofs with competing provers. In *Proceedings of the 20th IEEE Conference on Computational Complexity (CCC'05)*, pages 334–343, 2005.

[GW05] Gus Gutoski and John Watrous. Quantum interactive proofs with competing provers. In *Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science (STACS'05)*, volume 3404 of *Lecture Notes in Computer Science*, pages 605–616. Springer, 2005. arXiv:cs/0412102v1 [cs.CC].

[GW07] Gus Gutoski and John Watrous. Toward a general theory of quantum games. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC 2007)*, pages 565–574, 2007. arXiv:quant-ph/0611234v2.

[GW10] Gus Gutoski and Xiaodi Wu. Short quantum games characterize PSPACE. arXiv:1011.2787v1.

[JJUW10] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP=PSPACE. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 573–582, 2010. arXiv:0907.4737v2 [quant-ph].

[JUW09] Rahul Jain, Sarvagya Upadhyay, and John Watrous. Two-message quantum interactive proofs are in PSPACE. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 534–543, 2009. arXiv:0905.1300v1 [quant-ph].

[JW09] Rahul Jain and John Watrous. Parallel approximation of non-interactive zero-sum quantum games. In *Proceedings of the 24th IEEE Conference on Computational Complexity (CCC 2009)*, pages 243–253, 2009. arXiv:0808.2775v1 [quant-ph].

[Kal07] Satyen Kale. *Efficient algorithms using the multiplicative weights update method.* PhD thesis, Princeton University, 2007.

[Kit02] Alexei Kitaev. Quantum coin-flipping. Presentation at the 6th Workshop on *Quantum Information Processing* (QIP 2003), 2002.

[KM92] Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992.

[KMvS94] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the 26th ACM Symposium on Theory of Computing (STOC 1994)*, pages 750–759, 1994.

[KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof system. In *Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC 2000)*, pages 608–617, 2000.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.

[LN93] Michael Luby and Noam Nisan. A parallel approximation algorithm for positive linear programming. In *Proceedings of the 25th ACM Symposium on Theory of Computing (STOC 1993)*, pages 448–457, 1993.

[Meg92] Nimrod Megiddo. A note on approximate linear programming. *Information Processing Letters*, 42(1):53, 1992.

[MW05] Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005. arXiv:cs/0506068v1 [cs.CC].

[NC00] Michael Nielsen and Issac Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[Ser91] Maria Serna. Approximating linear programming is log-space complete for P. *Information Processing Letters*, 37(4):233–236, 1991.

[Sha92] Adi Shamir. IP = PSPACE. *Journal of the ACM*, 39(4):869–877, 1992.

[TX98] Luca Trevisan and Fatos Xhafa. The parallel complexity of positive linear programming. *Parallel Processing Letters*, 8(4):527–533, 1998.

[vN28] J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1928):295–320, 1928

[vzG93] J. von zur Gathen. Parallel linear algebra. *In J. Reif, editor, Synthesis of Parallel Algorithms*, 1993.

[WK06] Manfred Warmuth and Dima Kuzmin. Online variance minimization. In *Proceedings of the 19th Annual Conference on Learning Theory*, volume 4005 of *Lecture Notes in Computer Science*, pages 514–528. Springer, 2006.

[Wat08] John Watrous. *Lecture Notes on Theory of Quantum Information*. 2008.

[Wu10a] Xiaodi Wu. Equilibrium value method for the proof of QIP=PSPACE. arXiv:1004.0264v4 [quant-ph], 2010.

[Wu10b] Xiaodi Wu. Parallized solutions to semidefinite programmings in quantum complexity theory. arXiv:1009.2211v1 [quant-ph], 2010.