



Operator Scaling: Theory and Applications

Ankit Garg¹ · Leonid Gurvits² · Rafael Oliveira³ · Avi Wigderson⁴

Received: 7 March 2018 / Revised: 29 January 2019 / Accepted: 5 February 2019 / Published online: 6 May 2019
© SFoCM 2019

Abstract

In this paper, we present a deterministic polynomial time algorithm for testing whether a symbolic matrix in *non-commuting* variables over \mathbb{Q} is invertible or not. The analogous question for commuting variables is the celebrated polynomial identity testing (PIT) for symbolic determinants. In contrast to the commutative case, which has an efficient probabilistic algorithm, the best previous algorithm for the non-commutative setting required exponential time (Ivanyos et al. in *Comput Complex* 26(3):717–763, 2017) (whether or not randomization is allowed). The algorithm efficiently solves the “word problem” for the free skew field, and the identity testing problem for arithmetic formulae with division over non-commuting variables, two problems which had only exponential time algorithms prior to this work. The main contribution of this paper is a complexity analysis of an existing algorithm due to Gurvits (*J Comput Syst Sci* 69(3):448–484, 2004), who proved it was polynomial time for certain classes of inputs. We prove it always runs in polynomial time. The main component of our analysis is a simple (given the necessary known tools) lower bound on central notion of *capacity* of operators (introduced by Gurvits 2004). We extend the algorithm to actually approximate capacity to any accuracy in polynomial time, and use this analysis to give quantitative bounds on the continuity of capacity (the latter is used in a subsequent paper on Brascamp–Lieb inequalities). We also extend the algorithm to compute not only singularity, but actually the (non-commutative) rank of a symbolic matrix, yielding a factor 2 approximation of the commutative rank. This naturally raises a relaxation of the commutative PIT problem to achieving better deterministic approxi-

Communicated by Peter Bürgisser.

Ankit Garg: This research was done when the author was a student at Princeton University and his research was partially supported by Mark Braverman’s NSF Grant CCF-1149888, Simons Collaboration on Algorithms and Geometry, Simons Fellowship in Theoretical Computer Science and a Siebel Scholarship. Rafael Oliveira: This research was done when the author was a student at Princeton University and the research supported by NSF Career award (1451191) and by CCF-1523816 award. Avi Wigderson: This research was partially supported by NSF Grant CCF-1412958.

✉ Ankit Garg
garga@microsoft.com

Extended author information available on the last page of the article

mation of the commutative rank. Symbolic matrices in non-commuting variables, and the related structural and algorithmic questions, have a remarkable number of diverse origins and motivations. They arise independently in (commutative) invariant theory and representation theory, linear algebra, optimization, linear system theory, quantum information theory, approximation of the permanent and naturally in non-commutative algebra. We provide a detailed account of some of these sources and their interconnections. In particular, we explain how some of these sources played an important role in the development of Gurvits' algorithm and in our analysis of it here.

Keywords Derandomization · Non-commutative algebra · Polynomial identity testing · Invariant theory · Quantum operators

Mathematics Subject Classification 68Q10 · 68Q25

Contents

1	Introduction	225
1.1	The Free Skew Field	228
1.2	Word Problems and Identity Tests	230
1.3	Commutative and Non-commutative Rank of Symbolic Matrices	232
1.4	Compression Spaces, Optimization, and Gurvits' Algorithm G	233
1.5	Permanents, Quantum Operators, Origins and Nature of Algorithm G	235
1.6	Invariant Theory, Left–Right Action, Capacity and the Analysis of Algorithm G	237
1.7	Computation and Continuity of Capacity	241
	Subsequent Work	242
1.8	Organization	243
2	Quantum Operators and Analysis of Algorithm G	243
2.1	Completely Positive Operators and Capacity	243
2.2	Algorithm G and Its Convergence Rate	245
2.3	Main Theorem: The Lower Bound on Capacity	249
3	Properties of Capacity	253
3.1	Capacity of Almost Doubly Stochastic Operators	253
3.2	Characterization of Approximate Optimizers for Capacity	257
3.3	Multiplicativity of Capacity	259
4	Bit Complexity Analysis of Algorithm G and Continuity of Capacity	261
4.1	Evolution of the Operator Through Scaling	264
4.2	Bit Complexity Analysis of Algorithm G	267
4.3	Continuity of Capacity	269
5	Computing the Capacity of a Quantum Operator	274
5.1	Computing Capacity with Truncation	277
6	Conclusion and Open Problems	279
A	Symbolic Matrices with Polynomial Entries and Non-commutative Rank	281
A.1	Higman's Trick	281
A.2	Classical Reduction	284
A.3	The Quantum Reduction	286
	References	287

1 Introduction

This introduction will be unusually long, due to the unusual number of settings in which the problems we study appear, and their interconnections and history.

The main object of study in this paper are *symbolic* matrices (also called linear pencils) whose entries are linear forms in variables $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$ over a field¹ \mathbb{F} . Any such matrix can be expressed as a linear combination of the variables with matrix coefficients

$$L = x_1 A_1 + x_2 A_2 + \cdots + x_m A_m$$

where A_1, A_2, \dots, A_m are $n \times n$ matrices² over \mathbb{F} .

The main computational problem we will be concerned with in this paper (which we call SINGULAR) is determining whether such a symbolic matrix is invertible or not (over the field of fractions in the given variables). This problem has a dual life, depending on whether the variables commute or do not commute. In the *commutative* case, this problem has an illustrious history and significance. It was first explicitly stated by Edmonds [26] and shown to have a randomized polynomial time algorithm by Lovasz [61]. The completeness of determinant for arithmetic formulas by Valiant [80] means that singularity captures the celebrated polynomial identity testing (PIT) problem, and so in the commutative setting, we will refer to it as PIT. The derandomization of the latter probabilistic algorithm for PIT (namely, proving $\text{PIT} \in \mathcal{P}$) became all-important overnight when Kabanets and Impagliazzo [53] showed it would imply nontrivial arithmetic or Boolean lower bounds well beyond current reach.

On the other hand, in the *non-commutative* case even the meaning of this problem SINGULAR is unclear. It took decades to fully define and understand the related notion of a “field of fractions” for non-commutative polynomials, namely the *free skew field* over which we (attempt to) invert the matrix.³ But as we will explain below, this non-commutative SINGULAR problem has many intuitive and clean equivalent formulations (some entirely commutative!). It captures a non-commutative version of identity testing for polynomials and rational functions and provides a possible avenue to attack the notorious commutative PIT version, and quite surprisingly, its different formulations arise naturally in diverse areas of mathematics, revealing surprising connections between them. We only note that unlike the commutative PIT, it is not even clear from its definition that the problem SINGULAR is decidable. It requires some of the very nontrivial characterizations above, and other important results in commutative algebra, to prove a deterministic *exponential time* upper bound on its complexity

¹ Our main results will be for the rationals \mathbb{Q} (and will hold for \mathbb{R} and \mathbb{C} as well) but not for finite fields. However, many of the questions are interesting for any field.

² For all purposes, we may assume that the matrices A_i are linearly independent, namely span a space of matrices of dimension exactly m .

³ For now, the reader may think of the elements of this “free skew field” simply as containing all expressions (formulas) built from the variables and constants using the arithmetic operations of addition, multiplication and division (we define it more formally a bit later). We note that while this is syntactically the same definition one can use for commuting variables, the skew field is vastly more complex, and in particular its elements cannot be described canonically as ratios of polynomials.

(see two very different proofs in [17,51]), the best known before this work. No better bound was known even allowing randomness.

The main result of this paper is a *deterministic* polynomial time algorithm for this problem!

Theorem 1.1 *For non-commutative variables over \mathbb{Q} , SINGULAR $\in \mathcal{P}$. More specifically, there is a deterministic algorithm which, given m $n \times n$ integer matrices A_1, \dots, A_m with entries of bit size b , decides in time $\text{poly}(n, m, b)$ if the matrix $L = \sum_{i=1}^m x_i A_i$ is invertible over the free skew field.*

We now formulate the problem SINGULAR for non-commuting variables in its various contexts and incarnations. We will keep throughout an analogy with the commutative case. To avoid confusion, we will switch to calling commuting variables y , and keep x to denote non-commuting variables. As is common let $\mathbb{F}[\mathbf{y}]$ denote the algebra of polynomials in commuting variables $\mathbf{y} = \{y_1, y_2, \dots, y_m\}$ over a field \mathbb{F} , and $\mathbb{F}(\mathbf{y})$ the field of rational functions over these variables (every element of this field is a ratio of two polynomials in $\mathbb{F}[\mathbf{y}]$). For commuting variables, there are many simple equivalent formulations of the problem.

Fact 1.2 (PIT) *The following are equivalent for a matrix $L = \sum_{i=1}^m y_i A_i$ in commuting variables \mathbf{y} . Some of these equivalences only hold for large enough fields \mathbb{F} .*

- (1) L is singular (namely not invertible) over $\mathbb{F}(\mathbf{y})$.
- (2) L has nontrivial factors in $\mathbb{F}(\mathbf{y})$. Namely, there exist $r < n$, matrices K, M of dimensions $n \times r, r \times n$ (respectively) with entries in $\mathbb{F}(\mathbf{y})$ such that $L = KM$. The smallest r with this property is called the (commutative) rank of L , and is denoted $\text{rank}(L)$.
- (3) The linear space defined by L is singular, namely contains no non-singular matrix over \mathbb{F} . In other words, for every choice of constants $\beta_i \in \mathbb{F}$, the matrix $\sum_{i=1}^m \beta_i A_i$ is singular over \mathbb{F} .
- (4) $\text{Det}(\sum_{i=1}^m y_i A_i) \equiv 0$ as a polynomial over $\mathbb{F}[\mathbf{y}]$, where Det is the determinant polynomial.

Now consider the non-commutative situation. The algebra of polynomials over non-commutative variables $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$ is familiar enough, and denoted $\mathbb{F}\langle \mathbf{x} \rangle$. Its (skew) “field of fractions”,⁴ which we have yet to define, is denoted $\mathbb{F}\langle\langle \mathbf{x} \rangle\rangle$. As in the commutative case, its elements include all inverses of nonzero polynomials in $\mathbb{F}\langle \mathbf{x} \rangle$ (but as mentioned above, these elements may be far more complex than ratios of polynomials in $\mathbb{F}\langle \mathbf{x} \rangle$). One can formally define, as we will, the non-commutative SINGULAR problem, exactly as (1) in Fact 1.2 for the commutative case. But then most analogs of these equivalences seem to break down if we take them literally. For example, (3) clearly does not make sense (as it will not distinguish commuting from non-commuting elements), and (4) is ill-defined due to the plethora of non-commutative variants of the determinant polynomial (see, e.g., [36]). However, the beauty of the non-commutative setting is revealed in many more equivalences, which we collect in the following theorem (which in particular has interesting analogs of

⁴ Actually there are many, but only one “universal field of fractions”.

(2), (3), (4)). These equivalences illustrate some of the connections to optimization, (commutative) invariant theory, non-commutative algebra and quantum information theory. All of these equivalences will be discussed in this introduction, and elaborated on more formally later. We note that some of them are highly nontrivial theorems we shall later reference individually. To compare between the commutative and non-commutative worlds, and appreciate better the equivalences, it is instructive to have in mind the following example.

Example 1.3 The following 3×3 skew-symmetric matrix is singular over $\mathbb{F}(z, w)$ but not singular over $\mathbb{F}\langle z, w \rangle$.

$$\begin{bmatrix} 0 & z & w \\ -z & 0 & 1 \\ -w & -1 & 0 \end{bmatrix}$$

Theorem 1.4 (SINGULAR) *The following are equivalent for a matrix $L = \sum_{i=1}^m x_i A_i$ in non-commuting variables \mathbf{x} . Again some of these equivalences only hold for large enough fields \mathbb{F} and some only over reals or complex numbers. But for concreteness, we will fix the field to be \mathbb{C} .*

1. L is singular (namely not invertible) over $\mathbb{F}\langle \mathbf{x} \rangle$.
2. L has nontrivial polynomial factors in $\mathbb{F}\langle \mathbf{x} \rangle$. Namely, there exist $r < n$, matrices K, M of dimensions $n \times r, r \times n$ (respectively) with entries in $\mathbb{F}\langle \mathbf{x} \rangle$ ⁵ such that $L = KM$. The smallest r with this property is called the non-commutative rank (and sometimes inner rank) of L , and is denoted $\text{nc-rank}(L)$.
3. The linear space defined by L with matrix coefficients is singular. Namely, for every integer k and every tuple of $k \times k$ matrices (B_1, B_2, \dots, B_m) over \mathbb{F} , the matrix $\sum_{i=1}^m B_i \otimes A_i$ is singular over \mathbb{F} .
4. For every integer k , we have $\text{Det}(\sum_{i=1}^m X_i \otimes A_i) \equiv 0$, where the X_i are $k \times k$ matrices of commutative variables, Det is still the commutative determinant polynomial, acting here on $kn \times kn$ matrices.
5. L has a shrunk subspace. Namely, there are subspaces U, W of \mathbb{F}^n , with $\dim W < \dim U$, such that for all $i, A_i U \subseteq W$.
6. L has a Hall-blocker (in analogy with the Hall theorem for perfect matchings). Namely, there exist nonsingular matrices B, C over \mathbb{F} such that the symbolic matrix BLC has an all-zeros minor of size $i \times j$, with $i + j > n$.
7. The completely positive quantum operator (or map) associated with L is rank-decreasing. Namely, there exists a positive semi-definite matrix P such that $\text{rank}(\sum_{i=1}^m A_i P A_i^\dagger) < \text{rank}(P)$.⁶
8. The tuple of matrices (A_1, A_2, \dots, A_m) is in the null-cone of the left–right action of $(SL_n(\mathbb{F}))^2$ on $m \times n \times n$ matrices, namely they evaluate to zero on every non-constant homogeneous **invariant** polynomial under this action.⁷

⁵ Moreover, the polynomial entries of K, M in such a minimal decomposition can actually be taken to be polynomials of degree at most 1, namely affine combinations of the variables.

⁶ Here A^\dagger denotes the conjugate-transpose of a complex matrix A .

⁷ The left–right action and its invariant polynomials are defined as follows. Consider mn^2 commuting variables which are arranged in m matrices (Y_1, Y_2, \dots, Y_m) , and consider polynomials in these variables.

In the rest of this introduction, we will try to explain how and why the problem SINGULAR arises in these different contexts, how are they related. We will discuss the algorithm that solves SINGULAR (which already appears in [37]), its complexity analysis (which is the main result of this paper), and how they arise from these (and other!) sources. We will also discuss the extension of this algorithm to non-commutative rank computation, and the implications of these algorithms to the various settings. We also highlight the recurring analogs and connections between the commutative and non-commutative worlds. There are probably many ways to weave this meandering story due to the multiple connections between the topics; we chose one (some accounts of subsets of these connections appear, e.g., in [37,45,51]). Note that most descriptions will be informal or semi-formal; formal definitions for everything that is actually needed will be given in the technical sections, and for the rest, we provide references.

1.1 The Free Skew Field

Polynomials, both in commuting and non-commuting variables, can always be uniquely defined by their coefficients. And so, while their computational complexity is of course interesting, it is not part of their description. For rational functions, this persists in the commutative setting, as they are simply ratios of polynomials. It is perhaps amusing that in the non-commutative setting, the first definition of the *free skew field* of rational functions was computational, namely one whose objects are *described* by their computation via a sequence of arithmetic operations. Both this description and the subsequent discovery of a syntactic representation play an important role in our story.

Books on history and construction of the skew field include [16,74], and a thorough discussion from the computational perspective is given in [45]. Here, we will be relatively brief, highlighting the key roles played by matrices and symbolic matrices in these constructions of the free skew field.

Fix a field \mathbb{F} and a set of non-commuting variables $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$. As before, let $\mathbb{F}\langle\mathbf{x}\rangle$ denote the (free) algebra of non-commutative polynomials. A major objective is to construct a (minimal) “skew field of fractions”, which contains this algebra and is a division ring, namely every nonzero element is invertible. As it happens, there can be many non-isomorphic ways of doing that, but only one that is *universal*,⁸ the *free skew field* $\mathbb{F}\langle\langle\mathbf{x}\rangle\rangle$, first defined by Amitsur [4] as follows (see a more detailed and precise description in [54]).

A *rational expression* $r(\mathbf{x})$ is simply a formula whose inputs are variables and constants from \mathbb{F} , over the standard arithmetic operations of addition, multiplication and division, e.g., $x_1^{-1} - (x_1 + x_2)^{-1}$ and $(x_1 + (x_1 x_2^{-1}) x_1)^{-1}$ (note that inversions

Footnote 7 continued

Every pair B, C of determinant 1 matrices over \mathbb{F} defines a linear map of these variables by sending this tuple to $(BY_1C, BY_2C, \dots, BY_mC)$. A polynomial in these variables is *invariant* if it remains unchanged by this action for every such pair B, C .

⁸ This is a technical term which we will not define here.

can be arbitrarily nested⁹). Essentially, the elements of the field will consist of these expressions, modulo an equivalence relation that we now define. Note that substituting $d \times d$ matrices for the variables of an expression r it becomes a (partial) function.¹⁰ Two expressions r, s are *equivalent* if they agree (output the same matrix) for *all* inputs for which they are both defined (where we go over all tuples of matrices from $M_d(\mathbb{F})$ for all finite d).

Theorem 1.5 [4] *The rational expressions under the above equivalence relation comprise the (universal) free skew field $\mathbb{F}\langle\mathbf{x}\rangle$.*

One important complication (or perhaps a fountain of interesting problems) in the non-commutative setting is that a nonzero rational expression (or even a polynomial expression, namely one with no inverses) can be an identically zero function over the entire $M_d(\mathbb{F})$ for some finite d . The simple example $x_1x_2 - x_2x_1$ is a *polynomial identity* for $d = 1$, and more generally the *standard polynomial* $\sum_{\pi \in S_{2d}} (-1)^{\text{sgn}(\pi)} x_{\pi(1)}x_{\pi(2)} \cdots x_{\pi(2d)}$ is a polynomial identity for $M_d(\mathbb{F})$. By the celebrated Amitsur–Levitsky theorem, this example is tight in terms of the relation between the degree of the polynomial identity and dimension of matrices.

Theorem 1.6 [3] *For a fixed $d > \lfloor k/2 \rfloor$, a degree k polynomial cannot vanish on all inputs in $M_d(\mathbb{F})$.*

This theorem immediately implies a polynomial identity test in probabilistic polynomial time (just like in the commutative case) mentioned in the next subsection—simply plug random matrices of appropriate size for the variables. However, for rational expressions, which is our main interest here, such upper bounds on the sufficient dimension to exclude *rational identities* are much harder to prove, and what is known is much weaker. The best bounds (and only for some fields) are exponential, which seems to only provide probabilistic exponential time rational identity tests. The last subsection of this introduction shows how these exponential dimension bounds arise and obtained from invariant theory, and how we use them to nonetheless derive a *deterministic, polynomial* time identity test promised in our main theorem.

A second construction of the free skew field which is somehow more concrete was developed by Cohn. In the notation, we have already established of symbolic matrices Cohn proved:

Theorem 1.7 [13] *Every element of the free skew field $\mathbb{F}\langle\mathbf{x}\rangle$ is an entry of the inverse of some (invertible) symbolic matrix whose entries are polynomials in $\mathbb{F}\langle\mathbf{x}\rangle$.*

For this definition not to be self-referential (and other reasons), Cohn proved some of the important characterizations of invertible matrices, including items (2) and (5) in Theorem 1.4. It is clear that (each entry of) the inverse of such a matrix can be given by a rational expression, and the question of whether a matrix is invertible again arises naturally, and it turns out to capture rational identity testing.

⁹ In general *inversion height*, the minimum amount of nesting needed, can be arbitrarily high. This is an important theorem of Reutenauer [73]. However, in the example above, the nested inversion can be eliminated, and in fact, the two expressions are equal (a simple fact which the reader might try to prove)! This equality is called Hua’s identity [48], underlying the fundamental theorem of projective geometry.

¹⁰ when the expression attempts to invert a singular matrix it is undefined on that input. The domain of an expression is simply all input tuples on which it is defined.

1.2 Word Problems and Identity Tests

Word problems and their complexity are central throughout mathematics, arising whenever mathematical objects in a certain class have several representations. In such cases, a basic problem is whether two such representations describe the same object. Often, indeed, some problems of this form have served as early examples of decidable and undecidable problems.¹¹

For us, the objects in question are polynomials and rational functions in commuting and non-commuting variables. Their standard representations will be arithmetic formulas and circuits, which take the variables (and constants in \mathbb{F}) as inputs, and use plus, times and *division* gates. An excellent exposition of arithmetic complexity models and the state of the art in the subject is [77] (which discusses at length the polynomial identity testing problem¹²). We stress that we will specifically allow division gates in our computational models, both because we want to consider rational functions, and because unlike for computing polynomials, in the non-commutative case there is no known analog of Strassen's theorem [79] that divisions can be efficiently eliminated. This issue and partial results, many relevant to this paper, are discussed in [45]. The word problem in our arithmetic context is equivalent to an identity test, checking whether a given representation of a polynomial or rational function describes the trivial element, the zero polynomial.

As is well known, the main reason why PIT is so important in the commutative setting is that it captures the polynomial and rational function identity test for formulas, as proved by Valiant [80]. Combining it with the equivalences of Fact 1.2, we have

Theorem 1.8 [80] *There is an efficient algorithm which converts every arithmetic formula $\phi(\mathbf{y})$ in commuting variables \mathbf{y} of size s to a symbolic matrix L_ϕ of size $\text{poly}(s)$, such that ϕ is identically zero if and only if $L_\phi \in \text{PIT}$ (i.e., L_ϕ is singular).*

Theorem 1.7 of the previous subsection, showing that in the non-commutative setting as well SINGULAR captures polynomial and rational function identity test, is the analog Valiant's completeness theorem. It was proved even earlier by Cohn (see also Malcolmson's version [64]), using similar linear algebraic ideas. Here is a more precise statement which gives the analogy.

Theorem 1.9 [13] *There is an efficient algorithm which converts every arithmetic formula $\phi(\mathbf{x})$ in non-commuting variables \mathbf{x} of size s to a symbolic matrix L_ϕ of size $\text{poly}(s)$, such that the rational expression computed by ϕ is identically zero if and only if $L_\phi \in \text{SINGULAR}$.*

As mentioned, the structure of the free skew field is so complex that unlike the commutative case, even decidability of SINGULAR (and rational identity testing) is

¹¹ For example, deciding whether two knots diagrams describe the same knot was proved decidable by Haken [42], and deciding whether two presentations with generators and relations describe the same group were proved undecidable by Rabin [69].

¹² While this paper focuses on identity testing, we note that our interest is partly (and indirectly) motivated by the more basic problem of proving lower bounds for non-commutative circuits. We refer the reader to the papers [46,47,59,66] and their references for existing lower bounds on weaker models, some completeness results, and possible approaches to proving stronger lower bounds.

far from obvious. The first to prove decidability was Cohn in [14,15]. The first explicit bound on time was given by Cohn and Reutenauer in [17], reducing it to a system of commutative polynomial equations using characterization (5) of SINGULAR, proved earlier by Cohn, which puts it in $\mathcal{PSPAC}\mathcal{E}$. The best upper bound before this work was singly exponential time, obtained by [51], and of course yields the same bound for rational identity testing.

From our Theorem 1.1, we conclude a *deterministic, polynomial* identity test for non-commuting rational expressions.

Corollary 1.10 *The non-commutative rational identity testing problem is in \mathcal{P} . Namely, there is an algorithm which for any non-commutative formula over \mathbb{Q} of size s and bit complexity b determines in $\text{poly}(s, b)$ steps if it is identically zero.*

It is important to stress that unlike the commutative case, where symbolic matrix inversion can be simulated by quasi-polynomial sized formulas, in the non-commutative case symbolic matrix inversion is exponentially more powerful than formulas (and so Theorem 1.1 is far more powerful than its corollary above). We state these two results formally below for contrast.¹³ Note that when saying “computing the inverse”, we mean computing any entry in the inverse matrix.

Theorem 1.11 [8,49] *The inverse of a commutative $n \times n$ symbolic matrix in $\mathbb{F}(\mathbf{y})$ can be computed by formulas of size $n^{O(\log n)}$.*

Theorem 1.12 [45] *The inverse of a non-commutative $n \times n$ symbolic matrix in $\mathbb{F}\langle\mathbf{x}\rangle$ requires formulas of size $2^{\Omega(n)}$.*

Our algorithm of Theorem 1.1 and Corollary 1.10 is a “white-box” identity test (namely one which uses the description of the given matrix or formula), in contrast to a “black-box” algorithm which only has input–output access to the function computed by them. The best-known “black-box” identity test for these formulas, even if randomization is allowed(!), requires exponential time. It is a very interesting question to find a faster black-box algorithm.

When division is *not* allowed, efficient deterministic identity tests of non-commutative *polynomials* were known in several models. The strongest is for arithmetic branching programs (ABPs). As this model is easily simulable by matrix inversion (see Theorem 6.5 in [45], our algorithm provides an alternative (and very different) proof of the following theorem of Raz and Shpilka [71].

Theorem 1.13 [71] *There is a deterministic polynomial time white-box identity testing algorithm for non-commutative ABPs.*

For certain classes computing only polynomials, even efficient black-box algorithms are known, and we mention two below: the first deterministic, for non-commutative ABPs, and the second probabilistic (but for circuits).

Theorem 1.14 [31] *There is a deterministic quasi-polynomial time black-box identity testing algorithm for non-commutative ABPs.*

¹³ Replacing formulas by circuits there is no contrast—in both the commutative and non-commutative setting matrix inverse has a polynomial size circuit (with division of course) [45].

Theorem 1.15 [3,10] *There is a probabilistic polynomial time black-box identity testing algorithm for non-commutative circuits.*

1.3 Commutative and Non-commutative Rank of Symbolic Matrices

There are many sources, motivations and results regarding invertibility, and more generally rank, of *commutative* symbolic matrices, which are much older than the complexity theory interest in it. These mathematical papers, like the ones in the computational complexity literature, regard it as a difficult problem and attempt to understand a variety of special cases (of a different nature than restrictions of the computational power of the model). Some of the many references to this body of work can be found in the papers [33,35]. Often, the *non-commutative* rank (explicitly or implicitly) is used to give upper bound on the commutative rank, and their relationship becomes of interest. We focus on this connection here, and explain how our main result implies a deterministic *approximation algorithm* to the commutative rank.

In this section, we assume that the underlying field \mathbb{F} is infinite. We will use the same notation L for both a symbolic matrix, as well as the subspace of matrices spanned by it (when fixing the variables to constants in the field). We repeat the definitions and elaborate on what is known regarding the commutative and non-commutative ranks, from now on denoted by $\text{rank}(L)$ and $\text{nc-rank}(L)$. Note that the characterizations allow thinking about both without reference to the respective fields of fractions (over which they are most naturally defined), but only to polynomials.

Fact 1.16 *The following are equivalent for a matrix of linear forms over commutative variables, $L(\mathbf{y})$.*

- $\text{rank}(L(\mathbf{y})) = r$ over $\mathbb{F}(\mathbf{y})$
- r is the maximal rank of any matrix in the subspace L over \mathbb{F} .

While the characterization above is simple, the one below is very substantial, mostly developed by Cohn for his construction of the free skew field. The first characterization is due to him [16] and the second is due to Fortin and Reutenauer [33] who heavily use Cohn's techniques.

Theorem 1.17 *The following are equivalent for a matrix of linear forms over non-commutative variables, $L(\mathbf{x})$.*

- $\text{nc-rank}(L(\mathbf{x})) = r$ over $\mathbb{F}\langle\langle\mathbf{x}\rangle\rangle$.
- The inner rank of L over $\mathbb{F}\langle\langle\mathbf{x}\rangle\rangle$ is r . Namely, r is the minimal number such there exists matrices K, M of dimensions $n \times r, r \times n$ (respectively) with polynomial entries in $\mathbb{F}\langle\langle\mathbf{x}\rangle\rangle$ such that $L = KM$. Moreover, in this decomposition, the factors K, M can be assumed to have affine linear entries.
- The space L is r -decomposable. Namely, r is the minimal number such that there exists invertible matrices B, C over \mathbb{F} such that BLC has a minor of zeros of size $i \times j$ with $i + j = 2n - r$.

We can extend our main Theorem 1.1 from testing singularity to efficiently computing the non-commutative rank over \mathbb{Q} , this is done in Theorem A.4.

Theorem 1.18 *There is a deterministic algorithm which, given m $n \times n$ integer matrices A_1, \dots, A_m with entries of bit size b , computes $\text{nc-rank}(L)$ in time $\text{poly}(n, m, b)$ (where $L = \sum_{i=1}^m x_i A_i$).*

Remark 1.19 We note here that in some formulations of the problem, L is represented in *affine form*, namely where an additional constant matrix A_0 is added to L . (This may be viewed as a non-commutative analog of the rank completion problem.) However, the algorithm above works in this case as well, since $\text{nc-rank}(L)$ remains unchanged if an additional variable x_0 was added as well. So, we will stick with the linear formulation.

It is not hard to see from the definitions that for every L we have $\text{rank}(L) \leq \text{nc-rank}(L)$. We have already seen that they can be different in Example 1.3. Taking many copies of that 3×3 matrix we see that there can be a factor $3/2$ gap between the two: for any r there are matrices L with $\text{rank}(L) = 2r$ and $\text{nc-rank}(L) = 3r$. However, Fortin and Reutenauer [33] proved that this gap is never more than a factor of 2.

Theorem 1.20 [33] *For every L , we have $\text{nc-rank}(L) \leq 2\text{rank}(L)$.*

An immediate corollary of our main result is thus a factor 2 approximation of commutative rank.

Corollary 1.21 *There is a polynomial time algorithm which for every symbolic matrix in commuting variables over \mathbb{Q} approximates $\text{rank}(L)$ to within a factor of 2.*

We find the question of obtaining a better approximation ratio efficiently a very interesting problem, a different relaxation of the commutative PIT problem that as far as we are aware was not studied till now.

1.4 Compression Spaces, Optimization, and Gurvits’ Algorithm G

An important class of spaces of matrices L , which is studied in many of the papers mentioned above, is the class of *compression spaces* (this notation was apparently introduced by Eisenbud and Harris [28]). They are defined as those spaces L for which L is $\text{rank}(L)$ -decomposable. A simpler definition follows the characterization of [33] above, namely a space L is a compression space iff $\text{rank}(L) = \text{nc-rank}(L)$. The importance of compression spaces and their many origins will be surveyed below.

A deterministic polynomial time *commutative* PIT for compression spaces (over \mathbb{Q}) was discovered by Gurvits [37], a paper which serves as the starting point for this work.¹⁴ Indeed, Gurvits’ algorithm, which we will denote by Algorithm G , is *the same* one we use here for our main Theorem 1.1; our main contribution here is the *analysis* of its performance for any input L , not necessarily a compression space. We will return to the algorithm and its analysis, but first let us discuss its extension from testing singularity to rank computation. Before that, we state Gurvits’ theorem, which actually proves more about Algorithm G .

¹⁴ It is interesting to note that most recent progress on deterministic PIT algorithms (e.g., [25,31,55,57,75] among many others) are for polynomials computed by a variety of restricted classes of arithmetic circuits. Algorithm G seems to differ from all of them in solving PIT for a very different class of polynomials, which we do not know how to classify in arithmetic complexity terms.

Theorem 1.22 [37] *There is a deterministic polynomial time algorithm, Algorithm G , which for every n and every $n \times n$ matrix L given by a set of integer matrices (A_1, A_2, \dots, A_m) , outputs “singular” or “invertible”, and its output is guaranteed to be correct¹⁵ when either $\text{rank}(L) = n$ or $\text{nc-rank}(L) < n$.*

In particular, the algorithm always gives the correct answer for compression spaces. Our result on efficient computation of the non-commutative rank implies that for compression spaces we can determine the *commutative rank*.

Many natural spaces are compression spaces, and these arise from a variety of motivations and sources in linear algebra, geometry and optimization. One such source is the attempt to characterize singular spaces (namely when $\text{rank}(L) < n$) by relating the rank to the dimension of the space matrix L defines, denoted $\text{dim}(L)$. Perhaps the earliest result of this type is due to Dieudonné [23] who proved that if L is singular and $\text{dim}(L) = n^2 - n$, then for some nonsingular matrices B, C over \mathbb{F} the matrix BLC has an all-zero row or all-zero column. In other words, it has a Hall-blocker as in condition (6) in Theorem 1.4, and so is in particular decomposable. Many other similar results appear in, e.g., [5–7,28,35], which prove decomposability under more general conditions on the dimension, and study a variety of other properties of singular and rank-bounded spaces.

Yet another source, elaborated on in [37], is geometric duality theorems, such as the ones for matroid intersection and matroid parity problems. These sometimes give rise to compression spaces, with one prototypical example being spaces where the generating matrices A_i all have rank-1; this follows from the Edmonds–Rado [27, 70] duality theorem for matroid intersection.¹⁶ Another, simpler example are spaces generated by upper-triangular matrices. Other examples of compression spaces are given in [37]. Following his paper, we note that our rank algorithm above can solve such optimization problems *in the dark*, namely, when the given subspace has such *structured* spanning generators, but the generators A_i actually given to the algorithm may be *arbitrary*. This is quite surprising and, in general, cannot work by uncovering the original structured generators from the given ones: an efficient algorithm is not known for the problem of deciding if a given space of matrices is spanned by rank-1 matrices and we believe it is hard. It would be interesting to explore which other optimization problems can be encoded as the rank of a compression space, and whether the fact that it can be computed “in the dark” has any applications.

In recent related work, Ivanyos et al. [50] give a different “in the dark” algorithm for computing the rank of certain compression spaces (including the ones spanned by rank-1 matrices) using so-called second Wong sequences, a linear algebraic analog of “augmenting paths for matchings”. Its main advantage is that it works not only over subfields of \mathbb{C} as our algorithm, but also over large enough finite fields. This resolves an open problem in [37].

¹⁵ For both the commutative and non-commutative definitions.

¹⁶ It is an interesting question whether a compression space naturally arises from matroid parity duality of Lovasz [62,63].

1.5 Permanents, Quantum Operators, Origins and Nature of Algorithm G

We will describe the algorithm formally in Sect. 2.2. Here we give an informal description, which makes it easy to explain its origins and nature. We find these aspects particularly interesting. One (which has very few parallels) is that while the problem SINGULAR solved by Algorithm *G* in Theorem 1.1 is purely *algebraic*, the algorithm itself is purely *analytic*; it generates from the input a sequence of complex-valued matrices and attempts to discover if it is convergent. Another is that the algorithm arises as a quantum analog of another algorithm with very different motivation that we now discuss.

To give more insight to the working of algorithm *G*, let us describe another algorithm (for a different problem) which inspired it, which we call Algorithm *S*. This *matrix scaling* algorithm was developed by Sinkhorn [78] for applications in numerical analysis, and has since found many other applications (see survey and references in [60], who used it as a basis for their deterministic algorithm for approximating the permanent of nonnegative matrices). Two different analyses of Sinkhorn’s algorithm *S*, one of [60] and the other in the unpublished [41] inspire the analysis of Algorithm *G* in [37].

We describe the [60] analysis for Algorithm *S*. We need a few definitions. For a nonnegative matrix *A*, let *R*(*A*) denote the diagonal matrix whose (*i*, *i*)-entry is the inverse of the *L*₁ norm of row *i* (which here is simply the sum of its entries as *A* is nonnegative). Similarly, *C*(*A*) is defined for the columns.¹⁷

Algorithm *S* gets as input a nonnegative integer matrix *A*. For a fixed polynomial (in the input size) number of iterations, it repeats the following two steps

- Normalize rows: $A \leftarrow R(A) \cdot A$
- Normalize columns: $A \leftarrow A \cdot C(A)$

What does this algorithm do? It is clear that in alternate steps either *R*(*A*) = *I* or *C*(*A*) = *I*, where *I* is the identity matrix. Thus *A* itself alternates being row-stochastic and column-stochastic. The question is whether both converge to *I* together, namely, if this process converts *A* to a *doubly stochastic* matrix. In [60] it is proved that this happens if and only if *Per*(*A*) > 0, where *Per* is the permanent polynomial. Moreover, convergence is easy to detect after a few iterations! If we define *ds*(*A*) = ||*R*(*A*) – *I*||² + ||*C*(*A*) – *I*||² as a notion of distance between *A* and the doubly stochastic matrices, then the convergence test is simply whether *ds*(*A*) < 1/*n*. If it is that small at the end of the algorithm, then *Per*(*A*) > 0; otherwise, *Per*(*A*) = 0.

The analysis of convergence of Algorithm *S* in [60] is extremely simple, using the permanent itself as a progress measure (or potential function). It has the usual three parts which makes a potential function useful:

1. The input size provides an exponential lower bound on the starting value of *Per*(*A*),
2. The arithmetic-geometric mean inequality guarantees that it grows by a factor of 1 + 1/*n* at every iteration, and
3. The permanent of any stochastic matrix is upper bounded by 1.

¹⁷ A “non-triviality” assumption is that no row or column in *A* is all zero.

We shall return to this analysis of Algorithm S soon.

As it happens, Algorithm G is a *quantum* analog of Algorithm S ! In quantum analogs of classical situations two things typically happen, diagonal matrices (which commute) become general matrices (which do not), and the L_1 norm is replaced by L_2 . This happens here as well, and we do so almost syntactically, referring the reader to [37] for their quantum information theoretic intuition and meaning of all notions we mention.

The input to algorithm G is a symbolic matrix $L = \sum_i x_i A_i$, given by the $n \times n$ integer matrices (A_1, A_2, \dots, A_m) . Briefly, L is viewed as a *completely positive (quantum) operator*, or map, on psd matrices, mapping such a (complex valued) matrix P to $L(P) = \sum_i A_i P A_i^\dagger$ (P is typically a “density matrix” describing a quantum state, namely a psd matrix with unit trace, and the operator L will typically preserve trace or at least not increase it). The dual operator L^* acts (as you would expect) by $L^*(P) = \sum_i A_i^\dagger P A_i$. The analog “normalizing factors” for L , named $R(L)$ and $C(L)$ are defined¹⁸ by $R(L) = (\sum_i A_i A_i^\dagger)^{-\frac{1}{2}}$, and $C(L) = (\sum_i A_i^\dagger A_i)^{-\frac{1}{2}}$. Note that $R(L) = L(I)^{-\frac{1}{2}}$ and $C(L) = L^*(I)^{-\frac{1}{2}}$.

On input (A_1, A_2, \dots, A_m) Algorithm G repeats, for a fixed polynomial (in the input size) number of iterations, the following analogous two steps

- Normalize rows: $L \leftarrow R(L) \cdot L$
- Normalize columns: $L \leftarrow L \cdot C(L)$

So again, row and column operations are performed alternately, simultaneously on all matrices A_i . It is clear, as above, that after each step either $R(L) = I$ or $C(L) = I$. It is natural to define the case when both occur as “doubly stochastic”, and wonder under what conditions does this sequence converge, namely both $R(L)$ and $C(L)$ simultaneously approach I , and alternatively the limiting L simultaneously fixes the identity matrix I . A natural guess would be that it has to do with a “quantum permanent”. Indeed, Gurvits [37] defines such a polynomial (in the entries of the A_i) $\text{QPer}(L)$, and proves several properties and characterizations (for example, it is always nonnegative like the permanent, and moreover specializes to the permanent when the operator L is actually a “classical” operator described by a single nonnegative matrix A).

One can similarly define in an analogous way to the classical setting a “distance from double stochastic” by $\text{ds}(L) = \|R(L) - I\|^2 + \|C(L) - I\|^2$, and test (after polynomially many iterations) if $\text{ds} < 1/n$. This is precisely what Algorithm G does. It is not hard to see that if L (with non-commuting variables) is singular, then there is no convergence, the test above fails (and indeed $\text{QPer}(L) = 0$). However, in contrast to Algorithm S , the analysis in [37] falls short of proving that, otherwise, we have convergence (and $\text{QPer}(L) > 0$). It does so *only* under the strong (commutative) assumption $\text{Det}(L) \neq 0$, namely when the symbolic matrix, viewed with *commuting* variables, is nonsingular. This result was stated above as Theorem 1.22. The three-step complexity analysis proceeds exactly as in [60] for the classical Algorithm S described above, using the quantum permanent as a progress measure. However, if $\text{Det}(L) = 0$

¹⁸ Again using a “non-triviality” assumption these matrices are invertible.

then $\text{QPer}(L) = 0$, and lower bounding the initial quantum permanent from below in part (1) of that analysis fails.

To prove our main theorem, showing that convergence happens *if and only if* L is non-singular over the non-commutative skew field, we use another ingredient from Gurvits' paper. He proposes another progress measure, called *capacity* and denoted $\text{cap}(L)$. Capacity (defined in the next section) is a quantum analog of another classical progress measure, based on KL-divergence, which was used in [41] for the analysis of Algorithm S . The advantage of capacity, as pointed out in [37], is that it does not necessarily vanish when $\text{Det}(L) = 0$. Indeed, it is positive whenever L is non-singular in the non-commutative sense! However, Gurvits could only bound it sufficiently well from below as a function of the input size if $\text{Det}(L) \neq 0$. Thus, for a general input L , positive capacity only guarantees that Algorithm G converges in a finite number of steps, without providing an upper bound on that number.

Our main contribution, which leads to a polynomial upper bound, is proving an explicit capacity lower bound for *every* L that is nonsingular over the skew field. The source of the proof turns out to be commutative algebra and invariant theory, as we discuss next.

1.6 Invariant Theory, Left–Right Action, Capacity and the Analysis of Algorithm G

Invariant theory is a vast field; we will exposit here only the minimal background that is necessary for this paper. The books [18,20,58] provide expositions on the general theory. More focused discussions toward our applications appear in the appendix of [45] and Section 1.2 of [30].

Invariant theory deals with understanding the symmetries of mathematical objects, namely transformations of the underlying space which leave an object unchanged or *invariant*. Such a set of transformations always form a group (and every group arises this way). One major question in this field is, given a group acting on a space, characterize the objects left invariant under all elements of the group. Here we will only discuss very specific space and actions: polynomials (with commuting variables!) that are left invariant under certain linear transformations of the variables. The invariant polynomials under such action clearly form a ring (called the *invariant ring*). Important for us is the *null-cone* of the action, which is simply all assignments to variables which vanish on *all* non-constant homogeneous invariant polynomials. Namely, the null cone is the affine algebraic variety defined by the ideal generated by homogeneous invariant polynomials of positive degree.

Here are two motivating examples surely familiar to the reader. The first example is given by polynomials in n variables y_1, \dots, y_n invariant under the full group of permutations S_n ; here the invariant polynomials are (naturally) all *symmetric* polynomials. While this is an infinite set, containing polynomials of arbitrarily high degrees, note that this set of polynomials is generated (as an algebra) by the $n + 1$ *elementary symmetric polynomials*, with the largest degree being n . This finite generation is no coincidence! A general, important theorem of Hilbert [44] assures us that for “such” linear actions there is always a finite generating set (and hence a finite upper bound on their maximum degree). Obtaining upper bounds on the degree of generating sets,

finding descriptions of minimal generating sets for natural actions are the classical goals of this area, and a more modern one¹⁹ is obtaining succinct descriptions and efficient computation of these invariants (e.g., see [30,65]). The case of the action of the symmetric group is an excellent example of a perfect understanding of the invariant polynomial ring. Note that for this action the null-cone is simply the all-zero vector.

The second example, much closer to our interest here, is the “left–right” action of the group $(SL_n(\mathbb{F}))^2$ acting on n^2 variables y_{ij} arranged in an $n \times n$ matrix Y . Specifically, two $n \times n$ matrices (of determinant 1) (B, C) take Y to BYC , changing the basis on the left and right.²⁰ It is not hard to see that all invariant polynomials under this action are generated by one—the determinant $\text{Det}(Y)$ (which again has degree n). Consequently, the null-cone of this action is all singular matrices over the field.

What we consider here is the left–right action on $m \times n$ matrices, where a pair (B, C) as above takes (Y_1, Y_2, \dots, Y_m) to $(BY_1C, BY_2C, \dots, BY_mC)$. The study of the invariant ring of polynomials (in the mn^2 variables sitting in the entries of these matrices) for this action was achieved²¹ by [1,22,24,76], and will look very familiar from condition (4) in Theorem 1.4, as well as from Amitsur’s definition of the free skew field.²²

Theorem 1.23 [1,22,24,76] *Over algebraically closed fields, the invariant ring of polynomials of the left–right action above is generated by all polynomials of the form $\text{Det}(\sum_i D_i \otimes Y_i)$, for all d and all $d \times d$ matrices D_i .*

It is worthwhile to stress the connection forged between the commutative and non-commutative worlds by this theorem when combined with Amitsur’s and Cohn’s constructions of the skew field. A set of matrices (A_1, A_2, \dots, A_m) is in the null-cone of the left–right action if and only if the symbolic matrix $L = \sum_i x_i A_i$ is not invertible in the free skew field!²³ In other words, the non-commutative SINGULAR problem (and thus rational identity testing, and the word problem in the skew field) arises completely naturally in commutative algebra. Of course, invertibility itself is invariant under the left–right action (indeed, even by any invertible matrices B, C , not necessarily of determinant 1), so one expects a connection in at least one direction. At any rate, one may hope now that commutative algebraic geometric tools will aid in solving these non-commutative problems, and they do!

The generating set of the invariant ring given in the theorem is still an infinite set, and as mentioned above, Hilbert proved that some finite subset of it is already generating.

Theorem 1.24 [44] *The invariant ring of polynomials of the action of a reductive group on a finite-dimensional vector space is finitely generated.*

¹⁹ Arising in particular in the GCT program of Mulmuley and Sohoni.

²⁰ Note that for it to be a group action in the strict sense, one should study the action which takes Y to BYC^{-1} or BYC^T but for simplicity, we will avoid this distinction.

²¹ We note that this is part of the larger project of understanding *quiver representations*, started by the works of Procesi, Razmysolov, and Formanek [32,68,72].

²² Note though that the roles of which matrices in the tensor product are variable, and which are constant, has switched!

²³ This also follows from [11,56].

Hilbert proved his theorem for the actions of $SL_n(\mathbb{C})$ but his proof can be easily seen to be extendable to arbitrary reductive groups. Since $(SL_n(\mathbb{C}))^2$ is reductive, Hilbert’s theorem implies that the invariant ring of polynomials of the left–right action is finitely generated. A natural definition to make now is $\beta(n)$, the smallest integer such that taking only matrices with $d \leq \beta(n)$ generates the invariant ring.²⁴ Hilbert did not give an explicit bound. The first to do so was Popov [67], and this bound was significantly improved by Derksen [19]. A related quantity $\sigma(n)$ is the smallest integer such that taking only matrices with $d \leq \sigma(n)$ suffice for testing membership in the null-cone. Derksen [19] proved that $\sigma(n)$ and $\beta(n)$ are polynomially related (over algebraically closed fields of characteristic zero) and that (for algebraically closed fields of characteristic zero) $\sigma(n) \leq n^2 4^{n^2}$. All bounds above hold for general reductive groups. For the specific left–right action further progress was made! First, Derksen’s bound was further improved to a smaller exponential by [51], and extended to finite fields.

Theorem 1.25 [51] $\sigma(n) \leq (n + 1)!$

This exponential upper bound was the state of the art when we wrote our paper. Subsequent to it, this bound was greatly improved to linear!

Theorem 1.26 [21,52]

$$\sigma(n) \leq n - 1$$

This much better bound gave rise to further important developments discussed below under “Subsequent Work”, that highlight even more how invariant theory informs algorithmic efficiency in certain problems. We conclude this section describing the idea behind our proof, which uncovers this connection.

We plan to use the exact same three-step complexity analysis of Algorithm G as the one for Algorithm S described in Sect. 1.5, and which Gurvits [37] used to analyze Algorithm G for a special subset of inputs. As mentioned, the only missing ingredient is step (1), namely an explicit lower bound on the *capacity* of an input operator L . Just like in the classical case of Algorithm S , an exponential lower bound (in terms of the size of L) is required to establish a polynomial convergence. This is a proper time to finally define this important parameter of positive operators.²⁵

Given $L = (A_1, A_2, \dots, A_m)$, recall that $L(P) = \sum_i A_i P A_i^\dagger$ and define $\text{cap}(L)$ by

$$\text{cap}(L) = \inf \text{Det}(L(P))$$

where the infimum is taken over all psd matrices P with $\text{Det}(P) \geq 1$.

First, observe that capacity is the solution to a non-convex optimization problem (over a convex domain²⁶). We will see in the next section that we can not only

²⁴ This bound may a-priori depend on m , the number of matrices, but we already noted that $m \leq n^2$.

²⁵ We note that this notion of capacity seems to have nothing to do with the usual capacity of a quantum channel.

²⁶ Recall that log determinant is a concave function over the domain of positive definite matrices.

bound it, but actually compute it efficiently! Next, observe that capacity is an invariant (though not a polynomial invariant) of the left–right action. What we know is that $\text{cap}(L) \neq 0$. From this we infer that some polynomial invariant of the form above, $\text{Det}(\sum_i A_i \otimes D_i)$ is nonzero, where the D_i are $d \times d$ matrices for some d .

This expression naturally suggests considering several new positive operators, and relating their capacities. The first, in dimension d , is the operator, $M = (D_1, D_2, \dots, D_m)$. Two others, in nd dimensions are $L_M = (A_1 \otimes D_1, A_2 \otimes D_2, \dots, A_m \otimes D_m)$, and $L \otimes M = (A_1 \otimes D_1, \dots, A_i \otimes D_j, \dots, A_m \otimes D_m)$. The non-vanishing of the determinant above proves that both $\text{cap}(L_M) > 0$ and $\text{cap}(L \otimes M) > 0$. And it is easy to obtain exponential lower bounds on both quantities, which however depend not only on the description size of L (which is given), but also on that of M which could be much larger.

To facilitate our original capacity lower bound, which only used $d \leq \sigma(n) \leq \text{exp}(\text{poly}(n))$, we proved that the normalized²⁷ capacity is multiplicative, a fact of independent interest (see details in Sect. 3.3):

$$\text{cap}(L \otimes M)^{1/nd} = \text{cap}(L)^{1/n} \cdot \text{cap}(M)^{1/d}$$

Using only the easier direction, namely that the left hand side is at most the right hand side, it suggests that our lower bound on $\text{cap}(L)$ will follow from a lower bound on $\text{cap}(L \otimes M)$ and an upper bound on $\text{cap}(M)$. Both of these in turn follow from an upper bound on the entry sizes in M , namely in the D_i , which easily follow from the given exponential upper bound on d . A key thing to note is that even when d is exponential in n , and so these size bounds are doubly exponential, it is the d 'th root of this bound that gives the required lower bound $\text{cap}(L)$.

For the new lower bound, presented in Sect. 2.3, we bound $\text{cap}(L)$ using directly a bound on $\text{cap}(L \otimes M)$. Using much more careful size bounds on the entries of M , which follow from Alon's Nullstellensatz theorem, and tightening the reduction between the two capacities, we are able to derive the desired capacity lower bound using any finite bound on d !

This new proof eliminates the need for good degree upper bounds for the purpose of analyzing this algorithm (if one does not care about efficiency beyond having a polynomial bound). However, good degree upper bounds (that are anyway an important goal in invariant theory) were essential for the original proof, and for general actions were useful for other algorithmic problems which arose in other contexts in computational complexity, for example in geometric complexity theory (GCT). Further, we believe they will be important for studying other problems in algebraic complexity and optimization. Indeed, in a work in progress, we need degree bounds for invariant rings whose generating sets are not understood as well as for the left–right action, but Derksen's general exponential degree upper bound still holds and is useful.

The capacity lower bound above implies, as mentioned, a polynomial upper bound on the *number of iterations* of Algorithm G , assuming we can perform exact arithmetic! To actually bound the running time the algorithm has to be refined, truncating numbers so that they can be represented by polynomially many bits, and so a careful analysis of

²⁷ Taking the dimension's root of capacity.

the bit complexity is required to actually prove its running time. This issue naturally leads to the next section, in which we discuss how the algorithm can be extended to actually *compute* capacity of the input operator, as well as provide bounds on its continuity. These results in turn are essential in a follow-up work [34] we have done on the Brascamp–Lieb inequalities, where capacity arises naturally via a simple reduction to operator scaling.

1.7 Computation and Continuity of Capacity

We also show that Algorithm G can be modified to compute the capacity of a completely positive operator L . Recall that the capacity of an operator L is defined to be $\text{cap}(L) = \inf \text{Det}(L(P))$, where the infimum is taken over all psd matrices P with $\text{Det}(P) = 1$. The following theorem is proven as Theorem 5.3 in Sect. 5.

Theorem 1.27 *There is a $\text{poly}(n, b, 1/\epsilon)$ time algorithm for computing a $(1 + \epsilon)$ -multiplicative approximation to $\text{cap}(L)$. Here n is the dimension on which L acts and b denotes the bit sizes involved in description of L .*

This is quite fascinating since we don’t know of a convex formulation for $\text{cap}(L)$. The idea for the computation of capacity is quite simple. We know that the capacity of a doubly stochastic operator is exactly 1. We also know how operator scaling changes the capacity. If $L_{B,C}$ is a scaling of L , i.e.,

$$L_{B,C}(X) = B \cdot L(C \cdot X \cdot C^\dagger) \cdot B^\dagger$$

then

$$\text{cap}(L) = |\text{Det}(B)|^{-2} \cdot |\text{Det}(C)|^{-2} \cdot \text{cap}(L_{B,C}) = |\text{Det}(B)|^{-2} \cdot |\text{Det}(C)|^{-2}$$

if $L_{B,C}$ is doubly stochastic. Thus if we can find a doubly stochastic-scaling of L , then we can compute $\text{cap}(L)$ exactly. Algorithm G helps in finding an approximately doubly stochastic-scaling of L , which results in an approximation scheme for computing $\text{cap}(L)$. It is a very intriguing open problem if a $(1 + \epsilon)$ approximation to $\text{cap}(L)$ can be computed in time $\text{poly}(n, b, \log(1/\epsilon))$. In the aforementioned upcoming paper, we use this result (via a black-box reduction!) to approximating optimal constants in Brascamp–Lieb inequalities.

Although, we do not know of a convex formulation for capacity, it does have some nice properties. Let us look at the Lagrangian for $\log(\text{cap}(L))$:

$$f(X, \lambda) = \log(\text{Det}(L(X))) + \lambda \cdot \log(\text{Det}(X))$$

A critical point C , i.e., where the gradient $\nabla f(C, \lambda) = 0$ should satisfy $L^*(L(C)^{-1}) = C^{-1}$. This follows from

$$\nabla f(C, \lambda) = \left(L^* \left(L(C)^{-1} \right) + \lambda C^{-1}, \log(\text{Det}(C)) \right)$$

and the normalization condition

$$\text{tr} \left[C \cdot L^* \left(L(C)^{-1} \right) \right] = \text{tr} \left[L(C) \cdot L(C)^{-1} \right] = n$$

This is implied by Lemma 3.9 that

$$\text{cap}(L) \geq \frac{\text{Det}(L(C))}{\text{Det}(C)}$$

and hence any critical point is in fact a global minimum for $\text{cap}(L)$! It would be interesting to see if our techniques can help solve other non-convex optimization problems which share the same property.

Algorithm G can also be thought of as an alternating minimization algorithm for computing capacity. First let us write capacity in an alternative way.

$$n \cdot \text{cap}(L)^{1/n} = \inf \text{tr}[L(X) \cdot Y] = \inf \text{tr}[X \cdot L^*(Y)] \text{ s.t.} \\ \text{Det}(X) \geq 1, \text{Det}(Y) \geq 1, X, Y > 0$$

While the constraints are convex (by log-concavity of determinant), the objective is quadratic. However, if we fix X or Y , this program is convex. In fact, if we fix X , the optimum Y is $\text{Det}(L(X))^{1/n} \cdot L(X)^{-1}$. Similarly, if we fix Y , the optimum X is $\text{Det}(L^*(Y))^{1/n} \cdot L^*(Y)^{-1}$. Starting from $X_0 = I$, if we do alternate minimization, we get exactly Algorithm G , and our results imply that this process converges in polynomial number of steps!

From the description of capacity, it is not clear if it is continuous. The reason is that a priori, the optimizing matrices could be radically different for $\text{cap}(L_1)$ and $\text{cap}(L_2)$ even if L_1 and L_2 are quite close. We show that this not the case, essentially because $\text{cap}(L)$ (as well as the optimizer) can be approximated by a simple iterative process namely Algorithm G ! And Algorithm G is clearly continuous in L . We prove this formally in Theorem 4.5 of Sect. 4.3. In the aforementioned upcoming paper, we use this result to prove continuity of Brascamp–Lieb constants.

Remark 1.28 The continuity of capacity can also be proven via other methods and is already mentioned in [37]. But here we manage to get explicit bounds on the continuity parameter which we do not know how to get by methods other than analyzing Algorithm G .

Subsequent Work

After our work, Derksen and Makam [21] obtained polynomial degree bounds for the left–right action over any field! Using this bound, Ivanyos et al. [52] designed a completely different deterministic polynomial time algorithm for SINGULAR that works over all fields. This algorithm has a combinatorial/linear algebraic flavor, and has several important advantages over our algorithm, but also some limitations in comparison. One clear advantage is of course that it works over every field. Another is that this algorithm can *certify* the non-commutative singularity of an input L over

$\mathbb{F}\langle\langle \mathbf{x} \rangle\rangle$ by outputting a shrunk subspace when one exists! On the other hand, over the rationals, our algorithm computes capacity and finds an operator scaling which is doubly stochastic.

Using some of the techniques developed in [51,52], Bläser, Jindal and Pandey [9] designed a deterministic PTAS for the commutative rank, i.e., a $(1 + \epsilon)$ -approximation algorithm which runs in deterministic $n^{O(1/\epsilon)}$ time. This improves the factor-2 approximation algorithm given in this paper.

1.8 Organization

In Sect. 2, we describe Algorithm G for testing if quantum operators are rank-decreasing. We explain the notion of capacity of quantum operators, prove an explicit lower bound on the capacity of rank non-decreasing operators, and explain how it is used in the analysis to prove that Algorithm G runs in polynomial time. In Sect. 3, we study various properties of capacity, some of which are used in other places in the paper. In Sect. 4, we analyze the bit complexity of Algorithm G and also prove an explicit bound on the continuity of capacity. In Sect. 5, we will show how a modification of Algorithm G can be used to approximate capacity. In appendix Sect. A, we show how to compute the non-commutative rank. We conclude in Sect. 6 with a short discussion and open problems.

2 Quantum Operators and Analysis of Algorithm G

This section is devoted to Algorithm G and its analysis. We start with preliminaries about completely positive operators, their properties, and basic quantities associated with them. We then formally describe Algorithm G, and proceed to give a full analysis of its running time, proving the main theorem of this paper. Section 2.1 contains definitions and properties of completely positive operators and their capacity. Section 2.2 describes the Algorithm G and its analysis assuming an explicit lower bound on the capacity. Section 2.3 contains the main theorem concerning the lower bound on capacity of rank non-decreasing operators.

2.1 Completely Positive Operators and Capacity

Given a complex matrix A , we will use A^\dagger to denote the conjugate-transpose of A . For matrices with real entries, this will just be A^T . For a matrix A , $\text{Im}(A)$ will denote the image of A , i.e., $\{v \in \mathbb{C}^n \mid v = Au \text{ for some } u \in \mathbb{C}^n\}$.

Definition 2.1 (*Completely positive operators*) An operator (or map) $T : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ is called completely positive if there are $n \times n$ complex matrices A_1, \dots, A_m s.t. $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$. The matrices A_1, \dots, A_m are called Kraus operators of T (and they are not unique). T is called completely positive trace preserving (cptp) if $T^*(I) = \sum_{i=1}^m A_i^\dagger A_i = I$.

Remark 2.2 The above is actually not the usual definition of completely positive operators. T is defined to be positive if $T(X) \geq 0$ whenever $X \geq 0$. T is completely positive if $I_k \otimes T$ is positive for all $k \geq 1$. Choi [12] proved that an operator is completely positive iff it is of the form stated above.

Definition 2.3 (*Tensor products of operators*) Given operators $T_1 : M_{d_1}(\mathbb{C}) \rightarrow M_{d_1}(\mathbb{C})$ and $T_2 : M_{d_2}(\mathbb{C}) \rightarrow M_{d_2}(\mathbb{C})$, we define their tensor product $T_1 \otimes T_2 : M_{d_1 d_2}(\mathbb{C}) \rightarrow M_{d_1 d_2}(\mathbb{C})$ in the natural way

$$(T_1 \otimes T_2)(X \otimes Y) = T_1(X) \otimes T_2(Y)$$

and extend by linearity to the whole of $M_{d_1 d_2}(\mathbb{C})$.

Definition 2.4 If $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$ is a completely positive operator, we define its dual T^* by $T^*(X) = \sum_{i=1}^m A_i^\dagger X A_i$. If both T and T^* are trace preserving, namely $T(I) = T^*(I) = I$, then we call T (and T^*) *doubly stochastic*.

Definition 2.5 (*Rank-Decreasing Operators*) A completely positive operator T is said to be rank-decreasing if there exists an $X \geq 0$ s.t. $\text{rank}(T(X)) < \text{rank}(X)$. T is said to be c -rank-decreasing if there exists an $X \geq 0$ s.t. $\text{rank}(T(X)) \leq \text{rank}(X) - c$. We will sometimes refer to operators that are not rank-decreasing as rank non-decreasing.

Now that we defined completely positive operators, we define their *capacity*, which is a very important complexity measure of such operators suggested in [37]. Its evolution will be central to the complexity analysis of our algorithm.

Definition 2.6 (*Capacity* [37]) The capacity of a completely positive operator T , denoted by $\text{cap}(T)$, is defined as

$$\text{cap}(T) = \inf\{\text{Det}(T(X)) : X \succ 0, \text{Det}(X) = 1\}$$

This notion of capacity has a very interesting history. Some special cases of capacity were defined in [39,40]. It was then extended to hyperbolic polynomials and this also led to a resolution (and extremely elegant proofs) of the Van der Waerden Conjecture for Mixed Discriminants [38].

The next proposition shows how capacity changes when linear transformations are applied (as in the algorithm) to the completely positive operator.

Proposition 2.7 [37] *Let T be the operator defined by A_1, \dots, A_m and let $T_{B,C}$ be the operator defined by BA_1C, \dots, BA_mC , where B, C are invertible matrices. Then*

$$\text{cap}(T_{B,C}) = |\text{Det}(B)|^2 |\text{Det}(C)|^2 \text{cap}(T)$$

Proof

$$\begin{aligned}
 \text{cap}(T_{B,C}) &= \inf \left\{ \text{Det} \left(\sum_{i=1}^m B A_i C X C^\dagger A_i^\dagger B^\dagger \right) : X \succ 0, \text{Det}(X) = 1 \right\} \\
 &= |\text{Det}(B)|^2 \cdot \inf \left\{ \text{Det} \left(\sum_{i=1}^m A_i C X C^\dagger A_i^\dagger \right) : X \succ 0, \text{Det}(X) = 1 \right\} \\
 &= |\text{Det}(B)|^2 \cdot \inf \left\{ \text{Det} \left(\sum_{i=1}^m A_i X A_i^\dagger \right) : X \succ 0, \text{Det}(X) = |\text{Det}(C)|^2 \right\} \\
 &= |\text{Det}(B)|^2 \cdot |\text{Det}(C)|^2 \cdot \text{cap}(T) \quad \square
 \end{aligned}$$

The next proposition gives a useful upper bound on the capacity of trace-preserving completely positive operators; this will be used for the convergence analysis of the algorithm.

Proposition 2.8 (Capacity Upper Bound) *Let T be a completely positive operator with Kraus operators A_1, \dots, A_m (which are $n \times n$ complex matrices). Also suppose that either $\sum_{i=1}^m A_i^\dagger A_i = I$ or $\sum_{i=1}^m A_i A_i^\dagger = I$. Then $\text{cap}(T) \leq 1$.*

Proof Note that $\text{tr}(T(I)) = n$ in either case.

$$\begin{aligned}
 \text{cap}(T) &\leq \text{Det}(T(I)) \leq (\text{tr}(T(I))/n)^n \\
 &= 1
 \end{aligned}$$

The second inequality follows from the AM–GM inequality. □

2.2 Algorithm G and Its Convergence Rate

We now describe Algorithm G to test whether a completely positive operator T (given in terms of Kraus operators A_1, \dots, A_m) is rank non-decreasing (equivalently properties (1)–(8) in Theorem 1.4 one of which is that A_1, \dots, A_m admit no shrunk subspace). We will then analyze its convergence rate, namely the number of scaling iterations needed as a function of the input size. In Sect. 4, we will continue with the finer analysis of the bit complexity of this algorithm.

Since the property of A_1, \dots, A_m having a shrunk subspace or not remains invariant if we replace A_1, \dots, A_m by a basis spanning the subspace spanned by A_1, \dots, A_m , we can always assume wlog that $m \leq n^2$. Suppose the maximum bit size of the entries of A_1, \dots, A_m is b . Since scaling the matrices A_1, \dots, A_m does not change the problem, we can assume that A_1, \dots, A_m have integer entries of magnitude at most $M = 2^{O(b)}$.

Algorithm G below is essentially Gurvits’ algorithm [37] for Edmonds’ problem for subspaces of matrices having Edmonds–Rado property. It is a non-commutative generalization of the Sinkhorn scaling procedure to scale nonnegative matrices to doubly stochastic matrices (see, for example, [60] and references therein). The algorithm alternates applying a “normalizing” basis change from the left and right to

the given matrices, so as to alternately make the operator or its dual trace preserving. The idea is that this process will converge to a doubly stochastic operator iff it is not rank-decreasing, and furthermore, we can bound the number t of iterations by $\text{poly}(n, \log(M))$. We will use the following to measure of our operator to being doubly stochastic.

Definition 2.9

$$\text{ds}(T) = \text{tr} \left[(T(I) - I)^2 \right] + \text{tr} \left[(T^*(I) - I)^2 \right]$$

Definition 2.10 Given a completely positive operator T , define its *right normalization* T_R as follows:

$$T_R(X) = T \left(T^*(I)^{-1/2} X T^*(I)^{-1/2} \right)$$

Note that $T_R^*(I) = I$.

Definition 2.11 Given a completely positive operator T , define its *left normalization* T_L as follows:

$$T_L(X) = T(I)^{-1/2} T(X) T(I)^{-1/2}$$

Note that $T_L(I) = I$. These operations are referred to as row and column operations in [37].

We next prove that if a completely positive operator T is doubly stochastic, or close to being one in this measure, then it is rank non-decreasing. This will be important for the termination condition of Algorithm G .

Theorem 2.12 [37] *If T is a completely positive operator which is right (or left) normalized satisfying $\text{ds}(T) \leq 1/(n + 1)$, then T is rank non-decreasing.*

Proof Wlog assume that $T(I) = I$ and

$$\text{tr} \left[(T^*(I) - I)^2 \right] \leq 1/(n + 1)$$

Suppose $X \geq 0$ is a psd matrix s.t. $\text{Rank}(X) = r$. We would like to prove that $\text{Rank}(T(X)) \geq r$. Let

$$X = \sum_{i=1}^r \lambda_i v_i v_i^\dagger$$

be the eigenvalue decomposition of X where v_i 's are orthonormal and $\lambda_i > 0$. Then

$$T(X) = \sum_{i=1}^r \lambda_i T \left(v_i v_i^\dagger \right)$$

Let us denote $T(v_i v_i^\dagger)$ by R_i . Since $T(I) = I$, we get that $\sum_{i=1}^r R_i \preceq T(I) = I$. Also note that

$$\text{Rank} \left(\sum_{i=1}^r \lambda_i R_i \right) = \text{Rank} \left(\sum_{i=1}^r R_i \right)$$

This is because R_i 's are psd matrices and hence a vector is in the kernel of $\sum_{i=1}^r \lambda_i R_i$ iff it is in the kernel of all the R_i 's. Because of $\sum_{i=1}^r R_i \preceq I$, we get that

$$\text{Rank}(T(X)) = \text{Rank} \left(\sum_{i=1}^r R_i \right) \geq \text{tr} \left[\sum_{i=1}^r R_i \right]$$

Suppose $T^*(I) = I + \Delta$. We know that $\text{tr}[\Delta^2] \leq 1/(n + 1)$. Then

$$\text{tr}[R_i] = \text{tr} \left[T^*(I) v_i v_i^\dagger \right] = 1 + \text{tr} \left[\Delta v_i v_i^\dagger \right]$$

Let P denote the projection onto $\{v_1, \dots, v_r\}$. Adding the above equations for all $i \in \{1, \dots, r\}$, we get

$$\begin{aligned} \text{tr} \left[\sum_{i=1}^r R_i \right] &= r + \text{tr}[\Delta P] \\ &\geq r - \text{tr}[\Delta^2]^{1/2} \text{tr}[P^2]^{1/2} \\ &\geq r - \sqrt{\frac{r}{n + 1}} \end{aligned}$$

Since $r < n + 1$, this gives us

$$\text{Rank}(T(X)) > r - 1$$

which completes the proof since $\text{Rank}(T(X))$ is an integer. □

We are now ready to present the algorithm.

Input: Completely positive operator T given in terms of Kraus operators $A_1, \dots, A_m \in \mathbb{Z}^{n \times n}$. Each entry of A_j has absolute value at most M .

Output: Is T rank non-decreasing?

1. Check whether $T(I)$ and $T^*(I)$ are singular. If any one of them is singular, then output that the operator is rank-decreasing; otherwise, proceed to step 2.
2. Perform right and left normalizations on T alternatively for t steps. Let T_j be the operator after j steps. Also let $\epsilon_j = \text{ds}(T_j)$.
3. Check if $\min\{\epsilon_j : 1 \leq j \leq t\} \leq 1/(6n)$. If yes, then output that the operator is rank non-decreasing otherwise output rank-decreasing.

Algorithm 1: Algorithm G

Remark 2.13 In algorithm G , the operators are maintained in terms of the Kraus operators and it is easy to see the effect of right and left normalizations on the Kraus operators. In fact, they are named such since for a right (left) normalization, the Kraus operators multiplied on the right (left) by $T^*(I)^{-1/2}$ ($T(I)^{-1/2}$).

The main objective is to analyze the minimum number of steps t for which this algorithm terminates with the correct answer. The following theorem of [37] gives the time analysis of the algorithm *assuming* an initial lower bound on the capacity of the input completely positive operator. In the next subsection, we will prove our main result, an appropriate lower bound, which shows that the algorithm terminates with the correct answer in polynomial time.

Theorem 2.14 [37] *Let T be a completely positive operator that is rank non-decreasing, with Kraus operators A_1, \dots, A_m , which are $n \times n$ integer matrices such that each entry of A_i has absolute value at most M . Suppose T_R is the right normalization of T . If $\text{cap}(T_R) \geq f(n, M)$, then Algorithm G when applied for at least $t = 2 + 36n \cdot \log(1/f(n, M))$ steps is correct.*

For completeness sake, we provide a full proof of this theorem. Again, this analysis follows similar ones for the classical Sinkhorn iterations, e.g., as in [41,60]. Basically, capacity increases by a factor roughly $1 + 1/36n$ per iteration as long as it is not too close to 1.

Proof If either $T(I)$ or $T^*(I)$ is singular, then T is rank-decreasing. When $T(I)$ is singular, T decreases the rank of I . When $T^*(I)$ is singular, any vector in the kernel of $T^*(I)$ lies in the kernels of all the A_i 's. If $T(I)$ and $T^*(I)$ are both non-singular, it is easy to verify that $T_j(I)$ and $T_j^*(I)$ will remain non-singular for all j and hence step 2 is well defined. Also using Theorem 2.12 and the fact that right and left normalizations do not change the property of being rank-decreasing, Algorithm G will always output rank-decreasing if T is rank-decreasing.

So what is left to prove is if T is rank non-decreasing, then $\min\{\epsilon_j : 1 \leq j \leq t\} \leq 1/6n$. Assume to the contrary that it is not. Denote by cap_j to be the capacity of the operator T_j . By Lemma 5.2 below (which essentially is a robust version of the AM–GM inequality), if $\epsilon_j > 1/6n$, then $\text{cap}_{j+1} \geq \exp(1/36n) \cdot \text{cap}_j$. From the assumption of the theorem, we know that $\text{cap}_1 \geq f(n, M)$. Also it is easy to see that $\text{cap}_j \leq 1$ for all j (Proposition 2.8). However by the assumption that $\min\{\epsilon_j : 1 \leq j \leq t\} > 1/6n$ and the increase in capacity per iteration seen above, we get that

$$1 \geq \text{cap}_t \geq \exp\left(\frac{t-1}{36n}\right) \cdot \text{cap}_1 = \exp\left(\frac{t-1}{36n}\right) \cdot f(n, M)$$

Plugging in $t = 2 + 36n \cdot \log(1/f(n, M))$ gives us the required contradiction. □

In the main Theorem 2.20, we will prove that the quantity $f(n, M)$ used in the statement of Theorem 2.14 is $\geq \exp(-4n \log(Mmn))$, which will prove that the number of iterations needed in Algorithm G is $O(n^2 \log(Mmn))$. But this alone does

not guarantee that the algorithm is actually polynomial time, since the bit sizes of numbers involved might get exponential. As it happens, simple truncation suffices to overcome this problem, as shown in [37], and we reproduce this analysis in Sect. 4.2 for completeness.

2.3 Main Theorem: The Lower Bound on Capacity

In this subsection, we prove our main theorem, a lower bound on capacity of an operator in terms of its description size, in Theorem 2.20. Before diving into the proof, we provide a high level plan. We will first prove that if a completely positive operator with integer entries is rank non-decreasing (that is has positive capacity), then the capacity is actually non-negligible. Our starting point (Theorem 1.4) is the statement of the equivalence between the rank non-decreasing property and an algebraic condition (non-vanishing of a certain determinant). Using Alon’s Combinatorial Nullstellensatz, we can ensure small coefficients in the algebraic condition above. We then state and prove (Lemma 2.17) a Cauchy–Schwarz inequality for matrices needed next. The main result (Theorem 2.18) proceeds by applying an appropriate scaling to the original integral operator, which reduces it to one that preserves the identity matrix. The latter property (which provides a trace bound), the integrality of the original operator (which provides an integral determinantal lower bound) and the multiplicativity of capacity under scaling combine (via the above inequality) to give the desired bound. We now proceed with the details of this plan.

We will need the fact that (nonzero) polynomials of degree d cannot vanish on all points with nonnegative integer coordinates with sum $\leq d$. This follows from Alon’s combinatorial nullstellensatz [2, Theorem 1.2].

Lemma 2.15 [2] *If $p(x_1, \dots, x_n) \in \mathbb{C}[x_1, \dots, x_n]$ is a (nonzero) polynomial of degree d , where the individual degree of the variable x_i is at most d_i , then there exists $(a_1, \dots, a_n) \in \mathbb{Z}_{\geq 0}^n$ such that $\sum_{i=1}^n a_i \leq d$ and $a_i \leq d_i$, for which $p(a_1, \dots, a_n) \neq 0$.*

As a corollary of Theorem 1.4 and of Lemma 2.15, we get the following:

Corollary 2.16 *Let A_1, \dots, A_m be $n \times n$ complex matrices s.t. the completely positive operator T_A defined by A_1, \dots, A_m is rank non-decreasing. Then there exist matrices $D_1, \dots, D_m \in \mathcal{M}_d(\mathbb{C})$ for some d such that*

$$\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m) \neq 0.$$

Furthermore, D_1, \dots, D_m can be chosen to be matrices with nonnegative integer entries s.t.

$$\sum_{j=1}^m \sum_{k,l=1}^d D_j(k, l)^2 \leq n^2 d$$

Proof By equivalences (3) and (7) of Theorem 1.4, we know that there exist matrices $F_1, \dots, F_m \in \mathcal{M}_d(\mathbb{C})$, for some d , such that $\text{Det}(F_1 \otimes A_1 + \dots + F_m \otimes A_m) \neq 0$. This implies that the polynomial $p(X_1, \dots, X_m) = \text{Det}(X_1 \otimes A_1 + \dots + X_m \otimes A_m)$, where X_i are symbolic $d \times d$ matrices, is nonzero. Hence, by Lemma 2.15, we know that there exist matrices of positive integers D_1, \dots, D_m such that $\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m) \neq 0$ and that

$$\sum_{j=1}^m \sum_{k,l=1}^d D_j(k, l) \leq \text{deg}(p) = nd$$

$$D_j(k, l) \leq n$$

This is because the total degree of p is nd and the individual degree of each variable $X_j(k, l)$ in p is at most n . This implies the desired upper bound. \square

We will also need the following Cauchy–Schwarz like inequality:

Lemma 2.17 *Let C_1, \dots, C_m and D_1, \dots, D_m be d_1 and d_2 dimensional complex matrices, respectively. Then*

$$\text{tr} \left[\left(\sum_{i=1}^m C_i \otimes D_i \right) \left(\sum_{j=1}^m C_j^\dagger \otimes D_j^\dagger \right) \right] \leq \text{tr} \left[\sum_{i=1}^m C_i C_i^\dagger \right] \cdot \text{tr} \left[\sum_{j=1}^m D_j D_j^\dagger \right]$$

Proof We start with the inequality:

$$\begin{aligned} 0 &\leq \text{tr} \left[(C_i \otimes D_j - C_j \otimes D_i) (C_i^\dagger \otimes D_j^\dagger - C_j^\dagger \otimes D_i^\dagger) \right] \\ &= \text{tr} \left[C_i C_i^\dagger \otimes D_j D_j^\dagger + C_j C_j^\dagger \otimes D_i D_i^\dagger \right] - \text{tr} \left[C_i C_j^\dagger \otimes D_j D_i^\dagger + C_j C_i^\dagger \otimes D_i D_j^\dagger \right] \\ &= \text{tr} \left[C_i C_i^\dagger \right] \cdot \text{tr} \left[D_j D_j^\dagger \right] + \text{tr} \left[C_j C_j^\dagger \right] \cdot \text{tr} \left[D_i D_i^\dagger \right] - \text{tr} \left[C_i C_j^\dagger \otimes D_i D_j^\dagger \right] \\ &\quad - \text{tr} \left[C_j C_i^\dagger \otimes D_j D_i^\dagger \right] \end{aligned}$$

The first inequality is just $\text{tr} [YY^\dagger] \geq 0$. The second equality follows from linearity of trace, $\text{tr} [Y \otimes Z] = \text{tr} [Y] \cdot \text{tr} [Z]$ and $\text{tr} [XY^\dagger] = \text{tr} [YX^\dagger]$. Rearranging the above, we get the following inequality:

$$\begin{aligned} \text{tr} \left[C_i C_j^\dagger \otimes D_i D_j^\dagger \right] + \text{tr} \left[C_j C_i^\dagger \otimes D_j D_i^\dagger \right] &\leq \text{tr} \left[C_i C_i^\dagger \right] \cdot \text{tr} \left[D_j D_j^\dagger \right] \\ &+ \text{tr} \left[C_j C_j^\dagger \right] \cdot \text{tr} \left[D_i D_i^\dagger \right] \end{aligned} \tag{1}$$

Summing Eq. (1) over pairs (i, j) with $i < j$, we obtain the following:

$$\sum_{i \neq j} \text{tr} \left[C_i C_j^\dagger \otimes D_i D_j^\dagger \right] \leq \sum_{i \neq j} \text{tr} \left[C_i C_i^\dagger \right] \cdot \text{tr} \left[D_j D_j^\dagger \right]$$

Adding $\sum_i \text{tr} [C_i C_i^\dagger] \cdot \text{tr} [D_i D_i^\dagger]$ to both sides completes the proof (due to linearity of trace). \square

Now we are ready to prove our main theorem for operators with integral Kraus operators.

Theorem 2.18 (Capacity of Square Operators) *Suppose T_A is a completely positive operator which is rank non-decreasing and has Kraus operators $A_1, \dots, A_m \in \mathcal{M}_n(\mathbb{Z})$. In this case:*

$$\text{cap}(T_A) \geq \frac{1}{n^{2n}}$$

Proof Since T_A is rank non-decreasing, by Corollary 2.16, there exist nonnegative integer matrices D_1, \dots, D_m of some dimension d s.t. $\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m) \neq 0$ and also

$$\sum_{j=1}^m \sum_{k,l=1}^d D_j(k, l)^2 \leq n^2 d \tag{2}$$

Since A_1, \dots, A_m are also integer matrices, $\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m)$ is a nonzero integer and hence

$$|\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m)| \geq 1 \tag{3}$$

Let $X > 0$ be such that $\text{Det}(X) = 1$. Consider the matrices $C_i = T_A(X)^{-1/2} A_i X^{1/2}$. This is a scaling intended so that the operator T_C defined by the matrices C_i preserves identity, i.e., satisfies

$$\begin{aligned} T_C(I) &= \sum_{i=1}^m C_i C_i^\dagger = T_A(X)^{-1/2} \left(\sum_{i=1}^m A_i X^{1/2} X^{1/2} A_i^\dagger \right) T_A(X)^{-1/2} \\ &= I_n \end{aligned}$$

Hence

$$\text{tr} \left[\sum_{i=1}^m C_i C_i^\dagger \right] = n \tag{4}$$

Let Y denote the matrix $D_1 \otimes C_1 + \dots + D_m \otimes C_m$. Then

$$|\text{Det}(Y)| = |\text{Det}(D_1 \otimes C_1 + \dots + D_m \otimes C_m)| \tag{5}$$

$$\begin{aligned} &= |\text{Det}(D_1 \otimes A_1 + \dots + D_m \otimes A_m)| \cdot \text{Det}(X)^{d/2} \cdot \text{Det}(T_A(X))^{-d/2} \\ &\geq \text{Det}(T_A(X))^{-d/2} \tag{6} \end{aligned}$$

The equality follows from multiplicativity of determinant and the fact that $\text{Det}(I_d \otimes Y) = \text{Det}(Y)^d$. The inequality follows from $\text{Det}(X) = 1$ and Eq. (3). Hence we obtain

$$\text{Det}(YY^\dagger) = |\text{Det}(Y)|^2 \geq \text{Det}(T_A(X))^{-d} \tag{7}$$

On the other hand, by the AM–GM inequality

$$\text{Det}(YY^\dagger) \leq \left(\frac{\text{tr}[YY^\dagger]}{nd} \right)^{nd} \tag{8}$$

Now consider

$$\begin{aligned} \text{tr}[YY^\dagger] &= \text{tr} \left[\left(\sum_{i=1}^m D_i \otimes C_i \right) \left(\sum_{j=1}^m D_j^\dagger \otimes C_j^\dagger \right) \right] \\ &\leq \text{tr} \left[\sum_{i=1}^m D_i D_i^\dagger \right] \cdot \text{tr} \left[\sum_{j=1}^m C_j C_j^\dagger \right] \\ &= \left(\sum_{i=1}^m \sum_{k,l=1}^d D_i(k,l)^2 \right) \cdot n \\ &\leq n^3 d \end{aligned} \tag{9}$$

The first inequality follows from Lemma 2.17. The second equality follows from Eq. (4). The second inequality follows from Eq. (2). Combining Eqs. (7)–(9), we obtain:

$$\text{Det}(T_A(X)) \geq \frac{1}{n^{2n}}$$

Since X was an arbitrary matrix s.t. $X \succ 0$ and $\text{Det}(X) = 1$, we obtain the desired lower bound on $\text{cap}(T_A)$. □

Remark 2.19 A slightly stronger bound can be obtained by using quantum permanents [37] in the case when the D_i ’s are of constant dimension.

Theorem 2.18 immediately implies the following capacity lower bound that we need.

Theorem 2.20 *Suppose T is a cptp map which is rank non-decreasing and is obtained by right normalization of an operator with Kraus operators A_1, \dots, A_m , which are $n \times n$ integer matrices such that each entry of A_i has absolute value at most M . Then $\text{cap}(T) \geq \exp(-4n \log(Mmn))$.*

Proof Let T_A be the operator defined by A_1, \dots, A_m . Since T is the right normalization of T_A , by Proposition 2.7, we get that

$$\text{cap}(T) = \frac{\text{cap}(T_A)}{\text{Det}(T^*(I))} \tag{10}$$

Note that $T^*(I)$ is a psd matrix. Also

$$\begin{aligned} \text{tr}[T^*(I)] &= \sum_{i=1}^m \text{tr}[A_i^\dagger A_i] \\ &= \sum_{i=1}^m \sum_{k,l=1}^n |A_i(k,l)|^2 \\ &\leq M^2 mn^2 \end{aligned}$$

This implies (via the AM–GM inequality) that

$$\begin{aligned} \text{Det}(T^*(I)) &\leq \left(\frac{\text{tr}[T^*(I)]}{n} \right)^n \\ &\leq (M^2 mn)^n \end{aligned} \tag{11}$$

Combining Eqs. (10) and (11), and Theorem 2.18 gives the desired bound. □

3 Properties of Capacity

In this section, we prove some interesting properties of capacity. In Sect. 3.1, we prove that the capacity of operators which are close to doubly stochastic is close to 1. This is used in the proof of Theorem 3.11. In Sect. 3.2, we characterize the almost optimal points of capacity in terms of approximate fixed points of an operator. This will be used later in Sect. 4 in the proof of continuity of capacity (Theorem 4.5). In Sect. 3.3, we prove a multiplicativity property of capacity under tensor products of operators. This property was used in the lower bound on capacity in the previous version of this paper. While it is no longer needed for this purpose now, it is an interesting and intriguing fact by itself, so we include a proof.

3.1 Capacity of Almost Doubly Stochastic Operators

The following definition of capacity for nonnegative matrices is due to [41]. They used it to analyze Sinkhorn’s algorithm for matrix scaling.

Definition 3.1 (*Capacity of nonnegative matrices*) Given a nonnegative matrix A , its capacity is defined as follows:

$$\text{cap}(A) = \inf \left\{ \prod_{i=1}^n (Ax)_i : \prod_{i=1}^n x_i = 1, x > 0 \right\}$$

The next lemma states that the capacity of almost doubly stochastic matrices is close to 1.

Lemma 3.2 *Let A be a nonnegative $n \times n$ matrix s.t. the row sums of A are 1 and*

$$\sum_{i=1}^n (c_i - 1)^2 \leq \epsilon$$

where c_1, \dots, c_n are the column sums of A . Then

$$\text{cap}(A) \geq (1 - \sqrt{n\epsilon})^n$$

Proof We start by proving that for doubly stochastic matrices B , $\text{cap}(B) \geq 1$. This will follow from concavity of log. Indeed, let $x \in \mathbb{R}^n$ s.t. $\prod_{i=1}^n x_i = 1$ and $x > 0$. Then

$$\begin{aligned} \sum_{i=1}^n \log((Bx)_i) &= \sum_{i=1}^n \log\left(\sum_{j=1}^n B_{i,j}x_j\right) \\ &\geq \sum_{i=1}^n \sum_{j=1}^n B_{i,j} \log(x_j) \\ &= \sum_{j=1}^n \log(x_j) \sum_{i=1}^n B_{i,j} \\ &= \sum_{j=1}^n \log(x_j) \\ &= 0 \end{aligned}$$

For the inequality, we used concavity of log and the fact that row sums of B are 1. For the third equality, we used that the column sums of B are 1.

Now we need to move on to A which is almost doubly stochastic. A direct argument like for the doubly stochastic case does not work. We need to first prove the following claim [60]:

Claim 3.3 *There exists a decomposition of $A = \lambda B + Z$, where B is doubly stochastic, Z is nonnegative and $\lambda \geq 1 - \sqrt{n\epsilon}$.*

Given the claim, it is easy to complete the proof of the lemma.

$$\text{cap}(A) \geq \lambda^n \text{cap}(B) \geq (1 - \sqrt{n\epsilon})^n$$

So we end with a proof of the claim from [60]. We will first prove that there is a decomposition $A = D + Z$ where D is multiple of a doubly stochastic matrix and $\text{per}(Z) = 0$. Initially, we start with $D = 0$ and $Z = A$. As long as $\text{per}(Z) > 0$, there is an $\alpha > 0$ and a permutation matrix P s.t. $Z' = Z - \alpha P$ is nonnegative and the number of nonzero entries in Z' is strictly less than that of Z . Replace D by $D' = D + \alpha P$ and Z by Z' . Note that D' is also a multiple of a doubly stochastic matrix. After at most n^2 iterations, we will find a decomposition $A = \lambda B + Z$, where $\text{per}(Z) = 0$ (and B is doubly stochastic). We will now prove that $\lambda \geq 1 - \sqrt{n\epsilon}$. If $\lambda = 1$, then we are already done, so assume $\lambda < 1$. In this case, consider the matrix $C = \frac{Z}{1-\lambda}$. Row sums of C are 1 and the i th column sum $c'_i = \frac{c_i - \lambda}{1-\lambda}$. Then

$$\sum_{i=1}^n (c'_i - 1)^2 = \frac{1}{(1-\lambda)^2} \sum_{i=1}^n (c_i - 1)^2 \leq \frac{\epsilon}{(1-\lambda)^2}$$

Since $\text{per}(C) = 0$ and row sums are 1, it follows that (see Lemma 5.2 in [60])

$$\sum_{i=1}^n (c'_i - 1)^2 \geq 1/n$$

and hence

$$\frac{\epsilon}{(1-\lambda)^2} \geq 1/n$$

which implies $\lambda \geq 1 - \sqrt{n\epsilon}$ □

The next lemma says that capacity of almost doubly stochastic operators is close to 1. The proof will proceed by reducing the operator case to the nonnegative matrix case established above.

Lemma 3.4 *Let T be a positive operator acting on $n \times n$ matrices such that $\text{tr}[(T(I) - I)^2] \leq \epsilon$ and $T^*(I) = I$ (equivalently T is trace preserving). Then $\text{cap}(T) \geq (1 - \sqrt{n\epsilon})^n$.*

Remark 3.5 The proof can be adapted to obtain a similar statement when both $T(I), T^*(I) \approx I$.

Proof Let X be a positive definite matrix s.t. $\text{Det}(X) = 1$. Let

$$X = \sum_{j=1}^n \lambda_j v_j v_j^\dagger$$

be an eigenvalue decomposition of X with v_1, \dots, v_n forming a orthonormal basis. Then

$$T(X) = \sum_{j=1}^n \lambda_j T(v_j v_j^\dagger)$$

Let us denote $T(v_j v_j^\dagger)$ by R_j . Then since T is trace preserving, we have that $\text{tr}(R_j) = 1$ for all j . Also let

$$T(X) = \sum_{i=1}^n \sigma_i u_i u_i^\dagger$$

be an eigenvalue decomposition for $T(X)$ with u_1, \dots, u_n forming an orthonormal basis. It follows that for all i ,

$$\sigma_i = \sum_{j=1}^n \lambda_j u_i^\dagger R_j u_i$$

Let A denote the nonnegative matrix s.t. $A_{i,j} = u_i^\dagger R_j u_i$. Then $\sigma = A\lambda$. Also the column sums of A are all 1 since $\text{tr}(R_j) = 1$ for all j . We will also prove that

$$\sum_{i=1}^n (r_i - 1)^2 \leq \epsilon$$

where r_1, \dots, r_n denote the column sums of A . Then the lemma follows from Lemma 3.2 (applied with row sums switched with column sums) and the facts that $\text{Det}(T(X)) = \prod_{i=1}^n \sigma_i$ and $\text{Det}(X) = \prod_{i=1}^n \lambda_i$. Now note that

$$\begin{aligned} r_i &= \sum_{j=1}^n A_{i,j} = \sum_{j=1}^n u_i^\dagger R_j u_i = u_i^\dagger \left(\sum_{j=1}^n T(v_j v_j^\dagger) \right) u_i \\ &= u_i^\dagger T \left(\sum_{j=1}^n v_j v_j^\dagger \right) u_i = u_i^\dagger T(I) u_i \end{aligned}$$

Hence

$$\begin{aligned} \sum_{i=1}^n (r_i - 1)^2 &= \sum_{i=1}^n \left(u_i^\dagger T(I) u_i - 1 \right)^2 = \sum_{i=1}^n \left(u_i^\dagger (T(I) - I) u_i \right)^2 \\ &\leq \sum_{i=1}^n u_i^\dagger (T(I) - I)^2 u_i \\ &= \text{tr}[(T(I) - I)^2] \leq \epsilon \end{aligned}$$

The first inequality can be proved via convexity of square. □

We also prove an easy lemma for the other direction: Operators with capacity almost 1 are almost doubly stochastic.

Lemma 3.6 *Let T be a positive operator acting on $n \times n$ matrices s.t. $\text{cap}(T) \geq \exp(-\delta)$, $\delta \leq 1/6$ and also $T^*(I) = I$. Then $\text{ds}(T) = \text{tr}[(T(I) - I)^2] \leq 6\delta$.*

Proof $\text{cap}(T) \geq \exp(-\delta)$ implies that $\text{Det}(T(I)) \geq \exp(-\delta)$. Also $\text{tr}[T(I)] = \text{tr}[T^*(I)] = n$. Suppose $\text{tr}[(T(I) - I)^2] = \alpha$. Then Lemma 5.1 implies that

$$\text{Det}(T(I)) \leq \exp(-\alpha/6)$$

This implies that $\alpha \leq 6\delta$. □

Remark 3.7 Note that the parameters in Lemmas 3.4 and 3.6 do not match and that is because capacity and the distance to doubly stochasticity do not characterize each other exactly.

3.2 Characterization of Approximate Optimizers for Capacity

The following definition will be useful.

Definition 3.8 (Fixed points) Let T be a completely positive operator. Define $\text{Fixed}(T, \epsilon)$ to be the set of Hermitian positive-definite matrices which are ϵ -approximate fixed point of the operator $X \rightarrow T^*(T(X)^{-1})^{-1}$, i.e., $C \in \text{Fixed}(T, \epsilon)$ if the following holds:

$$\text{tr} \left[\left(C \cdot T^* \left(T(C)^{-1} \right) - I \right)^2 \right] \leq \epsilon$$

The following lemma essentially says that the elements of $\text{Fixed}(T, \epsilon)$ are approximate minimizers for $\text{cap}(T)$.

Lemma 3.9 *Suppose $C \in \text{Fixed}(T, \epsilon)$. If $\epsilon \leq 1/(n+1)$, then T is rank non-decreasing. Furthermore,*

$$(1 - \sqrt{n\epsilon})^n \cdot \frac{\text{Det}(T(C))}{\text{Det}(C)} \leq \text{cap}(T) \leq \frac{\text{Det}(T(C))}{\text{Det}(C)}$$

Here n is the size of matrices on which T acts. Similar statement also holds for the operator $X \rightarrow T(T^*(X)^{-1})^{-1}$.

We will use a $C \in \text{Fixed}(T, \epsilon)$ to find a scaling of T which is almost doubly stochastic and then apply Theorem 2.12 (saying that almost doubly stochastic operators are rank non-decreasing) and Lemma 3.4 (giving a lower bound on capacity of almost doubly stochastic operators).

Proof Consider the operator

$$Z(X) = T(C)^{-1/2} \cdot T \left(C^{1/2} \cdot X \cdot C^{1/2} \right) \cdot T(C)^{-1/2}$$

Then $Z(I) = I$ and also

$$\begin{aligned} \text{tr} \left[(Z^*(I) - I)^2 \right] &= \text{tr} \left[\left(C^{1/2} \cdot T^* \left(T(C)^{-1} \right) \cdot C^{1/2} - I \right)^2 \right] \\ &= \text{tr} \left[\left(C \cdot T^* \left(T(C)^{-1} \right) - I \right)^2 \right] \leq \epsilon \end{aligned}$$

When $\epsilon \leq 1/(n + 1)$, by Theorem 2.12, Z is rank non-decreasing and hence T is rank non-decreasing. Also by Lemma 3.4,

$$\text{cap}(Z) = \text{cap}(Z^*) \geq (1 - \sqrt{n\epsilon})^n$$

Then by Proposition 2.7

$$\text{cap}(T) = \text{cap}(Z) \cdot \frac{\text{Det}(T(C))}{\text{Det}(C)} \geq (1 - \sqrt{n\epsilon})^n \cdot \frac{\text{Det}(T(C))}{\text{Det}(C)}$$

The upper bound on $\text{cap}(T)$ is clear from the definition of capacity. □

Next we prove the other direction but again, not with matching parameters.

Lemma 3.10 *Suppose $C \succ 0$ is s.t.*

$$\text{cap}(T) \geq \exp(-\delta) \cdot \frac{\text{Det}(T(C))}{\text{Det}(C)}$$

where $\delta \leq 1/6$. Then $C \in \text{Fixed}(T, 6\delta)$.

Proof As in the proof of Lemma 3.9, define the operator

$$Z(X) = T(C)^{-1/2} \cdot T \left(C^{1/2} \cdot X \cdot C^{1/2} \right) \cdot T(C)^{-1/2}$$

Then $Z(I) = I$ and by Proposition 2.7

$$\text{cap}(Z^*) = \text{cap}(Z) \geq 1 - \delta$$

Now by Lemma 3.6 (applied to Z^*),

$$\text{ds}(Z) = \text{ds}(Z^*) = \text{tr}[(Z^*(I) - I)^2] \leq 6\delta$$

This implies that $C \in \text{Fixed}(T, 6\delta)$ because of the way operator Z is set up. □

3.3 Multiplicativity of Capacity

In this subsection, we prove a curious multiplicativity property of capacity of completely positive operators. This is not required for the lower bound on capacity presented in this version, but the easy direction of this theorem was required in the lower bound on capacity in the previous version.

Theorem 3.11 *Let T_1 and T_2 be completely positive operators where T_1 acts on matrices of dimension d_1 and T_2 acts on matrices of dimension d_2 . Then*

$$\text{cap}(T_1 \otimes T_2)^{1/(d_1 d_2)} = \text{cap}(T_1)^{1/d_1} \cdot \text{cap}(T_2)^{1/d_2}$$

We first prove the easy \leq direction of the above theorem. If X and Y are the minimizers for $\text{cap}(T_1)$ and $\text{cap}(T_2)$, then looking at how much $T_1 \otimes T_2$ shrinks the determinant of $X \otimes Y$ will give us the required statement. The infimum for $\text{cap}(T_1)$ and $\text{cap}(T_2)$ might not be achieved but that is fine.

Lemma 3.12 *Let T_1 and T_2 be completely positive operators where T_1 acts on matrices of dimension d_1 and T_2 acts on matrices of dimension d_2 . Then*

$$\text{cap}(T_1 \otimes T_2)^{1/(d_1 d_2)} \leq \text{cap}(T_1)^{1/d_1} \cdot \text{cap}(T_2)^{1/d_2}$$

Proof Suppose X, Y be d_1 and d_2 dimensional matrices respectively s.t. $X, Y \succ 0$ and $\text{Det}(X), \text{Det}(Y) = 1$. Then $X \otimes Y \succ 0$ and $\text{Det}(X \otimes Y) = 1$ as well. Also

$$\text{Det}((T_1 \otimes T_2)(X \otimes Y)) = \text{Det}(T_1(X) \otimes T_2(Y)) = \text{Det}(T_1(X))^{d_2} \cdot \text{Det}(T_2(Y))^{d_1}$$

This proves that $\text{cap}(T_1)^{d_2} \cdot \text{cap}(T_2)^{d_1} \geq \text{cap}(T_1 \otimes T_2)$ (by taking X and Y to be sequences of matrices which approach $\text{cap}(T_1)$ and $\text{cap}(T_2)$, respectively). Taking $1/(d_1 d_2)$ powers on both sides completes the proof of the easy direction. \square

We are now ready to prove Theorem 3.11. The main ingredient is the duality between capacity being 0 and the existence of a scaling to almost doubly stochastic position.

Proof of Theorem 3.11 The easier \leq direction was proved above in Lemma 3.12. Now we prove the harder \geq direction. We can assume wlog that $\text{cap}(T_1) > 0$ and $\text{cap}(T_2) > 0$ (otherwise \geq direction is trivial). Now by the analysis of Algorithm G (Theorem 2.14), we know that for any arbitrary $\epsilon > 0$, there exist invertible matrices $B, C \in M_{d_1}(\mathbb{C})$ and $D, E \in M_{d_2}(\mathbb{C})$ s.t. the scaled operators S_1 and S_2 , defined below, are almost doubly stochastic.

$$\begin{aligned} S_1(X) &= B \cdot T_1(C \cdot X \cdot C^\dagger) \cdot B^\dagger \\ S_2(Y) &= D \cdot T_2(E \cdot Y \cdot E^\dagger) \cdot D^\dagger \end{aligned}$$

More formally, $S_1^*(I) = I, S_2^*(I) = I$ and $\text{tr}[(S_1(I) - I)^2] \leq \epsilon, \text{tr}[(S_2(I) - I)^2] \leq \epsilon$. Now consider the operator $S = S_1 \otimes S_2$. Note that $S(I) = S_1(I) \otimes S_2(I)$ and

$S^*(I) = S_1^*(I) \otimes S_2^*(I)$. Hence $S^*(I) = I$ and $\text{tr}[(S(I) - I)^2] \leq \epsilon' = 2(d_1 + d_2 + \epsilon)\epsilon$. This can be seen from the following chain of inequalities (and equalities):

$$\begin{aligned} \text{tr}[(S(I) - I)^2] &= \text{tr}[(S_1(I) \otimes S_2(I) - I)^2] \\ &= \text{tr}[(S_1(I) \otimes S_2(I) - S_1(I) \otimes I + S_1(I) \otimes I - I)^2] \\ &\leq 2\text{tr}[(S_1(I) \otimes (S_2(I) - I))^2] + 2\text{tr}[(S_1(I) - I \otimes I)^2] \\ &\leq 2(d_1 + \epsilon)\epsilon + 2\epsilon d_2 \\ &= 2(d_1 + d_2 + \epsilon) \cdot \epsilon \end{aligned}$$

The first inequality follows from the following Cauchy–Schwarz inequality for psd matrices:

$$\text{tr}[(M_1 + M_2)^2] \leq 2\text{tr}[M_1^2] + 2\text{tr}[M_2^2]$$

The second inequality follows from $\text{tr}[S_1(I)^2] \leq d_1 + \epsilon$, which follows from the fact that $\text{tr}[S_1(I)] = \text{tr}[S_1^*(I)] = \text{tr}[I_{d_1}] = d_1$ and that $\text{tr}[(S_1(I) - I)^2] \leq \epsilon$, as well as $\text{tr}[(S_2(I) - I)^2] \leq \epsilon$.

Note that the operator S is explicitly given as follows:

$$S(Z) = (B \otimes D)(T_1 \otimes T_2) \left((C \otimes E)Z(C^\dagger \otimes E^\dagger) \right) (B^\dagger \otimes D^\dagger)$$

Hence by Proposition 2.7,

$$\begin{aligned} \text{cap}(S) &= |\text{Det}(B \otimes D)|^2 \cdot |\text{Det}(C \otimes E)|^2 \cdot \text{cap}(T_1 \otimes T_2) \\ &= |\text{Det}(B)|^{2d_2} \cdot |\text{Det}(D)|^{2d_1} \cdot |\text{Det}(C)|^{2d_2} \\ &\quad \cdot |\text{Det}(E)|^{2d_1} \cdot \text{cap}(T_1 \otimes T_2) \end{aligned}$$

By Lemma 3.4, we have that $\text{cap}(S) \geq (1 - \sqrt{d_1 d_2 \epsilon'})^{d_1 d_2}$ and hence

$$\begin{aligned} \text{cap}(T_1 \otimes T_2) &\geq |\text{Det}(B)|^{-2d_2} \cdot |\text{Det}(D)|^{-2d_1} \cdot |\text{Det}(C)|^{-2d_2} \\ &\quad \cdot |\text{Det}(E)|^{-2d_1} \cdot (1 - \sqrt{d_1 d_2 \epsilon'})^{d_1 d_2} \end{aligned} \tag{12}$$

Combining the fact that $S_1^*(I) = I$ and $S_2^*(I) = I$ along with Proposition 2.8, we get that $\text{cap}(S_1), \text{cap}(S_2) \leq 1$. Also by Proposition 2.7, we have that

$$\text{cap}(T_1) = |\text{Det}(B)|^{-2} \cdot |\text{Det}(C)|^{-2} \cdot \text{cap}(S_1) \leq |\text{Det}(B)|^{-2} \cdot |\text{Det}(C)|^{-2} \tag{13}$$

$$\text{cap}(T_2) = |\text{Det}(D)|^{-2} \cdot |\text{Det}(E)|^{-2} \cdot \text{cap}(S_2) \leq |\text{Det}(D)|^{-2} \cdot |\text{Det}(E)|^{-2} \tag{14}$$

Combining Eqs. (12)–(14), we get that

$$\text{cap}(T_1 \otimes T_2)^{1/d_1 d_2} \geq \text{cap}(T_1)^{1/d_1} \cdot \text{cap}(T_2)^{1/d_2} \cdot (1 - \sqrt{d_1 d_2 \epsilon'})$$

Since $\epsilon' = 2(d_1 + d_2 + \epsilon)\epsilon$ can be taken to be arbitrarily close to 0, this completes the proof. \square

4 Bit Complexity Analysis of Algorithm G and Continuity of Capacity

In this section, we will provide the bit complexity analysis of Algorithm G and also provide explicit bounds on the continuity of capacity by a sensitivity analysis of Algorithm G. We start by analyzing a natural iterative sequence associated with an operator T that will be very useful, both in the bit complexity and continuity analysis. Section 4.1 describes how the operators T_j in Algorithm G evolve with respect this iterative sequence. Section 4.2 provides the bit complexity analysis of Algorithm G and Sect. 4.3 provides explicit bounds for continuity of capacity.

Consider the sequence of matrices given by $S_0 = T^*(I)$ (I is of size $n \times n$), and

$$S_j = \begin{cases} T(S_{j-1}^{-1}) & j \text{ odd, } j \geq 1 \\ T^*(S_{j-1}^{-1}) & j \text{ even, } j \geq 1. \end{cases} \tag{15}$$

The next proposition studies the stability properties of this sequence of matrices.

Proposition 4.1 *For a real symmetric positive definite matrix X , let $l(X)$ denote its smallest eigenvalue and $u(X)$ its largest.*

1. *Suppose A is an $n \times n$ real symmetric positive definite matrix with integer entries s.t. the magnitude of any entry of A is at most M . Then $l(A) \geq \frac{1}{(Mn)^{n-1}}$ and $u(A) \leq Mn$.*
2. *Define $\alpha = (M^2n^2m)^{n-1}$. Let T be a completely positive operator whose Kraus operators A_1, \dots, A_m are $n \times n$ integer matrices and each entry of A_i is of magnitude at most M . Also assume that $T(I), T^*(I)$ are both non-singular. Then*

$$(M^2n^2m)I \succeq T(I), T^*(I) \succeq \frac{1}{(M^2n^2m)^{n-1}}I$$

3. *Let T, M, n, m be as before. Let X be a real symmetric matrix. Then $\|T(X)\| \leq \alpha\|X\|$. Also if X is real symmetric positive definite, then $l(T(X)) \geq \alpha^{-1} \cdot l(X)$ and $u(T(X)) \leq \alpha \cdot u(X)$. Similarly, $\|T^*(X)\| \leq \alpha\|X\|, l(T^*(X)) \geq \alpha^{-1} \cdot l(X)$ and $u(T^*(X)) \leq \alpha \cdot u(X)$.*
4. *S_j are real symmetric positive definite matrices for all $0 \leq j \leq t$. Also $l(S_j) \geq \alpha^{-(j+1)}$ and $u(S_j) \leq \alpha^{j+1}$ for all $0 \leq j \leq t$.*
5. *For all $1 \leq k, l \leq n, |S_j(k, l)| \leq \alpha^{j+1}$.*

Proof 1. Let us first prove that $u(A) \leq Mn$.

$$\begin{aligned}
 u(A) &= \max_{v \text{ st. } \|v\|_2=1} \|Av\|_2 \\
 &= \max_{v \text{ st. } \|v\|_2=1} \sqrt{\sum_{k=1}^n \left(\sum_{l=1}^n A_{k,l} v_l\right)^2} \\
 &\leq \max_{v \text{ st. } \|v\|_2=1} \sqrt{\sum_{k=1}^n \left(\sum_{l=1}^n M|v_l|\right)^2} \\
 &= \max_{v \text{ st. } \|v\|_2=1} \sqrt{n} \cdot M \cdot \left(\sum_{l=1}^n |v_l|\right) \\
 &\leq nM
 \end{aligned}$$

The last inequality follows from Cauchy–Schwarz. Now since A is positive definite and the determinant is an integer, $\text{Det}(A) \geq 1$. Then

$$1 \leq \text{Det}(A) \leq u(A)^{n-1}l(A) \leq (Mn)^{n-1}l(A)$$

from which it follows that $l(A) \geq \frac{1}{(Mn)^{n-1}}$.

2. Follows from previous point by noting that each entry of $T(I)$, $T^*(I)$ is an integer of magnitude at most M^2nm and $T(I)$, $T^*(I)$ are both symmetric matrices.
3. We know that

$$\alpha I \succeq T(I), T^*(I) \succeq \alpha^{-1}I$$

Suppose X is a real symmetric matrix satisfying $\|X\| = \beta$. Then

$$\beta I \succeq X \succeq -\beta I$$

Then

$$T(X + \beta I) \succeq 0$$

since T is completely positive, and hence it maps psd matrices to psd matrices. Thus we get

$$T(X) + \beta T(I) = T(X + \beta I) \succeq 0$$

This follows from linearity of T . Then

$$T(X) \succeq -\beta T(I) \succeq -\beta \alpha I$$

Similarly one can prove that $\beta \alpha I \succeq T(X)$ which would imply $\|T(X)\| \leq \alpha \beta = \alpha \|X\|$. Other statements can be proven in a similar manner.

4. We will prove this via induction on j . It is true for S_0 by point 2 in the proposition. Suppose the statement holds for S_{j-1} . Then $S_j = T(S_{j-1}^{-1})$ or $S_j = T^*(S_{j-1}^{-1})$. It does not really matter which case it is for the purpose of this proof. Lets assume the former. Then

$$\begin{aligned} l(S_j) &= l(T(S_{j-1}^{-1})) \\ &\geq \alpha^{-1} \cdot l(S_{j-1}^{-1}) \\ &= \alpha^{-1} \cdot u(S_{j-1})^{-1} \\ &\geq \alpha^{-1} \cdot \alpha^{-j} \\ &= \alpha^{-(j+1)} \end{aligned}$$

The first inequality is by point 3 in the proposition. The second inequality is by induction hypothesis. Similarly, we can also prove that $u(S_j) \leq \alpha^{j+1}$, which would complete the induction step.

- 5.

$$\begin{aligned} |S_j(k, l)| &= |e_k^T S_j e_l| \\ &\leq \|e_k\|_2 \|S_j e_l\|_2 \\ &\leq \alpha^{j+1} \end{aligned}$$

Here e_k and e_l are the standard basis vectors.

□

Define another sequence of matrices defined by $U_0 = S_0$ and

$$U_j = \begin{cases} T(U_{j-1}^{-1}) + \Delta_j & j \text{ odd, } j \geq 1 \\ T^*(U_{j-1}^{-1}) + \Delta_j & j \text{ even, } j \geq 1 \end{cases}$$

where Δ_j 's are small perturbations. We now study the closeness of the sequences $\{U_j\}$ and $\{S_j\}$.

Lemma 4.2 *If $\|\Delta_j\| \leq \frac{1}{2^j \cdot \alpha^{j+1}}$ for all $j \geq 1$, then $l(U_j) \geq \frac{1}{2^j \cdot \alpha^{j+1}}$ and $u(U_j) \leq 2^j \cdot \alpha^{j+1}$, for all $j \geq 0$. Here $\alpha = (M^2 n^2 m)^{n-1}$.*

Proof We will prove this using induction on j . It is true for $j = 0$ since $U_0 = S_0$. Assume that $l(U_j) \geq \frac{1}{2^j \cdot \alpha^{j+1}}$ and $u(U_j) \leq 2^j \cdot \alpha^{j+1}$. Then $U_{j+1} = T(U_j^{-1}) + \Delta_{j+1}$ or $U_{j+1} = T^*(U_j^{-1}) + \Delta_{j+1}$. Suppose it is $T(U_j^{-1}) + \Delta_{j+1}$. The other case is similar.

$$\begin{aligned} l(U_{j+1}) &\geq l(T(U_j^{-1})) - \|\Delta_{j+1}\| \\ &\geq \alpha^{-1} \cdot l(U_j^{-1}) - \|\Delta_{j+1}\| \\ &= \alpha^{-1} \cdot u(U_j)^{-1} - \|\Delta_{j+1}\| \end{aligned}$$

$$\begin{aligned} &\geq \alpha^{-1} \cdot \frac{1}{2^j \alpha^{j+1}} - \frac{1}{2^{j+1} \alpha^{j+2}} \\ &= \frac{1}{2^{j+1} \alpha^{j+2}} \end{aligned}$$

Also

$$\begin{aligned} u(U_{j+1}) &\leq u(T(U_j^{-1})) + \|\Delta_{j+1}\| \\ &\leq \alpha \cdot u(U_j^{-1}) + \|\Delta_{j+1}\| \\ &= \alpha \cdot l(U_j)^{-1} + \|\Delta_{j+1}\| \\ &\leq \alpha \cdot 2^j \alpha^{j+1} + \frac{1}{2^{j+1} \alpha^{j+2}} \\ &\leq 2^{j+1} \cdot \alpha^{j+2} \end{aligned}$$

Note that there is a lot of slack in this part, but it does not matter for our purposes. \square

Lemma 4.3 *The following holds for every integer t : suppose $\|\Delta_j\| \leq \delta$ for all $1 \leq j \leq t$, where $\delta \leq \frac{1}{(2\alpha)^{t+1}}$. Then*

$$\|S_j - U_j\| \leq (2\alpha)^{(2t+1) \cdot (t+1)} \delta$$

for all $0 \leq j \leq t$. Here $\alpha = (M^2 n^2 m)^{n-1}$.

Proof

$$\begin{aligned} \|S_j - U_j\| &= \|T(S_{j-1}^{-1}) - T(U_{j-1}^{-1}) - \Delta_j\| \\ &\leq \|T(S_{j-1}^{-1} - U_{j-1}^{-1})\| + \delta \\ &\leq \alpha \cdot \|S_{j-1}^{-1} - U_{j-1}^{-1}\| + \delta \\ &= \alpha \cdot \|S_{j-1}^{-1}(U_{j-1} - S_{j-1})U_{j-1}^{-1}\| + \delta \\ &\leq \alpha \cdot \|S_{j-1}^{-1}\| \cdot \|S_{j-1} - U_{j-1}\| \cdot \|U_{j-1}^{-1}\| + \delta \\ &\leq 2^{j-1} \alpha^{2j+1} \cdot \|S_{j-1} - U_{j-1}\| + \delta \\ &\leq (2\alpha)^{2t+1} \cdot \|S_{j-1} - U_{j-1}\| + \delta \end{aligned}$$

For the fourth inequality, we used Proposition 4.1 and Lemma 4.2. Note that value of δ is small enough so that conditions of Lemma 4.2 are satisfied. \square

4.1 Evolution of the Operator Through Scaling

In this subsection, we compute succinct expressions for the operators T_j which appear in Algorithm G , together with expressions for their capacity and distance to doubly stochastic. These expressions will involve the matrices S_j defined in the previous subsection. These operators T_j are the intermediary operators in the scaling procedure,

and the succinct expressions will be important for us when approximating capacity, or proving the stability of the capacity of a quantum operator.

Starting from a completely positive operator T which satisfies $T(I) = I$, denote the sequence of operators obtained after right and left normalizations by $\{T_j\}$, i.e., T_j is the operator obtained after j iterations of right and left normalizations (as in Algorithm G). Note that

$$T_j(X) = \begin{cases} T_{j-1} \left(T_{j-1}^*(I)^{-1/2} \cdot X \cdot T_{j-1}^*(I)^{-1/2} \right) & j \text{ odd, } j \geq 1 \\ T_{j-1}(I)^{-1/2} \cdot T_{j-1}(X) \cdot T_{j-1}(I)^{-1/2} & j \text{ even, } j \geq 1 \end{cases}$$

For each j , T_j is an operator scaling of T . So

$$T_j(X) = C_j \cdot T \left(D_j^\dagger \cdot X \cdot D_j \right) \cdot C_j^\dagger \tag{16}$$

for some non-singular matrices C_j and D_j . Let us denote $C_j^\dagger C_j$ by P_j and $D_j^\dagger D_j$ by Q_j . It turns out P_j and Q_j are the matrices that really matter for the purpose of analyzing Algorithm G , and we will see next how these evolve.

When j is odd, $T_j^*(I) = I$. Note that

$$T_j^*(X) = D_j \cdot T^*(C_j^\dagger \cdot X \cdot C_j) \cdot D_j^\dagger \tag{17}$$

So the condition $T_j^*(I) = I$ gives us that $Q_j = T^*(P_j)^{-1}$. Also note that when j is odd, $P_j = P_{j-1}$. When j is even, we have, $T_j(I) = I$, which implies $P_j = T(Q_j)^{-1}$ and also $Q_j = Q_{j-1}$ holds. This can be summarized as follows:

$$(P_j, Q_j) = \begin{cases} (P_{j-1}, T^*(P_{j-1})^{-1}) & j \text{ odd, } j \geq 1 \\ (T(Q_{j-1})^{-1}, Q_{j-1}) & j \text{ even, } j \geq 1 \end{cases}$$

along with $T_0 = T$ and $P_0, Q_0 = I$. Thus we can see that

$$(P_j, Q_j) = \begin{cases} (S_{j-2}^{-1}, S_{j-1}^{-1}) & j \text{ odd, } j \geq 1 \\ (S_{j-1}^{-1}, S_{j-2}^{-1}) & j \text{ even, } j \geq 1 \end{cases} \tag{18}$$

with the convention $S_{-1} = I$. For Algorithm G , what matters is

$$\epsilon_j = \begin{cases} \text{tr} \left[(T_j(I) - I)^2 \right] & j \text{ odd, } j \geq 1 \\ \text{tr} \left[(T_j^*(I) - I)^2 \right] & j \text{ even, } j \geq 1 \end{cases}$$

Since

$$T_j(I) = C_j \cdot T \left(D_j^\dagger D_j \right) \cdot C_j^\dagger = \left(C_j^\dagger \right)^{-1} \cdot P_j \cdot T(Q_j) \cdot C_j^\dagger$$

and

$$T_j^*(I) = D_j \cdot T^*(C_j^\dagger C_j) \cdot D_j^\dagger = \left(D_j^\dagger\right)^{-1} \cdot Q_j \cdot T^*(P_j) \cdot D_j^\dagger$$

we get (because similar matrices have the same trace) that

$$\epsilon_j = \begin{cases} \text{tr} \left[(P_j \cdot T(Q_j) - I)^2 \right] & j \text{ odd, } j \geq 1 \\ \text{tr} \left[(Q_j \cdot T^*(P_j) - I)^2 \right] & j \text{ even, } j \geq 1 \end{cases}$$

and hence

$$\epsilon_j = \begin{cases} \text{tr} \left[(S_{j-2}^{-1} \cdot T(S_{j-1}^{-1}) - I)^2 \right] = \text{tr} \left[(S_{j-2}^{-1} \cdot S_j - I)^2 \right] & j \text{ odd, } j \geq 1 \\ \text{tr} \left[(S_{j-2}^{-1} \cdot T^*(S_{j-1}^{-1}) - I)^2 \right] = \text{tr} \left[(S_{j-2}^{-1} \cdot S_j - I)^2 \right] & j \text{ even, } j \geq 1 \end{cases} \tag{19a}$$

$$\tag{19b}$$

For the computation of capacity, the determinants that matter are

$$d_j = \begin{cases} \text{Det}(T_{j-1}^*(I))^{-1} & j \text{ odd, } j \geq 1 \\ \text{Det}(T_{j-1}(I))^{-1} & j \text{ even, } j \geq 1 \end{cases}$$

Then by Proposition 2.7, we have that

$$\text{cap}(T_r) = \text{cap}(T_0) \cdot \prod_{j=1}^r d_j = \text{cap}(T) \cdot \prod_{j=1}^r d_j$$

By the discussion above and the fact that similar matrices have the same determinants, we get that

$$d_j = \begin{cases} \text{Det} \left(Q_{j-1}^{-1} \right) \cdot \text{Det} \left(T^*(P_{j-1}) \right)^{-1} = \text{Det}(S_{j-3}) \cdot \text{Det}(S_{j-1})^{-1} & j \text{ odd, } j \geq 1 \\ \text{Det} \left(P_{j-1}^{-1} \right) \cdot \text{Det} \left(T(Q_{j-1}) \right)^{-1} = \text{Det}(S_{j-3}) \cdot \text{Det}(S_{j-1})^{-1} & j \text{ even, } j \geq 1 \end{cases}$$

with the convention $S_{-2}, S_{-1} = I$. Thus we get

$$\begin{aligned} \text{cap}(T_r) &= \text{cap}(T) \cdot \prod_{j=1}^r d_j \\ &= \text{cap}(T) \cdot \prod_{j=1}^r \text{Det}(S_{j-3}) \cdot \text{Det}(S_{j-1})^{-1} \\ &= \text{cap}(T) \cdot \text{Det}(S_{-2}) \cdot \text{Det}(S_{-1}) \cdot \text{Det}(S_{r-2})^{-1} \cdot \text{Det}(S_{r-1})^{-1} \\ &= \text{cap}(T) \cdot \text{Det}(S_{r-2})^{-1} \cdot \text{Det}(S_{r-1})^{-1} \end{aligned} \tag{20}$$

4.2 Bit Complexity Analysis of Algorithm G

We are now ready to analyze the bit complexity of Algorithm G . We will prove that if one runs Algorithm G while truncating the numbers to an appropriate polynomial number of bits there is essentially no change in the convergence and required number of iterations. We will call the algorithm working with truncated inputs Algorithm G' .

Given a matrix A , let $\text{Trn}(A)$ be the matrix obtained by truncating the entries of A up to $P(n, \log(M))$ bits after the decimal point. $P(n, \log(M))$ is a polynomial which we will specify later. Note that $\|A - \text{Trn}(A)\|_\infty \leq 2^{-P(n, \log(M))}$. We now describe Algorithm G' , which is the variant of Algorithm G with truncation.

Input: Completely positive operator T given in terms of Kraus operators $A_1, \dots, A_m \in \mathbb{Z}^{n \times n}$. Each entry of A_i has absolute value at most M .

Output: Is T rank non-decreasing?

1. Check whether $T(I)$ and $T^*(I)$ are singular. If any one of them are singular, then output that the operator is rank-decreasing, otherwise proceed to step 2.
2. Let $U_{-1} = S_{-1} = I$ and $U_0 = S_0 = T^*(I)$.
For($j = 1$ to t):

$$U_j = \begin{cases} \text{Trn}(T(U_{j-1}^{-1})), & \text{if } j \text{ odd and} \\ \text{Trn}(T^*(U_{j-1}^{-1})), & \text{if } j \text{ even.} \end{cases}$$

Let $\tilde{\epsilon}_j = \text{tr} \left[\left(U_{j-2}^{-1} \cdot U_j - I \right)^2 \right]$.
3. Check whether $\min\{\tilde{\epsilon}_j : 1 \leq j \leq t\} \leq 1/6n$. If yes, then output that the operator is rank non-decreasing; otherwise, output rank-decreasing.

Algorithm 2: Algorithm G' (Algorithm G with truncation)

The parameter t will be chosen as before: $t = 2 + 36n \cdot \log(1/f(n, M))$. We now show that throughout the iterations, distances to double stochasticity is essentially the same for the original and truncated algorithms G and G' .

Lemma 4.4 *For an appropriate choice of $P(n, \log(M))$, $|\epsilon_j - \tilde{\epsilon}_j| \leq 1/12n$, for $1 \leq j \leq t = 2 + 36n \cdot \log(1/f(n, M))$.*

Here $\epsilon_j = \text{ds}(T_j)$ as defined in Algorithm G . Also note that by Eqs. (19a) and (19b),

$$\epsilon_j = \text{tr} \left[\left(S_{j-2}^{-1} \cdot S_j - I \right)^2 \right]$$

Let us first prove the correctness of Algorithm G' assuming Lemma 4.4. As before, we can assume that $T(I)$ and $T^*(I)$ are both non-singular. If $\min\{\tilde{\epsilon}_j : 1 \leq j \leq t\} \leq 1/6n$, then by Lemma 4.4,

$$\min\{\epsilon_j : 1 \leq j \leq t\} \leq 1/6n + 1/12n \leq 1/3n$$

Hence by Theorem 2.12, T is rank non-decreasing. Now assume, for the reverse direction, that T is rank non-decreasing. Then by the analysis in Sect. 2.2, $\min\{\epsilon_j : 1 \leq j \leq t\} \leq 1/12n$. Hence by Lemma 4.4

$$\min\{\tilde{\epsilon}_j : 1 \leq j \leq t\} \leq 1/12n + 1/12n = 1/6n$$

This proves the correctness of Algorithm G' .

Proof of Lemma 4.4 $\epsilon_j = \text{tr} \left[\left(S_{j-2}^{-1} \cdot S_j - I \right)^2 \right]$ and $\tilde{\epsilon}_j = \text{tr} \left[\left(U_{j-2}^{-1} \cdot U_j - I \right)^2 \right]$.

Let $M_j = S_{j-2}^{-1} \cdot S_j$ and $N_j = U_{j-2}^{-1} \cdot U_j$. Then

$$\|M_j\|_\infty = \|S_{j-2}^{-1} \cdot S_j\|_\infty \leq \|S_{j-2}^{-1}\| \cdot \|S_j\| \leq \alpha^{2j}$$

where the last inequality follows from Proposition 4.1. Also

$$\|N_j\|_\infty = \|U_{j-2}^{-1} \cdot U_j\|_\infty \leq \|U_{j-2}^{-1}\| \cdot \|U_j\| \leq 2^{2j-2} \cdot \alpha^{2j}$$

where the last inequality follows from Lemma 4.2 and the fact that $P(n, \log(M))$ will be chosen to be a large enough polynomial so that the perturbations Δ_j in the lemma satisfy the condition $\|\Delta_j\| \leq 2^{-P(n, \log(M))} \leq \frac{1}{2^j \cdot \alpha^{j+1}}$. Now

$$\begin{aligned} |\epsilon_j - \tilde{\epsilon}_j| &\leq |\text{tr}(M_j^2 - N_j^2)| + 2|\text{tr}(M_j - N_j)| \\ &= |\text{tr}[(M_j + N_j)(M_j - N_j)]| + 2|\text{tr}(M_j - N_j)| \\ &\leq (2\alpha)^{2j} \cdot \sum_{k,l=1}^n |(M_j - N_j)(k, l)| + 2n\|M_j - N_j\| \\ &\leq (2\alpha)^{2j} \cdot n \cdot \sqrt{\sum_{k,l=1}^n |(M_j - N_j)(k, l)|^2} + 2n\|M_j - N_j\| \\ &\leq (2\alpha)^{2j} \cdot n \cdot \sqrt{n} \cdot \|M_j - N_j\| + 2n\|M_j - N_j\| \end{aligned}$$

In the first equality, we used the fact that $\text{tr}(M_j N_j) = \text{tr}(N_j M_j)$. In the second inequality, we used the fact that the maximum magnitude of any entry in $M_j + N_j$ is bounded by $(2\alpha)^{2j}$ and that $|\text{tr}(M_j - N_j)|$ is upper bounded by $n\|M_j - N_j\|$. The third inequality is just Cauchy–Schwarz, and the fourth is the fact that Frobenius norm of a matrix is upper bounded by $n\|M_j - N_j\|^2$. Let us try to upper bound $\|M_j - N_j\|$ now.

$$\begin{aligned} \|M_j - N_j\| &= \|S_{j-2}^{-1} \cdot S_j - U_{j-2}^{-1} \cdot U_j\| \\ &= \|S_{j-2}^{-1} \cdot S_j - S_{j-2}^{-1} \cdot U_j + S_{j-2}^{-1} \cdot U_j + U_{j-2}^{-1} \cdot U_j\| \\ &\leq \|S_{j-2}^{-1} \cdot S_j - S_{j-2}^{-1} \cdot U_j\| + \|S_{j-2}^{-1} \cdot U_j + U_{j-2}^{-1} \cdot U_j\| \\ &\leq \|S_{j-2}^{-1}\| \cdot \|S_j - U_j\| + \|S_{j-2}^{-1} - U_{j-2}^{-1}\| \cdot \|U_j\| \\ &= \|S_{j-2}^{-1}\| \cdot \|S_j - U_j\| + \|S_{j-2}^{-1} (U_{j-2} - S_{j-2}) U_{j-2}^{-1}\| \cdot \|U_j\| \\ &\leq \|S_{j-2}^{-1}\| \cdot \|S_j - U_j\| + \|S_{j-2}^{-1}\| \cdot \|U_{j-2} - S_{j-2}\| \cdot \|U_{j-2}^{-1}\| \cdot \|U_j\| \\ &\leq \alpha^{j-1} \cdot (2\alpha)^{(2t+1)\cdot(t+1)}\delta + \alpha^{j-1} \cdot (2\alpha)^{(2t+1)\cdot(t+1)}\delta \\ &\quad \cdot 2^{j-2}\alpha^{j-1} \cdot 2^j\alpha^{j+1} \\ &\leq (2\alpha)^{10t^2} \cdot \delta \end{aligned}$$

The second last inequality follows by application of Proposition 4.1 and Lemmas 4.2 and 4.3. Also δ here is at most $n \cdot 2^{-P(n, \log(M))}$ since $\|\Delta_j\|_\infty \leq 2^{-P(n, \log(M))}$.

Now $\alpha = (M^2 n^2 m)^{n-1} = \exp(\Theta(n \log(n) \log(M)))$ (since $m \leq n^2$). Hence

$$\begin{aligned} (2\alpha)^{10t^2} &= \exp(\Theta(n \log(n) \log(M) \cdot t^2)) \\ &= \exp\left(\Theta\left(n^3 \log(n) \log(M) \log^2(1/f(n, M))\right)\right) \end{aligned}$$

Hence choosing

$$P(n, \log(M)) = n^4 \log(M) \log^2(1/f(n, M))$$

suffices to get

$$|\epsilon_j - \tilde{\epsilon}_j| \leq 1/12n$$

We proved in Theorem 2.20 that $\log(1/f(n, M))$ is poly($n, \log(M)$), so $P(n, \log(M))$ is also a polynomial in n and $\log(M)$. □

4.3 Continuity of Capacity

In this section, we prove the continuity of capacity. We prove the following theorem:

Theorem 4.5 *Suppose A_1, \dots, A_m and B_1, \dots, B_m are two tuples of $n \times n$ matrices such that the bit complexity of elements of A_i 's is b and $\|A_i - B_i\| \leq \delta$ for all i . Let T_A*

be the operator defined by A_1, \dots, A_m and T_B be the operator defined by B_1, \dots, B_m . Then there exists a polynomial $P(n, b, \log(m))$ s.t. if $\delta \leq \exp(-P(n, b, \log(m)))$, then $\text{cap}(T_A) > 0$ implies $\text{cap}(T_B) > 0$. Furthermore,

$$\left(1 - \frac{P(n, b, \log(m))}{\log(1/\delta)^{1/3}}\right) \leq \frac{\text{cap}(T_A)}{\text{cap}(T_B)} \leq \left(1 + \frac{P(n, b, \log(m))}{\log(1/\delta)^{1/3}}\right)$$

The fact that capacity is continuous is already mentioned in [37] and can be proved by other methods. But here we provide explicit bounds on the continuity parameters. Recall that $\text{Fixed}(T, \epsilon)$ (defined in Sect. 3.2) is the set of Hermitian positive-definite matrices C which are ϵ -approximate fixed points of the operator $X \rightarrow T(T^*(X)^{-1})^{-1}$, i.e., satisfy the following condition:

$$\text{tr} \left[\left(C \cdot T^* \left(T(C)^{-1} \right) - I \right)^2 \right] \leq \epsilon$$

The main insight is that by the analysis of Algorithm G , since $\text{cap}(T_A) > 0$, there exists a $C \in \text{Fixed}(T_A, \epsilon)$ with low spectral norm. Since T_B is close to T_A , $C \in \text{Fixed}(T_B, \epsilon')$ for ϵ' close to ϵ and then the proof is finished by applying Lemma 3.9.

Proof of Theorem 4.5 The proof of Theorem 2.14 can be modified to prove the following: there exists a polynomial $Q(n, b)$ s.t. for all $\epsilon > 0$, if we run $t = Q(n, b)/\epsilon$ iterations of Algorithm G starting from $T_0 = T_A$ satisfying $\text{cap}(T_A) > 0$, then for some $1 \leq j \leq t$, $\epsilon_j = \text{ds}(T_j) \leq \epsilon$. Essentially, at each step, capacity increases by roughly $\exp(\Omega(\epsilon))$ if $\epsilon_j > \epsilon$, capacity is lower bounded by $\exp(-Q(n, b))$ initially and upper bounded by 1 always.

By Eqs. (19a) and (19b), we know that

$$\epsilon_j = \text{tr} \left[\left(S_{j-2}^{-1} S_j - I \right)^2 \right]$$

where $\{S_i\}$ is the sequence of matrices given by $S_0 = T_A^*(I)$, and

$$S_i = \begin{cases} T_A(S_{i-1}^{-1}) & i \text{ odd}, i \geq 1 \\ T_A^*(S_{i-1}^{-1}) & i \text{ even}, i \geq 1 \end{cases}$$

Suppose $1 \leq r \leq t$ be such that $\epsilon_r \leq \epsilon$. Wlog assume that r is odd. Then $S_r = T_A \left(T_A^* \left(S_{r-2}^{-1} \right)^{-1} \right)$. $\epsilon_r \leq \epsilon$ implies that S_{r-2}^{-1} is an ϵ -approximate fixed point of the operator $X \rightarrow T_A \left(T_A^*(X)^{-1} \right)^{-1}$. Hence by Lemma 3.9,

$$(1 - \sqrt{n\epsilon})^n \cdot \frac{\text{Det} \left(T_A \left(S_{r-2}^{-1} \right) \right)}{\text{Det} \left(S_{r-2}^{-1} \right)} \leq \text{cap}(T_A) \leq \frac{\text{Det} \left(T_A \left(S_{r-2}^{-1} \right) \right)}{\text{Det} \left(S_{r-2}^{-1} \right)} \tag{21}$$

We will prove now that S_{r-2}^{-1} is also an ϵ' -approximate fixed point of the operator $X \rightarrow T_B (T_B^*(X)^{-1})^{-1}$ for an appropriate choice of ϵ' . Let us denote S_{r-2}^{-1} by C . By an application of Proposition 4.1, it follows that the lowest and highest eigenvalues of C , $l(C)$ and $u(C)$ satisfy

$$\frac{1}{l(C)}, u(C) \leq \exp(O(n \cdot b \cdot r \cdot \log(nm))) \leq \exp(Q_1(n, b, \log(m))/\epsilon) \tag{22}$$

where $Q_1(n, b, \log(m))$ is another polynomial s.t. $Q_1(n, b, \log(m)) = O(Q(n, b) \cdot n \cdot b \cdot \log(nm))$. Let D be an arbitrary matrix. Then

$$\begin{aligned} \|T_A(D) - T_B(D)\| &= \left\| \sum_{i=1}^m A_i D A_i^\dagger - \sum_{i=1}^m B_i D B_i^\dagger \right\| \\ &\leq \sum_{i=1}^m \|A_i D A_i^\dagger - B_i D B_i^\dagger\| \\ &\leq \sum_{i=1}^m \left(\|A_i D A_i^\dagger - A_i D B_i^\dagger\| + \|A_i D B_i^\dagger - B_i D B_i^\dagger\| \right) \\ &\leq \sum_{i=1}^m \left(\|A_i\| \cdot \|D\| \cdot \|A_i^\dagger - B_i^\dagger\| + \|A_i - B_i\| \cdot \|D\| \cdot \|B_i^\dagger\| \right) \\ &\leq 2m \cdot (n \cdot \exp(b) + \delta) \cdot \|D\| \cdot \delta \\ &= \exp(Q_2(n, b, \log(m))) \cdot \|D\| \cdot \delta \end{aligned} \tag{23}$$

The first two inequalities are just the triangle inequality. The third inequality follows from submultiplicativity of the spectral norm. The fourth inequality follows from the fact that

$$\|A_i\| \leq n \cdot \|A_i\|_\infty \leq n \cdot \exp(b)$$

and that

$$\|B_i^\dagger\| = \|B_i\| \leq \|A_i\| + \|A_i - B_i\| \leq \|A_i\| + \delta$$

Similarly, we have that

$$\|T_A^*(D) - T_B^*(D)\| \leq \exp(Q_2(n, b, \log(m))) \cdot \|D\| \cdot \delta \tag{24}$$

We will now upper bound $\|T_A (T_A^*(C)^{-1}) - T_B (T_B^*(C)^{-1})\|$:

$$\begin{aligned} &\|T_A (T_A^*(C)^{-1}) - T_B (T_B^*(C)^{-1})\| \\ &\leq \|T_A (T_A^*(C)^{-1}) - T_B (T_A^*(C)^{-1})\| + \|T_B (T_A^*(C)^{-1}) - T_B (T_B^*(C)^{-1})\| \end{aligned}$$

$$\begin{aligned} &\leq \|T_A (T_A^*(C)^{-1}) - T_B (T_A^*(C)^{-1})\| + \|T_A (T_A^*(C)^{-1} - T_B^*(C)^{-1}) \\ &\quad - T_B (T_A^*(C)^{-1} - T_B^*(C)^{-1})\| \\ &\quad + \|T_A (T_A^*(C)^{-1} - T_B^*(C)^{-1})\| \end{aligned} \tag{25}$$

Equation (22) along with Proposition 4.1 implies that the lowest and highest eigenvalues of $T_A^*(C)$ satisfy the following:

$$\frac{1}{l(T_A^*(C))}, u(T_A^*(C)) \leq \exp(Q_4(n, b, \log(m))/\epsilon) \tag{26}$$

Now applying equation (23) with $D = T_A^*(C)^{-1}$ along with equation (26) gives us the following:

$$\|T_A (T_A^*(C)^{-1}) - T_B (T_A^*(C)^{-1})\| \leq \exp(Q_5(n, b, \log(m))/\epsilon) \cdot \delta \tag{27}$$

Now we will upper bound $\|T_A^*(C)^{-1} - T_B^*(C)^{-1}\|$.

$$\begin{aligned} \|T_A^*(C)^{-1} - T_B^*(C)^{-1}\| &= \|T_A^*(C)^{-1}T_B^*(C)^{-1} (T_A^*(C) - T_B^*(C))\| \\ &\leq \|T_A^*(C)^{-1}\| \cdot \|T_B^*(C)^{-1}\| \cdot \|T_A^*(C) - T_B^*(C)\| \\ &\leq \exp(Q_6(n, b, \log(m))/\epsilon) \cdot \delta \end{aligned} \tag{28}$$

The last inequality follows from Eqs. (26) and (24). Note that we need $\delta \leq \exp(-P(n, b, \log(m)))$ for a sufficiently large polynomial P here to upper bound $\|T_B^*(C)^{-1}\|$ via Eqs. (26) and (24). Now applying equation (23) with $D = T_A^*(C)^{-1} - T_B^*(C)^{-1}$ along with equation (28) gives us that

$$\begin{aligned} &\|T_A (T_A^*(C)^{-1} - T_B^*(C)^{-1}) - T_B (T_A^*(C)^{-1} - T_B^*(C)^{-1})\| \\ &\leq \exp(Q_7(n, b, \log(m))/\epsilon) \cdot \delta^2 \end{aligned} \tag{29}$$

We are left to upper bound $\|T_A (T_A^*(C)^{-1} - T_B^*(C)^{-1})\|$. This follows from Proposition 4.1 and equation (28).

$$\|T_A (T_A^*(C)^{-1} - T_B^*(C)^{-1})\| \leq \exp(Q_8(n, b, \log(m))/\epsilon) \cdot \delta \tag{30}$$

Combining Eqs. (25), (27), (29), and (30), we get the following

$$\|T_A (T_A^*(C)^{-1}) - T_B (T_B^*(C)^{-1})\| \leq \exp(Q_9(n, b, \log(m))/\epsilon) \cdot \delta \tag{31}$$

Let us denote the matrix $D_1 = C \cdot T_A (T_A^*(C)^{-1})$ and $D_2 = C \cdot T_B (T_B^*(C)^{-1})$. D_1 determines whether C is an approximate fixed point of T_A and D_2 determines whether C is an approximate fixed point of T_B .

$$\begin{aligned} \|D_1 - D_2\| &\leq \|C\| \cdot \|T_A(T_A^*(C)^{-1}) - T_B(T_B^*(C)^{-1})\| \\ &\leq \exp(Q_{10}(n, b, \log(m))/\epsilon) \cdot \delta \end{aligned} \tag{32}$$

The second inequality follows from Eqs. (22) and (31). We also have the following elementary inequality:

$$\text{tr}[(D_1 - I)^2] - \text{tr}[(D_2 - I)^2] \leq (\|D_1\| + \|D_2\| + 2n) \cdot \|D_1 - D_2\| \tag{33}$$

The above inequality implies that C is an ϵ' -approximate fixed point of T_B for

$$\epsilon' = \epsilon + \exp(Q_{11}(n, b, \log(m))/\epsilon) \cdot \delta$$

We can now choose $\epsilon = \frac{1}{2(n+1)}$. Then as long as $\delta \leq \frac{\exp(-Q_{11}(n, b, \log(m))/\epsilon)}{2(n+1)}$, then C is a $1/(n + 1)$ -approximate fixed point of the operator $X \rightarrow T_B(T_B^*(X)^{-1})^{-1}$ and by Lemma 3.9, T_B is rank non-decreasing and hence $\text{cap}(T_B) > 0$. This proves first part of the theorem.

For the second part of the theorem, since C is an ϵ' -approximate fixed point of T_B for

$$\epsilon' = \epsilon + \exp(Q_{11}(n, b, \log(m))/\epsilon) \cdot \delta$$

by Lemma 3.9, we have that

$$(1 - \sqrt{n\epsilon'})^n \cdot \frac{\text{Det}(T_B(C))}{\text{Det}(C)} \leq \text{cap}(T_B) \leq \frac{\text{Det}(T_B(C))}{\text{Det}(C)} \tag{34}$$

Also note that

$$\begin{aligned} \|T_A(C) \cdot T_B(C)^{-1} - I\| &= \|(T_A(C) - T_B(C)) \cdot T_B(C)^{-1}\| \\ &\leq \|T_A(C) - T_B(C)\| \cdot \|T_B(C)^{-1}\| \\ &\leq \exp(Q_{12}(n, b, \log(m))/\epsilon) \cdot \delta \end{aligned}$$

by Eq. (23) and Proposition 4.1. Hence

$$\begin{aligned} (1 - \exp(Q_{12}(n, b, \log(m))/\epsilon) \cdot \delta)^n &\leq \frac{\text{Det}(T_A(C))}{\text{Det}(T_B(C))} \\ &\leq (1 + \exp(Q_{12}(n, b, \log(m))/\epsilon) \cdot \delta)^n \end{aligned} \tag{35}$$

Now, combining Eqs. (21) and (21), we get that

$$(1 - \sqrt{n\epsilon})^n \cdot \frac{\text{Det}(T_A(C))}{\text{Det}(T_B(C))} \leq \frac{\text{cap}(T_A)}{\text{cap}(T_B)} \leq \frac{1}{(1 - \sqrt{n\epsilon'})^n} \cdot \frac{\text{Det}(T_A(C))}{\text{Det}(T_B(C))}$$

Combining this with Eq. (35) gives us

$$\begin{aligned} (1 - \sqrt{n\epsilon})^n \cdot (1 - \exp(Q_{12}(n, b, \log(m))/\epsilon) \cdot \delta)^n &\leq \frac{\text{cap}(T_A)}{\text{cap}(T_B)} \\ &\leq \frac{1}{(1 - \sqrt{n\epsilon'})^n} \cdot (1 + \exp(Q_{12}(n, b, \log(m))/\epsilon) \cdot \delta)^n \end{aligned}$$

Now choose $\epsilon = 2 \cdot \max \{Q_{11}(n, b, \log(m)), Q_{12}(n, b, \log(m))\} / \log(1/\delta)$. This ensures that $\epsilon' \leq 2\epsilon$ and elementary calculations can then finish the proof of the theorem. □

5 Computing the Capacity of a Quantum Operator

In this section, we show how algorithm G can be used to compute an approximation to the capacity of any quantum operator. For simplicity of exposition, we will present in this section an analysis of convergence of algorithm G without truncation. Afterward, in Sect. 5.1, we show how to adapt the analysis of algorithm G to handle the truncation. This corresponds to the analysis of algorithm G' in the previous section.

We begin with the following lemma, which is an adaptation of Lemma 3.10 from [60].

Lemma 5.1 *Let x_1, \dots, x_n be positive real numbers such that $\sum_{i=1}^n x_i = n$ and $\sum_{i=1}^n (x_i - 1)^2 = \alpha$. Then*

$$\prod_{i=1}^n x_i \leq \begin{cases} \exp(-\alpha/6), & \text{if } \alpha \leq 1, \\ \exp\left(\frac{-1}{6}\right), & \text{otherwise.} \end{cases}$$

Proof We have two cases to analyze:

Case 1: $\alpha \leq 1$.

In this case, by using the inequality $1 + t \leq \exp\left(t - \frac{t^2}{2} + \frac{t^3}{3}\right)$, which holds for all $t \in \mathbb{R}$, we have:

$$\begin{aligned} \prod_{i=1}^n x_i &= \prod_{i=1}^n [1 + (x_i - 1)] \\ &\leq \prod_{i=1}^n \exp\left((x_i - 1) - \frac{(x_i - 1)^2}{2} + \frac{(x_i - 1)^3}{3}\right) \end{aligned}$$

$$\begin{aligned}
 &= \exp\left(\sum_{i=1}^n (x_i - 1) - \frac{1}{2} \cdot \sum_{i=1}^n (x_i - 1)^2 + \frac{1}{3} \cdot \sum_{i=1}^n (x_i - 1)^3\right) \\
 &\leq \exp\left(-\frac{1}{2} \cdot \alpha + \frac{1}{3} \cdot \alpha^{3/2}\right) \leq \exp\left(-\frac{\alpha}{6}\right)
 \end{aligned}$$

Where in the last inequalities we used the fact that $\sum_{i=1}^n (x_i - 1)^3 \leq \left(\sum_{i=1}^n (x_i - 1)^2\right)^{3/2}$ and $\alpha^{3/2} \leq \alpha$.

Case 2: $\alpha > 1$.

Consider the function $f(\lambda) = \prod_{i=1}^n (1 + \lambda(x_i - 1))$. We will prove that f is a decreasing function of λ when $\lambda \in [0, 1]$. In that case

$$\prod_{i=1}^n x_i = \prod_{i=1}^n [1 + (x_i - 1)] = f(1) \leq f\left(\frac{1}{\sqrt{\alpha}}\right) = \prod_{i=1}^n \left[1 + \frac{x_i - 1}{\sqrt{\alpha}}\right] \leq \exp\left(-\frac{1}{6}\right)$$

where the last inequality follows from Case 1. Now let us prove that f is decreasing.

$$f'(\lambda) = f(\lambda) \cdot \left(\sum_{j=1}^n \frac{x_j - 1}{1 + \lambda(x_j - 1)}\right) \leq f(\lambda) \cdot \left(\sum_{j=1}^n (x_j - 1)\right) = 0$$

This completes the proof. □

As a corollary of Lemma 5.1, we obtain the following quantitative progress measure toward computing capacity:

Lemma 5.2 (Quantitative Progress) *Let T be a right (left) normalized quantum operator such that $ds(T) = \alpha$. Additionally, let \tilde{T} be the left (right) normalization of operator T . Then,*

$$\text{cap}(\tilde{T}) \geq \begin{cases} \text{cap}(T) \cdot \exp(\alpha/6), & \text{if } \alpha \leq 1, \\ \text{cap}(T) \cdot \exp\left(\frac{1}{6}\right), & \text{otherwise.} \end{cases}$$

Proof Suppose T is right normalized and \tilde{T} is the left normalization of T . Proposition 2.7 tells us that

$$\text{cap}(\tilde{T}) = \det(T(I))^{-1} \cdot \text{cap}(T).$$

Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of $T(I)$. As $n = \text{tr}(T(I)) = \sum_{i=1}^n \lambda_i$ and $\alpha =$

$ds(T) = \text{tr}[(T(I) - I)^2] = \sum_{i=1}^n (\lambda_i - 1)^2$, the conditions of Lemma 5.1 apply and we have

$$\det(T(I)) = \prod_{i=1}^n \lambda_i \leq \begin{cases} \exp(-\alpha/6), & \text{if } \alpha \leq 1, \\ \exp\left(\frac{-1}{6}\right), & \text{otherwise.} \end{cases}$$

This implies the desired lower bounds on $\text{cap}(\tilde{T})$. Since the case where T is left normalized is analogous, we omit the argument. □

We now state a slight modification of Algorithm G, with a view toward computing the capacity of a quantum operator.

Input: Quantum operator T given in terms of Kraus operators $A_1, \dots, A_m \in \mathbb{Z}^{n \times n}$ and approximation parameter $\epsilon > 0$. Each entry of A_i has absolute value at most M .

Output: $\text{cap}(T)$ with multiplicative error of $(1 \pm \epsilon)$.

1. Check if $T(I)$ and $T^*(I)$ are singular. If any one of them is singular, then output $\text{cap}(T) = 0$, otherwise proceed to step 2.
2. Alternately perform right and left normalizations on $T = T_0$ for t steps. Let T_j be the operator after j steps. Also let $\epsilon_j = \text{ds}(T_j)$. Go to step 3.
3. If $\epsilon_j \leq \frac{\epsilon^2}{n^3}$ for any $0 \leq j \leq t$, go to step 4. Otherwise, output $\text{cap}(T) = 0$.
4. For the smallest $0 \leq j \leq t$ such that $\epsilon_j \leq \frac{\epsilon^2}{n^3}$, output

$$\text{cap}(T) = \prod_{i=0}^{j-1} \det(R_i), \text{ where } R_i = \begin{cases} T_i(I), & \text{if } i \text{ is odd,} \\ T_i^*(I), & \text{otherwise.} \end{cases}$$

Algorithm 3: Algorithm G, computing capacity

Theorem 5.3 *Let T be a completely positive operator, whose Kraus operators are given by $n \times n$ integer matrices A_1, \dots, A_m , such that each entry of A_i has absolute value at most M . Algorithm G when applied for $t = \frac{n^3}{\epsilon^2} \cdot (1 + 10n^2 \log(Mn))$ steps approximates $\text{cap}(T)$ within a multiplicative factor of $1 \pm \epsilon$.*

Proof If either $T(I)$ or $T^*(I)$ is singular, then T decreases the rank of I . When $T(I)$ is singular, rank-decreasing follows by definition. When $T^*(I)$ is singular, one way to see it is that $\text{Im}(A_i) \subseteq \text{Im}(T^*(I))$ for all i . Since $\text{Im}(A_i A_i^\dagger) = \text{Im}(A_i)$, we get that $\text{Im}(T(I)) \subseteq \text{Im}(T^*(I))$ and hence $T(I)$ is singular. Therefore, the algorithm is correct on step 1, by outputting $\text{cap}(T) = 0$.

If $T(I)$ and $T^*(I)$ are both non-singular, it is easy to verify that $T_j(I)$ and $T_j^*(I)$ will remain non-singular for all j and hence step 3 is well defined.

We now divide the proof into two cases:

Case 1: T is rank-decreasing.

In this case, since right and left normalizations do not change the property of being rank-decreasing, we have $\text{cap}(T_j) = 0$ for all $0 \leq j \leq t$. Hence, Lemma 3.4 implies that $\text{ds}(T_j) > \frac{\epsilon^2}{n^3}$ for all $0 \leq j \leq t$. In this case, step 3 of Algorithm G will correctly output $\text{cap}(T) = 0$.

Case 2: T is rank non-decreasing.

In this case, we will show that there must exist $0 \leq j \leq t$ such that $\epsilon_j \leq \frac{\epsilon^2}{n^3}$. Assume the contrary, for the sake of contradiction. By Theorem 2.20, we know that $\text{cap}(T_1) \geq \exp(-10n^2 \log(Mn))$. Also Proposition 2.8 implies that $\text{cap}(T_j) \leq 1$ for all j . However by the assumption that $\epsilon_j > \frac{\epsilon^2}{n^3}$, Lemma 5.2 implies that $\text{cap}(T_{j+1}) \geq \exp(\epsilon^2/n^3) \cdot \text{cap}(T_j)$ for all $0 \leq j \leq t$. Hence, we obtain:

$$1 \geq \text{cap}(T_{t+1}) \geq \exp\left(\frac{t\epsilon^2}{n^3}\right) \cdot \text{cap}(T_1) \geq \exp\left(\frac{t\epsilon^2}{n^3}\right) \cdot \exp(-10n^2 \log(Mn))$$

Plugging in $t = \frac{n^3}{\epsilon^2} \cdot (1 + 10n^2 \log(Mn))$ gives us the required contradiction.

Now that we showed the existence of $0 \leq j \leq t$ such that $\epsilon_j \leq \frac{\epsilon^2}{n^3}$, we will show that step 4 indeed computes a good approximation to capacity. For the first ϵ_j such that $\epsilon_j \leq \frac{\epsilon^2}{n^3}$, Lemma 3.4 implies that $\text{cap}(T_j) \geq (1 - \sqrt{n\epsilon_j})^n \geq (1 - \epsilon/n)^n \geq 1 - \epsilon$.

Since $\text{cap}(T_j) = \text{cap}(T) \cdot \left(\prod_{i=0}^{j-1} \det(R_i)\right)^{-1}$, we have

$$\text{cap}(T) = \text{cap}(T_j) \cdot \prod_{i=0}^{j-1} \det(R_i).$$

As $1 - \epsilon \leq \text{cap}(T_j) \leq 1$, we obtain the correct approximation. □

5.1 Computing Capacity with Truncation

In this subsection, we analyze the computation of capacity when we truncate the input matrices. This analysis will be similar to the one in Sect. 4.2. We begin with some intuition on why truncation works.

Notice that to approximate the capacity, all we need is to compute the determinants of the matrices U_i in Algorithm 2. The U_i 's are the truncations of the matrices S_i from Eq. (15), the latter matrices being important as they describe the scaled operator T_j in terms of the original operator T , see Eqs. (16–18). The reason why truncating the input works is because the eigenvalues of U_i are very similar to the eigenvalues of S_i . Therefore, we can show that $\det(S_i) \approx \det(U_i)$. This will imply that the truncated capacity is a good approximation to the actual capacity. The analysis will rely mainly on Lemma 4.3, which gives a good bound on the spectral norm of $S_i - U_i$. Now we state the truncated algorithm.

Given a matrix A , let $\text{Trn}(A)$ be the matrix obtained by truncating the entries of A up to $P(n, 1/\epsilon, \log(M)) = \frac{1}{\epsilon} \cdot (n^{12} \log^4(Mn)) \cdot \log(n^4/\epsilon^2)$ bits after the decimal point.

Input: Quantum operator T given in terms of Kraus operators $A_1, \dots, A_m \in \mathbb{Z}^{n \times n}$ and approximation parameter $\epsilon > 0$. Each entry of A_i has absolute value at most M .

Output: $\text{cap}(T)$ with multiplicative error of $(1 \pm \epsilon)$.

1. Check whether $T(I)$ and $T^*(I)$ are singular. If any one of them is singular, then output $\text{cap}(T) = 0$, otherwise proceed to step 2.
2. Let $U_{-1} = S_{-1} = I_n$ and $U_0 = S_0 = T^*(I)$. Additionally, for $1 \leq j \leq t$, let

$$U_j = \begin{cases} \text{Trn}(T(U_{j-1}^{-1})), & \text{if } i \text{ is odd,} \\ \text{Trn}(T^*(U_{j-1}^{-1})), & \text{otherwise.} \end{cases}$$

Also let $\tilde{\epsilon}_j = \text{tr}[(U_{j-2}^{-1} \cdot U_j - I_n)^2]$. Go to step 3.

3. If $\tilde{\epsilon}_j \leq \frac{\epsilon^2}{4n^3}$ for any $0 \leq j \leq t$, go to step 4. Otherwise, output $\text{cap}(T) = 0$.
4. For the first $0 \leq j \leq t$ such that $\tilde{\epsilon}_j \leq \frac{\epsilon^2}{4n^3}$, output

$$\text{cap}(T) = \det(U_{j-1}) \cdot \det(U_{j-2}).$$

Algorithm 4: Algorithm G with truncation, computing capacity

We now proceed to the analysis of Algorithm 4. In Theorem 5.3, we proved the correctness of Algorithm G without truncation. Thus, to prove correctness of Algorithm 4, it is enough to prove two statements:

1. if $\tilde{\epsilon}_j \leq \epsilon^2/4n^3$, then the operators T_j will satisfy the $\epsilon_j \leq \epsilon^2/n^3$ bounds
3. $\|U_i - S_i\| \leq 2^{-P(n, 1/\epsilon, \log(M))/2}$.

The first item implies that steps 1 to 3 of the algorithm above are correct, and the second item will tell us that step 4 indeed computes an $1 \pm \epsilon$ approximation to capacity. More formally, we have the following theorem.

Theorem 5.4 *Let T be a completely positive operator, whose Kraus operators are given by $n \times n$ matrices $A_1, \dots, A_m \in \mathbb{Z}^{n \times n}$, such that each entry of A_i has absolute value at most M . Algorithm 4, with truncation parameter $P(n, 1/\epsilon, \log(M)) = \frac{1}{\epsilon} \cdot (n^{12} \log^4(Mn)) \cdot \log(n^4/\epsilon^2)$ when applied for $t = \frac{4n^3}{\epsilon^2} \cdot (1 + 10n^2 \log(Mn))$ steps approximates $\text{cap}(T)$ within a multiplicative factor of $1 \pm \epsilon$.*

Proof By applying Lemma 4.4 with parameters t and $P(n, 1/\epsilon, \log(M))$ as above, we get that

$$|\tilde{\epsilon}_i - \epsilon_i| < \frac{\epsilon^2}{n^4} \text{ for all } 0 \leq i \leq t.$$

Therefore, steps 1 to 3 of Algorithm 4 work just as if we had not done any truncation (as in Algorithm 3). This implies that we will always output $\text{cap}(T) = 0$ whenever the operator T is rank-decreasing.

We are now left with the computation of capacity when T is rank non-decreasing, which is done in step 4. By applying Lemma 4.3 with parameter $\delta = 2^{-P(n, 1/\epsilon, \log(M))}$, we get

$$\|U_i - S_i\| \leq 2^{-P(n,1/\epsilon,\log(M))/2}, \text{ for all } 0 \leq i \leq t.$$

Let $0 \leq \mu_{i1} \leq \mu_{i2} \leq \dots \leq \mu_{in}$ be the eigenvalues of U_i and $0 \leq \lambda_{i1} \leq \lambda_{i2} \leq \dots \leq \lambda_{in}$ be the eigenvalues of S_i . From $\|U_i - S_i\| \leq 2^{-P(n,1/\epsilon,\log(M))/2}$ and Lemma 4.2, we have

$$|\mu_{i\ell} - \lambda_{i\ell}| \leq 2^{-P(n,1/\epsilon,\log(M))/2} \leq \frac{\epsilon\mu_{i\ell}}{4n} \Rightarrow \mu_{i\ell} \left(1 - \frac{\epsilon}{4n}\right) \leq \lambda_{i\ell} \leq \mu_{i\ell} \left(1 + \frac{\epsilon}{4n}\right).$$

Hence, we have that

$$\begin{aligned} & \det(S_{j-1}) \cdot \det(S_{j-2}) \\ &= \prod_{i=j-2}^{j-1} \prod_{\ell=1}^n \lambda_{i,\ell} \geq \prod_{i=j-2}^{j-1} \prod_{\ell=1}^n \mu_{i,\ell} \left(1 - \frac{\epsilon}{4nt}\right) \\ &\geq \left(1 - \frac{\epsilon}{4n}\right)^{2n} \cdot \prod_{i=j-2}^{j-1} \prod_{\ell=1}^n \mu_{i,\ell} \\ &\geq (1 - \epsilon/2) \cdot \prod_{i=j-2}^{j-1} \prod_{\ell=1}^n \mu_{i,\ell} = (1 - \epsilon/2) \cdot \det(U_{j-1}) \cdot \det(U_{j-2}) \end{aligned}$$

Similarly, we have that $\det(S_{j-1}) \cdot \det(S_{j-2}) \leq (1 + \epsilon/2) \cdot \det(U_{j-1}) \cdot \det(U_{j-2})$.

As $\tilde{\epsilon}_j \leq \frac{\epsilon^2}{4n^3}$ implies that $\epsilon_j \leq \frac{\epsilon^2}{4n^3} + \frac{\epsilon^2}{n^4}$, by Lemma 3.4 we have $\text{cap}(T_j) \in [(1 - \epsilon/2), 1]$. Thus, Eq. (20) yields

$$\begin{aligned} \text{cap}(T) &= \text{cap}(T_j) \det(S_{j-1}) \det(S_{j-2}) \\ &\Rightarrow (1 - \epsilon/2) \det(S_{j-1}) \det(S_{j-2}) \leq \text{cap}(T) \leq \det(S_{j-1}) \det(S_{j-2}) \\ &\Rightarrow (1 - \epsilon/2)^2 \det(U_{j-1}) \det(U_{j-2}) \leq \text{cap}(T) \\ &\leq (1 + \epsilon/2) \det(U_{j-1}) \det(U_{j-2}). \end{aligned}$$

The inequalities above imply that $\det(U_{j-1}) \cdot \det(U_{j-2})$, that is, the output of Algorithm 4, lies in the interval $[(1 - \epsilon)\text{cap}(T), (1 + \epsilon)\text{cap}(T)]$. \square

6 Conclusion and Open Problems

In this paper, we gave a polynomial time algorithm for computing the non-commutative rank of a symbolic matrix over any subfield of the complex numbers. We stated its different incarnations and implications to the many different areas in which this problem arises (indeed we feel that expositing these many connections, some essential to the present result, may yield better future interaction between them with possible more benefits). We note that our algorithm and the analysis bypasses the need to use any degree bounds at all. We further note again that despite the purely algebraic nature

of the problem our algorithm is purely analytic, generating a sequence of complex matrices and testing its convergence.

We collect now the most obvious directions for future research, some of them already mentioned in the paper.

- Find more applications of this algorithm to optimization problems.
- Can we use these techniques to design an efficient deterministic algorithm for the orbit-closure intersection problem for the left–right action? In terms of invariants, this is equivalent to asking if two tuples of matrices can be separated by invariants of the left–right action (over algebraically closed fields of characteristic 0). More formally given two tuples of matrices, (A_1, \dots, A_m) and (B_1, \dots, B_m) , check whether for all (T_1, \dots, T_m) of arbitrary dimension,

$$\text{Det}(T_1 \otimes A_1 + \dots T_m \otimes A_m) = \text{Det}(T_1 \otimes B_1 + \dots T_m \otimes B_m)$$

The results of [21] give a randomized polynomial time algorithm for this problem (over algebraically closed fields of characteristic 0): just plug in random (T_1, \dots, T_m) of polynomial dimension.

- Find a black-box algorithm for SINGULAR. That is, efficiently produce (deterministically) a polynomial size set \mathcal{S} of tuples of polynomial dimension matrices s.t. for all $L = \sum_{i=1}^m x_i A_i$ s.t. L is non-singular, it holds that for some $(T_1, \dots, T_m) \in \mathcal{S}$,

$$\text{Det}(T_1 \otimes A_1 + \dots T_m \otimes A_m) \neq 0$$

Due to the recent polynomial dimension bounds of [21], it can be proved that a random set \mathcal{S} works. The challenge is to produce it deterministically. As a special case, this captures deterministic parallel algorithms for the decision version of bipartite perfect matching (when A_1, \dots, A_m are elementary matrices $E_{i,j}$ representing the edges of a bipartite graph). So perhaps, techniques from the recent breakthrough work [29] can be useful.

- Explore further the connection between commutative and non-commutative PIT problems. We feel that beyond the many connections between commutative and non-commutative settings that arise here, this different angle of looking at the commutative PIT problem, relating it to its non-commutative cousin, may help in the major quest of finding an efficient deterministic algorithm for it. As mentioned above, this viewpoint has already resulted in a deterministic PTAS for the commutative rank [9].
- We design an efficient algorithm for checking if a completely positive operator is rank-decreasing. Can we do the same for positive operators? Algorithm G in fact works for positive operators as well and all that is needed is to prove an effective lower bound on the capacity $\text{cap}(T)$ of a positive operator T which is rank non-decreasing (similar to Theorem 2.20). It was already proven in [37] that $\text{cap}(T) > 0$ for a positive operator T which is rank non-decreasing.
- Design a strongly polynomial time algorithm for operator scaling. Strongly polynomial time algorithms for matrix scaling were given by [60]. Can they be extended to the operator case?

- Can we compute $(1 + \epsilon)$ approximation to $\text{cap}(T)$ in time $\text{poly}(n, b, \log(1/\epsilon))$? For computing capacity of nonnegative matrices, such algorithms exist. One of the algorithms in [60] has this stronger convergence rate. Also for matrices, capacity can be formulated as a convex program and hence the Ellipsoid algorithm also gives this stronger convergence rate [41].
- Can we design efficient algorithms for testing the null-cone of general quivers? There is reduction from general quivers to Kronecker-quiver or the left–right action (e.g., see [21]) but the reduction is not always efficient. What about the general problem of testing the null-cone of actions of reductive groups?

Acknowledgements We would like to thank Harm Derksen, Pavel Hrubes, Louis Rowen and K. V. Subrahmanyam for helpful discussions. We would also like to thank Oded Regev for suggesting us that operator scaling could be used for approximating capacity. Finally, we thank the anonymous reviewers for a comprehensive reading of the paper and pointing several typographical errors and minor bugs.

A Symbolic Matrices with Polynomial Entries and Non-commutative Rank

In this section, we show how to compute the non-commutative rank of any (not necessarily square) matrix with linear entries over the free skew field $\mathbb{Q}\langle\mathbf{x}\rangle$. This will be achieved in two ways: the first, in Sect. A.2, by reducing this problem to testing singularity of a certain square matrix with linear entries, and the second, in Sect. A.3, by a purely quantum approach which in a sense mimics the reduction from maximum matching to perfect matching.

In fact, we solve a more general problem. Section A.1 starts with a reduction of computing nc-rank of a matrix with *polynomial* entries (given by formulae), to the problem of computing the nc-rank of a matrix with linear entries, via the so-called “Higman’s trick” (Proposition A.2). We give the simple quantitative analysis of this reduction, which as far as we know does not appear in the literature and may be useful elsewhere. This reduction, with the two above, allow computing the non-commutative rank of any matrix in time polynomial in the description of its entries.

A.1 Higman’s Trick

Before stating the full version of the effective Higman trick, we need to define the bit complexity of a formula computing a non-commutative polynomial.

Definition A.1 (*Bit Complexity of a Formula*) Let Φ be a non-commutative formula without divisions such that each of its gates computes a polynomial in $\mathbb{Q}\langle\mathbf{x}\rangle$ (i.e., the inputs to the formula are either rational numbers or non-commutative variables). The *bit complexity* of Φ is the maximum bit complexity of any rational input appearing in the formula Φ .

With this definition in hand, we can state and prove Higman’s trick, which first appeared in [43]. In the proposition below, it will be useful to have the following notation to denote the direct sum of two matrices A and B :

$$A \oplus B = \begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix},$$

where the zero matrices in the top right and bottom left corners are of appropriate dimensions. Before stating and proving Higman’s trick, let us work through a small example which showcases the essence of the trick.

Suppose we want to know the nc-rank of matrix $\begin{pmatrix} 1 & x \\ y & z + xy \end{pmatrix}$. The problem here is that this matrix is not linear, and we need to have a linear matrix. How can we convert this matrix into a linear matrix while preserving the rank, or the complement of the rank? To do this, we need to remove the multiplication happening in $z + xy$.

Notice that the complement of its rank does not change after the following transformation:

$$\begin{pmatrix} 1 & x \\ y & z + xy \end{pmatrix} \mapsto \begin{pmatrix} 1 & x & 0 \\ y & z + xy & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Since the complement of the rank does not change after we perform elementary row or column operations, we can first add $x \cdot$ (third row) to the second row, and then subtract (third column) $\cdot y$ to the second column, to obtain:

$$\begin{pmatrix} 1 & x & 0 \\ y & z + xy & 0 \\ 0 & 0 & 1 \end{pmatrix} \mapsto \begin{pmatrix} 1 & x & 0 \\ y & z + xy & x \\ 0 & 0 & 1 \end{pmatrix} \mapsto \begin{pmatrix} 1 & x & 0 \\ y & z & x \\ 0 & -y & 1 \end{pmatrix}$$

The complement of the rank of this last matrix is the same as the complement of the rank of our original matrix! In particular, if this last matrix is full rank, it implies that our original matrix is also full rank. This is the essence of Higman’s trick. We now proceed to its full version.

Proposition A.2 (Effective Higman’s Trick) *Let $A \in \mathbb{Q}(\mathbf{x})^{m \times n}$ be a matrix where each entry a_{ij} is computed by a non-commutative formula of size $\leq s$ and bit complexity $\leq b$ without divisions. Let k be the total number of multiplication gates used in the computation of the entries of A . There exist matrices $P \in \mathbb{GL}_{m+k}(\mathbb{Q}(\mathbf{x}))$, $Q \in \mathbb{GL}_{n+k}(\mathbb{Q}(\mathbf{x}))$ such that $P(A \oplus I_k)Q$ is a matrix with linear entries and coefficients with bit complexity bounded by b . Moreover, given access to the formulas computing the entries, one can construct P and Q efficiently in time $\text{poly}(m, n, s, b)$. Since P and Q are non-singular matrices, the co-rank and the co-nc-rank of $P(A \oplus I_k)Q$ are the same as the co-rank and the co-nc-rank of A .*

Proof Let $\text{Mult}(a_{ij})$ be the number of multiplication gates in the formula computing entry a_{ij} and

$$T = \sum_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} \text{Mult}(a_{ij}).$$

That is, T is the total number of multiplication gates used to compute all entries of the matrix A .

We prove this proposition by induction on T , for matrices of all dimensions. The base case, when $T = 0$, is trivial, as in this case A itself has linear entries. Suppose now that the proposition is true for all matrices (regardless of their dimensions) which can be computed by formulas using $< T$ multiplication gates.

Let A be our matrix, which can be computed using T multiplications. W.l.o.g., we can assume that $\text{Mult}(a_{mn}) \geq 1$. Then, by finding a multiplication gate in the formula for a_{mn} that has no other multiplication gate as an ancestor, we can write a_{mn} in the form $a_{mn} = a + b \cdot c$, where

$$\text{Mult}(a_{mn}) = \text{Mult}(a) + \text{Mult}(b) + \text{Mult}(c) + 1.$$

Hence, the matrix

$$B = \left(I_{m-1} \oplus \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \right) (A \oplus 1) \left(I_{n-1} \oplus \begin{pmatrix} 1 & 0 \\ -c & 1 \end{pmatrix} \right)$$

is such that

$$b_{ij} = \begin{cases} a_{ij}, & \text{if } i \leq m, j \leq n \text{ and } (i, j) \neq (m, n) \\ a, & \text{if } (i, j) = (m, n) \\ b, & \text{if } (i, j) = (m, n + 1) \\ -c, & \text{if } (i, j) = (m + 1, n) \\ 1, & \text{if } (i, j) = (m + 1, n + 1) \\ 0, & \text{otherwise} \end{cases}$$

Therefore, the number of multiplications needed to compute B is given by

$$\begin{aligned} \sum_{\substack{1 \leq i \leq m+1 \\ 1 \leq j \leq n+1}} \text{Mult}(b_{ij}) &= \left(\sum_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} \text{Mult}(a_{ij}) \right) - \text{Mult}(a_{mn}) \\ &\quad + \text{Mult}(a) + \text{Mult}(b) + \text{Mult}(c) \\ &= T - \text{Mult}(a_{mn}) + \text{Mult}(a) + \text{Mult}(b) + \text{Mult}(c) \\ &= T - 1 \end{aligned}$$

Since B is an $(m + 1) \times (n + 1)$ matrix which can be computed by using a total of $T - 1$ multiplication gates, by the induction hypothesis, there exist $P' \in \text{GL}_{m+1+(T-1)}(\mathbb{Q}\langle \mathbf{x} \rangle) = \text{GL}_{m+T}(\mathbb{Q}\langle \mathbf{x} \rangle)$ and $Q' \in \text{GL}_{n+1+(T-1)}(\mathbb{Q}\langle \mathbf{x} \rangle) = \text{GL}_{n+T}(\mathbb{Q}\langle \mathbf{x} \rangle)$ such that $P'(B \oplus I_{T-1})Q'$ is a linear matrix. Since

$$\begin{aligned} B \oplus I_{T-1} &= \left(I_{m-1} \oplus \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \oplus I_{T-1} \right) (A \oplus I_T) \left(I_{n-1} \oplus \begin{pmatrix} 1 & 0 \\ -c & 1 \end{pmatrix} \oplus I_{T-1} \right) \\ &= R(A \oplus I_T)S, \end{aligned}$$

where $R = \left(I_{m-1} \oplus \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \oplus I_{T-1} \right) \in \text{GL}_{m+T}(\mathbb{Q}\langle \mathbf{x} \rangle)$ and $S = \left(I_{n-1} \oplus \begin{pmatrix} 1 & 0 \\ -c & 1 \end{pmatrix} \oplus I_{T-1} \right) \in \text{GL}_{n+T}(\mathbb{Q}\langle \mathbf{x} \rangle)$, we have that

$$P'(B \oplus I_{T-1})Q' = (P'R)(A \oplus I_T)(SQ').$$

Setting $P = P'R$ and $Q = SQ'$ proves the inductive step and completes the proof. Since we only use subformulas of the formulas computing the entries of A , the bound on the bit complexity does not change. \square

A.2 Classical Reduction

Having shown the effective version of Higman’s trick, we can now compute the nc-rank of a matrix over $\mathbb{Q}\langle \mathbf{x} \rangle$. We begin with a lemma which will tell us that we can reduce the problem of computing the nc-rank of a matrix by testing fullness of a smaller matrix with polynomial entries.

Lemma A.3 (Reduction to Fullness Testing) *Let $M \in \mathbb{F}\langle \mathbf{x} \rangle^{m \times n}$ be any matrix. In addition, let $U = (u_{ij})$ and $V = (v_{ij})$ be generic matrices in new, non-commuting variables u_{ij} and v_{ij} , of dimensions $r \times m$ and $n \times r$, respectively. Then, $\text{nc-rank}(M) \geq r$ iff the matrix UMV is full.*

Proof Since $\text{nc-rank}(M) \geq r$, there exists an $r \times r$ minor of M of full rank. Let Q be such a minor of M . W.l.o.g.,²⁸ we can assume that Q is the $[r] \times [r]$ principal minor of M . Hence, we have that

$$UMV = \begin{pmatrix} U_1 & U_2 \end{pmatrix} \begin{pmatrix} Q & M_2 \\ M_3 & M_4 \end{pmatrix} \begin{pmatrix} V_1 \\ V_2 \end{pmatrix},$$

where U_1 and V_1 are $r \times r$ matrices and the others are matrices with the proper dimensions.

Letting $U' = (I_r \ 0)$ and $V' = \begin{pmatrix} I_r \\ 0 \end{pmatrix}$, the equality above becomes:

$$U'MV' = Q.$$

As

$$r \geq \text{nc-rank}(UMV) \geq \text{nc-rank}(U'MV') = \text{nc-rank}(Q) = r,$$

we obtain that UMV is full, as we wanted. Notice that the second inequality comes from the fact that rank does not increase after restrictions of the new variables. \square

²⁸ Notice that we can make the following assumption just to simplify notation. In actuality, we do not know where the full rank minor is located in M .

Notice that we do not know the rank $\text{nc-rank}(M)$ a priori. Therefore, our algorithm will try all possible values of $r \in [n]$ and output the maximum value of r for which we find a full matrix.

For each $r \times r$ matrix UMV , we can use the effective Higman’s trick to convert UMV into a $s \times s$ matrix with linear entries. With this matrix, we can use the truncated Gurvits’ algorithm to check whether the matrix we just obtained is full. Since we have this test, we will be able to output the correct rank. Algorithm 5 is the precise formulation of the procedure just described.

Input: $M \in \mathbb{Q}(\mathbf{x})^{m \times n}$ s.t. each entry of M is a polynomial computed by a formula of size bounded by s and bit complexity bounded by b .
 Output: $\text{nc-rank}(M)$

1. For $1 \leq r \leq \min(m, n)$, let U_r and V_r be $r \times m$ and $n \times r$ generic matrices in new, non-commuting variables $u_{ij}^{(r)}, v_{ij}^{(r)}$.
2. Let $M_r = U_r M V_r$.
3. Apply the effective Higman’s trick to M_r to obtain a matrix N_r with linear entries on the variables x_1, \dots, x_m and $u_{ij}^{(r)}, v_{ij}^{(r)}$.
4. Use Algorithm G' to test whether N_r is full rank.
5. Output the maximum value of r for which N_r is full rank.

Algorithm 5: Non-commutative Rank Algorithm

Theorem A.4 *Let $M \in \mathbb{Q}(\mathbf{x})^{m \times n}$ be s.t. each entry of M is a polynomial computed by a formula of size bounded by s and bit complexity bounded by b . There exists a deterministic algorithm that finds the non-commutative rank of M in time $\text{poly}(m, n, s, b)$.*

Proof To prove this theorem, it is enough to show that Algorithm 5 is correct and it runs with the desired runtime.

Without loss of generality, we can assume that $n \leq m$. Therefore we have that $\text{nc-rank}(M) \leq n$. By Lemma A.3, if $r \leq \text{nc-rank}(M)$, then matrix M_r will be of full rank (and therefore will not have a shrunk subspace, by Theorem 1.4). Since $M_r = U_r M V_r$, from the formulas computing the entries of M , we obtain formulas of size at most $2smn$ computing the entries of M_r . Moreover, the bit complexities of these formulas will still be bounded by b , as multiplication by generic matrices do not mix any of the polynomials of M .

By Proposition A.2 and the fact that the size of the formulas computing the entries of M_r are bounded by $2smn$, we have that N_r is a linear matrix of dimensions $(k + r) \times (k + r)$, where $k \leq 2s(mn)^2$ and the bit complexity of the coefficients bounded by b . Moreover, $N_r = P(M_r \oplus I_k)Q$ implies that N_r is full if, and only if, M_r is full, which is true if, and only if, $\text{nc-rank}(M) \geq r$.

Now, by Theorem 1.1, we have a deterministic polynomial time algorithm to determine whether N_r is full rank. If $r \leq \text{nc-rank}(M)$, N_r will be full, and the maximum such r will be exactly when $r = \text{nc-rank}(M)$. Therefore, by outputting the maximum r for which N_r we compute $\text{nc-rank}(M)$. This proves that our algorithm is correct. Notice that the runtime is polynomial in the input size, as we perform at most n applications of the Higman trick and of Algorithm G' . This completes the proof. \square

A.3 The Quantum Reduction

Here we present a different reduction from computing non-commutative rank to fullness testing from a quantum viewpoint. We will only work with square matrices though. As we saw, by Higman’s trick, we can assume the matrices to be linear. So we are given a matrix $L = \sum_{i=1}^m x_i A_i \in M_n(\mathbb{F}\langle \mathbf{x} \rangle)$. A combination of Theorems 1.4 and 1.17 shows that $\text{nc-rank}(L) \leq r$ iff the operator defined by A_1, \dots, A_m is $n - r$ -rank-decreasing. So we just want to check whether a completely positive operator is c -rank-decreasing, and we will do this by using an algorithm for checking if an operator is rank-decreasing as a black box using the following lemma:

Lemma A.5 *Let $T : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ be a completely positive operator. Define an operator $\bar{T} : M_{n+c-1}(\mathbb{C}) \rightarrow M_{n+c-1}(\mathbb{C})$ as follows:*

$$\bar{T} \left(\begin{bmatrix} X_{1,1} & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix} \right) = \begin{bmatrix} T(X_{1,1}) + \text{tr}(X_{2,2})I_n & 0 \\ 0 & \text{tr}(X_{1,1})I_{c-1} \end{bmatrix}$$

Here $X_{1,1}, X_{1,2}, X_{2,1}, X_{2,2}$ are $n \times n, n \times c - 1, c - 1 \times n, c - 1 \times c - 1$ matrices, respectively. Then \bar{T} is completely positive and T is c -rank-decreasing iff \bar{T} is rank-decreasing. Note that we are considering $c \leq n$.

Proof A well known characterization due to Choi [12] states that \bar{T} is completely positive iff $\sum_{i,j=1}^{n+c-1} E_{i,j} \otimes \bar{T}(E_{i,j})$ is psd. Here $E_{i,j}$ is the matrix with 1 at i, j position and 0 everywhere else. Now

$$\begin{aligned} & \bar{T}(E_{i,j}) \\ &= \begin{cases} \begin{bmatrix} T(E_{i,j}) & 0 \\ 0 & I_{c-1} \end{bmatrix} & 1 \leq i = j \leq n \\ \begin{bmatrix} T(E_{i,j}) & 0 \\ 0 & 0 \end{bmatrix} & 1 \leq i, j \leq n, i \neq j \\ \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} & 1 \leq i \leq n, n+1 \leq j \leq n+c-1 \text{ or } n+1 \leq i \leq n+c-1, 1 \leq j \leq n \\ \begin{bmatrix} I_n & 0 \\ 0 & 0 \end{bmatrix} & n+1 \leq i = j \leq n+c-1 \\ \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} & n+1 \leq i, j \leq n+c-1, i \neq j \end{cases} \end{aligned}$$

From here, it is easy to verify that $\sum_{i,j=1}^{n+c-1} E_{i,j} \otimes \bar{T}(E_{i,j})$ is psd given that $\sum_{i,j=1}^n E_{i,j} \otimes T(E_{i,j})$ is psd. Now suppose that \bar{T} is rank-decreasing. This can only happen if $X_{1,1} = 0$ or $X_{2,2} = 0$, otherwise

$$\bar{T} \left(\begin{bmatrix} X_{1,1} & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix} \right) = \begin{bmatrix} T(X_{1,1}) + \text{tr}(X_{2,2})I_n & 0 \\ 0 & \text{tr}(X_{1,1})I_{c-1} \end{bmatrix}$$

is full rank. If $X_{1,1} = 0$, then

$$\begin{bmatrix} 0 & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix}$$

can be psd (and Hermitian) only if $X_{1,2} = X_{2,1} = 0$. In this case a $c - 1$ ranked matrix is mapped to rank n matrix. So $X_{2,2}$ has to be zero. Then again by the psd condition $X_{1,2} = X_{2,1} = 0$. So

$$\bar{T} \left(\begin{bmatrix} X_{1,1} & 0 \\ 0 & 0 \end{bmatrix} \right) = \begin{bmatrix} T(X_{1,1}) & 0 \\ 0 & \text{tr}(X_{1,1})I_{c-1} \end{bmatrix}$$

and $X_{1,1} \neq 0$ and

$$\text{Rank} \left(\begin{bmatrix} X_{1,1} & 0 \\ 0 & 0 \end{bmatrix} \right) > \text{Rank} \left(\begin{bmatrix} T(X_{1,1}) & 0 \\ 0 & \text{tr}(X_{1,1})I_{c-1} \end{bmatrix} \right)$$

Hence $\text{Rank}(T(X_{1,1})) \leq \text{Rank}(X_{1,1}) - c$. This proves one direction. Now suppose that T is c -rank-decreasing and $\text{Rank}(T(X)) \leq \text{Rank}(X) - c$, then

$$\text{Rank} \left(\bar{T} \left(\begin{bmatrix} X & 0 \\ 0 & 0 \end{bmatrix} \right) \right) = \text{Rank} \left(\begin{bmatrix} T(X) & 0 \\ 0 & \text{tr}(X)I_{c-1} \end{bmatrix} \right) < \text{Rank} \left(\begin{bmatrix} X & 0 \\ 0 & 0 \end{bmatrix} \right)$$

This proves the lemma. □

Remark A.6 This seems to be the “quantum” analogue of obtaining a maximum matching oracle based on a perfect matching oracle: add $c-1$ dummy vertices to both sides of the bipartite graph and connect them to everything. Then the new graph has a perfect matching iff the original graph had a matching of size $\geq n - c + 1$.

Remark A.7 Here we didn’t specify a set of Kraus operators for the operator \bar{T} which seem to be needed to run Algorithms 1 and 2 but Kraus operators can be obtained by looking at the eigenvectors of $\sum_{i,j=1}^{n+c-1} E_{i,j} \otimes \bar{T}(E_{i,j})$. Alternatively Algorithms 1 and 2 can also be interpreted as acting directly on the Choi-Jamiolkowski state of \bar{T} , i.e., $\sum_{i,j=1}^{n+c-1} E_{i,j} \otimes \bar{T}(E_{i,j})$.

References

1. B. Adsul, S. Nayak, and K. V. Subrahmanyam. A geometric approach to the Kronecker problem ii : rectangular shapes, invariants of $n \times n$ matrices, and a generalization of the Artin-Procesi theorem. *Manuscript, available at <http://www.cmi.ac.in/~kv/ANS10.pdf>*, 2010.
2. N. Alon. Combinatorial Nullstellensatz. *Combinatorics, Probability and Computing*, 8(1-2):7–29, 1999.

3. S. A. Amistur and J. Levitzki. Minimal identities for algebras. *Proceedings of the American Mathematical Society*, 1:449–463, 1950.
4. S. Amitsur. Rational identities and applications to algebra and geometry. *Journal of Algebra*, 3:304–359, 1966.
5. M. D. Atkinson. Spaces of matrices with several zero eigenvalues. *Bulletin of the London Mathematical Society*, 12(89–95), 1980.
6. M. D. Atkinson and S. Lloyd. Large spaces of matrices of bounded rank. *Quarterly Journal of Math. Oxford*, 31:253–262, 1980.
7. L. B. Beasley. Nullspaces of spaces of matrices of bounded rank. *Current trends in matrix theory*, 1987.
8. S. J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.
9. M. Bläser, G. Jindal, and A. Pandey. Greedy strikes again: A deterministic PTAS for commutative rank of matrix spaces. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 79. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
10. A. Bogdanov and H. Wee. More on noncommutative polynomial identity testing. *Computational Complexity*, pages 92–99, 2005.
11. M. Bürgin and J. Draisma. The hilbert null-cone on tuples of matrices and bilinear forms. *Math Z*, 254(4):785–809, 2006.
12. M. Choi. Completely positive linear maps on complex matrices. *Linear Algebra and Its Applications*, pages 285–290, 1975.
13. P. M. Cohn. The embedding of firs in skew fields. *Proceedings of the London Mathematical Society*, 23:193–213, 1971.
14. P. M. Cohn. The word problem for free fields. *The Journal of Symbolic Logic*, 38(2):309–314, 1973.
15. P. M. Cohn. The word problem for free fields: A correction and an addendum. *Journal of Symbolic Logic*, 40(1):69–74, 1975.
16. P. M. Cohn. *Skew Fields, Theory of General Division Rings*. Cambridge University Press, 1995.
17. P. M. Cohn and C. Reutenauer. On the construction of the free field. *International journal of Algebra and Computation*, 9(3):307–323, 1999.
18. D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties, and Algorithms*. Undergraduate Texts in Mathematics. Springer, New York, third edition edition, 2007.
19. H. Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.
20. H. Derksen and G. Kemper. *Computational Invariant Theory*, volume 130. Springer-Verlag, Berlin, 2002.
21. H. Derksen and V. Makam. Polynomial degree bounds for matrix semi-invariants. *Advances in Mathematics*, 310:44–63, 2017.
22. H. Derksen and J. Weyman. Semi-invariants of quivers and saturation for Littlewood-Richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000.
23. J. Dieudonné. Sur une généralisation du groupe orthogonal à quatre variables. *Arch. Math.*, 1:282–287, 1949.
24. M. Domokos and A. N. Zubkov. Semi-invariants of quivers as determinants. *Transformation Groups*, 6(1):9–24, 2001.
25. Z. Dvir and A. Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. *SIAM J. Comput.*, 2006.
26. J. Edmonds. Systems of distinct representatives and linear algebra. *Journal of research of the National Bureau of Standards*, 71(241–245), 1967.
27. J. Edmonds. Submodular functions, matroids, and certain polyhedra. *Lectures, Calgary International Symposium on Combinatorial Structures*, 1969.
28. D. Eisenbud and J. Harris. Vector spaces of matrices of low rank. *Advances in Math*, 70:135–155, 1988.
29. S. A. Fenner, R. Gurjar, and T. Thierauf. Bipartite perfect matching is in quasi-NC. *STOC*, 2016.
30. M. Forbes and A. Shpilka. Explicit noether normalization for simultaneous conjugation via polynomial identity testing. *RANDOM*, 2013.
31. M. Forbes and A. Shpilka. Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. *FOCS*, pages 243–252, 2013.
32. E. Formanek. Generating the ring of matrix invariants. *Ring Theory*, pages 73–82, 1986.

33. M. Fortin and C. Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. 2004.
34. A. Garg, L. Gurvits, R. Oliveira, and A. Wigderson. Algorithmic and optimization aspects of Brascamp-Lieb inequalities, via operator scaling. *Geometric and Functional Analysis*, 28(1):100–145, 2018.
35. B. Gelbord and R. Meshulam. Spaces of singular matrices and matroid parity. *European Journal of Combinatorics*, 23(4):389–397, 2002.
36. I. Gelfand, S. Gelfand, V. Retakh, and R. Wilson. Quasideterminants. [arXiv:math/0208146](https://arxiv.org/abs/math/0208146), 2002.
37. L. Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004.
38. L. Gurvits. Hyperbolic polynomials approach to Van der Waerden/Schrijver-Valiant like conjectures: sharper bounds, simpler proofs and algorithmic applications. *STOC*, pages 417–426, 2006.
39. L. Gurvits and A. Samorodnitsky. A deterministic polynomial-time algorithm for approximating mixed discriminant and mixed volume. *STOC*, 2000.
40. L. Gurvits and A. Samorodnitsky. A deterministic algorithm approximating the mixed discriminant and mixed volume, and a combinatorial corollary. *Discrete Computational Geometry*, 27(531-550), 2002.
41. L. Gurvits and P. N. Yianilos. The deflation-inflation method for certain semidefinite programming and maximum determinant completion problems. *Technical Report, NECI*, 1998.
42. W. Haken. Theorie der Normalflächen. *Acta Math*, 105:245–375, 1961.
43. G. Higman. *Units in group rings*. PhD thesis, Balliol College, 1940.
44. D. Hilbert. Über die vollen Invariantensysteme. *Math. Ann.*, 42:313–370, 1893.
45. P. Hrubes and A. Wigderson. Non-commutative arithmetic circuits with division. *ITCS*, 2014.
46. P. Hrubes, A. Wigderson, and A. Yehudayoff. Relationless completeness and separations. In *Computational Complexity (CCC), 2010 IEEE 25th Annual Conference on*, pages 280–290. IEEE, 2010.
47. P. Hrubes, A. Wigderson, and A. Yehudayoff. Non-commutative circuits and the sum-of-squares problem. *Journal of the American Mathematical Society*, 24(3):871–898, 2011.
48. L.-K. Hua. Some properties of a sfield. *Proceedings of National Academy of Sciences USA*, 35:533–537, 1949.
49. L. Hyafil. On the parallel evaluation of multivariate polynomials. *SIAM Journal on Computing*, 8(2):120–123, 1979.
50. G. Ivanyos, M. Karpinski, Y. Qiao, and M. Santha. Generalized Wong sequences and their applications to Edmonds’ problems. *Journal of Computer and System Sciences*, 81(7):1373–1386, 2015.
51. G. Ivanyos, Y. Qiao, and K. Subrahmanyam. Non-commutative Edmonds’ problem and matrix semi-invariants. *Computational Complexity*, 26(3):717–763, 2017.
52. G. Ivanyos, Y. Qiao, and K. V. Subrahmanyam. Constructive noncommutative rank computation in deterministic polynomial time over fields of arbitrary characteristics. *Computational Complexity*, 27(4):561–593, December 2018.
53. V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13:1–46, 2004.
54. D. S. Kaliuzhnyi-Verbovetskyi and V. Vinnikov. Noncommutative rational functions, their difference-differential calculus and realizations. *Multidimensional Systems and Signal Processing*, 23(1-2):49–77, 2010.
55. N. Kayal and N. Saxena. Polynomial identity testing for depth 3 circuits. *Computational Complexity*, 2007.
56. A. D. King. Moduli of representations of finite dimensional algebras. *The Quarterly Journal of Mathematics*, 45(4):515–530, 1994.
57. A. Klivans and D. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Annual STOC*, 2001.
58. H. Kraft and C. Procesi. Classical invariant theory, a primer. <https://math.unibas.ch/uploads/x4epersdb/files/primernew.pdf>, 1996.
59. N. Limaye, G. Malod, and S. Srinivasan. Lower bounds for non-commutative skew circuits. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 22, page 22, 2015.
60. N. Linial, A. Samorodnitsky, and A. Wigderson. A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. *STOC*, pages 644–652, 1998.
61. L. Lovasz. On determinants, matchings, and random algorithms. *Fundamentals of Computation Theory*, pages 565–574, 1979.

62. L. Lovasz. Selecting independent lines from a family of lines in a space. *Acta Sci. Math.*, 42(121–131), 1980.
63. L. Lovasz. Singular spaces of matrices and their application in combinatorics. *Bulletin of the Brazilian Mathematical Society*, 20:87–99, 1989.
64. P. Malcolmson. A prime matrix ideal yields a skew field. *Journal of the London Mathematical Society*, 18(221–233), 1978.
65. K. Mulmuley. Geometric complexity theory V: Equivalence between blackbox derandomization of polynomial identity testing and derandomization of noether’s normalization lemma. *FOCS*, pages 629–638, 2012.
66. N. Nisan. Lower bounds for non-commutative computation. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 410–418. ACM, 1991.
67. V. L. Popov. The constructive theory of invariants. *Izvestiya: Mathematics*, 19(2):359–376, 1982.
68. C. Procesi. The invariant theory of $n \times n$ matrices. *Advances in Mathematics*, 19:306–381, 1976.
69. M. O. Rabin. Recursive unsolvability of group theoretic problems. *Annals of Mathematics*, 67(172–194), 1958.
70. R. Rado. A theorem on independence relations. *Quarterly Journal of Math. Oxford*, 13:83–89, 1942.
71. R. Raz and A. Shpilka. Deterministic polynomial identity testing in non commutative models. *Computational Complexity*, 14:1–19, 2005.
72. J. P. Razmyslov. Trace identities of full matrix algebras over a field of characteristic zero. *Mathematics of the USSR-Izvestiya*, 8(4):727, 1974.
73. C. Reutenauer. Inversion height in free fields. *Selecta Mathematica*, 2(1):93–109, 1996.
74. L. H. Rowen. *Polynomial identities in ring theory*. Academic Press, New York, 1980.
75. S. Saraf and I. Volkovich. Black-box identity testing of depth 4 multilinear circuits. In *Proceedings of the 43rd annual STOC*, 2011.
76. A. Schofield and M. V. den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001.
77. A. Shpilka and A. Yehudayoff. *Arithmetic Circuits: A Survey of Recent Results and Open Questions*, volume 5. NOW, Foundations and Trends in Theoretical Computer Science, 2010.
78. R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35:876–879, 1964.
79. V. Strassen. Vermeidung von Divisionen. *Journal für Reine Angew. Math*, 264:182–202, 1973.
80. L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Ankit Garg¹ · Leonid Gurvits² · Rafael Oliveira³ · Avi Wigderson⁴

Leonid Gurvits
l.n.gurvits@gmail.com

Rafael Oliveira
rafael@cs.toronto.edu

Avi Wigderson
avi@math.ias.edu

¹ Microsoft Research India, Bengaluru, India

² Department of Computer Science, The City College of New York, New York, USA

³ Simons Institute, Berkeley, USA

⁴ Institute for Advanced Study, Princeton, USA