Convergence of Markovian Semigroups with Applications to Quantum Information Theory

Daniel Stilck França

Zentrum Mathematik M5

Technische Universität München

A thesis submitted for the degree of

Doctor rerum naturalium

July 2018

TECHNISCHE UNIVERSITÄT MÜNCHEN
ZENTRUM MATHEMATIK

Lehrstuhl für mathematische Physik

Convergence of Markovian Semigroups with Applications to Quantum Information Theory

Daniel Stilck França

Vollständiger Abdruck der von der Fakultät für Mathematik der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigten Dissertation.

Vorsitzender: Prof. Dr. Christian Kühn

Prüfer der Dissertation:

1. Prof. Dr. Michael M. Wolf
2. Prof. Dr. Holger Boche
3. Prof. Dr. Matthias Christandl
   (schriftliche Beurteilung)

Die Dissertation wurde am 07.06.2018 bei der Technischen Universität München eingereicht und durch die Fakultät für Mathematik am 04.11.2018 angenommen.

Repito que el menos importante de sus recuerdos era más minucios y más vivo que nuestra percepción de un goce físico o de un tormento físico. Hacia el Este, en un trecho no amanzanado, había casas nuevas, desconocidas. Funes las imaginaba negras, compactas, hechas de tiniebla homogénea; en esa dirección volvía la cara para dormir. También solía imaginarse en el fondo del río, mecido y anulado por la corriente. Había aprendido sin esfuerzo el inglés, el francés, el portugués, el latín. Sospecho, sin embargo, que no era muy capaz de pensar. Pensar es olvidar diferencias, es generalizar, abstraer. En el abarrotado mundo de Funes no había sino detalles, casi inmediatos.
– *Jorge Luis Borges, Funes El Memorioso*

De pensamento em chamas
Inspiração
Arte de criar o saber
Arte, descoberta, invenção
Theoría em grego quer dizer
O ser em contemplação
Cântico dos cânticos
Quântico dos quânticos
Sei que a arte é irmã da ciência
Ambas filhas de um deus fugaz
Que faz num momento e no mesmo momento desfaz
Esse vago deus por trás do mundo
Por detrás do detrás
– *Gilberto Gil, Quanta*

Acknowledgements

I would like to start by thanking my advisor, Prof. Dr. Michael M. Wolf. Thank you for supporting me all these years and guiding me through my journey in quantum information theory!

I would like to thank Prof. Dr. Boche and Prof. Dr. Christandl for being on my committee. It is an honor to be evaluated by such distinguished researchers. Thank you for taking the time to read my work!

I would also like to thank Prof. Dr. Michael Wolf and Prof. Dr. Robert König for creating such a great work environment. I always felt very free and encouraged to pursue my ideas and interests. I made many great friends during my years at M5 and will have many fond memories of this period. Furthermore, I have to thank all my colleagues at M5 during the last years for being so fun to be around, whether we were discussing some new paper or naked mole rats. Special thanks go to Anna. Sharing an office with you certainly made my routine much more fun. I believe the time we got an email asking why we were laughing so much summarizes the experience of all these years pretty well. I would like to thank my coauthors Anna, Andreas, Alex and Michael. I know it took a lot of patience to check my morally correct proofs again and again, but I still hope that you had at least half the fun I had while working on our projects. I also have to thank Stefan for proofreading this thesis and finding many typos!

I am grateful to Prof. Brandão for letting me visit his group at Caltech for a couple of months. It was an exciting experience!

I should also thank all people that helped me to navigate the bureaucracy during my Ph.D., with a special thanks to the TopMath team, Carl-Friedrich and Agnieszka, and Wilma. I am also eternally grateful to all the people of the DAAD, who gave me the fantastic opportunity to come to Germany and study here.

I cannot thank my family enough. You were incredibly supportive during my whole life. Even having an ocean between us for the past seven years, I felt your encouragement and warmth daily! To my parents, Heitor and Veronika, I want to say that I highly value the education you gave me and how you always encouraged me to pursue my goals, be the best version of myself and all your love and understanding. I could not have wished for better parents! To my sister, Nina, I want to say it is a great pleasure to see that my little sister has become a strong, incredible woman who is much wiser than I am and to whom I can always turn for advice, as I did many times while writing this. I must thank all my uncles and aunts. Your support during the last years was incredibly valuable, and I feel so lucky for having a family that encouraged my intellectual pursuits since I was a kid. I would also like to thank all my friends. I am fortunate enough to be able to say that I have uncountably many good friends and, as we all know, trying to list them all is pointless. I will always forget at least one. But I do believe that they all know how important they are to me and how valuable their company was during good and bad times.

Finally, I would like to dedicate this thesis to all the teachers in my life, not necessarily only the ones I met in a classroom. Being the fourth generation of teachers in my family, I highly value the fact that many of you took an active interest in my development as a thinker and fostered my curiosity and inquisitive spirit. Thank you very much!

List of contributed articles

Core articles as principal author

I) A. Müller-Hermes and D. Stilck França.
Sandwiched Renyi Convergence for Quantum Evolutions.
Quantum, 2, 55, 2018.

II) D. Stilck França.
Perfect Sampling for Quantum Gibbs States.
Quantum Information and Computation, 18(5), 2018.

Further articles as principal author under review

III) A. Bluhm and D. Stilck França.
Dimensionality reduction of SDPs through sketching.
arXiv:1707.09863, 2017.
Submitted to *Linear Algebra and its Applications* August 2017.

IV) A.K. Hashagen and D. Stilck França.
Approximate Randomized Benchmarking for Finite Groups.
arXiv:1803.03621, 2018.
Submitted to *Journal of Physics A: Mathematical and Theoretical* March 2018.

Articles as co-author

V) A. Müller-Hermes, D. Stilck França, and M. M. Wolf.
Relative entropy convergence for depolarizing channels.
Journal of Mathematical Physics, 57(2), 2016.

The author of the present dissertation would like to reiterate that the core articles of this dissertation are I and II and that he is the principal author of articles I, II, III and IV. Articles III and IV are still being reviewed at the moment of the submission of this dissertation. Article V is included due to its close connection to the other articles included in this thesis, in particular article I, which was inspired by it. Alexander Müller-Hermes is the principal author of that article.

Contents

1 Introduction **1**

1.1 Summary and Discussion of Results . . . 2

2 Notation and Preliminaries **7**

2.1 Finite Dimensional Quantum Mechanics . . . 7

2.1.1 States and Measurements . . . 7

2.1.2 Time Evolutions . . . 8

3 Quantum Dynamical Semigroups **11**

3.1 Convergence of Semigroups . . . 11

3.1.1 The Phase Estimation Algorithm . . . 13

3.1.2 Important Examples . . . 14

3.1.3 Covariant Channels and Twirling . . . 16

3.1.4 Classical Markov Chains . . . 17

4 Distance Measures and Convergence of Quantum Markov Chains **21**

4.1 Standard Distance Measures . . . 21

4.1.1 Noncommutative lp Spaces . . . 22

4.2 Entropic Measures . . . 23

4.2.1 Sandwiched Rényi divergences . . . 23

5 Functional Inequalities and Convergence of Semigroups **25**

5.1 Convergence in Variance: Spectral Gap . . . 25

5.2 Convergence in Information-Theoretic Divergence . . . 26

5.2.1 Convergence of Rényi Divergences . . . 27

5.2.2 Logarithmic Sobolev Inequalities and Hypercontractivity . . . 28

5.2.3 Relations between the constants . . . 29

5.2.4 Tensorization . . . 31

6 Applications of Convergence Bounds **33**

6.1 Dissipative preparation of states . . . 33

6.2 Entropic inequalities and Capacity Bounds . . . 33

6.3 Randomized Benchmarking . . . 35

Bibliography **39**

A Articles as co-author **41**

A.1 Relative entropy convergence for depolarizing channels . . . 41

A.1.1 Main result . . . 42

B Core Articles 61
B.1 Sandwiched Renyi Convergence for Quantum Evolutions . . . 61
B.1.1 Main Results . . . 62
B.1.2 Applications . . . 62
B.1.3 Individual Contribution . . . 63
B.2 Perfect Sampling for quantum Gibbs States . . . 101
B.2.1 The Algorithm . . . 102
B.2.2 Stability Results . . . 102
B.2.3 Individual Contribution . . . 104

C Further articles as principal author under review 135
C.1 Dimensionality reduction of SDPs through sketching . . . 135
C.1.1 SDPs and Quantum Information . . . 136
C.1.2 Sketching the HS scalar product . . . 136
C.1.3 Sketching SDPs to approximate the optimum value . . . 137
C.1.4 Linear Matrix Inequality feasibility problems . . . 138
C.1.5 Complexity and Memory Considerations . . . 139
C.1.6 Individual Contribution . . . 139
C.2 Approximate Randomized Benchmarking for Finite Groups . . . 172
C.2.1 Main Results . . . 173
C.2.2 Applications . . . 174
C.2.3 Individual Contribution . . . 174

1

Introduction

Quantum computers are poised to revolutionize the way we process information. There is now a great scientific effort to build them and understand how to harness their power to solve different problems. One of the big obstacles to building large-scale quantum computer is the fact that any physical system is subject to noise. Therefore, understanding how to characterize, quantify and suppress noise is of central importance to the development of quantum technologies. On the other hand, in some situations, noise and randomness can actually be helpful to solve some computational task. The main example are Monte Carlo algorithms which have a wide applicability. Thus, understanding how we can develop and adapt such techniques to quantum computers is a very promising avenue for developing new quantum algorithms that profit from noise. This thesis is concerned with the two sides of this coin. On one hand, we develop several different techniques to quantify how noisy a given quantum system is for a given task. On the other hand, we discuss how noisy processes can be used to develop algorithms to sample from quantum states that are analogous to Markov chain Monte Carlo techniques in the classical setting. The main technical tool we use are quantum dynamical semigroups and we study their convergence in detail. These mathematical objects can be used to model quantum systems under memoryless noise and provide a good approximate description for many physical systems of relevance.

We start by giving a brief summary of the two core contributed articles and of the individual contributions of the author. We then proceed to give brief summaries of the articles of which the present is the principal author, but are still under review and one for which the present author is not the principal author. After that, we lay the mathematical and technical foundations related to the articles, besides setting our notation. We start by introducing the basic notions of finite-dimensional quantum mechanics, followed by some basic facts on and examples of quantum dynamical semigroups. We proceed by discussing several different distance measures used in quantum information theory and their operational interpretations. After that, we give a brief overview of how functional inequalities can be used to study the convergence of semigroups under several of the distance measures introduced before and finally discuss some applications of these techniques. Considering that this is an extensive area of research, we focus only on results and concepts from the existing literature which are particularly relevant for the present dissertation. This is followed by the contributed articles. We start with the article in which the author of this dissertation is a co-author, V [1]. This is justified by the fact that one of the core articles, I [2], is inspired by it. We then include the core published articles and later include the articles under review. Each article is preceded by a technical summary of its main results, which, although still short, is more involved than the one found in the next section, and a more elaborate account of the individual contributions of the present author. They are also accompanied by the authorization to include the the articles in this thesis in case they have already been published.

1.1 Summary and Discussion of Results

The articles included in the present dissertation can be classified loosely into three categories. The first one consists of those concerning quantifying *how noisy* some given quantum dynamical semigroup is. The figure of merit to quantify this depends on the specific task at hand. Articles I, IV and V [1, 2, 3] are concerned with this issue. The second category consists of those involved with the question of how to prepare quantum states on a quantum computer exploring quantum dynamical semigroups, ideally efficiently. This includes articles I and II [2, 4]. The last category, which is admittedly slightly disjoint thematically from the others, consists of article III [5], where we show how to explore the *compression* power of positive maps to solve semidefinite programs, a natural framework to formulate and solve many optimization problems in quantum information theory. Here we also study the limitations of such maps for these tasks.

Core articles as principal author

- *Article I [2]: Sandwiched Renyi Convergence for Quantum Evolutions*
  In this article, we develop the technical machinery necessary to study the convergence of quantum dynamical semigroups with full rank stationary states using the recently introduced sandwiched Rényi divergences entropies with $p \geq 1$ as a distance measure [6, 7]. Focusing on semigroups in continuous time, we show that the convergence is always exponential and equivalent to a functional inequality and define the optimal convergence rates. Using Pinsker's inequality, it is possible to derive rapid mixing times for such semigroups. We explore how these convergence rates connect to other essential constants in the study of the convergence of semigroups, such as the spectral gap and the logarithmic Sobolev constant. As expected, we find that the spectral gap of the semigroup gives an upper bound on the optimal convergence rates in the reversible case. Moreover, we also show that these convergence rates are lower-bounded by logarithmic Sobolev constants. Both of these proofs work by relating the functionals involved in the definition of such constants. By connecting the convergence rates with respect the sandwiched Rényi divergences and logarithmic Sobolev constants, we are able to derive mixing time bounds from them without resorting to technical assumptions such as l_p-regularity that were needed before [8]. We provide evidence that obtaining analytical bounds or even optimal values for these constants is much more feasible than using regular logarithmic Sobolev inequalities by computing the optimal constants for $p = 2$ in the case of depolarizing channels. The same computation for the $p = 1$ case, done in [1] is significantly more involved and was one of our inspirations to search for better techniques to compute convergence rates. This result implies a universal lower bound for the convergence rate in terms of the spectral gap and stationary state of the semigroup. We also briefly comment on how to use a similar approach to derive mixing times in discrete time. The techniques developed here are therefore natural candidates to show rapid mixing for some classes of quantum dynamical semigroups which are relevant for algorithmic applications. Finally, we explore the connection between sandwiched Rényi divergences and strong converse bounds on the classical capacity of quantum channels established in [7] and tensorization results for logarithmic Sobolev constants to obtain upper bounds on the classical capacity of quantum dynamical semigroups as a function of time, their spectral gap , and properties of the stationary state. Putting these together with bounds on the spectral gap of Davies generators available in the literature [9, 10], we obtain the first bound available on the classical capacity of stabilizer Hamiltonian under thermal noise as a function of time and temperature. These include widely studied models such as the $2D-$toric code. Moreover, these are bounds in the strong converse sense. These results clearly show that the techniques developed here are a powerful tool to study the classical capacity of quantum dynamical semigroups. I am the principal author of this article. The project's idea was motivated by discussions between Alexander Müller-Hermes and me after we finished [1]. I proved all main results and wrote all sections of it, except Appendix A, Theorem 4.1 and Theorem 4.3.

- *Article II [4]: Perfect Sampling for Quantum Gibbs States*
  In this article we propose an algorithm to overcome the need of having mixing time bounds to obtain certifiably good samples from measurements on quantum Gibbs states. That is, we develop an algorithm to obtain perfect samples from any measurement on a Gibbs state. We assume we can implement on a quantum computer a quantum dynamical semigroup satisfying certain conditions and a phase estimation routine for the Hamiltonian of interest. This is conceptually different to usual approaches to this problem which focus on obtaining approximate samples and usually do not provide any certificate on the quality of the samples. Moreover, the run-time analysis of these approaches is far from trivial, as the huge literature dedicated to Markov chain mixing attests. The algorithm works by adapting the coupling from the past algorithms developed by Propp and Wilson [11] to solve the analogous problem in the classical setting. The classical algorithm produces a perfect sample of the stationary distribution of a Markov chain given a black box which, fed some initial state, outputs a valid transition of the chain. Here we suppose we have access to a quantum computer that can run the phase estimation routine for the Hamiltonian of interest and also implement a quantum channel that drives the system to the desired Gibbs state and has some extra properties. These ensure that, fixing some eigenbasis of the Hamiltonian, the dynamics of this eigenbasis under the channel can be modelled as a classical Markov chain and that the transition probabilities between different eigenstates only depend on their energies. One example of such a channel is the one proposed in the quantum Metropolis algorithm of [12]. The algorithm then works by feeding transitions of the chain in the eigenbasis to a classical computer running an adapted version of the CFTP algorithm of Propp and Wilson. The runtime of the algorithm is probabilistic and its mean time depends strongly on the degeneracy of the Hamiltonian. Given that each eigenspace of the Hamiltonian in a $d-$dimensional Hilbert space has dimension at least $dr(d)^{-1}$ for some function r, then the runtime is of order $\mathcal{O}(r(d)^2\log(r(d))t_{\text{mix}})$, where t_{mix} is the mixing time of the channel. This is efficient for Hamiltonians with an extremely degenerate spectrum, i.e. $r(d)=\log(d)^m$, and in the worst case of nondegenerate spectra, i.e. $r(d)=1$, this runtime is the same as the classical one, up to a logarithmic factor. We also investigate the stability of the algorithm with respect to different sources of noise. We start by showing that the algorithm is stable against perturbations of the channel. We also show that the algorithm is stable against faulty phase estimation. Roughly speaking, the algorithm will be more stable against faulty phase estimation the further different eigenvalues are apart and if eigenvalues that are close have eigenspaces whose dimensions are of the same order of magnitude. Moreover, we show how to adapt other variations of perfect sampling algorithms to the quantum setting. I am solely responsible for all the writing and results of this article.

Further articles as principal author under review

- *Article III [5]: Dimensionality reduction of SDPs through sketching*
  Although most of the articles discussed before are concerned with quantifying or exploring noise in quantum systems, here we consider how to compress observables using positive maps and apply our results to the solution of semidefinite programs (SDPs). SDPs provide a natural framework to formulate and solve many optimization problems in quantum information theory. Although they are solvable in polynomial time under mild assumptions, their practical application is however restricted by the prohibitive amount of memory required to solve problems of even moderate dimension. Therefore, we develop an algorithm based on positive maps to approximately solve them using less memory. We start by showing how to use Johnson-Lindenstrauss transforms to obtain completely positive maps that approximately preserve the Hilbert-Schmidt scalar product between two hermitian matrices, but have a much smaller output dimension. We then apply this result to improve both complexity and storage space requirements to solve SDPs given in a certain universal form. This works by applying the positive maps we obtain to the matrices that define the SDP constraints and target functional to obtain a smaller SDP,

which we call the sketched SDP. This new SDP then has a much smaller dimension, and as the maps approximately preserve feasibility and the value of the original SDP, solving this sketched SDP yields information about the original one. One of the main advantages of this approach is that one can use already available solvers or algorithms for SDPs to solve the sketched version. These techniques work best for problems in which the Schatten 1-norm of the matrices specifying the SDP and of a solution to the problem is constant in the problem size. Moreover, we show how to apply similar ideas to probe the feasibility of certain linear matrix inequalities. We clarify the limitations of this and other slightly more general approaches to approximating the value of SDPs by showing some no-go results. We show that it is not possible to compress all SDPs nontrivially using linear maps by relating this problem to sketching the operator norm, which is known not to be sketchable [13]. Furthermore, we show that our results concerning positive maps that approximately preserve the Hilbert-Schmidt scalar product cannot be improved significantly, a result which might be of independent interest to the quantum information community and complements those of [14]. The project's idea was motivated by discussions between Andreas Bluhm and me. I proved the majority of the statements of the article and wrote most of the sections, although the discussions with Andreas Bluhm were central in the process.

- *Article IV [3]: Approximate Randomized Benchmarking for Finite Groups*
  Randomized benchmarking [15] is an experimental protocol to estimate the average gate fidelity of a set of quantum gates efficiently. It is usually used to estimate the average fidelity of Clifford gates and works under the assumption that the quantum channel that describes the noise in the implementation is independent of the gate and constant in time. The protocol works by exploring the properties of twirled and covariant channels and by estimating how fast a covariant quantum channel converges. From the information on the speed of the convergence of the quantum channel, it is then possible to infer the average gate fidelity of the channel. In this article, we generalize the randomized benchmarking protocol in three different ways. First, we show how to apply the protocol to estimate the average gate fidelity of an arbitrary representation of a finite group, not necessarily Cliffords. How to interpret the experimental data and extract the average fidelity from it depends on properties of the representation at hand. In the usual setting of randomized benchmarking, it is also assumed that we have access to samples from the Haar distribution of the group. As it may not be possible or straightforward to obtain these efficiently, we show that obtaining samples that are approximately Haar also suffices for the implementation of the protocol. These results allow us to apply Markov chain Monte Carlo techniques to obtain approximate samples efficiently and are also a stability result for randomized benchmarking protocols. Finally, we show how to implement the randomized benchmarking protocol only having to implement a set of gates that generate the group and that is closed under taking inverses, and one arbitrary element of the group. These results simplify the implementation of the protocol and makes it easier to justify the error model, as usually gates have to be broken down into generators. The price to pay is that we need to assume that the quantum channel that describes the model is close to a covariant channel. We apply our methods to the subgroup of unitary matrices that is formed by permutation matrices multiplied by diagonal unitaries whose entries are roots of unity. These are interesting candidates for the application of our methods because one is only required to estimate two parameters when performing randomized benchmarking with them, the scaling of multiplying and inverting elements of this group is not too prohibitive, and gates from this set allow for universal quantum computation together with the Cliffords. We also apply our methods of approximate samples and generator benchmarking to the Clifford group with success. This project started after Anna-Lena K. Hashagen asked me to review a paper of hers on a similar topic. I realized that many techniques could be generalized and we followed this path. I proved, formulated and wrote most of the statements of the article.

Articles as co-author

- *Article V [1]: Relative entropy convergence for depolarizing channels*
  This article investigates how fast states converge under depolarizing semigroups which have full rank fixed points in terms of the relative entropy. This convergence is always exponential in time, and we compute the optimal exponents as a function of the state. Therefore, these exponents quantify how noisy these semigroups are. Moreover, these are the first known results of optimal constants for convergence in the relative entropy. To arrive at this result, we explore properties of quasi-convex functions and show how to simplify the computation of these constants using Birkhoff's Theorem on doubly-stochastic matrices. By reformulating the relative entropy in terms of the von Neumann entropy, we also show how our results imply improved concavity bounds for it. We compare it to other estimates available in the literature, such as the one by Kim et al. [16], and provide numerical evidence that they are not comparable. We prove a quantum version of Shearer's inequality by again exploring the connections between the relative entropy and the von Neumann entropy. This inequality is then used to establish a uniform bound on the convergence exponent of the relative entropy under tensor powers of the depolarizing semigroup going to the maximally mixed state. Finally, we prove an optimal version of Pinsker's inequality for a fixed second argument of the relative entropy. Alexander Müller-Hermes is the principal author of this article.

2

Notation and Preliminaries

Throughout this thesis $\mathcal{M}_{d,d'}$ will denote the space of $d \times d'$ complex matrices. We will write $\mathcal{M}_d$ for short in case $d = d'$. Given a matrix $A \in \mathcal{M}_{d,d'}$, we will express its adjoint by $A^\dagger$. Moreover, for linear operators $T : \mathcal{M}_d \to \mathcal{M}_{d'}$, we will denote their adjoint w.r.t. the Hilbert-Schmidt scalar product by T^*. By $\mathcal{M}_d^+$ we denote the set of positive definite matrices. We will use the symbol $\mathbb{1}_d$ for the identity matrix on $\mathcal{M}_d$ and will denote the identity map from $\mathcal{M}_d \to \mathcal{M}_d$ by $\mathrm{id}_d : \mathcal{M}_d \to \mathcal{M}_d$. Sometimes we will drop the d index if the dimension is clear from context. We will denote the set of orthogonal projections in $\mathcal{M}_d$ by $\mathcal{P}_d$. Given vector spaces V, W, we will denote the set of bounded linear maps from V to W by $\mathcal{B}(V, W)$ or just $\mathcal{B}(V)$ in case $V = W$. We adopt the bra-ket notation to denote vectors, matrices and their duals. We will denote the group of $d-$dimensional unitary matrices by $U(d)$. Given some unitary $U_g \in U(d)$, we will denote the unitary conjugation with U_g by $\mathcal{U}_g$. That is, we define the map $\mathcal{U}_g : \mathcal{M}_d \to \mathcal{M}_d$ as

$$\mathcal{U}_g(\cdot) = U_g \cdot U_g^\dagger .$$

We will sometimes call self-adjoint matrices $A \in \mathcal{M}_d$ observables or Hamiltonians.

2.1 Finite Dimensional Quantum Mechanics

Throughout this thesis we will only deal with finite dimensional systems. This will allow us to bypass many of the technicalities involved when studying infinite dimensional systems. We will briefly introduce most of the concepts of quantum mechanics we will need throughout this thesis. We will adopt a quantum information perspective on quantum mechanics and refer to [17] for more on these basic concepts from a similar point of view.

2.1.1 States and Measurements

The state of a d-dimensional quantum system is described by a density matrix $\rho \in \mathcal{M}_d$, which are positive semi-definite matrices of trace 1. We will denote the set of d-dimensional quantum states by $\mathcal{D}_d$ and by $\mathcal{D}_d^+ = \mathcal{M}_d^+ \cap \mathcal{D}_d$ the set of full rank states. A measurement of the system corresponds to a positive operator valued measure (POVM) on $\mathbb{C}^d$. These are positive semi-definite operators $\{E_i\}_{i=1}^k \subset \mathcal{M}_d^+$ such that

$$\sum_{i=1}^k E_i = \mathbb{1}. \tag{2.1}$$

We will call each E_i a POVM element and k the number of outcomes.

These conditions on states and POVMs assure that they induce a probability distribution $p \in \mathbb{R}^k$ through $p(i) = \mathrm{Tr}(E_i \rho)$. This probability distribution describes the probability of observing outcome i when measuring the POVM on a system described by the state ρ. Given two systems A, B of dimensions d_A, d_B, respectively, the composite system AB is described by the tensor product of the individual systems. That is, a state ρ_{AB} of the composite system AB is an element of $\mathcal{M}_{d_A} \otimes \mathcal{M}_{d_B} \simeq \mathcal{M}_{d_A d_B}$. Given a $X_{AB} \in \mathcal{M}_{d_A} \otimes \mathcal{M}_{d_B}$ we will denote the partial trace over B by $\operatorname{Tr}_B(\cdot)$. When it is clear from context which system is of primary interest and which one is the auxiliary system, we will denote the partial trace over the auxiliary system by $\operatorname{Tr}_2(\cdot)$. Given an inverse temperature $\beta > 0$ and a Hamiltonian $H \in \mathcal{M}_d$, we define its Gibbs state to be given by $\frac{e^{-\beta H}}{\operatorname{Tr}(e^{-\beta H})}$. The function $\beta \mapsto \operatorname{Tr}\left(e^{-\beta H}\right)$ is called the partition function and we will denote it by $\mathcal{Z}_\beta$.

2.1.2 Time Evolutions

The most general way to describe the time evolution of states of a quantum system is through quantum channels:

Definition 2.1.1 (Quantum Channel)**.** *We call a linear map $T : \mathcal{M}_d \to \mathcal{M}_{d'}$ a quantum channel if it is trace preserving and completely positive.*

Recall that a linear map $T : \mathcal{M}_d \to \mathcal{M}_{d'}$ is trace preserving if for all $X \in \mathcal{M}_d$ we have $\operatorname{Tr}(X) = \operatorname{Tr}(T(X))$ or equivalently $T^*(\mathbb{1}) = \mathbb{1}$. A map is completely positive if we have for all $D \in \mathbb{N}$ that for $X \in (\mathcal{M}_d \otimes \mathcal{M}_D)^+$ that $T \otimes \mathrm{id}_D(X) \geq 0$. The trace preserving property and the complete positivity ensure that states are mapped to states under T, even when considering composite systems. One can alternatively characterize and describe quantum channels $T\mathcal{M}_d \to \mathcal{M}_{d'}$ in terms of their duals with respect to the Hilbert-Schmidt scalar product, $T^* : \mathcal{M}_{d'} \to \mathcal{M}_d$. The map T is usually referred to as the channel in the Schrödinger picture and the map is the channel in the Heisenberg picture. T^* can be seen as describing the evolution of observables under the channel. One can easily see that T being a quantum channel in the Schrödinger picture is equivalent to T^* being completely positive and unital, i.e. $T^*(\mathbb{1}) = \mathbb{1}$.

There are many different equivalent characterizations of quantum channels or more generally completely positive maps. Here we recall some of them. We again refer to [17] for more details on this and proofs. A very useful characterization of completely positive maps is given through the Choi-Jamiolkowski isomorphism. Define the maximally entangled state $|\Omega\rangle\langle\Omega| \in \mathcal{M}_d \otimes \mathcal{M}_d$ to be given by

$$|\Omega\rangle\langle\Omega| = \frac{1}{d}\sum_{i,j}^{d} |i\rangle\langle j| \otimes |i\rangle\langle j| \,.$$

Here $\{|i\rangle\}_{i=1}^d$ is an orthonormal basis of $\mathbb{C}^d$. The Choi-Jamiolkowski matrix is then given by

$$T \otimes \mathrm{id}\left(|\Omega\rangle\langle\Omega|\right). \tag{2.2}$$

One can infer many properties of the map T from properties of the state defined in Equation (2.2). For instance, T is completely positive if and only if the Choi matrix is positive semi-definite. Another important characterization of completely positive maps is given through the Kraus decomposition.

Theorem 2.1.2 (Kraus decomposition)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_{d'}$ be a linear map. Then T is completely positive if and only if, there exist $K_i \in \mathcal{M}_{d',d}$, $1 \leq i \leq r$, s.t.*

$$T(X) = \sum_{i=1}^{r} K_i X K_i^\dagger.$$

Moreover, T is trace preserving if and only if, $\sum_{i=1}^{r} K_i^\dagger K_i = \mathbb{1}$.

Note that this decomposition is not unique. We call the minimal r s.t. such a decomposition exists the Kraus rank of the completely positive map. One can show that this corresponds to the rank of the Choi matrix. The evolution of closed systems is described by the conjugation with unitary operators in quantum mechanics. The connection between the dynamics of closed system and dynamics described by quantum channels is made clear by the Stinespring dilation, which gives yet another characterization of quantum channels:

Theorem 2.1.3 (Stinespring Dilation)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_{d'}$ be a quantum channel. Then there exists a $D \leq d^2$, a unitary operator U on $\mathbb{C}^d \otimes \mathbb{C}^D$ and a state $\rho \in \mathcal{D}_D$ such that*

$$T(X) = \mathrm{Tr}_2 \left(UX \otimes \rho U^\dagger \right) \tag{2.3}$$

for all $X \in \mathcal{M}_d$.

Equation (2.3) can be interpreted in the following way: the evolution under T of any state $\sigma \in \mathcal{D}_d$ can be implemented by preparing the initially uncorrelated state $\sigma \otimes \rho$, evolving it with the unitary U and then tracing out the auxiliary system. The quantum channel can thus be interpreted as describing the evolution of a state that interacts with a larger, closed system, and is initially uncorrelated with it.

3

Quantum Dynamical Semigroups

In this chapter, we will introduce some of the basic concepts related to quantum dynamical semigroups, which are semigroups of quantum channels, and their convergence properties. We will mostly work in the Schrödinger picture here.

Definition 3.0.1 (Quantum dynamical semigroup). *Let $I \in \{\mathbb{R}^+, \mathbb{N}\}$. A function $f : I \to \mathcal{B}(\mathcal{M}_d)$ is called a quantum dynamical semigroup if*

1. *$f(0) = id_d$*
2. *$f(t+s) = f(t)f(s)$ for all $t, s \in I$*
3. *$f(t)$ is a quantum channel for all $t \in I$.*

In case $I = \mathbb{R}^+$ we also demand that $f(t)$ depends continuously on t. We will refer to the case $I = \mathbb{R}^+$ as a semigroup in continuous time and $I = \mathbb{N}$ as a semigroup in discrete time.

We will usually denote a quantum dynamical semigroup by T_t, where $T_t = f(t)$ and will usually denote $f(I)$ by $\{T_t\}_{t\in I}$. One should think of T_t as the time evolution of the system in a time interval $[0, t]$. The semigroup structure ensures that the evolution is both homogeneous and memoryless over time. It is easy to see that in case we have a semigroup in discrete time there is a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ such that $T_t = T^t$, that is, we just iterate a quantum channel t times. We call this quantum channel the generator of the chain. In the case of continuous time we may also find generators of the semigroup and they have a richer structure:

Theorem 3.0.2 (Generators of semigroups in continuous time). *Let $\{T_t\}_{t\in\mathbb{R}^+}$ be a quantum dynamical semigroup. Then there exists a $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ s.t. $T_t = e^{t\mathcal{L}}$. Moreover, $\mathcal{L}$ may be written in any of these equivalent forms:*

$$\mathcal{L}(X) = \Phi(X) - X\kappa - \kappa^\dagger X \tag{3.1}$$

$$\mathcal{L}(X) = i[X, H] + \sum_j L_j X L_j^\dagger - \frac{1}{2}\{L_j^\dagger L_j, X\}, \tag{3.2}$$

where $\Phi : \mathcal{M}_d \to \mathcal{M}_d$ is completely positive and satisfies $\Phi^*(\mathbb{1}) = \kappa + \kappa^\dagger$, $\kappa, H, L_j \in \mathcal{M}_d$ and $H = H^\dagger$. We will refer to any operator satisfying Equation (3.1) as a Liouvillian.

We refer to e.g. [17] for more details on this and proofs.

3.1 Convergence of Semigroups

In this section, we will discuss some properties of quantum channels or Liouvillians that simplify the study of their convergence as $t \to \infty$. Given a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$, we will call the states $\rho \in \mathcal{M}_d$ such that $T(\rho) = \rho$ stationary states. Analogously, for Liouvillians

$\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ we call states such that $\mathcal{L}(\rho) = 0$ stationary states, as these are then stationary states of the semigroups they generate for all $t \in \mathbb{R}^+$. One of the most important concepts in this thesis is that of a primitive channel or Liouvillian.

Definition 3.1.1 (Primitive generator)**.** *We call a semigroup $\{T_t\}_{t\in I}$ primitive if there is a unique full-rank state $\sigma \in \mathcal{D}_d^+$ such that for all $\rho \in \mathcal{D}_d$*

$$\lim_{t\to\infty} T_t(\rho) = \sigma.$$

Analogously, we call a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ primitive if there is a unique full-rank state $\sigma \in \mathcal{D}_d^+$ such that $\mathcal{L}(\sigma) = 0$.

We will refer to the state σ as the stationary state of the semigroup. Another concept which is closely related to the convergence of semigroups is that of irreducibility:

Definition 3.1.2 (Irreducible quantum channel)**.** *A quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is called irreducible if for a hermitian projection $P \in \mathcal{M}_d$ we have that $T(P\mathcal{M}_dP) = P\mathcal{M}_dP$ implies that $P \in \{0, \mathbb{1}\}$.*

With some abuse of terminology, we will call a Liouvillian primitive or irreducible if the semigroup it generates is primitive for some t. There are many different equivalent characterizations of primitive or irreducible quantum channels. We will now collect some of them which are central to the other sections and refer to e.g. [18] for more details and proofs. Here we also differentiate between continuous and discrete time, as it is easier to characterize primitive or irreducible semigroups in continuous time.

Theorem 3.1.3 (Generators of primitive semigroups)**.** *The semigroup generated by a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ is primitive and irreducible if one of the following equivalent conditions holds:*

1. *There is a $t_0 > 0$ such that T_{t_0} is irreducible.*
2. *T_t is irreducible for all $t > 0$.*
3. *T_t is primitive for all $t > 0$.*
4. $\ker \mathcal{L} =$ *span*$\{\sigma\}$

That is, primitivity and irreducibility are equivalent in continuous time and one just has to check the kernel of $\mathcal{L}$ to determine whether it is primitive or not. In discrete time the situation is more subtle. One can show that any quantum channel satisfies $\|T\|_\infty = 1$ and there is always a positive semi-definite $X \in \mathcal{M}_d$ such that $T(X) = X$. We call eigenvalues $\lambda \in \mathbb{C}$ of a quantum channel of the peripheral spectrum if $|\lambda| = 1$ and call the peripheral spectrum of the channel trivial if it only contains 1 with multiplicity 1. We then have

Theorem 3.1.4 (Primitive quantum channels)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel. Then the following are equivalent:*

1. *T is primitive.*
2. *T has a trivial peripheral spectrum.*
3. *there is an $n \in \mathbb{N}$ such that for all $\rho \in \mathcal{D}_d$ $T^n(\rho) > 0$.*

Theorem 3.1.4 allows us to characterize primitive quantum channels in terms of their spectrum, which is usually more readily accessible. Another property of semigroups that significantly simplifies the study of their convergence is that of detailed balance, as it is the case for classical Markov chains. There are many different generalizations of this condition to the quantum setting which are not equivalent (see e.g. [19]). The following is the most appropriate for our purposes:

Definition 3.1.5 (Detailed balance/reversibility)**.** *A semigroup $\{T_t\}_{i\in I}$ is said to satisfy detailed balance with respect to a state $\sigma \in \mathcal{D}_d^+$ if for all $X \in \mathcal{M}_d$ and $t \in I$*

$$T_t\left(\sigma^{\frac{1}{2}} X \sigma^{\frac{1}{2}}\right) = \sigma^{\frac{1}{2}} T_t^*(X) \sigma^{\frac{1}{2}}. \tag{3.3}$$

We will also sometimes call semigroups that satisfy detailed balance reversible.

We will see later in Subsection 4.1.1 that the condition in Equation (3.3) is equivalent to the generator of the semigroup being self-adjoint with respect to a certain scalar product.

3.1.1 The Phase Estimation Algorithm

The quantum phase estimation algorithm and its variations, first developed in [20], certainly is one of the most important quantum algorithms. Although not strictly related to quantum dynamical semigroups, at first sight, this algorithm is a vital subroutine to implement primitive quantum channels that converge to quantum Gibbs states, such as the quantum Metropolis algorithm proposed in [12]. We will discuss how exactly this is used in quantum Gibbs sampling later and discuss the basics of this subroutine here. The perfect phase estimation procedure or algorithm has as an input a Hamiltonian $H \in \mathcal{M}_d$ and produces a unitary $U \in \mathcal{M}_d \otimes \mathcal{M}_{2^m}$ that acts as follows. Given an eigenstate $|\psi_i\rangle$ of H such that $H|\psi_i\rangle = E_i |E_i\rangle$, we have that

$$U|\psi_i\rangle \otimes |0\rangle = |\psi_i\rangle \otimes |E_i\rangle .$$

Here we have assumed that E_i can be expressed exactly with m binary digits. This will not generally be the case for generic Hamiltonians and most phase estimation procedures only produce an output that peaks around $|\psi_i\rangle \otimes |E_i\rangle$. Phase estimation algorithms are still subject of current research, but we will briefly discuss here the probability of making errors for a simple implementation discussed in [21]. To implement the phase estimation algorithm we assume we may implement the semigroups of unitaries generated by H, $V_t = e^{itH} \in U(d)$. We also assume we may implement Hadamard gates $Ha \in U(2)$, where

$$Ha = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

on each of the qubits of the second register. We also need to apply a controlled V_t gate between the first system and any of the m qubits on the second register, CV_j, where j is the qubit. The unitary acts as

$$CV_j |x\rangle |0\rangle \mapsto |x\rangle |0\rangle , \quad CV_j |x\rangle |1\rangle \mapsto V_{2^{j-1}} |x\rangle |1\rangle$$

and as the identity on the other qubits. Finally, we also need to implement the inverse discrete Fourier transform, $FT^\dagger$. Following the notation of [21], for a binary sequence $j_1, j_2, \ldots, j_m$ we set $0.j_l, j_{l+1}, \ldots, j_n$ to be the binary fraction

$$0.j_l, j_{l+1}, \ldots, j_m = \frac{j_l}{2} + \frac{j_{l+1}}{4} + \ldots + \frac{j_m}{2^{n-l+1}}.$$

The Fourier transform gate FT then acts as follows in the computational basis:

$$|j_1 j_2 \ldots j_n\rangle \mapsto \frac{1}{2^{m/2}} \left(|0\rangle + e^{2\pi i 0.j_m} |1\rangle\right) \left(|0\rangle + e^{2\pi i 0.j_{m-1} j_m} |1\rangle\right) \ldots \left(|0\rangle + e^{2\pi i 0.j_1 j_2 \ldots j_m} |1\rangle\right) .$$

The phase estimation procedure is given by applying the sequence of gates

$$(\mathbb{1}_d \otimes FT^\dagger)(CV_m CV_{m-1} \ldots CV_1)(\mathbb{1}_d \otimes Ha^{\otimes m}).$$

It is possible to check that if E_i can be expressed exactly with m binary digits, we will indeed observe E_i if we measure the second register in the computational basis, but in general it will only peak around the best approximation of E_i using m binary digits. Given that the eigenstate of the system is $|\psi_i\rangle$, we have that the probability that the observed outcome is E is bounded by

$$|2^m(E_i - E) \mod 2^m|^{-2}.$$

Using this it is then possible to show that if we use

$$m \geq n + \log\left(2 + (2\delta)^{-1}\right)$$

qubits to perform phase estimation, then we obtain E_i accurate to n bits with probability at least $1 - \delta$.

The relevance of discussing faulty phase estimation is that it might lead us to implement quantum channels that do not have the desired state as their stationary state and it is therefore always important to also show the stability of the quantum channels with respect to such errors.

3.1.2 Important Examples

Here we discuss some examples of semigroups that are important in the theory of quantum dynamical semigroups.

Depolarizing Channels:
The depolarizing channels may be considered the simplest example of a quantum dynamical semigroup. Given a full rank state $\sigma \in \mathcal{D}_d^+$, we define the semigroup generated by

$$\mathcal{L}_\sigma(X) = \operatorname{Tr}(X)\,\sigma - X \tag{3.4}$$

to be the $\sigma-$depolarizing semigroup. It is easy to see that the channels generated, the depolarizing channels, are of the form

$$T_{t,\sigma}(\rho) = e^{-t}\sigma + (1 - e^{-t})\rho \tag{3.5}$$

for a state $\rho \in \mathcal{D}_d$. As was made clear by [1], the depolarizing channels play a central role in the theory of functional inequalities for semigroups, as they are a very simple example of a primitive semigroup for any initial state $\sigma \in \mathcal{D}_d^+$.

Davies Generators:
Davies generators describe a system weakly coupled to a thermal bath under an appropriate approximation [22]. Here we will only review their most basic properties and refer to [23, 24, 25] for more details.

Suppose a $d-$dimensional system is weakly coupled to a thermal bath of dimension d_B at inverse inverse temperature $\beta > 0$. Consider a Hamiltonian $H_{\text{tot}} \in \mathcal{M}_d \otimes \mathcal{M}_{d_B}$ of the system and the bath of the form

$$H_{\text{tot}} = H \otimes \mathbb{1}_{d_B} + \mathbb{1}_S \otimes H_B + H_I,$$

where $H \in \mathcal{M}_d$ is the Hamiltonian of the system, $H_B \in \mathcal{M}_{d_B}$ of the bath and

$$H_I = \sum_\alpha S^\alpha \otimes B^\alpha \in \mathcal{M}_d \otimes \mathcal{M}_{d_B} \tag{3.6}$$

describes the interaction between the system and the bath. Here the operators S^α and B^α are self-adjoint. Let $\{\lambda_k\}_{k\in[d]}$ be the spectrum of the Hamiltonian H. We then define the Bohr-frequencies $\omega_{i,j}$ to be given by the differences of eigenvalues of H, that is, $\omega_{i,j} = \lambda_i - \lambda_j$ for different values of λ. We will drop the indices on ω from now on to avoid cumbersome notation, as is usually done. Moreover, we introduce operators $S^\alpha(\omega)$ which are the Fourier components of the coupling operators S^α and satisfy

$$e^{iHt}S^\alpha e^{-iHt} = \sum_\omega S^\alpha(\omega)e^{i\omega t}.$$

The canonical form of the Davies generator at inverse temperature $\beta > 0$ in the Heisenberg picture, $\mathcal{L}^*_\beta$, is then given by

$$\mathcal{L}^*_\beta(X) = i[H,X] + \sum_{\omega,\alpha} \mathcal{L}^*_{\omega,\alpha}(X),$$

where

$$\mathcal{L}^*_{\omega,\alpha}(X) = G^\alpha(\omega)\left(S^\alpha(\omega)^\dagger X S^\alpha(\omega) - \frac{1}{2}\{S^\alpha(\omega)^\dagger S^\alpha(\omega), X\}\right).$$

Here $\{X,Y\} = XY + YX$ is the anticommutator and $G^\alpha : \mathbb{R} \to \mathbb{R}$ are the transition rate functions. Their form depends on the choice of the bath model [24]. For our purposes, it will be enough to assume that these are functions that satisfy the KMS condition [26], that is, $G^\alpha(-\omega) = G^\alpha(\omega)e^{-\beta\omega}$. Note that under some assumptions on the operators $S^\alpha(\omega)$ [10, 27] and on the transition rate functions, the semigroup generated by $\mathcal{L}_\beta$ converges to the Gibbs state $\frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})}$ and is reversible [26]. In the examples considered here, this will always be the case.

Their relevance comes from the fact that they are one of the standard ways of modelling quantum systems under thermal noise [10].

Quantum Metropolis:

Although the last two classes of examples are of semigroups in continuous time and most of the methods we will discuss later are more suited for semigroups in continuous time, an example of a semigroup in discrete time which is particularly relevant is the one described in the quantum Metropolis algorithm introduced in [12]. It can be seen as a generalization of the classical Metropolis algorithm to sample from classical Gibbs states. As before, the goal is to obtain a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ which will be primitive and whose stationary state will be $\frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})}$ for some given Hamiltonian $H \in \mathcal{M}_d$ and inverse temperature $\beta > 0$. The algorithm works with four registers, $ABCD$. The first register, A, will be of dimension d, the second and the third one, BC, will be of dimension 2^m and the fourth, D, will be of dimension 2. We will assume we are able to perform phase estimation for the Hamiltonian H exactly. We refer to [12] for more details on the performance of the algorithms with faulty phase estimation. Moreover, we will assume we may sample from a probability distribution μ on $U(d)$ with the property that $d\mu(C) = d\mu(C^\dagger)$. The convergence of the quantum channel will typically depend on how we choose these unitaries. Assuming for simplicity that the support of μ is finite, the only thing necessary to ensure convergence is that for any eigenstate $|\psi_i\rangle$ of H, there is another eigenstate $|\psi_j\rangle$ and $C \in U(d)$ such that

$$\mu(C) > 0, |\langle\psi_i| C |\psi_j\rangle|^2 > 0.$$

If $d = 2^k$ for some $k \in \mathbb{N}$ we may take e.g. the uniform measure on the Clifford group on k qubits. The algorithm works as follows. We first start by preparing some fixed initial state, say $|0\rangle|0\rangle|0\rangle|0\rangle$. We will then follow a combination of these steps:

1. Step 1: reinitialize registers BCD to $|0\rangle|0\rangle|0\rangle$. Perform phase estimation between registers A and B and measure B in the computational basis. The state of the system now is

   $$|\psi_i\rangle|E_i\rangle|0\rangle|0\rangle$$

   for some eigenstate $|\psi_i\rangle$ such that $H|\psi_i\rangle = E_i|\psi_i\rangle$.

2. Step 2: apply a random unitary drawn according to μ, V, to the first register followed by a phase estimation routine between registers A and C.

3. Step 3: define the unitary $W(E_i, E_j) \in U(2)$ with

   $$W(E_i,E_j) = \begin{bmatrix} \sqrt{1-f(E_i,E_j)} & \sqrt{f(E_i,E_j)} \\ \sqrt{f(E_i,E_j)} & -\sqrt{1-f(E_i,E_j)} \end{bmatrix},$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ is defined as $f(E_i, E_j) = \min\{1, e^{-\beta(E_j - E_i)}\}$. Apply $W(E_i, E_j)$ to register D conditioned on registers BC. Denote this unitary on BCD by W.

4. Step 4: measure register D in the computational basis. If the outcome is 1, go back to Step 1. If the outcome is 0, apply the unitary W^*, followed by V^* and the inverse of phase estimation on the system. Call this sequence of unitaries U^*. Define the following POVM elements a POVM $Q = \{Q_0, Q_1\}$:

$$Q_0 = U^* \left(\mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1} \otimes |0\rangle\langle 0|\right) U,$$
$$Q_1 = U^* \left(\mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1} \otimes |1\rangle\langle 1|\right) U.$$

Moreover, also define the following POVM $R = \{R_0, R_1\}$:

$$R_0 = \sum_i P_i \otimes |E_i\rangle\langle E_i| \otimes \mathbb{1} \otimes \mathbb{1},$$
$$R_1 = \mathbb{1} - P_0.$$

Here the sum is over all eigenvalues of H and P_i corresponds to the orthogonal projection onto the eigenspace of H corresponding to the eigenvalue E_i. We note that the POVM R may easily be implemented using phase estimation. We then measure the POVM R. If the outcome is 0, then we go to step 1. If not, for a fixed number of tries, say $k \in \mathbb{N}$, we measure the POVM Q followed by another measurement of R. If at any point we observe the outcome 0 after measuring R we go back to step 1 and we abort the procedure if this is not the case after k steps.

As shown in [12], the probability of observing 1 after measuring R decreases exponentially in the number of tries. We can interpret the unitary C as implementing the proposed moves, as in the classical Metropolis. The choice of $W(E_i, E_k)$ is also inspired by the function used in the classical Metropolis algorithm for the probability of accepting or rejecting a proposed move. Whether we accept or reject a proposed move is decided by the measurement of register D, with an outcome of 1 being accepting the move and 0 rejecting. Therefore, we have to reverse the move at step 4 in that case. It is easy to see that the algorithm just implements a classical random walk on the eigenspaces of H that is very similar to the classical Metropolis algorithm. Let $T_\beta : \mathcal{M}_d \to \mathcal{M}_d$ be the quantum channel which describes one complete step of the algorithm for some inverse temperature $\beta > 0$ on register A. Our assumptions on the distribution of unitaries μ ensures that this quantum channel is irreducible, as we may observe any eigenstate of H with positive probability by iterating T_β. Moreover, with the choice of $W(E_i, E_k)$ it is possible to show that the quantum channel T_β satisfies detailed balance with respect to the Gibbs state $\frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})}$ and therefore it is a stationary state. As shown in Section 3.1, this ensures that the implemented channel is indeed primitive.

3.1.3 Covariant Channels and Twirling

A very useful tool to study symmetry properties or the structure of quantum dynamical semigroups is that of covariance or twirling of quantum channels. Given a compact or finite group G and a representation $g \mapsto U_g \in U(d)$ of it, we say that a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is covariant with respect to this representation if we have that for all $X \in \mathcal{M}_d$ and $g \in G$

$$T(U_g X U_g^\dagger) = U_g(X)U_g^\dagger.$$

One could in principle also consider the case in which we have two different representations of the group acting on each side of the equation or more general classes of groups, but this goes beyond our purposes. Moreover, one can also define the twirl of a channel with respect to this representation, $\mathcal{T}(T)$, given by

$$\mathcal{T}(T) = \int_G U_g^\dagger T(U_g X U_g^\dagger) U_g dU_g,$$

where we integrate over the Haar measure on G. It is easy to check that $\mathcal{T}(T)$ is covariant with respect to the representation. This concept is useful in many settings in quantum information, as many properties of a channel are invariant under twirls. One important example is that of the depolarizing channels converging to $\frac{\mathbb{1}}{d}$, $T_{t,\frac{\mathbb{1}}{d}}$, as defined in Equation (3.5) as for them we have

$$T_{t,\frac{\mathbb{1}}{d}}(UXU^\dagger) = UT_{t,\frac{\mathbb{1}}{d}}(X)U^\dagger$$

for all $U \in U(d)$.

3.1.4 Classical Markov Chains

As we often use tools and draw inspiration from the classical analogues of quantum channels, Markov chains, we will review some of their basic properties and fix our notation. We refer to [28] for more details on these topics.

Definition 3.1.6 (Markov chain)**.** *A sequence $X_0, X_1, X_2, \dots$ of random variables taking values in a (finite) set S, referred to as the state space, is called a Markov chain if we have*

$$P(X_{n+1} = j | X_n = i) = \pi(i,j)$$

for a $|S| \times |S|$ matrix π. π is called the transition matrix of the chain.

As in the case of quantum dynamical semigroups, we will be interested in the behaviour of the Markov chain as $n \to \infty$. It is easy to see that if the probability distribution of X_0 is $\nu \in \mathbb{R}^{|S|}$, then the distribution of X_n is given by $\pi^n \nu$. We briefly discuss the translation of the definitions and results mentioned in Section 3.1 to the classical setting.

A probability distribution μ on S is called stationary if we have that $\pi\mu = \mu$. A transition matrix is called primitive if there exists a stationary distribution μ such that for any other distribution ν on S we have that $\lim\limits_{n\to\infty} \pi^n \nu = \mu$ and μ is strictly positive on all elements of S. A Markov chain is said to be irreducible if

$$\forall i, j \in S \ \exists n : \pi^n(i,j) > 0.$$

It is aperiodic if

$$\forall i \in S : \gcd\{n \in \backslash\{0\} : \pi^n(i,i) > 0\} = 1.$$

The importance of these concepts comes from the fact that together they imply that the chain is primitive. Analogously to the quantum case, we say that the transition matrix π satisfies detailed balance with respect to μ if

$$\pi(i,j)\mu(i) = \mu(j)\pi(j,i).$$

If μ satisfies detailed balance it again implies that it is stationary.

There is a standard way of embedding a Markov chain into a quantum channel. Let $d = |S|$ and $\{|i\rangle\}_{i=1}^d$ be the computational basis. Given π, we define the quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ as

$$T(X) = \sum_{i,j=1}^{d} \operatorname{Tr}\left(|i\rangle\langle i| \, X\right) \pi(i,j) \, |j\rangle\langle j| \, .$$

If the distribution of X_0 is ν and defining the state

$$\rho = \sum_{i=1}^{d} \nu(i) \, |i\rangle\langle i| \, ,$$

it is easy to see that $\nu_n' \in \mathbb{R}^d$ defined as

$$\nu_n'(i) = \operatorname{Tr}\left(T^n(\rho)\, |i\rangle\langle i|\right)$$

is a probability distribution which coincides with the one of X_n. One example of a Markov chain which will be of importance is that of a random walk on a (finite) group. Given a finite group G and a set $A \subset G$, we define the sequence of random variables $X_0, X_1, X_2, \ldots$ with

$$X_0 = e, \quad X_{n+1} = Y_{n+1}X_n.$$

Here $e \in G$ is the group's identity and the Y_i are i.i.d. random variables distributed according to some probability measure μ on A. It is easy to check that the X_i are indeed a Markov chain and that the transition matrix is given by

$$\pi(g_1, g_2) = \mu(g_2 g_1^{-1}).$$

One case of particular importance is that of a set A that is closed under inversions, i.e.

$$g \in A \implies g^{-1} \in A,$$

and such that A generates G. If we pick μ as the uniform measure on A it is not difficult to see that the resulting Markov chain is primitive, converges to the Haar measure on the group and satisfies detailed balance.

Another concept of Markov chains which will be important for our purposes will be that of a lumpable chain, sometimes also called a projected chain. Given an equivalence relation $\sim$ on the state space of a Markov chain $X_0, X_1, X_2, \ldots$ it is possible to define new random variables from the Markov chain. The state space of these new random variables is given by the equivalence classes of the equivalence relation and they are defined as follows. Define the function $f : S \to S\backslash \sim$ which maps a state to its equivalence class. The random variables are then given by $f(X_0), (X_1), f(X_2), \ldots$. If this stochastic process is again a Markov chain for all possible initial probability distributions on S, the chain is said to be lumpable with respect to this equivalence relation. The next Theorem gives necessary and sufficient conditions for lumpability.

Theorem 3.1.7 (Lumpable Chain)**.** *A necessary and sufficient condition for a Markov chain with state space S to be lumpable with respect to an equivalence relation $\sim$ on S is that for every pair $S_l, S_k \in S\backslash \sim$ we have for all $l, l' \in S_l$*

$$\sum_{k \in S_k} \pi(l, k) = \sum_{k \in S_k} \pi(l', k) \tag{3.7}$$

Moreover, the transition probability between S_l and S_k in the lumpable chain is given by Equation (3.7)*.*

Proof. We refer to [29, Theorem 6.3.2] for a proof. □

Lumpable chains are a natural setting when we can only access partial information about the current state of the Markov chain or as a technique to reduce the size of the state space of a given Markov chain. We extend the concept of a lumpable chain to a lumpable channel.

Definition 3.1.8. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel,*

$$H = \sum_{i=1}^{d'} E_i P_i \in \mathcal{M}_d$$

be a Hamiltonian and $P_i \in \mathcal{P}_d$ the projections onto its eigenspaces. The quantum channel T is said to be lumpable with respect to the Gibbs state $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ if $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ is a stationary state, it is primtive and there exists a function $f : \mathbb{R} \to \mathbb{R}$ s.t.

$$\operatorname{Tr}\left(T\left(|\psi_i\rangle\langle\psi_i|\right)|\psi_j\rangle\langle\psi_j|\right) = f(E_i, E_j)$$

for all eigenstates s.t. $H\left|\psi_i\right\rangle = E_i\left|\psi_i\right\rangle, H\left|\psi_j\right\rangle = E_j\left|\psi_j\right\rangle$*. Moreover, we demand that*

$$[T(P_i), P_j] = 0$$

for all eigenprojectors of H*.*

The reason to consider lumpable channels is that they induce a classical Markov chain when we fix an eigenbasis of H that is lumpable with respect to the equivalence relation given by different elements of the basis corresponding to the same eigenvalue.

4

Distance Measures and Convergence of Quantum Markov Chains

There are many different distance measures that can be used to quantify how close two quantum states are or how fast a quantum dynamical semigroup converges. Of course, the right distance measure depends on the particular application one has in mind. We will now review the main distance measures used in this work and comment on their interpretation. Note that most of the times we use the concept of a distance measure loosely, that is, any way of quantifying how close states are, and not necessarily in the mathematical definition of a metric.

4.1 Standard Distance Measures

One of the most standard distance measures used in quantum information theory is that induced by Schatten norms.

Definition 4.1.1 (Schatten $p-$Norm). *For $X \in \mathcal{M}_d$ and $p \in [1, \infty)$ we define the Schatten $p-$Norm of X, denoted by $\| \cdot \|_p$, to be given by*

$$\|X\|_p^p = \mathrm{Tr}\left(|X|^p\right),$$

with $|X| = (XX^*)^{\frac{1}{2}}$. Moreover, for $p = \infty$ we set $\|X\|_\infty = \lim\limits_{p\to\infty} \|X\|_p$.

It is easy to see that if X has singular values $\{s_i\}_{i=1}^d$, then $\|X\|_p^p = \sum_{i=1}^d s_i^p$ and $\|X\|_\infty = \max\{s_i\}_{i=1}^d$. These norms have similar properties as the usual l_p norms in $\mathbb{R}^n$. The norm for $p = 2$, also called the Hilbert-Schmidt norm, is induced by a scalar product on $\mathcal{M}_d$ given by $\langle X|Y\rangle = \mathrm{Tr}\left(X^\dagger Y\right)$ and the space of matrices equipped with this scalar product is a Hilbert space. Moreover, these norms satisfy a Hölder inequality and are monotonically decreasing in p.

Given a linear operator $\Phi : \mathcal{M}_d \to \mathcal{M}_D$ we may define the $p \to q$ norm which is induced by the Schatten norms.

Definition 4.1.2 ($p \to q$ norm). *Let $\Phi : \mathcal{M}_d \to \mathcal{M}_D$ be a linear operator and $p, q \in [1, \infty]$. We define the $p \to q$ norm of Φ, denoted by $\| \cdot \|_{p\to q}$, to be given by*

$$\|\Phi\|_{p\to q} = \sup_{X\in\mathcal{M}_d, X\neq 0} \frac{\|\Phi(X)\|_q}{\|X\|_p}.$$

The Schatten 1−norm, also called trace norm, will be of particular importance to quantify the convergence of Quantum Markov chains. Given two states $\rho, \sigma \in \mathcal{D}_d$, we have that

$$\|\rho - \sigma\|_1 = 2 \sup_{P \in \mathcal{P}_d} \operatorname{Tr}\left(P\left(\rho - \sigma\right)\right), \tag{4.1}$$

where again $\mathcal{P}_d$ is the set of $d-$dimensional orthogonal projections. We refer to e.g. [21, p. 404] for a proof of this claim. Equation 4.1 has a clear operational interpretation: the trace norm between two quantum states gives the optimal probability of distinguishing two quantum states from one measurement. Similarly, we have that the $1 \to 1$ norm between two quantum channels gives twice the optimal probability of distinguishing them by performing a measurement on the output state. Thus, it is natural to define the l_1-mixing time of a quantum Markov chain to be given by

Definition 4.1.3 (l_1- mixing time). *Given a quantum Markov chain $\{T_t\}_{t \in I}$ with stationary state $\sigma \in \mathcal{D}_d$ we define the l_1-mixing time for some $\epsilon > 0$, denoted by $t_1(\epsilon)$, to be given by*

$$t_1(\epsilon) = \inf\{t \in I | \forall \rho \in \mathcal{D}_d : \|T_t(\rho) - \sigma\|_1 \leq \epsilon\}.$$

The interpretation of this quantity is clear. It tells us how long the quantum dynamical semigroup has to run so that we may not distinguish the current state from the stationary state with probability greater than ϵ regardless of the initial state. Another standard measure is the fidelity F between two quantum states $\rho, \sigma \in \mathcal{D}_d$, which is given by

$$F(\rho, \sigma) = \|\sqrt{\rho}\sqrt{\sigma}\|_1. \tag{4.2}$$

It is easy to see that $0 \leq F(\rho, \sigma) \leq 1$ and $F(\sigma, \rho) = 1$ if and only if, $\rho = \sigma$. The fidelity is one of the most used distance measures in quantum information theory, as it is easily accessible experimentally. It also has an operational interpretation, although it is not as natural as for the trace distance. Given two quantum states $\sigma, \rho \in \mathcal{D}_d$, their fidelity is just

$$\begin{aligned} &\sup_{T, |\psi\rangle, |\phi\rangle} |\langle\psi|\phi\rangle|^2 \\ &\text{s.t. } T : \mathcal{M}_D :\to \mathcal{M}_d \text{ is a quantum channel and} \\ &T(|\psi\rangle\langle\psi|) = \sigma, \quad T(|\phi\rangle\langle\phi|) = \rho. \end{aligned}$$

As in the case of Schatten norms, one can also lift the fidelity, a distance measure between states, to the average fidelity, a distance measure for quantum channels.

Definition 4.1.4 (Average fidelity). *Let $T_1, T_2 : \mathcal{M}_d \to \mathcal{M}_d$ be quantum channels. The average fidelity between T_1 and T_2, $\mathcal{F}(T_1, T_2)$, is defined to be*

$$\mathcal{F}(T_1, T_2) = \int \|\sqrt{T_1(|\psi\rangle\langle\psi|)}\sqrt{T_2(|\psi\rangle\langle\psi|)}\|_1 d\psi,$$

where we integrate over the Haar measure on quantum states.

4.1.1 Noncommutative lp Spaces

Given a full rank state $\sigma \in \mathcal{D}_d^+$, one can define the following generalization of the classical (l^p, μ) spaces, with μ a probability measure.

Definition 4.1.5 (σ p-Norm). *Let $\sigma \in \mathcal{D}_d^+$ and $p \in [1, +\infty)$. We define the noncommutative $p-$norm for $X \in \mathcal{M}_d$ to be given by*

$$\|X\|_{p,\sigma}^p = \operatorname{Tr}\left(|\sigma^{\frac{1}{2p}} X \sigma^{\frac{1}{2p}}|^p\right). \tag{4.3}$$

For $p = +\infty$ we set $\|X\|_{\infty,\sigma} = \|X\|_\infty$.

The motivation to choose this particular generalization of the norm comes from interpolation theory [30]. Given these definitions, we may define $p \to q$ norms of linear operators between these spaces in an analogous way to those of Schatten norms (see Definition 4.1.2). If we fix $\|\cdot\|_{1,\sigma}$ to be given by Equation (4.3) and $\|\cdot\|_{\infty,\sigma}$ to just be the regular operator norm, then the definition of the other p-norms follows from the construction based on complex interpolation. This gives the norms many desirable features, such as the fact that they satisfy a Riesz-Thorin interpolation theorem:

Theorem 4.1.6 (Riesz-Thorin Interpolation Theorem). *Let $L : \mathcal{M}_d \to \mathcal{M}_d$ be a linear map, $1 \leq p_0 \leq p_1 \leq +\infty$ and $1 \leq q_0 \leq q_1 \leq +\infty$. For $\theta \in [0,1]$ define p_θ to satisfy*

$$\frac{1}{p_\theta} = \frac{\theta}{p_0} + \frac{1-\theta}{p_1}$$

and q_θ analogously. Then for $\sigma \in \mathcal{D}_d^+$ we have:

$$\|L\|_{p_\theta \to q_\theta,\sigma} \leq \|L\|^{\theta}_{p_0 \to q_0,\sigma} \|L\|^{1-\theta}_{p_1 \to q_1,\sigma}$$

As we will see later, this Theorem implies many fundamental inequalities for quantum Renyi divergences, such as the data processing inequality.

To obtain more accessible expressions when working with these spaces it is often useful to define the power operator $\Gamma_p : \mathcal{M}_d \to \mathcal{M}_d$ for some $p \in \mathbb{R}$. This operator acts as $\Gamma^p(X) = \sigma^{\frac{p}{2}} X \sigma^{\frac{p}{2}}$. We also adopt the convention $\Gamma^1 = \Gamma$.

As the Schatten norms, they also satisfy a Hölder inequality and $p = 2$ has a Hilbert space structure. The scalar product,$\langle \cdot | \cdot \rangle_\sigma$ is given by

$$\langle X | Y \rangle_\sigma = \operatorname{Tr}\left(X^\dagger \Gamma_\sigma(Y) \right).$$

It is then easy to check that a semigroup satisfies detailed balance w.r.t. σ if it is self-adjoint w.r.t. this scalar product.

However, the ordering of the noncommutative p-norms is reversed. That is, we have $\|X\|_{p,\sigma} \geq \|X\|_{q,\sigma}$ for $p \geq q$.

An important measure which might be used to quantify convergence of semigroups is the variance.

Definition 4.1.7 (Variance). *Let $X \in \mathcal{M}_d$ and $\sigma \in \mathcal{D}_d^+$. We define the variance of X w.r.t. to σ to be given by*

$$Var_\sigma(X) = \|X\|^2_{2,\sigma} - \|X\|^2_{1,\sigma}$$

Unlike the trace norm, this can be seen as a distance or convergence measure for observables instead of states. We define the l_2-mixing time of a semigroup, $t_2(\epsilon)$, to be given by

$$t_2(\epsilon) = \inf\{t \in I | \forall X \in \mathcal{M}_d \text{ s.t.} \|X\|_{1,\sigma} = 1 : \operatorname{Var}_\sigma(T_t^*(X)) \leq \epsilon\}.$$

A very useful concept in the study of convergence of semigroups is that of the relative density of a state w.r.t. another one. If $\sigma \in \mathcal{D}_d^+$ we may define the relative density of a state $\rho \in \mathcal{D}_d$ as $X_\rho = \Gamma^{-1}(\rho) = \sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}}$. Through relative densities, it is possible to upper-bound the trace distance between two states as follows. We have

$$\|\rho - \sigma\|_1^2 \leq \operatorname{Var}_\sigma(X_\rho). \tag{4.4}$$

4.2 Entropic Measures

4.2.1 Sandwiched Rényi divergences

In [6, 7] the following generalization of the Rényi divergence was proposed:

Definition 4.2.1 (Sandwiched p-Rényi divergence)**.** *Let $\rho, \sigma \in \mathcal{D}_d$. For $p \in (0,1) \cup (1, \infty)$, the* ***sandwiched p-Rényi divergence*** *is defined as:*

$$D_p\left(\rho \| \sigma\right) = \begin{cases} \frac{1}{p-1} \ln \left(\operatorname{Tr}\left(\left(\sigma^{\frac{1-p}{2p}} \rho \sigma^{\frac{1-p}{2p}} \right)^p \right) \right) & \textit{if } ker(\sigma) \subseteq ker(\rho) \textit{ or } p \in (0,1) \\ +\infty, & \textit{otherwise} \end{cases} \tag{4.5}$$

where $ker(\sigma)$ is the kernel of σ.

We may recover another important entropic quantity from taking limits of these entropies. One can show that

$$\lim_{p \to 1} D_p\left(\rho \| \sigma\right) = D\left(\rho \| \sigma\right) = \operatorname{Tr}\left(\rho\left(\log \rho - \log \sigma\right)\right),$$

where $D\left(\rho \| \sigma\right)$ is the quantum relative entropy [31].

They have been shown to satisfy a data processing inequality [32, 33], that is, they contract under quantum channels. Moreover, they have become a useful tool in quantum Shannon theory, as we will explain further in Chapter 6. They share a strong connection to the noncommutative l_p norms introduced in Definition 4.1.5. It is easy to see that

$$D_p\left(\rho \| \sigma\right) = \frac{\log(\|X\|_{p,\sigma}^p)}{p-1}.$$

This allows us to translate statements or inequalities for the $p-$norms to the sandwiched p-Rényi divergence. Observe that the conditions for a matrix $\rho \in \mathcal{M}_d$ to be a quantum state, $\operatorname{Tr}(\rho) = 1$ and $\rho \geq 0$, translate to a matrix $X \in \mathcal{M}_d$ being a relative density with respect to σ if we have $X \geq 0$ and $\|X\|_{1,\sigma} = 1$. From this one can easily prove that

$$\sup_{\rho \in \mathcal{D}_d} D_p\left(\rho \| \sigma\right) = \log\left(\|\sigma^{-1}\|_\infty\right). \tag{4.6}$$

For our purposes it will be important that we may upper-bound the trace norm through the sandwiched p-Rényi divergences using Pinsker's inequality:

Theorem 4.2.2 (Pinker's inequality)**.** *For the quantum Kullback-Leibler divergence we have*

$$\frac{1}{2}\|\sigma - \rho\|_1^2 \leq D\left(\rho \| \sigma\right) \leq D_p\left(\rho \| \sigma\right) \tag{4.7}$$

for any $p \geq 1$ and all $\rho, \sigma \in \mathcal{D}_d$. Moreover, Equation (4.7) *is optimal.*

Proof. See [34, Theorem 3.1] for a proof of Pinsker's inequality. The constant $\frac{1}{2}$ has been shown to be optimal in the classical case (see [35]), i.e. restricting to ρ that commute with σ, and is therefore also optimal here. □

Although not a distance measure, an important concept related to the Rényi divergences is that of Rényi entropies:

Definition 4.2.3 (Rényi entropy)**.** *Let $\rho \in \mathcal{D}_d$ and $p \in (1, \infty)$. We define the $p-$Rényi entropy, $S_p(\rho)$, to be given by*

$$S_p(\rho) = D_p\left(\rho \| \mathbb{1}\right) = \frac{\log\left(\operatorname{Tr}\left(\rho^{p-1}\right)\right)}{p-1}.$$

Moreover, for $p = 1$, we may take the limit $p \to 1$ and consistently define the von Neumann entropy to be given by

$$S_1(\rho) = S(\rho) = -\operatorname{Tr}\left(\rho \log\left(\rho\right)\right).$$

5

Functional Inequalities and Convergence of Semigroups

In this chapter, we will show how one can use functional inequalities to estimate the convergence of semigroups. We will mostly focus on continuous time. Most of the estimates will be based on the following simple observation: let $f(t,\rho) = d(T_t(\rho),\sigma)$ with f differentiable with respect to t, $\sigma \in \mathcal{D}_d^+$ a stationary state and $\rho \in \mathcal{D}_d$. For most distance measures, one expects that if the semigroup is primitive, then $T_t(\rho) \to \sigma$ exponentially fast, that is, we have

$$d(T_t(\rho),\sigma) \leq e^{-\mu t} d(\rho,\sigma), \tag{5.1}$$

where $\mu > 0$ is independent of ρ. It is then easy to see that an inequality like the one in Equation (5.1) is equivalent to the differential inequality

$$f(t,\rho) \leq -\mu \frac{\partial f(t,\rho)}{\partial t}. \tag{5.2}$$

Note that we assume that the inequality is valid for all ρ. One can then take the supremum over all μ that satisfy the inequality in Equation (5.2) and obtain the optimal exponential convergence rate for this particular distance measure under this semigroup. We will now discuss the exact form inequality (5.2) takes for different distance measures, what they imply for the convergence in the trace norm and how different convergence rates are related to each other. Although we formulated the general framework here for distance measures based on states, formulating everything in terms of relative densities will often lead to simpler expressions, as we will see. To this end, it is also useful to define the following

Definition 5.0.1 (Generator of the evolution of the relative density)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with stationary state $\sigma \in \mathcal{D}_d^+$. We define the generator of the evolution of the relative density, $\hat{\mathcal{L}}$, to be*

$$\hat{\mathcal{L}} = \Gamma^{-1} \circ \mathcal{L} \circ \Gamma.$$

As the name already suggests, it is easy to see that given some initial state $\rho \in \mathcal{D}_d$ and let $X = \Gamma_\sigma^{-1}(\rho)$ be its relative density w.r.t. σ. Then the relative density of $\rho_t = e^{t\mathcal{L}}(\rho)$ is given by $X_t = e^{t\hat{\mathcal{L}}}(\rho)$. That is, $e^{t\hat{\mathcal{L}}}$ describes the time evolution of the relative density. $\hat{\mathcal{L}}$ also has some other desirable properties, such as the fact that it has the same spectrum of $\mathcal{L}$, as it is just given by a similarity transformation on $\mathcal{L}$. Moreover if $\mathcal{L}$ satisfies detailed balance, we have that $\hat{\mathcal{L}} = \mathcal{L}^*$. Analogously to what we have done for the generator, we will denote the semigroup that gives the evolution of the relative density by $\hat{T}_t$.

5.1 Convergence in Variance: Spectral Gap

One of the simplest ways to bound the convergence of semigroups is through the spectral gap of the generator. It is what we get if we choose the variance (see Definition 4.1.7) as our distance measure. Let X be a relative density w.r.t. σ and T_t a primitive semigroup in continuous time. One can then show (see e.g. [8]) that

$$\frac{d}{dt}\mathrm{Var}_\sigma(X_t) = -\mathcal{E}_2^{\mathcal{L}}(X_t),$$

where $\mathcal{E}_2^{\mathcal{L}}(X) = -\frac{1}{2}\mathrm{Tr}\left(\hat{\mathcal{L}}(X)\sigma^{\frac{1}{2}}X\sigma^{\frac{1}{2}}\right)$ is the 2-Dirichlet Form. This motivates the following definition:

Definition 5.1.1 (Spectral Gap or Poincaré Inequality)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with stationary state $\sigma \in \mathcal{D}_d^+$. We define the spectral gap of $\mathcal{L}$, $\lambda(\mathcal{L})$, to be given by*

$$\lambda(\mathcal{L}) = \sup\{\lambda \in \mathbb{R}^+ | \forall X \in \mathcal{D}_{d,\sigma} : 2\lambda\, Var_\sigma(X) \leq \mathcal{E}_2^{\mathcal{L}}(X)\}.$$

From the previous discussion, it follows that

Theorem 5.1.2. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with stationary state $\sigma \in \mathcal{D}_d^+$ and spectral gap $\lambda > 0$. Then for all $X \geq 0$ we have*

$$Var_\sigma(e^{t\hat{\mathcal{L}}}X) \leq e^{-2\lambda t}\, Var_\sigma(X). \tag{5.3}$$

Moreover, we have the following bounds on mixing times for $\epsilon > 0$:

$$t_2(\epsilon) \leq \frac{1}{2\lambda}\log\left(\frac{\|\sigma^{-1}\|_\infty}{\epsilon}\right) \tag{5.4}$$

and

$$t_1(\epsilon) \leq \frac{1}{\lambda}\log\left(\frac{\|\sigma^{-1}\|_\infty}{\epsilon}\right) \tag{5.5}$$

Proof. As we have $\frac{d}{dt}\mathrm{Var}_\sigma(X_t) = -\mathcal{E}_2^{\mathcal{L}}(X_t)$, Equation (5.3) follows from the previous discussion. The bounds on $t_2(\epsilon)$ follows from (5.3) after noting that

$$\mathrm{Var}_\sigma(X) \leq \|\sigma^{-1}\|_\infty, \tag{5.6}$$

for $X \geq 0$ and $\|X\|_{1,\sigma} = 1$, from which it follows that

$$\mathrm{Var}_\sigma(e^{t\hat{\mathcal{L}}}X) \leq e^{-2\lambda t}\left(\|\sigma^{-1}\|_\infty\right). \tag{5.7}$$

Choosing t as in the r.h.s. of Equation 5.4 we obtain the claim. The statement in Equation (5.5) follows from the bound in Equation (4.4) and the previous discussion.

□

In case the Liouvillian $\mathcal{L}$ is reversible, it is not difficult to see that the spectral gap is nothing but the smallest nonzero eigenvalue of $-\mathcal{L}$ by the variational characterization of eigenvalues.

5.2 Convergence in Information-Theoretic Divergence

We will now discuss how to study the convergence of quantum dynamical semigroups under the sandwiched Rényi divergences introduced in Subsection 4.2.1 and how this convergence can be related to hypercontractive and logarithmic Sobolev inequalities.

5.2.1 Convergence of Rényi Divergences

If we take our distance measure d as in Equation (5.1) to be given by $d(\rho, \sigma) = D_p(\rho||\sigma)$, we may study its convergence using ideas similar to those introduced in Section 5.1. As in the case of the spectral gap, it is more convenient to express all quantities in terms of relative densities instead of states. In order to express derivatives of sandwiched Rényi divergences and later of l_p-norms it is convenient to introduce the following:

Definition 5.2.1 (Power Operator). *The power operator $I_{q,p} : \mathcal{M}_d \to \mathcal{M}_d$ for a state $\sigma \in \mathcal{D}_d$ and $q, p \in \mathbb{R}\backslash\{0\}$ is defined as*

$$I_{p,q}(X) = \Gamma_\sigma^{-\frac{1}{p}}\left(\left|\Gamma_\sigma^{\frac{1}{q}}(X)\right|^{\frac{q}{p}}\right).$$

Definition 5.2.2 (Dirichlet Form). *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For $p > 1$ we define the p-Dirichlet form of $X \in \mathcal{M}_d^+$ as*

$$\mathcal{E}_p^{\mathcal{L}}(X) = -\frac{p}{2(p-1)}\langle I_{q,p}(X)|\hat{\mathcal{L}}(X)\rangle_\sigma$$

where $\frac{1}{p} + \frac{1}{q} = 1$. For $p = 1$ we may take the limit $p \to 1$ and consistently define the 1-Dirichlet form by

$$\mathcal{E}_1^{\mathcal{L}}(X) = -\frac{1}{2}\mathrm{Tr}\left(\Gamma_\sigma\left(\hat{\mathcal{L}}(X)\right)(\log(\Gamma_\sigma(X)) - \log(\sigma))\right).$$

Note that both the power operator and the Dirichlet form depend on the state σ, which we omit from the notation as it will always be the stationary state of some semigroup which should be clear from context.

We also introduce the following functional:

Definition 5.2.3 (κ_p-Functional). *For any $p > 1$ we introduce the functional $\kappa_p : \mathcal{M}_d^+ \to \mathbb{R}$ as*

$$\kappa_p(X) = \frac{1}{p-1}\|X\|_{p,\sigma}^p \log\left(\frac{\|X\|_{p,\sigma}^p}{\|X\|_{1,\sigma}^p}\right) \tag{5.8}$$

for $X \in \mathcal{M}_d^+$. For $p = 1$ we may again take the limit $p \to 1$ and obtain $\kappa_1(X) := \lim_{p\to 1}\kappa_p(X) = \mathrm{Tr}(\Gamma_\sigma(X)(\log(\Gamma_\sigma(X)) - \log(\sigma)))$.

Note that κ_p is well-defined and non-negative as $\|X\|_{p,\sigma} \geq \|X\|_{1,\sigma}$ for $p \geq 1$. Strictly speaking, the definition also depends on a reference state $\sigma \in \mathcal{D}_d^+$, which we usually omit as it is always the fixed point of the primitive Liouvillian under consideration. The motivation to introduce this functional is the following , which we proved in [2, Theorem 3.1]:

Theorem 5.2.4 (Derivative of the sandwiched p-Rényi divergence). *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For any $\rho \in \mathcal{D}_d$ and $p > 1$ we have*

$$\frac{d}{dt}D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} = \frac{p}{p-1}\|X\|_{p,\sigma}^{-p}\langle I_{q,p}(X)|\hat{\mathcal{L}}(X)\rangle_\sigma \tag{5.9}$$

with $\frac{1}{p} + \frac{1}{q} = 1$.

We can see that the l.h.s. of Equation (5.9) is closely related to the Dirichlet form. Moreover, the κ_p-Functional is closely related to the sandwiched Rényi divergences , as one can see that

$$\frac{\frac{d}{dt}D_p(e^{t\mathcal{L}}(\rho)\|\sigma)}{D_p(e^{t\mathcal{L}}(\rho)\|\sigma)} = -\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)}.$$

For the relative density X associated to ρ. From the discussion preceding Equation (5.1) it follows that a lower bound on $\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)}$ for all relative densities X implies the exponential convergence the sandwiched Rényi divergences . We therefore define

Definition 5.2.5 (Entropic convergence constant for p-Rényi divergence)**.** *For any primitive Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ and $p \geq 1$ we define*

$$\beta_p(\mathcal{L}) = \sup\{\beta \in \mathbb{R}^+ : \beta\kappa_p(X) \leq \mathcal{E}_p^{\mathcal{L}}(X) \forall X \in \mathcal{D}_{d,\sigma}\}. \tag{5.10}$$

We then know that

$$D_p\left(e^{t\mathcal{L}}(\rho)\|\sigma\right) \leq e^{-2\beta_p(\mathcal{L})t} D_p\left(\rho\|\sigma\right). \tag{5.11}$$

From this, it is straightforward to derive mixing time bounds.

Theorem 5.2.6 (Mixing time from Rényi convergence)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$t_1(\epsilon) \leq \frac{1}{2\beta_p(\mathcal{L})} \log\left(\frac{2\log\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right).$$

Proof. From Pinsker's inequality (see Theorem 4.2.2) and the fact that $D_p\left(\rho\|\sigma\right) \leq \log\left(\ \|\sigma^{-1}\|_\infty\right)$ it follows that

$$\log\left(\|\sigma^{-1}\|_\infty\right) e^{-2\beta_p(\mathcal{L})t} \geq \frac{1}{2}\|e^{t\mathcal{L}}(\rho) - \sigma\|_1^2.$$

for any $\rho \in \mathcal{D}_d$. The claim follows after rearranging the terms. □

5.2.2 Logarithmic Sobolev Inequalities and Hypercontractivity

One of the most powerful frameworks to obtain convergence bounds for semigroups is that of Logarithmic Sobolev (LS) inequalities and their relation to hypercontractive inequalities. This framework was first introduced in the quantum setting in [36] and further developed from the quantum information perspective in [8]. However, based on the results of [2], we will take the approach of understanding them through sandwiched Rényi divergences . We will only show the results that establish the connection between convergence results in Section 5.2.3, but will introduce the basic definitions and give the intuition why studying hypercontractive inequalities is related to convergence already in this section. We begin by introducing one of the main functionals in the study of LS inequalities:

Definition 5.2.7 (p-relative entropy)**.** *For a full rank state $\sigma \in \mathcal{M}_d^+$ and $p > 1$ we define the p-relative entropy of $X \in \mathcal{M}_d^+$ as*

$$Ent_{p,\sigma}(X) = \langle I_{q,p}\left(X\right), S_p(X)\rangle_\sigma - \|X\|_{p,\sigma}^p \log\left(\|X\|_{p,\sigma}\right), \tag{5.12}$$

where $\frac{1}{q} + \frac{1}{p} = 1$. For $p = 1$ we can consistently define

$$Ent_{1,\sigma}(X) = \mathrm{Tr}\left(\Gamma_\sigma(X)\left(\log\left(\Gamma_\sigma(X)\right) - \log(\sigma)\right)\right).$$

by taking the limit $p \to 1$.

Although this is called a relative entropy in the literature, note that this is not a relative entropy in the information theoretic sense, although it is related to the usual relative entropy through

$$\mathrm{Ent}_{p,\sigma}\left(I_{p,1}\left(\Gamma_\sigma^{-1}\left(\rho\right)\right)\right) = \frac{1}{p} D\left(\rho\|\sigma\right).$$

We are now ready to introduce LS constants:

Definition 5.2.8 (Logarithmic Sobolev Inequalities)**.** *For a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and $p \geq 1$ the p-logarithmic Sobolev constant is defined as*

$$\alpha_p\left(\mathcal{L}\right) = \sup\{\alpha \in \mathbb{R}^+ : \alpha Ent_{p,\sigma}(X) \leq \mathcal{E}_p^{\mathcal{L}}(X) \text{ for all } X > 0\}. \tag{5.13}$$

At first sight, it is not clear how this inequality fits into the framework of differential inequalities we discussed before unless $p = 1$, as in this case $\mathrm{Ent}_{1,\sigma} = \kappa_1$ and we obtain $\alpha_1 = \beta_1$. That is, α_1 characterizes the optimal convergence rate in the usual relative entropy. The connection to differential inequalities comes from the intimate relationship between LS inequalities and hypercontractivity for $p \geq 2$:

Lemma 5.2.9. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and $p \geq 2$ and suppose that*

$$\|\hat{T}_t\|_{p \to p(t),\sigma} \leq 1 \tag{5.14}$$

with $p(t) = 1 + (p-1)e^{2\alpha t}$. Then $\alpha_p(\mathcal{L}) \geq \alpha$.

Proof. As it is shown in e.g. [8, Lemma 14], we have for any $X \in \mathcal{M}_d^+$ that

$$\frac{d}{dt} \log \left(\|\hat{T}_t(X)\|_{p,\sigma} \right) \big|_{t=0} = \frac{p'(0)}{p(0)\|X\|_{p,\sigma}} \left(\mathrm{Ent}_{p,\sigma}(X) - \frac{1}{\alpha} \mathcal{E}_p^{\mathcal{L}}(X) \right).$$

Define the function $f(t) = \|\hat{T}_t(X)\|_{p(t),\sigma}$. Note that we clearly have $f(0) = \|\hat{T}_0(X)\|_{p,\sigma}$ and from Equation (5.14) it follows that $\frac{d}{dt} \log(f(t)) \big|_{t=0} \leq 0$. To see this, note that $\frac{d}{dt} \log(f(t)) \big|_{t=0} > 0$ would imply that $\|\hat{T}_{t_0}\|_{p \to p(t),\sigma} > 1$ for t_0 sufficiently small, as log is a monotone function. This contradicts Equation (5.14). As X was arbitrary, it follows that

$$\mathrm{Ent}_{p,\sigma}(X) - \frac{1}{\alpha} \mathcal{E}_p^{\mathcal{L}}(X) \leq 0,$$

which is equivalent to $\alpha_p \geq \alpha$. □

From Lemma 5.2.9 we see that a LS inequality is implied by a hypercontractive inequality by looking at the derivative of $p \to q$ norms under semigroups. Thereofre, hypercontractive inequalities provide a framework to obtain lower bounds on LS constants, which in turn can be used to obtain bounds on β_p, as we will see later. We also note that under some technical assumptions (see e.g. [8, Theorem 15]) the converse to Lemma 5.2.9 also holds: one can derive hypercontractive inequalities from a LS inequality. At last, we would like to provide some intuition why hypercontractive inequalities are related to convergence in the first place, as it might seem elusive why one should consider them. The ratio $\|X\|_{p,\sigma}\|X\|_{q,\sigma}^{-1}$ for $p > q \geq 1$ provides a measure of how flat the spectrum of X is weighted by σ. To see this, consider X commuting with σ. Then we have

$$\|X\|_{p,\sigma}^p = \sum_{i=1}^d \sigma_i X_i^p,$$

where σ_i, X_i are the eigenvalues of σ and X, respectively. If the X_i are close to each other for large eigenvalues of σ, then $\|X\|_{p,\sigma} \simeq \|X\|_{q,\sigma}$ and they will differ significantly if the spectrum is not flat. Moreover, the larger the difference between p and q, the more accentuated this difference becomes. As for a primitive semigroup we have $\hat{T}_t(X) \to \mathrm{Tr}(X\sigma)\mathbb{1}$ as $t \to \infty$, all initial operators converge to an operator with a constant, and therefore very flat, spectrum. A hypercontractive inequality like that in Equation (5.14) therefore allows us to quantify how fast the spectrum of operators becomes flat under the semigroup. Therefore, it should be expected that this also gives information on the convergence of the semigroup.

5.2.3 Relations between the constants

All the constants introduced in this section are related to each other. We will only formulate the relation for semigroups that are reversible, as the general result is a bit more involved and technical, but remark that similar relations hold in general.

Theorem 5.2.10 (Relation between constants)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive, reversible Liouvillian with stationary state $\sigma \in \mathcal{D}_d^+$. Then, for all $p \geq 1$:*

$$\lambda\left(\mathcal{L}\right) \geq \beta_p(\mathcal{L}) \geq \frac{\alpha_p(\mathcal{L})}{p}. \tag{5.15}$$

Proof. As $\mathcal{L}$ is reversible, there is a s.a. $X \in \mathcal{M}_d$ such that $\mathcal{L}(X) = \lambda\left(\mathcal{L}\right) X$. For ϵ_0 small enough, $\mathbb{1} + \epsilon X$ is positive for all $\epsilon \in (0, \epsilon_0)$. One can then show that (see [2, Theorem 4.1])

$$\lim_{\epsilon \to 0} \frac{\mathcal{E}_p^{\mathcal{L}}(\mathbb{1} + \epsilon X)}{\kappa_p(\mathbb{1} + \epsilon X)} = \lambda\left(\mathcal{L}\right).$$

This implies the upper bound in Equation 5.15 by the variational definition of β_p. To show the lower bound, note that it holds that

$$\frac{\kappa_p(Y)}{p} \leq \mathrm{Ent}_{p,\sigma}(Y) \tag{5.16}$$

for all $Y \in \mathcal{M}_d^+$. We refer to [2, Lemma 4.2] for a proof of Equation (5.16). It then follows from Equation (5.16) that

$$\frac{\alpha_p(\mathcal{L})}{p}\kappa_p(Y) \leq \alpha_p(\mathcal{L})\mathrm{Ent}_{p,\sigma}(Y) \leq \mathcal{E}_p^{\mathcal{L}}(Y)$$

for all $Y \in \mathcal{M}_d^+$ and the claim follows from the variational definition of β_p. □

Theorem 5.2.10 has implications for convergence rates of semigroups. First of all, it shows that a LS inequality always implies a mixing time bound, as can easily be seen by combining it with Theorem B.1.3. This was only known to hold under further assumptions on the semigroup [8]. As remarked before, if the spectral gap and β_p are of the same order, then β_p gives an exponentially better mixing time, but the convergence rate cannot be larger than the one given by the spectral gap. The proof of the Theorem also illustrates that the convergence rate in the sandwiched Rényi divergences is governed by the spectral gap for operators that are close to the identity. One can also show inequalities in the other direction for $p = 2$. That is, lower bounds on β_p or α_p in terms of the spectral gap. But they must have a dimensional factor, as we will explain below. To illustrate these bounds, we show:

Theorem 5.2.11. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank stationary state $\sigma \in \mathcal{D}_d^+$. Then*

$$\beta_2\left(\mathcal{L}\right) \geq \lambda\left(\mathcal{L}\right) \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\log\left(\|\sigma^{-1}\|_\infty\right)}. \tag{5.17}$$

Moreover, this inequality is tight.

Proof. We have for all $X \in \mathcal{M}_d^+$ that

$$\frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\log\left(\|\sigma^{-1}\|_\infty\right)} \kappa_p(X) \leq \mathrm{Var}_\sigma(X).$$

To prove this, one can e.g. compute the optimal constant for the depolarizing semigroups. The claim then follows from the variational definition of β_p, as

$$\lambda\left(\mathcal{L}\right) \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\log\left(\|\sigma^{-1}\|_\infty\right)} \kappa_p(X) \leq \lambda\left(\mathcal{L}\right) \mathrm{Var}_\sigma(X) \leq \mathcal{E}_2^{\mathcal{L}}(X)$$

holds for all $X \in \mathcal{M}_d^+$. Moreover, for the Liouvillian $\mathcal{L}_\sigma(X) = \mathrm{Tr}\left(X\right)\sigma - X$, we have (see [2, Theorem 3.3]):

$$\beta_2(\mathcal{L}_\sigma) = \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\log\left(\|\sigma^{-1}\|_\infty\right)}, \quad \lambda\left(\mathcal{L}\right) = 1,$$

which shows that the inequality is tight. □

One can also show lower bounds on α_2 in terms of β_2 [2] for reversible Liouvillians and all bounds available have dimensional factors, although they are of order $\log(\log(\|\sigma^{-1}\|_\infty))$. It is an open question if this is indeed optimal.

5.2.4 Tensorization

As we will see later in Chapter 6, it is of central importance to applications to obtain functional inequalities that tensorize. That is, inequalities of the form

$$\mu(\mathcal{L}^{(n)}) \geq c, \tag{5.18}$$

where μ is one of the constants discussed in the last section, such as the spectral gap or the LS constant, and $c \in \mathbb{R}$ is a constant which is independent of $n \in \mathbb{N}$. It is easy to see that the spectral gap tensorizes, that is if we show a bound on the spectral gap of $\mathcal{L}$, it also holds for $\mathcal{L}^{(n)}$ for all n. It is known that LS inequalities tensorize for classical Markov chains [37] but this is unexpected to hold in the quantum case. By the equivalence between hypercontractivity and LS constants, the tensorization of a LS inequality would imply that

$$\|\hat{T}_t^{\otimes n}\|_{2\to q,\sigma} = \|\hat{T}_t\|_{2\to q,\sigma}^n \tag{5.19}$$

for small t. As shown in [38], the $p \to q$ norm of quantum channels is not multiplicative in general and it is therefore not expected that Equation 5.19 holds in general, although the multiplicativity was established in some cases [39, 40]. It is therefore not immediately clear that an inequality like that of Equation 5.18 can hold for the LS inequality or the β_p constants. It follows from the results of Section 5.2.3 that it is sufficient to obtain lower bounds on the LS-2 constant to also obtain lower bounds on other quantities of interest. Therefore, we focus on obtaining lower bounds on LS-2 constants and outline a strategy developed in [41] to accomplish this task, although other approaches exist, e.g. the one followed in [42]. Note that for all functional inequalities discussed here, only the Dirichlet form depends on the semigroup in question, the other functional involved depends only on the stationary state. Thus, if we can show a comparison inequality of the form

$$\mathcal{E}_2^{\mathcal{L}} \leq a\mathcal{E}_2^{\mathcal{L}'} \tag{5.20}$$

for some constant $a \in \mathbb{R}$ and $\mathcal{L}, \mathcal{L}' : \mathcal{M}_d \to \mathcal{M}_d$ primitive Liouvillians with the same stationary state, we may obtain lower bounds on many functional inequalities for $\mathcal{L}'$ from functional inequalities for $\mathcal{L}$. As shown in [41], inequalities like that in Equation 5.20 tensorize. We therefore pursue the following strategy to obtain functional inequalities that tensorize for a generator $\mathcal{L}'$:

1. Pick a generator $\mathcal{L}$ with a particularly simple structure and show a functional inequality that tensorizes, such as a LS-2 inequality.
2. Show an inequality like that of Equation 5.20 relating the Dirichlet forms.

With this general strategy, it is possible to obtain many different tensorization results, as illustrated in [2, 41]. In particular, one can show using this strategy with depolarizing semigroups that we have

$$\alpha_2(\mathcal{L}^{(n)}) \geq 2\frac{\lambda(\mathcal{L})}{2\left(\ln\left(d^4\|\sigma^{-1}\|_\infty\right) + 11\right)}$$

for all $n \in \mathbb{N}$.

6

Applications of Convergence Bounds

In this chapter, we will briefly discuss some of the applications of convergence estimates for quantum dynamical semigroups.

6.1 Dissipative preparation of states

Markov chain Monte Carlo methods are ubiquitous in the design of classical algorithms. It is expected that with the advent of the quantum computer we will be able to implement powerful algorithms that are based on sampling not from a classical distribution but rather quantum states. Therefore, one needs techniques to (approximately) prepare states of interests. One of the ways proposed to do this is through dissipative processes [43]. That is, to design a primitive quantum dynamical semigroup that drives the system to the state we want to sample from. The runtime of algorithms based on such ideas is given by the mixing time of the chain and it is, therefore, of central importance to develop tools to estimate this. Although one can perform universal quantum computation in this computational model [43], the dissipative preparation of one class of states is of particular importance, namely Gibbs states. This class is particularly important because these states describe quantum systems in thermal equilibrium and preparing them is therefore important to many questions of physical relevance. Moreover, in the many-body setting, understanding how the mixing time scales with the size of the system has deep physical implications. These range from statements on the correlations of the Gibbs state [44] and on the stability of these states and preparation processes to perturbations [45, 46]. Tools based on functional inequalities such as the ones described in Section 5 played a central role in the establishment of such connections in the classical setting [47], and one can expect that further developing them is of importance to translating these statements to the quantum setting. A further motivation to focus on Gibbs states are the recent quantum algorithms based on Gibbs sampling to solve semidefinite programs [48]. These algorithms are based on subroutines in which Gibbs states are prepared and their runtime is again given by how fast we can prepare a given Gibbs state of interest. Thus, it is no exaggeration to say that the preparation of Gibbs states is one of the central tasks of quantum computation.

6.2 Entropic inequalities and Capacity Bounds

It is possible to derive entropic inequalities and capacity bounds from the convergence bounds of the form

$$D_p\left(T_t(\rho)||\sigma\right) \leq e^{-\alpha t} D_p\left(\rho||\sigma\right)$$

for some of the divergences D_p discussed before. One example is the following:

Theorem 6.2.1 (Entropy production and Convergence)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with stationary state $\frac{\mathbb{1}}{d} \in \mathcal{D}_d^+$ and $p \in [1, \infty)$. Then $\beta_p(\mathcal{L}) \geq \beta$ is equivalent to*

$$S_p(T_t(\rho)) \geq (1 - e^{-\beta t}) \log(d) - e^{-\beta t} S_p(\rho) \tag{6.1}$$

for all $\rho \in \mathcal{D}_d$ and $t \geq 0$.

Proof. Note that we have

$$D_p\left(\rho \middle\| \frac{\mathbb{1}}{d}\right) = \log(d) - S_p(\rho).$$

As $\beta_p(\mathcal{L}) \geq \beta$ is equivalent to

$$D_p\left(T_t(\rho) \middle\| \frac{\mathbb{1}}{d}\right) \leq e^{-\beta t} D_p\left(\rho \middle\| \frac{\mathbb{1}}{d}\right)$$

for all $\rho \in \mathcal{D}_d$ and $t \geq 0$. The claim then follows by rearranging the terms. □

Statements like that of Theorem 6.2.1 are useful when estimating important parameters of quantum channels, such as its minimal output entropy [41]. Other examples of entropic inequalities that can be obtained this way include estimates for the concavity of the von Neumann entropy [1]. Another quantity we may bound from convergence inequalities is the information radius:

Definition 6.2.2 (p-information radius)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel and $p \in [1, \infty)$. The* p-information radius*, $K_p(T)$, of T is defined as*

$$K_p(T) = \frac{1}{\log(2)} \min_{\sigma \in \mathcal{D}_d} \max_{\rho \in \mathcal{D}_d} D_p(T(\rho) \| \sigma). \tag{6.2}$$

The connection between the information radius and convergence inequalities is given by the following theorem:

Theorem 6.2.3. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full-rank fixed point $\sigma \in \mathcal{D}_d^+$ such that for some $p \in (1, \infty)$ $\beta_p(\mathcal{L}) \geq \beta$. Then*

$$K_p(T_t) \leq e^{-\beta t} \log_2\left(\|\sigma^{-1}\|_\infty\right).$$

Proof. It is clear by the definition of the information radius in Equation 6.2 that $K_p(T_t) \leq \log(2)^{-1} \max\limits_{\rho \in \mathcal{D}_d} D_p(T_t(\rho) \| \sigma)$. Applying the convergence bound and Equation (4.6) the claim follows. □

The relevance of the information radius follows from its strong connection to the classical capacity of a quantum channel.

When classical information is encoded in a quantum state and sent through a quantum channel, the classical capacity is the supremum of transmission rates such that the probability for a decoding error goes to 0 as we allow arbitrarily many uses of the channel. In general, it is not possible to retrieve the information perfectly when it is sent over a finite number of channels, and the probability for successful decoding will be smaller than 1. Here we want to derive bounds on this probability for quantum dynamical semigroups. More specifically we are interested in strong converse bounds on the classical capacity. An upper bound on the capacity is called a strong converse bound if whenever a transmission rate exceeds the bound the probability of successful decoding goes to zero as the number of channel uses goes to infinity.

We refer to [21, Chapter 12] for the exact definition of the classical capacity $C(T)$ of a quantum channel T and to [7, 49, 50, 51, 52] for more details on strong converses and strong converse bounds.

We will call a coding scheme for the transmission of m classical bits via n uses of the channel T for which the probability of successful decoding is p (see again [21, Chapter 12] for an exact definition) a (m, n, p)-coding scheme for classical communication using a quantum channel T. The following theorem shown in [7, Section 6] relates the information radius and the probability of successful decoding and thus to classical capacities:

Theorem 6.2.4 (Bound on the success probability in terms of information radius)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, $n \in \mathbb{N}$ and $R \geq 0$. For any (nR, n, p_{succ})-coding scheme for classical communication via T we have*

$$p_{succ} \leq 2^{-n\left(\frac{p-1}{p}\right)\left(R - \frac{1}{n}K_p(T^{\otimes n})\right)}. \tag{6.3}$$

Putting these pieces together we obtain

Theorem 6.2.5. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full-rank fixed point $\sigma \in \mathcal{D}_d^+$ such that for some $p \in (1, \infty)$ there exists $\beta > 0$ fulfilling $\beta_p(\mathcal{L}^{(n)}) \geq \beta$ for all $n \in \mathbb{N}$. Then for any (nR, n, p_{succ})-coding scheme for classical communication via the quantum dynamical semigroup $T_t = e^{t\mathcal{L}}$ we have*

$$p_{succ} \leq 2^{-n\left(\frac{p-1}{p}\right)\left(R - e^{-2ct}\log_2(\|\sigma^{-1}\|_\infty)\right)}.$$

Proof. Combining Equation (C.7) and Equation (4.6) we have

$$K_p(T_t^{\otimes n}) \leq \frac{1}{\log(2)} \max_{\rho \in \mathcal{D}_{d^n}} D_p(T_t^{\otimes n}(\rho)\|\sigma) \leq \frac{n}{\log(2)} e^{-2\beta_p(\mathcal{L}^{(n)})t} \log_2(\|\sigma^{-1}\|_\infty).$$

Now Theorem 6.2.4 together with the assumption $\beta_p(\mathcal{L}^{(n)}) \geq \beta$ finishes the proof. □

Theorem 6.2.5 implies in particular that the classical capacity $C(T_t)$ is bounded by

$$e^{-\beta t} \log_2(\|\sigma^{-1}\|).$$

Exploring the relation between β_p and LS inequalities and tensorization results for LS inequalities discussed in Sections 5.2.2 and 5.2.4, respectively, it is then possible to derive bounds on the classical capacity in terms of the spectral gap. It is notoriously difficult to compute the classical capacity of a quantum channel [53, 54]. Through these methods, it is possible to obtain bounds on this quantity for new classes of quantum channels, such as the semigroups generated by Davies generators. These can then be used to study the lifetime of quantum memories under thermal noise.

6.3 Randomized Benchmarking

Another area in which semigroups of quantum channels play a central role is randomized benchmarking [15]. The goal of randomized benchmarking is to efficiently estimate the average gate fidelity of a set of quantum gates $\{U_g\}_{g \in G} \subset U(d)$ that are a representation of a (usually finite) group G. This is an important figure of merit when assessing the quality of a given implementation of the gates and how long we may perform computations. One makes the further assumption that all gates are affected by the same noise, that is, there is a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ such that for all $g \in G$ we actually implement $T_g = \mathcal{U}_g \circ T$. We then wish to estimate $\mathcal{F}(T_g, \mathcal{U}_g)$. One can show that

Lemma 6.3.1. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, $\{U_g\}_{g \in G} \subset U(d)$ a unitary representation of a finite group G and $T_g = \mathcal{U}_g \circ T$. Then*

$$\mathcal{F}(T_g, \mathcal{U}_g) = \frac{d\mathrm{Tr}\,(T) + d^2}{d^2(d+1)}.$$

Here we mean the trace of T as a linear operator between the vector spaces $\mathcal{M}_d$.

Lemma 6.3.1 implies that to estimate the average gate fidelity in the scenario described above, we only need to determine the trace of the channel T. The next crucial observation for the randomized benchmarking protocol is that twirling the channel with respect to the group G does not change the trace of a channel. Working with the twirled channel $\mathcal{T}(T)$, which is covariant with respect to the group, simplifies the estimation of its trace. We will not discuss the randomized benchmarking protocol in detail, but remark that by preparing random sequences of gates it allows us to estimate the quantity:

$$\operatorname{Tr}\left(\mathcal{T}(T)^m(\rho)E_i\right) \tag{6.4}$$

for some initial state $\rho \in \mathcal{D}_d$ and POVM element E_i. To make the analysis more concrete and as it is the simplest case, we will now restrict to the case in which the group G is given by the Clifford group. It plays a central role in the theory of quantum computation and error correction [55] and is the usual setting for randomized benchmarking, although we have extended the protocol to arbitrary representations. One of the reasons one usually considers the Clifford group is that we have

$$\mathcal{T}(T) = T_{t_0,\frac{\mathbb{1}}{d}}.$$

for some $t_0 \in \mathbb{R}^+$. We refer to [15] for a proof of this fact. Here $T_{t,\frac{\mathbb{1}}{d}}$ is again the depolarizing channel. It is easy to see that we have $\operatorname{Tr}\left(T_{t_0,\frac{\mathbb{1}}{d}}\right) = (d^2-1)e^{-t_0}+1$. Therefore, we only have to estimate the parameter t_0 to obtain an estimate on the average fidelity. This corresponds to estimating how fast the semigroup $\{T^n_{t_0,\frac{\mathbb{1}}{d}}\}_{n\in\mathbb{N}}$ converges. As we have access to the expectation values in Equation (6.4), we can plug in the exact form of $\mathcal{T}(T)$ and see that

$$\operatorname{Tr}\left(\mathcal{T}(T)^m(\rho)E_i\right) = \operatorname{Tr}\left(\frac{\mathbb{1}}{d}E_i\right) + e^{-t_0 m}\operatorname{Tr}\left(\left(\rho - \frac{\mathbb{1}}{d}\right)E_i\right).$$

By fitting the experimental data to a curve of the form $f(m) = A + Be^{-\lambda m}$ we may then estimate t_0, the convergence speed and the average fidelity.

Bibliography

[1] A. Müller-Hermes, D. Stilck França, and M. M. Wolf. **Relative entropy convergence for depolarizing channels**. *J. Math. Phys.*, **57**(2), 2016. 1, 2, 5, 14, 34

[2] A. Müller-Hermes and D. Stilck Franca. **Sandwiched Renyi Convergence for Quantum Evolutions**. *Quantum*, **2**:55, 2018. 1, 2, 27, 28, 30, 31, 42, 43

[3] Daniel Stilck França and Anna-Lena Hashagen. **Approximate Randomized Benchmarking for Finite Groups**. *arXiv preprint quant-ph/1803.03621*, 2018. 2, 4

[4] D. Stilck França. **Perfect Sampling for Quantum Gibbs States**. *Quantum Info. Comput.*, **18**(5), 2018. 2, 3

[5] A. Bluhm and D. Stilck Franca. **Dimensionality reduction of SDPs through sketching**. *arXiv preprint 1707.09863*, 2017. 2, 3

[6] M. Müller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel. **On quantum Rényi entropies: A new generalization and some properties**. *J. Math. Phys.*, **54**(12):122203, 2013. 2, 23, 62

[7] M. M. Wilde, A. Winter, and D. Yang. **Strong Converse for the Classical Capacity of Entanglement-Breaking and Hadamard Channels via a Sandwiched Rényi Relative Entropy**. *Comm. Math. Phys.*, **331**:593–622, 2014. 2, 23, 34, 35, 62, 63

[8] M. J. Kastoryano and K. Temme. **Quantum logarithmic Sobolev inequalities and rapid mixing**. *J. Math. Phys.*, **54**(5):052202, 2013. 2, 26, 28, 29, 30, 62

[9] K. Temme. **Lower bounds to the spectral gap of Davies generators**. *J. Math. Phys.*, **54**(12):122110, 2013. 2

[10] K. Temme. **Thermalization Time Bounds for Pauli Stabilizer Hamiltonians**. *Comm. Math. Phys.*, **350**(2):603–637, 2017. 2, 15, 63

[11] J. G. Propp and D. B. Wilson. **How to Get a Perfectly Random Sample from a Generic Markov Chain and Generate a Random Spanning Tree of a Directed Graph**. *Journal of Algorithms*, **27**(2):170 – 217, 1998. 3, 102

[12] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete. **Quantum Metropolis sampling**. *Nature*, **471**(7336):87–90, 2011. 3, 13, 15, 16, 102

[13] Y. Li and D. P Woodruff. **Tight bounds for sketching the operator norm, schatten norms, and subspace embeddings**. In *LIPIcs-Leibniz International Proceedings in Informatics*, **60**, 2016. 4

[14] A. W. Harrow, A. Montanaro, and A. J. Short. **Limitations on Quantum Dimensionality Reduction**. In *Automata, Languages and Programming*, 2011. 4, 137

[15] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. **Randomized benchmarking of quantum gates**. *Phys. Rev. A*, **77**:012307, 2008. 4, 35, 36, 173

[16] I. Kim and M. B. Ruskai. **Bounds on the concavity of quantum entropy**. *J. Math. Phys.*, **55**(9):–, 2014. 5

[17] T. Heinosaari and M. Ziman. *The mathematical language of quantum theory: from uncertainty to entanglement*. Cambridge University Press, 2011. 7, 8, 11

[18] D. Burgarth, G. Chiribella, V. Giovannetti, P. Perinotti, and K. Yuasa. **Ergodic and mixing quantum channels in finite dimensions**. *New J. Phys.*, **15**(7):073045, 2013. 12

[19] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete. **The χ^2-divergence and mixing times of quantum Markov processes**. *J. Math. Phys.*, **51**(12):122201, 2010. 12

[20] A. Y. Kitaev. **Quantum measurements and the Abelian Stabilizer Problem**. *arXiv preprint quant-ph/9511026*, 1995. 13

[21] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000. 13, 22, 34, 35

[22] H. Spohn and J. L. Lebowitz. **Irreversible thermodynamics for quantum systems weakly coupled to thermal reservoirs**. *Adv. Chem. Phys*, **38**:109–142, 1978. 14

[23] E.B. Davies. *Quantum Theory of Open Systems*. Academic Press, 1976. 14

[24] H.P. Breuer and F. Petruccione. *The Theory of Open Quantum Systems*. OUP Oxford, 2007. 14, 15

[25] E.B. Davies. **Generators of dynamical semigroups**. *J. Funct. Anal.*, **34**(3):421 – 432, 1979. 14

[26] A. Kossakowski, A. Frigerio, V. Gorini, and M. Verri. **Quantum detailed balance and KMS condition**. *Comm. Math. Phys.*, **57**(2):97–110, 1977. 15

[27] H. Spohn. **Entropy production for quantum dynamical semigroups**. *J. Math. Phys.*, **19**(5):1227–1230, 1978. 15

[28] D. A Levin and Y. Peres. *Markov chains and mixing times*, **107**. American Mathematical Soc., 2017. 17

[29] J.G. Kemény and J.L. Snell. *Finite markov chains*. University series in undergraduate mathematics. Van Nostrand, 1960. 18

[30] J. Bergh and J. Lofstrom. *Interpolation spaces: an introduction*, **223**. Springer, 2012. 23

[31] H. Umegaki. **Conditional expectation in an operator algebra. IV. Entropy and information**. *Kodai Math. Sem. Rep.*, **14**(2):59–85, 1962. 24

[32] S. Beigi. **Sandwiched Rényi divergence satisfies data processing inequality**. *J. Math. Phys.*, **54**(12):122202, 2013. 24

[33] R. L. Frank and E. H. Lieb. **Monotonicity of a relative Rényi entropy**. *J. Math. Phys.*, **54**(12):122201, 2013. 24

[34] F. Hiai, M. Ohya, and M. Tsukada. **Sufficiency, KMS condition and relative entropy in von Neumann algebras.** *Pacific J. Math.*, **96**(1):99–109, 1981. 24

[35] G. L. Gilardoni. **On Pinsker's and Vajda's Type Inequalities for Csiszar's f -Divergences**. *IEEE Trans. Inform. Theory*, **56**(11):5377–5386, 2010. 24

[36] R. Olkiewicz and B. Zegarlinski. **Hypercontractivity in Noncommutative Lp Spaces**. *Journal of Functional Analysis*, **161**(1):246 – 285, 1999. 28, 62

[37] P. Diaconis and L. Saloff-Coste. **Logarithmic Sobolev inequalities for finite Markov chains**. *Ann. Appl. Probab.*, **6**(3):695–750, 1996. 31

[38] M. B Hastings. **Superadditivity of communication capacity using entangled inputs**. *Nat. Phys.*, **5**(4):255, 2009. 31

[39] C. King. **Multiplicativity of Superoperator Norms for Some Entanglement Breaking Channels**. *Quantum Info. Comput.*, **14**(13-14):1203–1212, 2014. 31

[40] C. King. **Hypercontractivity for Semigroups of Unital Qubit Channels**. *Comm. Math. Phys.*, **328**(1):285–301, 2014. 31

[41] A. Müller-Hermes, D. Stilck França, and M. M. Wolf. **Entropy production of doubly stochastic quantum channels**. *J. Math. Phys.*, **57**(2), 2016. 31, 34

[42] T. Bodineau and B. Zegarlinski. **Hypercontractivity via spectral theory**. *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, **03**(01):15–31, 2000. 31

[43] F. Verstraete, M. M Wolf, and J. I. Cirac. **Quantum computation and quantum-state engineering driven by dissipation**. *Nat. Phys.*, **5**(9):633–636, 2009. 33

[44] M. J. Kastoryano and J. Eisert. **Rapid mixing implies exponential decay of correlations**. *J. Math. Phys.*, **54**(10):102201, 2013. 33

[45] T. S. Cubitt, A. Lucia, S. Michalakis, and D. Perez-Garcia. **Stability of Local Quantum Dissipative Systems**. *Comm. Math. Phys.*, **337**(3):1275–1315, 2015. 33

[46] A. Lucia, T. S Cubitt, S. Michalakis, and D. Pérez-García. **Rapid mixing and stability of quantum dissipative systems**. *Phys. Rev. A*, **91**(4):040302, 2015. 33

[47] D. W. Stroock and B. Zegarlinski. **The equivalence of the logarithmic Sobolev inequality and the Dobrushin-Shlosman mixing condition**. *Comm. Math. Phys.*, **144**(2):303–323, 1992. 33

[48] F. GSL Brandão and K. M. Svore. **Quantum Speed-Ups for Solving Semidefinite Programs**. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 415–426. IEEE, 2017. 33

[49] A. Winter. **Coding theorem and strong converse for quantum channels**. *IEEE Trans. Inform. Theory*, **45**(7):2481–2485, 1999. 34

[50] T. Ogawa and H. Nagaoka. **Strong converse to the quantum channel coding theorem**. *IEEE Trans. Inform. Theory*, **45**(7):2486–2489, 1999. 34

[51] R. König and S. Wehner. **A Strong Converse for Classical Channel Coding Using Entangled Inputs**. *Phys. Rev. Lett.*, **103**(7):070504, 2009. 34

[52] M. Tomamichel, M. M. Wilde, and A. Winter. **Strong converse rates for quantum communication**. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 2386–2390. IEEE, 2015. 34

[53] M. B Hastings. **A counterexample tao additivity of minimum output entropy**. *arXiv preprints 0809.3972*, 2008. 35

[54] Toby S. Cubitt, David Elkouss, William Matthews, Maris Ozols, David Pérez-García, and Sergii Strelchuk. **Unbounded number of channel uses may be required to detect quantum capacity**. *Nat Commun*, **6**, 2015. 35

[55] Daniel Gottesman. **Stabilizer codes and quantum error correction**. *arXiv preprint quant-ph/9705052*, 1997. 36

[56] Dorit Aharonov, Michael Ben-Or, Russell Impagliazzo, and Noam Nisan. **Limitations of noisy reversible computation**. *arXiv preprint quant-ph/9611028*, 1996. 43

[57] K. Temme, F. Pastawski, and M. J. Kastoryano. **Hypercontractivity of quasi-free quantum semigroups**. *J. Phys. A: Math. Theor.*, **47**(40):405303, 2014. 63

[58] X. Wang, W. Xie, and R. Duan. **Semidefinite programming converse bounds for classical communication over quantum channels**. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pages 1728–1732. IEEE, 2017. 136

[59] S. Boyd, P. Diaconis, and L. Xiao. **Fastest mixing Markov chain on a graph**. *SIAM review*, **46**(4):667–689, 2004. 136

[60] D. P. Woodruff. **Sketching as a Tool for Numerical Linear Algebra**. *Foundations and Trends in Theoretical Computer Science*, **10**(1–2):1–157, 2014. 136

Appendix A

Articles as co-author

A.1 Relative entropy convergence for depolarizing channels

Relative Entropy Convergence for Depolarizing Channels

A. Müller-Hermes, D. Stilck França and M.M. Wolf

We study the convergence of quantum Markovian time-evolutions generated by Liouvillians depolarizing to a full rank state in the relative entropy. These are Liouvillians $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ given by $\mathcal{L}_\sigma(\rho) = \mathrm{Tr}\,(\rho)\,\sigma - \rho$ for some quantum state $\sigma \in \mathcal{D}_d^+$. We compute the optimal convergence rate α such that

$$D(e^{t\mathcal{L}_\sigma}(\rho)\|\sigma) \leq e^{-2\alpha t} D(\rho\|\sigma) \tag{A.1}$$

holds for any $\rho \in \mathcal{D}_d$ and any $t \in \mathbb{R}^+$. This can be seen to be equal to the logarithmic-Sobolev constant $\alpha_1\left(\mathcal{L}_\sigma\right)$.

A.1.1 Main result

I would like to start by clarifying that A. Müller-Hermes is the principal author of this article. Our main result is the computation of $\alpha_1\left(\mathcal{L}_\sigma\right)$. This is the first class of semigroups for which we know the exact value of α_1. Here $D_2(x\|y)$ denotes the relative entropy between the binary distributions $(x, 1-x)$ and $(y, 1-y)$.

Theorem A.1.1. *Let $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ be the depolarizing Liouvillian with full-rank fixed point $\sigma \in \mathcal{D}_d$. Then we have*

$$\alpha_1\left(\mathcal{L}_\sigma\right) = \min_{x\in[0,1]} \frac{1}{2}\left(1 + \frac{D_2(s_{min}(\sigma)\|x)}{D_2(x\|s_{min}(\sigma))}\right),$$

where $s_{min}(\sigma)$ denotes the minimal eigenvalue of σ.

As a direct consequence of Equation (A.1), we have to solve the following optimization problem to show the theorem:

$$\alpha_1\left(\mathcal{L}_\sigma\right) = \inf_{\rho\in\mathcal{D}_d^+} \frac{1}{2}\Big(1 + \frac{D(\sigma\|\rho)}{D(\rho\|\sigma)}\Big).$$

We show that the quotient of relative entropies appearing in the optimization is a quasi-linear function in the entries of the doubly stochastic matrix $P_{ij} = |\langle v_i|w_j\rangle|^2$. Here $\{|v_i\rangle\}$ and $\{|w_j\rangle\}$ are eigenbasis of ρ and σ. As the infimum of a quasi-linear function over a convex set is attained at extreme points, Birkhoff's theorem implies that the optimal ρ has to commute with σ. This simplification allows us to apply Lagrange-multipliers to obtain the claimed expression for the minimum.

As one can imagine from the sketch of the proof, computing the α_1 constant is not a trivial task even for very simple semigroups like depolarizing channels. This motivated A. Müller-Hermes and me to find tools to obtain similar results in a simpler way than using these techniques in [2]. Moreover, in this article we also obtain tensorization results for the case in which $\sigma = \frac{\mathbb{1}}{d}$:

Theorem A.1.2 (Tensorization for the maximally mixed state)**.** *For n tensor powers of the semigroup generated by $\mathcal{L}_{\frac{\mathbb{1}}{d}}$ we have*

$$\alpha_1\left(\mathcal{L}_{\frac{\mathbb{1}}{d}}^{(n)}\right) \geq \frac{1}{2}.$$

This bound is a direct consequence of the following entropy-production estimate:

Theorem A.1.3. *For any $\sigma, \rho \in \mathcal{D}_d$ (not necessarily full rank) we have*

$$S((T_t^\sigma)^{\otimes n}(\rho)) \geq e^{-t}S(\rho) + (1-e^{-t})S(\sigma^{\otimes n}).$$

This theorem was first considered (though with a wrong proof) in [56] for the particular case $\sigma = \frac{\mathbb{1}_2}{2}$. We prove the above theorem using a quantum version of Shearer's inequality for entropies. This bound can also be used to obtain bounds on different capacities of this class of depolarizing channels. Again, this proof explores the structure of depolarizing channels crucially. Tensorization results like this one are central in many applications of LS inequalities in classical probability theory and are usually obtained through the connection between α_1 and hypercontractive inequalities. Therefore, it would be desirable to obtain similar results for general semigroups, but this seems to be technically challenging, as it is not known if α_1 is always related to hypercontractive properties of the semigroup in the quantum case. This missing piece necessary to adapt classical techniques motivated us to find new ways to obtain tensorization results for the convergence in relative entropy in [2]. I, therefore, also include this article in this dissertation. Although A. Müller-Hermes is the principal author of it, it is intimately related to [2] and served as an inspiration for it.

Permission to attach:

A. Müller-Hermes, D. Stilck França, and M. M. Wolf.
Relative entropy convergence for depolarizing channels.
Journal of Mathematical Physics, 57(2), 2016.

Daniel Stilck França <dsfranca13@gmail.com>

JMP Permission Request

1 message

JMP Editorial Office <jmpeo@aip.org> Wed, Apr 18, 2018 at 2:05 PM
To: "dsfranca13@gmail.com" <dsfranca13@gmail.com>

Dear Mr. Franca,

You have permission to include your work:
A. Müller-Hermes, D. Stilck França, and M. M. Wolf.
Relative entropy convergence for depolarizing channels.
Journal of Mathematical Physics, 57(2), 2016.
(https://aip.scitation.org/doi/abs/10.1063/1.4939560?journalCode=jmp)
in your PhD thesis.

Sincerely,

Karen Beverlin
Assistant Editor

JOURNAL OF MATHEMATICAL PHYSICS **57**, 022202 (2016)

Relative entropy convergence for depolarizing channels

Alexander Müller-Hermes,[1,2,a)] Daniel Stilck França,[1,b)] and Michael M. Wolf[1,c)]
[1]*Zentrum Mathematik, Technische Universität München, 85748 Garching, Germany*
[2]*Department of Mathematical Sciences, University of Copenhagen, 2100 Copenhagen ϕ, Denmark*

(Received 6 October 2015; accepted 20 December 2015; published online 20 January 2016)

We study the convergence of states under continuous-time depolarizing channels with full rank fixed points in terms of the relative entropy. The optimal exponent of an upper bound on the relative entropy in this case is given by the log-Sobolev-1 constant. Our main result is the computation of this constant. As an application, we use the log-Sobolev-1 constant of the depolarizing channels to improve the concavity inequality of the von Neumann entropy. This result is compared to similar bounds obtained recently by Kim and we show a version of Pinsker's inequality, which is optimal and tight if we fix the second argument of the relative entropy. Finally, we consider the log-Sobolev-1 constant of tensor-powers of the completely depolarizing channel and use a quantum version of Shearer's inequality to prove a uniform lower bound. © *2016 AIP Publishing LLC.* [http://dx.doi.org/10.1063/1.4939560]

I. INTRODUCTION

Let $\mathcal{M}_d$ denote the set of complex $d \times d$-matrices, $\mathcal{D}_d \subset \mathcal{M}_d$ the set of quantum states, i.e., positive matrices with trace equal to 1, and $\mathcal{D}_d^+$ the set of strictly positive states. The relative entropy (also called quantum Kullback-Leibler divergence) of $\rho, \sigma \in \mathcal{D}_d$ is defined as

$$D(\rho\|\sigma) := \begin{cases} \mathrm{tr}[\rho(\log\rho - \log\sigma)], & \text{if } \mathrm{supp}(\rho) \subset \mathrm{supp}(\sigma) \\ +\infty, & \text{otherwise} \end{cases}. \tag{1}$$

The relative entropy defines a natural distance measure to study the convergence of Markovian time-evolutions. For some state $\sigma \in \mathcal{D}_d$, consider the generalized depolarizing Liouvillian $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ defined as

$$\mathcal{L}_\sigma(\rho) := \mathrm{tr}[\rho]\,\sigma - \rho. \tag{2}$$

This Liouvillian generates the generalized depolarizing channel $T_t^\sigma : \mathcal{M}_d \to \mathcal{M}_d$ with $T_t^\sigma(\rho) := e^{t\mathcal{L}_\sigma}(\rho) = (1-e^{-t})\mathrm{tr}[\rho]\,\sigma + e^{-t}\rho$, where $t \in \mathbb{R}^+$ denotes a time parameter. As $T_t^\sigma(\rho) \to \sigma$ for $t \to \infty$ we can study the convergence speed of the depolarizing channel with a full rank fixed point $\sigma \in \mathcal{D}_d$ by determining the largest constant $\alpha \in \mathbb{R}^+$ such that

$$D(T_t^\sigma(\rho)\|\sigma) \le e^{-2\alpha t} D(\rho\|\sigma) \tag{3}$$

holds for any $\rho \in \mathcal{D}_d$ and any $t \in \mathbb{R}^+$. This constant is known as the logarithmic Sobolev-1 constant[1,2] of $\mathcal{L}_\sigma$, denoted by $\alpha_1(\mathcal{L}_\sigma)$. In the following, we will compute this constant and then use it to derive an improvement on the concavity of von Neumann entropy.

[a)] Electronic address: muellerh@posteo.net
[b)] Electronic address: dsfranca@mytum.de
[c)] Electronic address: m.wolf@tum.de

0022-2488/2016/57(2)/022202/15/$30.00

II. PRELIMINARIES AND NOTATION

Consider a primitive[21] Liouvillian $\mathcal{L}$ with full rank fixed point $\sigma \in \mathcal{D}_d$ and denote by $T_t := e^{t\mathcal{L}}$ the quantum dynamical semigroup generated by $\mathcal{L}$. Consider the function $f(t) := D\left(T_t(\rho)\middle\|\sigma\right)$ for some initial state $\rho \in \mathcal{D}_d$ and note that if

$$\frac{df}{dt} \leq -2\alpha f$$

holds for some $\alpha \in \mathbb{R}_+$, then it follows that $f(t) \leq e^{-2\alpha t} f(0)$. The time derivative of the relative entropy at $t = 0$, also called the entropy production,[3] is given by

$$\frac{d}{dt} D\left(T_t(\rho)\middle\|\sigma\right)\Big|_{t=0} = -\mathrm{tr}[\mathcal{L}(\rho)(\log(\sigma) - \log(\rho))] \tag{4}$$

as $\mathrm{tr}(\mathcal{L}(\rho)) = 0$ for any $\rho \in \mathcal{D}_d$. This motivates the following definition.

Definition 2.1 (log-Sobolev-1 constant,[1,2]). For a primitive Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d$, we define its **log-Sobolev-1 constant** as

$$\alpha_1(\mathcal{L}) := \sup\left\{\alpha \in \mathbb{R} : \mathrm{tr}[\mathcal{L}(\rho)(\log(\sigma) - \log(\rho))] \geq 2\alpha D\left(\rho\|\sigma\right), \forall \rho \in \mathcal{D}_d^+\right\} \tag{5}$$

For a primitive Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ the preceding discussion shows that (3) holds for any $\alpha \leq \alpha_1(\mathcal{L})$. Furthermore, $\alpha_1(\mathcal{L})$ is the optimal constant for which this inequality holds independent of $\rho \in \mathcal{D}_d$ (for states ρ not of full rank, this follows from a simple continuity argument).

In the following, we will need some functions defined as continuous extensions of quotients of relative entropies. We denote by $Q_\sigma : \mathcal{D}_d^+ \to \mathbb{R}$ the continuous extension of the function $\rho \mapsto \frac{D(\sigma\|\rho)}{D(\rho\|\sigma)}$ (see Appendix B) given by

$$Q_\sigma(\rho) := \begin{cases} \dfrac{D(\sigma\|\rho)}{D(\rho\|\sigma)}, & \rho \neq \sigma \\ 1, & \rho = \sigma \end{cases}. \tag{6}$$

Note that for $x \in [0,1]$ and $y \in (0,1)$, the binary relative entropy is defined as

$$D_2(x\|y) := x \log\left(\frac{x}{y}\right) + (1-x)\log\left(\frac{1-x}{1-y}\right). \tag{7}$$

This is the classical relative entropy of the probability distributions $(x, 1-x)$ and $(y, 1-y)$. For $y \in (0,1)$, we denote by $q_y : (0,1) \to \mathbb{R}$ the continuous extension of $x \mapsto \frac{D_2(y\|x)}{D_2(x\|y)}$ given by

$$q_y(x) := \begin{cases} \dfrac{D_2(y\|x)}{D_2(x\|y)}, & x \neq y \\ 1, & x = y \end{cases}. \tag{8}$$

III. LOG-SOBOLEV-1 CONSTANT FOR THE DEPOLARIZING LIOUVILLIAN

Note that for the depolarizing Liouvillian $\mathcal{L}_\sigma$ with $\sigma \in \mathcal{D}_d^+$ as defined in (2), we have

$$\mathrm{tr}[\mathcal{L}_\sigma(\rho)(\log(\sigma) - \log(\rho))] = D(\rho\|\sigma) + D(\sigma\|\rho).$$

Inserting this into Definition 2.1, we can write

$$\alpha_1(\mathcal{L}_\sigma) = \inf_{\rho \in \mathcal{D}_d^+} \frac{1}{2}\left(1 + Q_\sigma(\rho)\right). \tag{9}$$

Our main result is the following theorem.

Theorem 3.1. *Let $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ be the depolarizing Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d$ as defined in (2). Then, we have*

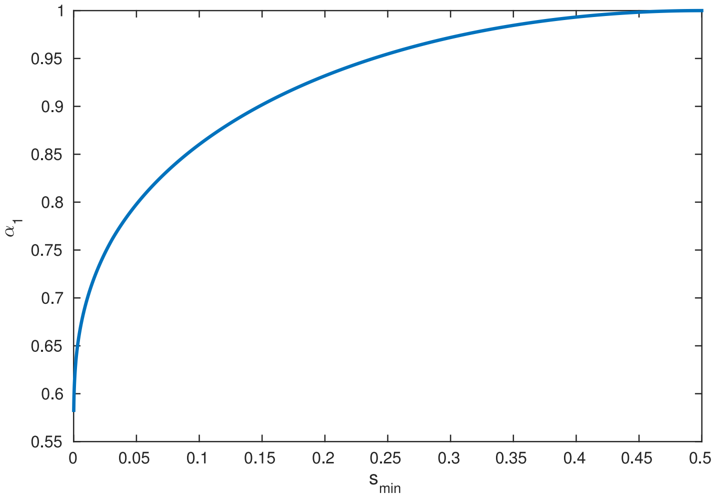

FIG. 1. $\alpha_1(\mathcal{L}_\sigma)$ for $s_{\min}(\sigma) \in [0,1]$.

$$\alpha_1(\mathcal{L}_\sigma) = \min_{x\in[0,1]} \frac{1}{2}\left(1 + q_{s_{min}(\sigma)}(x)\right),$$

where $s_{min}(\sigma)$ denotes the minimal eigenvalue of σ.

In Figure 1, the values of $\alpha_1(\mathcal{L}_\sigma)$ depending on $s_{\min}(\sigma) \in [0,1]$ are plotted. Note that by Theorem 3.1, we have $\alpha_1(\mathcal{L}_\sigma) \to 1/2$ in the limit $s_{\min}(\sigma) \to 0$ (as $D_2(s_{\min}(\sigma)\|0.5) \leq \log(2)$ and $D_2(0.5\|s_{\min}(\sigma)) \to \infty$ in this case).

Before we state the proof of Theorem 3.1 we need to make a technical comment. By (6), we have $Q_\sigma(\rho) \to +\infty$ as $\rho \to \partial\mathcal{D}_d$, i.e., as ρ converges to a rank-deficient state. Therefore, the infimum in (9) will be attained in a full rank state $\tilde{\rho} \in \mathcal{D}_d^+$ and we can restrict the optimization to the compact set $K_\sigma \subset \mathcal{D}_d$ (depending on σ) defined as

$$K_\sigma = \{\rho \in \mathcal{D}_d^+ : s_{\min}(\rho) \geq s_{\min}(\tilde{\rho}) - \epsilon\} \tag{10}$$

for some fixed $\epsilon \in (0, s_{\min}(\tilde{\rho}))$ and where $s_{\min}(\cdot)$ denotes the minimal eigenvalue. Note that the minimizing state $\tilde{\rho}$ is contained in the interior of K_σ. Now we have to solve the following optimization problem for fixed $\sigma \in \mathcal{D}_d^+$:

$$\inf_{\rho\in\mathcal{D}_d^+} Q_\sigma(\rho) = \inf_{\rho\in K_\sigma} Q_\sigma(\rho). \tag{11}$$

To prove Theorem 3.1, we will need the following lemma showing that the infimum in (11) is attained at states $\rho \in \mathcal{D}_d$ commuting with the fixed point σ.

Lemma 3.1. For any $\sigma \in \mathcal{D}_d^+$ we have

$$\inf_{\rho\in K_\sigma} Q_\sigma(\rho) = \inf_{\rho\in K_\sigma, [\rho,\sigma]=0} Q_\sigma(\rho)$$

where $Q_\sigma : D_d^+ \to \mathbb{R}$ denotes the continuous extension of $\rho \mapsto \frac{D(\sigma\|\rho)}{D(\rho\|\sigma)}$ (see (6)).

Proof. Consider the spectral decomposition $\sigma = \sum_{i=1}^d s_i |v_i\rangle\langle v_i|$ for $s \in \mathbb{R}_+^d$ and fix a vector $r \in \mathbb{R}_+^d$ which is not a permutation of s and fulfills $\min_i(r_i) \geq s_{\min}(\tilde{\rho}) - \epsilon$ (see (10)) and $\sum_{j=1}^d r_j = 1$. For some fixed orthonormal basis $\{|w_j\rangle\}_j$, consider $\rho := \sum_{j=1}^d r_j |w_j\rangle\langle w_j| \in K_\sigma$. Inserting ρ into Q_σ gives

$$Q_\sigma(\rho) = \frac{D(\sigma\|\rho)}{D(\rho\|\sigma)} = \frac{-S(\sigma) - \mathrm{tr}[\sigma\log(\rho)]}{-S(\rho) - \mathrm{tr}[\rho\log(\sigma)]} = \frac{-S(\sigma) - \langle s, P\log(r)\rangle}{-S(\rho) - \langle \log(s), Pr\rangle} =: F(P), \tag{12}$$

where we introduced $P \in \mathcal{M}_d$ given by $P_{ij} = |\langle v_i|w_j\rangle|^2$ and $\log(s), \log(r) \in \mathbb{R}^d$ are defined as $(\log(s))_i = \log(s_i)$ and $(\log(r))_j = \log(r_j)$. Note that P is a unistochastic matrix, i.e., a doubly stochastic matrix whose entries are squares of absolute values of the entries of a unitary matrix. We will show that the minimum of F over unistochastic matrices P is attained at a permutation matrix. By the definition of P, this shows that there exists a state $\rho' \in K_\sigma$ with spectrum r and commuting with σ, which fulfills $Q_\sigma(\rho') \leq Q_\sigma(\rho)$.

As the set of unistochastic matrices is in general not convex,[4] we want to consider the set of doubly stochastic matrices instead. By Birkhoff's theorem [Ref. 5, Theorem II.2.3], we can write any doubly stochastic $D \in \mathcal{M}_d$ as $D = \sum_{i=1}^k \lambda_i P_i$ for some $k \in \mathbb{N}$, numbers $\lambda_i \in [0,1]$ with $\sum_{i=1}^k \lambda_i = 1$ and permutation matrices P_i. Now we can write the denominator of $F(D)$ as

$$-S(\rho) - \langle \log(s), Dr\rangle = \sum_{i=1}^k \lambda_i \left(-S(\rho) - \langle \log(s), P_i r\rangle\right) = \sum_{i=1}^k \lambda_i D(\rho_i \| \sigma) > 0,$$

where ρ_i is the state obtained by permuting the eigenvectors of ρ with P_i. In the last step, we used Klein's inequality [Ref. 6, p. 511] together with the fact that $\rho_i \neq \sigma$ for any $1 \leq i \leq k$ as their spectra are different. The previous estimate shows that F is also well-defined on doubly stochastic matrices.

Any unistochastic matrix is also doubly stochastic and we have

$$\inf\left\{F(P) : P \in \mathcal{M}_d \text{ doubly stochastic}\right\} \leq \inf\left\{F(P) : P \in \mathcal{M}_d \text{ unistochastic}\right\}.$$

Note that $S(\sigma)$ and $S(\rho)$ in (12) only depend on $s \in \mathbb{R}^d_+$ and $r \in \mathbb{R}^d_+$ and thus the numerator and the denominator of F are positive affine functions in P. This shows that F is a quasi-linear function [Ref. 7, p. 97] on the set of doubly stochastic matrices. It can be shown (see Ref. 7) that the minimum of such a function over a compact and convex set is always attained in an extremal point of the set. By Birkhoff's theorem [Ref. 5, Theorem II.2.3], the extremal points of the compact and convex set of doubly stochastic matrices are the permutation matrices. As these are also unistochastic matrices, we have

$$\inf\left\{F(P) : P \in \mathcal{M}_d \text{ unistochastic}\right\} = \inf\left\{F(P) : P \in \mathcal{M}_d \text{ permutation matrix}\right\}.$$

This finishes the first part.

To prove the lemma, note that we have

$$\inf_{\rho \in K_\sigma} Q_\sigma(\rho) = Q_\sigma(\tilde{\rho})$$

for some minimizing full rank state $\tilde{\rho} \in \mathcal{D}^+_d$. Now consider some sequence $(\rho_n)_{n\in\mathbb{N}} \in K_\sigma^{\mathbb{N}}$ with $\rho_n \to \tilde{\rho}$ as $n \to \infty$ and such that the spectra of the ρ_n are no permutations of the spectrum of σ. By the first part of the proof, we find a sequence $(\rho'_n)_{n\in\mathbb{N}} \in K_\sigma^{\mathbb{N}}$ commuting with σ, such that

$$Q_\sigma(\tilde{\rho}) \leq Q_\sigma(\rho'_n) \leq Q_\sigma(\rho_n) \to Q_\sigma(\tilde{\rho})$$

as $n \to \infty$. Thus, $Q_\sigma(\rho'_n) \to Q_\sigma(\tilde{\rho})$ as $n \to \infty$. On the compact set K_σ, the sequence $(\rho'_n)_n$ has a converging subsequence $\left(\rho'_{n_k}\right)_{k\in\mathbb{N}}$ with $\rho'_{n_k} \to \rho' \in K_\sigma$ as $k \to \infty$. By continuity of Q_σ, we have $Q_\sigma(\rho') = Q_\sigma(\tilde{\rho}) = \inf_{\rho\in K_\sigma} Q_\sigma(\rho)$ and by continuity of the commutator $\rho \mapsto [\rho, \sigma]$ we have $[\rho', \sigma] = 0$. □

With this lemma we can prove our main result.

Proof of Theorem 3.1. By Lemma 3.1, we may restrict the optimization in (11) to states which commute with σ. Thus, we can repeat the construction of the compact set K_σ (see (10)) for a minimizer $\tilde{\rho} \in \mathcal{D}^+_d$ with $[\tilde{\rho}, \sigma] = 0$. By construction, $\tilde{\rho}$ lies in the interior of K_σ, which will be important for the following argument involving Lagrange-multipliers.

To find necessary conditions on the minimizers of (11), we abbreviate $C := \inf_{\rho\in K_\sigma} Q_\sigma(\rho)$ and note that $C > 0$. To see this, note that we may extend $Q_\sigma(\rho)$ continuously to 1 at σ, so there exists $\delta > 0$ s.t. for $\|\rho - \sigma\|_1 \leq \delta$, we have $Q_\sigma(\rho) \geq \frac{1}{2}$ and for ρ s.t. $\|\rho - \sigma\|_1 > \delta$, we have $Q_\sigma(\rho) \geq \frac{\delta^2}{2\log\left(s_{\min}(\sigma^{-1})\right)}$ using Pinsker's inequality and $D(\rho\|\sigma) \leq \log\left(s_{\min}(\sigma)\right)$. For any $\rho \in K_\sigma$ with $[\rho, \sigma] = 0$ and $\rho \neq \sigma$ have

$$\frac{D(\sigma\|\rho)}{D(\rho\|\sigma)} \geq C,$$

which is equivalent to

$$S(\sigma) \leq CS(\rho) + C\sum_{i=1}^{d} r_i \log(s_i) - \sum_{i=1}^{d} s_i \log(r_i). \tag{13}$$

Here, $\{r_i\}_{i=1}^{d}$ denote the eigenvalues of $\rho \in K_\sigma$ (see (10)) fulfilling $[\rho,\sigma] = 0$ and $\{s_i\}_{i=1}^{d}$ the eigenvalues of σ. As $\tilde{\rho}$ is a minimizer of (11) and commutes with σ its spectrum is a minimizer of the right-hand-side of (13) minimized over the set $\mathcal{S} := \{r \in \mathbb{R}^d : \min_i(r_i) \geq s_{\min}(\tilde{\rho}) - \epsilon\} \subset \mathbb{R}^d$ with ϵ chosen in the construction of K_σ (see (10)). We will now compute necessary conditions on the spectrum of $\tilde{\rho}$ using the formalism of Lagrange-multipliers (note that by construction the spectrum of $\tilde{\rho}$ lies in the interior of $\mathcal{S}$).

Consider the Lagrange function $F : \mathcal{S} \times \mathbb{R} \to \mathbb{R}$ given by

$$F(r_1,\ldots,r_d,\lambda) = CS(\rho) + C\sum_{i=1}^{d} r_i \log(s_i) - \sum_{i=1}^{d} s_i \log(r_i) + \lambda\left(\sum_{i=1}^{d} r_i - 1\right).$$

The gradient of F is given by

$$[\nabla F(r_1,\ldots,r_d,\lambda)]_j = \begin{cases} C(-\log(r_j) - 1 + \log(s_j)) - \dfrac{s_j}{r_j} + \lambda & 1 \leq j \leq d \\ \sum_{i=1}^{d} r_i - 1 & j = d+1 \end{cases}. \tag{14}$$

By the formalism of Lagrange-multipliers, any minimizer $r = (r_1,\ldots,r_d)$ of the right-hand-side of (13) in the interior of $\mathcal{S}$ has to fulfill $\nabla F(r_1,\ldots,r_d,\lambda) = 0$ for some $\lambda \in \mathbb{R}$. Summing up the first d of these equations (where the jth equation is multiplied with r_j) implies

$$\lambda = 1 + C(1 + D(\rho\|\sigma)).$$

Inserting this back into the equations $[\nabla F(r_1,\ldots,r_d,\lambda)]_j = 0$ and using $u_j = \frac{r_j}{s_j}$, we obtain

$$u_j(1 + CD(\rho\|\sigma)) - 1 = Cu_j \log(u_j) \tag{15}$$

for $1 \leq j \leq d$. For fixed $D(\rho\|\sigma)$, there are only two values for u_j solving the Equations (15), as an affine functions (the left-hand-side) can only intersect a strictly convex function (the right-hand-side) in at most two points. Thus, for a minimizer $\{r_i\}_{i=1}^{d}$ of the right-hand-side of (13) in the interior of $\mathcal{S}$ there are constants $c_1, c_2 \in \mathbb{R}^+$ such that for each $i \in \{1,\ldots,d\}$ either $r_i = c_1 s_i$ or $r_i = c_2 s_i$ holds.

We have obtained the following conditions on the spectrum of the minimizer $\tilde{\rho} \in K_\sigma$ (fulfilling $[\tilde{\rho},\sigma] = 0$) of (11): There exist constants $c_1, c_2 \in \mathbb{R}^+$ a permutation $\nu \in S_d$ (where S_d denotes the group of permutations on $\{1,\ldots,d\}$) and some $0 \leq n \leq d$ such that the spectrum $r \in \mathbb{R}^+_d$ of $\tilde{\rho}$ fulfills $r_{\nu(i)} = c_1 s_{\nu(i)}$ for any $1 \leq i \leq n$ and $r_{\nu(i)} = c_2 s_{\nu(i)}$ for any $n+1 \leq i \leq d$. Note that the cases $c_1 = c_2 = 1$, $n = 0$, and $n = d$, all correspond to the case $\rho = \sigma$ where we have $Q_\sigma(\sigma) = 1$. Thus, we can exclude the cases $n = 0$ and $n = d$ as long as we optimize over $c_1 = c_2 = 1$. Furthermore, note that we can use the normalization of $\tilde{\rho}$, i.e., $c_1\sum_{i=1}^{n} s_{\nu(i)} + c_2\sum_{i=n+1}^{d} s_{\nu(i)} = 1$ to eliminate c_2. Given a permutation $\nu \in S_d$ and $n \in \{1,\ldots,d\}$, we define $p(\nu,n) = \sum_{i=1}^{n} s_{\nu(i)}$. Inserting the above conditions in (11) and setting $c_1 = x$ and $0 < n < d$ yields

$$\inf_{\rho \in K_\sigma} Q_\sigma(\rho) = \inf_{\nu \in S_d} \inf_{1 \leq n \leq (d-1)} \inf_{x \in \left[0, p(\nu,n)^{-1}\right]} q_{p(\nu,n)}(xp(\nu,n)) \tag{16}$$

$$= \inf_{\nu \in S_d} \inf_{1 \leq n \leq (d-1)} \inf_{x \in [0,1]} q_{p(\nu,n)}(x), \tag{17}$$

where $q_y : [0,1] \to \mathbb{R}$ denotes the continuous extension of $x \mapsto \frac{D_2(y\|x)}{D_2(x\|y)}$ (see (8)). By Lemma A.1 in the Appendix, the function $y \mapsto q_y(x)$ is continuous and quasi-concave and hence the minimum over any convex and compact set is attained at the boundary. Thus, we have

$$q_{s_{\min}(\sigma)}(x) \geq \inf_{\nu \in S_d} \inf_{1 \leq n \leq (d-1)} q_{p(\nu,n)}(x) \geq \inf_{y \in [s_{\min}(\sigma), 1 - s_{\min}(\sigma)]} q_y(x) = q_{s_{\min}(\sigma)}(x)$$

using $q_{1-s_{\min}(\sigma)}(x) = q_{s_{\min}(\sigma)}(x)$ for any $x \in [0,1]$. Inserting this into (17) leads to

$$\inf_{\rho \in K_\sigma} Q_\sigma(\rho) = \inf_{x\in[0,1]} q_{s_{\min}(\sigma)}(x).$$

□

Lemma 3.1 implies that the log-Sobolev-1 constant of the depolarizing channels coincides with the classical one of the random walk on the complete graph with d vertices and distribution given by the spectrum of σ. This constant has been shown to imply other inequalities, such as in [Ref. 8, Proposition 3.13]. Using this result, Theorem 3.1 implies a refined transportation inequality on graphs.

Using the correspondence with the classical log-Sobolev-1 constant of a random walk on the complete graph, we may apply Ref. 9 [Example 3.10], which proves the following:

Corollary 3.1. Let $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ be the depolarizing Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d$ as defined in (2). Then we have

$$\alpha_1(\mathcal{L}_\sigma) \geq \frac{1}{2} + \sqrt{s_{\min}(\sigma)\,(1 - s_{\min}(\sigma))}$$

with equality iff $s_{\min}(\sigma) = \frac{1}{2}$. Again $s_{min}(\sigma)$ denotes the minimal eigenvalue of σ.

IV. APPLICATION: IMPROVED CONCAVITY OF VON NEUMANN ENTROPY

It is a well-known fact that the von Neumann entropy $S(\rho) = -\mathrm{tr}\,[\rho \log(\rho)]$ is concave in ρ. Using Theorem 3.1 we can improve the concavity inequality:

Theorem 4.1 (Improved concavity of the von Neumann entropy). *For $\rho, \sigma \in \mathcal{D}_d$ and $q \in [0,1]$, we have*

$$S((1-q)\sigma + q\rho) - (1-q)S(\sigma) - qS(\rho) \geq \max \begin{cases} q(1-q^{c(\sigma)})D(\rho\|\sigma) \\ (1-q)(1-(1-q)^{c(\rho)})D(\sigma\|\rho) \end{cases},$$

with

$$c(\sigma) = \min_{x\in[0,1]} \frac{D_2(s_{\min}(\sigma)\|x)}{D_2(x\|s_{\min}(\sigma))}$$

and $c(\rho)$ defined in the same way.

Note that this bound becomes trivial if both σ and ρ are not of full rank (as we have $c(\rho) = c(\sigma) = 0$ in this case). However, as long as $D(\rho\|\sigma)$ or $D(\sigma\|\rho) < \infty$, we may still get a bound by restricting both density matrices to the support of σ or ρ, respectively.

Proof. Note that for the Liouvillian $\mathcal{L} := -\log(q)\mathcal{L}_\sigma$, we have

$$e^{\mathcal{L}}(\rho) = q\rho + (1-q)\sigma.$$

By Theorem 3.1 and (3), we have

$$D\left(e^{\mathcal{L}}(\rho)\|\sigma\right) \leq e^{(1+c(\sigma))\log(q)} D\left(\rho\|\sigma\right) \tag{18}$$

Rearranging and expanding the terms in (18), we get

$$\begin{aligned} S(q\rho + (1-q)\sigma) &\geq (1-q)S(\sigma) - q\mathrm{tr}\,[\rho\log(\sigma)] + q^{1+c(\sigma)}D(\rho\|\sigma) \\ &= (1-q)S(\sigma) + qS(\rho) + q(1-q^{c(\sigma)})D(\rho\|\sigma). \end{aligned}$$

Interchanging the roles of ρ and σ in the above proof gives the second case under the maximum. □

In Ref. 10, another improvement on the concavity of the von Neumann entropy is shown,

$$S((1-q)\sigma + q\rho) - (1-q)S(\sigma) - qS(\rho) \geq \frac{q(1-q)}{(1-2q)^2} D(\rho_{\text{rev}} \| \rho_{\text{avg}}) \tag{19}$$

$$\geq \frac{1}{2} q(1-q) \|\rho - \sigma\|_1^2, \tag{20}$$

where $\rho_{\text{avg}} = (1-q)\sigma + q\rho$ and $\rho_{\text{rev}} = (1-q)\rho + q\sigma$. Note that this bound is valid for all states $\rho, \sigma \in \mathcal{D}_d$ while our bound in Theorem 4.1 becomes trivial unless the support ρ is contained in the support of σ or the other way around. We will therefore consider only full rank states in the following analysis.

By simple numerical experiments, our bound from Theorem 4.1 seems to be worse than (19). However, one can argue that (19) is not much simpler than the left-hand-side itself. In particular, the dependence on ρ and σ is only implicit via the relative entropy between ρ_{avg} and ρ_{rev}. Our bound from Theorem 4.1 depends on some spectral data (in terms of the smallest eigenvalues of ρ or σ), but whenever this is given, we have a bound for any $q \in [0,1]$ in terms of the relative entropies of ρ and σ.

Again, we can do simple numerical experiments to compare bounds (20) and Theorem 4.1. Recall that our bound is given in terms of the relative entropy and (20) in terms of the trace norm. In Figure 2, the bounds are compared for randomly generated quantum states in dimension $d = 10$. These plots show that the bounds are *not* comparable and depending of the choice of the states the bound from Theorem 4.1 will perform better than (20) or vice versa. Note that for q close to 0 or 1, our bound seems to perform better in both Figures. This is to be expected as $\alpha_1(\mathcal{L}_\sigma)$ is defined as the optimal constant α bounding the entropy production (4) (in $t = 0$) by $-2\alpha D(\rho\|\sigma)$. Therefore, Theorem 4.1 should be the optimal bound (in terms of relative entropy) for q near 0 or 1.

Note that by applying Pinsker's inequality,

$$D(\rho\|\sigma) \geq \frac{1}{2}\|\rho - \sigma\|_1^2 \tag{21}$$

for states $\rho, \sigma \in \mathcal{D}_d$ to our bound from Theorem 4.1 we can obtain an improvement on the concavity inequality in terms of the trace-distance similar to (20). Unfortunately, a simple computation shows that the resulting trace-norm bound is always worse than (20). In Sec. V, we will show that Pinsker's inequality can be improved in the case where the second argument in the relative entropy is fixed (which is the case in the bound from Theorem 4.1). This will lead to an additional improvement of the trace-norm bound obtained from Theorem 4.1, such that in some (but only very few) cases the bound becomes better than (20).

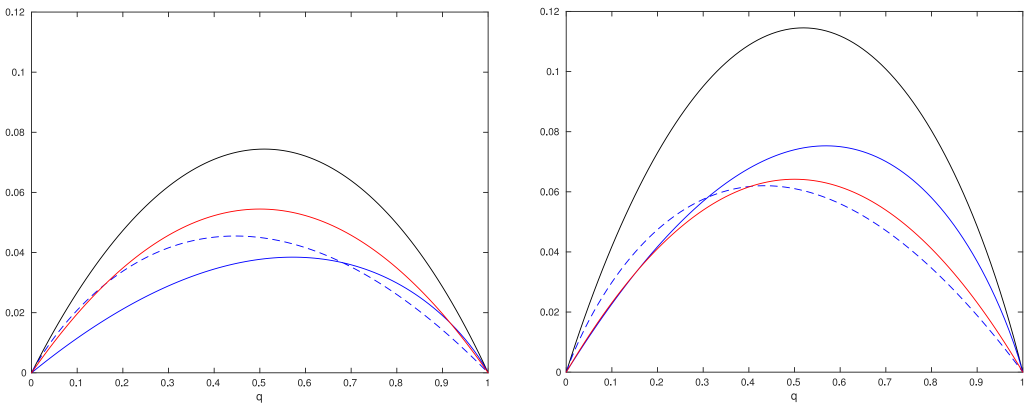

(a) (b)

FIG. 2. Comparison of bound (20) (red) and the bound from Theorem 4.1 (blue, where both choices of the ordering of ρ and σ are plotted) and the exact value $S((1-q)\sigma + q\rho) - (1-q)S(\sigma) - qS(\rho)$ (black) two pairs of randomly generated 10×10 quantum states and $q \in [0,1]$.

V. STATE-DEPENDENT OPTIMAL PINSKER'S INEQUALITY

Pinsker's inequality (21) can be applied to the bound in Theorem 4.1 to get an improvement of the concavity in terms of the trace distance of the two density matrices. It can also be applied to (3) to get a mixing time bound[2] for the depolarizing channel. Note that in both of these cases, the second argument of the relative entropy is fixed. Other improvements have been considered in the literature,[11] but here we will improve Pinsker's inequality in terms of the second argument of the relative entropy. More specifically, we compute the optimal constant $C(\sigma)$ (depending on σ) such that $D(\rho\|\sigma) \geq C(\sigma)\|\rho-\sigma\|_1^2$ holds when σ has full rank.

We will follow a strategy similar to the one pursued in Ref. 12 in proving this, where the analogous problem was considered for classical probability distributions. For a state $\rho \in \mathcal{D}_d$ let $s(\rho) = (s_1(\rho), \ldots, s_d(\rho))$ denote its vector of eigenvalues decreasingly ordered. We will adopt the convention that the minimum over an empty set is $+\infty$ and that $D_2(x\|y) = +\infty$ for $x, y \notin [0,1]$.

Lemma 5.1. Let $\sigma \in \mathcal{D}_d^+$ and for $A \subseteq \{1,\ldots,d\}$, define $P_\sigma(A) = \sum_{i\in A} s_i(\sigma)$. Then we have for $\epsilon > 0$:

$$\min_{\rho:\|\rho-\sigma\|_1\geq\epsilon} D(\rho\|\sigma) = \min_{A\subseteq\{1,\ldots,d\}} D_2\left(P_\sigma(A) + \frac{\epsilon}{2}\middle\| P_\sigma(A)\right) \tag{22}$$

Proof. Suppose that there exists $\rho \in \mathcal{D}_d$ such that $\|\rho-\sigma\|_1 = \delta$, with $\delta \geq \epsilon$. By Lidskii's theorem [Ref. 5, Corollary III.4.2], we have

$$s(\sigma-\rho) = s(\sigma) - Ls(\rho),$$

where L is a doubly stochastic matrix. Define ρ' to be the state which has eigenvalues $Ls(\rho)$ and commutes with σ. Then, we have

$$\|\rho-\sigma\|_1 = \|\rho'-\sigma\|_1.$$

By the operational interpretation for the 1− norm [Ref. 6, Theorem 9.1], there exist Hermitian projections $Q, Q' \in \mathcal{M}_n$ such that

$$2\mathrm{tr}[Q(\rho-\sigma)] = \|\rho-\sigma\|_1 = \|\rho'-\sigma\|_1 = 2\mathrm{tr}[Q'(\rho'-\sigma)]. \tag{23}$$

Now define the quantum channel $T : \mathcal{M}_d \to \mathcal{M}_2$ given by

$$T(\rho) = \mathrm{tr}[Q\rho]|0\rangle\langle 0| + \mathrm{tr}[(\mathbb{1}-Q)\rho]|1\rangle\langle 1|,$$

where $|0\rangle, |1\rangle$ is an orthonormal basis of $\mathbb{C}^2$. By the data processing inequality, we have

$$D(\rho\|\sigma) \geq D(T(\rho)\|T(\sigma)). \tag{24}$$

It is easy to see that the image of Q' must be spanned by eigenvectors of σ. Thus, we may associate a subset $A \subseteq \{1,\ldots,d\}$ to the projector Q' indicating the eigenvectors of σ spanning this subspace. Using (23) and the assumption that $\|\rho-\sigma\|_1 = \delta$ we have

$$\mathrm{tr}[Q'\rho'] = P_\sigma(A) + \frac{\delta}{2}.$$

Also observe that

$$D(T(\rho)\|T(\sigma)) = D_2\left(P_\sigma(A) + \frac{\delta}{2}\middle\| P_\sigma(A)\right) \geq D_2\left(P_\sigma(A) + \frac{\epsilon}{2}\middle\| P_\sigma(A)\right)$$

as the binary relative entropy is convex and $\delta \geq \epsilon$ was assumed. With (24), we have

$$\min_{\rho:\|\rho-\sigma\|_1\geq\epsilon} D(\rho\|\sigma) \geq \min_{A\subseteq\{1,\ldots,d\}} D_2\left(P_\sigma(A) + \frac{\epsilon}{2}\middle\| P_\sigma(A)\right) \tag{25}$$

Now given any $A \subseteq \{1, \ldots, d\}$ such that $P_\sigma(A) + \frac{\epsilon}{2} \leq 1$ (otherwise $D_2\big(P_\sigma(A) + \frac{\epsilon}{2} \| P_\sigma(A)\big) = +\infty$ by convention), define a state $\tau \in \mathcal{D}_d$ which commutes with σ and has spectrum:

$$s_i(\tau) = \begin{cases} \dfrac{(P_\sigma(A) + \epsilon/2)\, s_i(\sigma)}{P_\sigma(A)} & \text{for } i \in A \\ \dfrac{(1 - P_\sigma(A) - \epsilon/2)\, s_i(\sigma)}{1 - P_\sigma(A)} & \text{else} \end{cases}. \tag{26}$$

Note that $\|\sigma - \tau\|_1 = \epsilon$ and $D(\tau\|\sigma) = D_2\big(P_\sigma(A) + \frac{\epsilon}{2} \| P_\sigma(A)\big)$, i.e., the lower bound in (25) is attained.

If there does not exist $\rho \in \mathcal{D}_d$ such that $\|\rho - \sigma\|_1 = \delta$, with $\delta \geq \epsilon$, then the minimum on the l.h.s. of (22) is $+\infty$ by our convention. In this case, we also have $P_\sigma(A) + \frac{\epsilon}{2} > 1$ for all $A \subseteq \{1, \ldots, d\}$ and the minimum on r.h.s. of (22) is also $+\infty$ by convention. If we have $P_\sigma(A) + \frac{\epsilon}{2} < 1$ for some $A \subseteq \{1, \ldots, d\}$, we may construct a density matrix τ as in (26) s.t. $\|\tau - \sigma\|_1 = \delta$, a contradiction. □

We define the function $\phi : [0, \frac{1}{2}] \to \mathbb{R}$ as

$$\phi(p) = \frac{1}{1 - 2p} \log\left(\frac{1 - p}{p}\right), \tag{27}$$

extended continuously by $\phi\left(\frac{1}{2}\right) = 2$. Furthermore for any $\sigma \in \mathcal{D}_d$, we define

$$\pi(\sigma) = \max_{A \subseteq \{1, \ldots, d\}} \min\{P_\sigma(A), 1 - P_\sigma(A)\}. \tag{28}$$

With essentially the same proof as given in Ref. 12 for the classical case, we obtain the following improvement on Pinsker's inequality.

Theorem 5.1 (State-dependent Pinsker's Inequality). *For $\sigma, \rho \in \mathcal{D}_d$, we have*

$$D(\rho\|\sigma) \geq \frac{\phi(\pi(\sigma))}{4} \|\rho - \sigma\|_1^2 \tag{29}$$

with ϕ as in (27) and $\pi(\sigma)$ as in (28). Moreover, this inequality is tight.

Proof. For convenience, set $\|\rho - \sigma\|_1 = \delta$. Then we have

$$D(\rho\|\sigma) \geq \min_{\rho' : \|\rho' - \sigma\|_1 \geq \delta} D(\rho'\|\sigma) = \min_{A \subseteq \{1, \ldots, d\}} D_2\left(P_\sigma(A) + \frac{\delta}{2} \middle\| P_\sigma(A)\right) \tag{30}$$

using Theorem 5.1. By Ref. 12 [Proposition 2.2] for $p \in [0, \frac{1}{2}]$ and $\epsilon \geq 0$, we have

$$D_2(p + \epsilon \| p) \leq D_2(1 - p + \epsilon \| 1 - p)$$

so we may assume $P_\sigma(A) \leq \frac{1}{2}$ in (30). In Ref. 13 [Theorem 1] it is shown that for $p \in [0, \frac{1}{2}]$, we have

$$\inf_{\epsilon \in (0, 1 - p]} \frac{D_2(p + \epsilon \| p)}{\epsilon^2} = \phi(p) \tag{31}$$

which implies

$$\min_{A \subseteq \{1, \ldots, d\}} D_2\left(P_\sigma(A) + \frac{\delta}{2} \middle\| P_\sigma(A)\right) \geq \min_{A \subseteq \{1, \ldots, d\}} \frac{\phi(P_\sigma(A))}{4} \|\rho - \sigma\|_1^2.$$

By Ref. 12 [Proposition 2.4], the function ϕ is strictly decreasing. Thus, we have

$$\min_{A \subseteq \{1, \ldots, d\}} \frac{\phi(P_\sigma(A))}{4} = \frac{\phi(\pi(\sigma))}{4}$$

which, after combining the previous inequalities, finishes the proof of (29). To show that the inequality is tight, we may again follow the proof of Ref. 12 [Proposition 2.1]. Let $B \subseteq \{1, \ldots, d\}$

be a subset such that $\pi(\sigma) = P_\sigma(B) =: p$. Define a minimizing sequence $\{\epsilon_i\}_{i\in\mathbb{N}}$ with $\epsilon_i > 0$ for the infimum (with respect to p) in (31), i.e., such that

$$\lim_{i\to\infty} \frac{D_2(p+\epsilon_i\|p)}{\epsilon_i^2} = \phi(p).$$

Next, define a sequence of states ρ_i that commute with σ and have spectrum,

$$s_j(\rho_i) = \begin{cases} \dfrac{(p+\epsilon_i)\, s_i(\sigma)}{p} & \text{for } j \in B \\ \dfrac{(1-p-\epsilon_i)\, s_i(\sigma)}{1-p} & \text{else} \end{cases}.$$

One can check that $\|\rho_i - \sigma\|_1 = 2\epsilon_i$ and $D(\rho_i\|\sigma) = D_2(p+\epsilon_i\|p)$, from which we get:

$$\lim_{i\to\infty} \frac{D(\rho_i\|\sigma)}{\|\rho_i - \sigma\|^2} = \frac{\phi(\pi(\sigma))}{4}.$$

□

In some cases, the bound can be made more explicit, as illustrated in the next corollary.

Corollary 5.1. Let $\sigma, \rho \in \mathcal{D}_d$ be such that $\|\sigma\|_\infty \geq \frac{1}{2}$. Then:

$$D(\rho\|\sigma) \geq \frac{\phi(1-\|\sigma\|_\infty)}{4}\|\rho-\sigma\|_1^2 \tag{32}$$

Proof. In this , it is clear that $\pi(\sigma) = 1 - \|\sigma\|_\infty$. □

Note that we have $\phi(x) \to +\infty$ for $x \to 0$. Thus, there might be an arbitrary large improvement of (29) compared to usual Pinsker's inequality (21). This happens for instance in Corollary 5.1 when $\|\sigma\|_\infty \to 1$, i.e., when σ converges to a pure state.

By applying the improved inequality (29) to Theorem 4.1, we obtain for quantum states $\rho, \sigma \in \mathcal{D}_d$ and $q \in [0,1]$,

$$S((1-q)\sigma + q\rho) - (1-q)S(\sigma) - qS(\rho) \geq \max\begin{cases} q(1-q^{c(\sigma)})\dfrac{\phi(\pi(\sigma))}{4}\|\rho-\sigma\|^2 \\ (1-q)(1-(1-q)^{c(\rho)})\dfrac{\phi(\pi(\rho))}{4}\|\rho-\sigma\|^2 \end{cases}$$

with ϕ as in (27) and $\pi(\sigma)$ as in (28).

Even using this refinement of Pinsker's inequality, some numerical experiments indicate that (20) is stronger for randomly generated states. From Corollary 5.1, we can expect, our bound to perform well if σ has a large eigenvalue and the smallest eigenvalue is as large as possible. Such states have spectrum of the form $\left(p, \frac{1-p}{d-1}, \ldots, \frac{1-p}{d-1}\right)$. Indeed for $\sigma \in \mathcal{D}_5$ with spectrum (0.99, 0.0025, 0.0025, 0.0025, 0.0025) and $q < 0.2$ our bound performs better than (19) for randomly generated ρ. However, even in this case, the improvement is *not* significant.

Still we can expect that Theorem 5.1 will find more applications, for instance improving the mixing time bounds. Such bounds have been derived from log-Sobolev inequalities in Ref. 2. The next theorem can be used to improve these results.

Theorem 5.2. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with fixed point σ that satisfies*

$$D\left(e^{t\mathcal{L}}\rho\|\sigma\right) \leq e^{-2\alpha t} D(\rho\|\sigma) \tag{33}$$

for some $\alpha > 0$ and for all $\rho \in \mathcal{D}_d$ and $t \in \mathbb{R}^+$. Then we have

$$\|e^{t\mathcal{L}}(\rho) - \sigma\|_1 \leq 2e^{-\alpha t}\sqrt{\frac{\log(s_{\min}(\sigma))}{\phi(\pi(\sigma))}} \tag{34}$$

with ϕ as in (27), $\pi(\sigma)$ as in (28) and where $s_{\min}(\sigma)$ is the smallest eigenvalue of σ.

Proof. This is a direct consequence of (3) and (29). □

VI. TENSOR PRODUCTS OF DEPOLARIZING CHANNELS

For a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ generating the channel $\mathcal{T}_t = e^{t\mathcal{L}}$ and any $n \in \mathbb{N}$, we denote by $\mathcal{L}^{(n)} : \mathcal{M}_{d^n} \to \mathcal{M}_{d^n}$ the generator of the tensor-product semigroup $(T_t)^{\otimes n}$, i.e., $\mathcal{L}^{(n)} := \sum_{i=1}^n \mathrm{id}_d^{\otimes i-1} \otimes \mathcal{L} \otimes \mathrm{id}_d^{\otimes(n-i)}$.

Here we study $\alpha_1(\mathcal{L}_\sigma^n)$ in the special case where $\sigma = \frac{\mathbb{1}_d}{d}$. For simplicity, we denote the depolarizing Liouvillian onto $\sigma = \frac{\mathbb{1}_d}{d}$ by $\mathcal{L}_d := \mathcal{L}_{\frac{\mathbb{1}_d}{d}}$ and by $T_t^d = e^{t\mathcal{L}^d}$ the generated semigroup. In the case $d = 2$, it is known[2] that $\alpha_1\left(\mathcal{L}_2^{(n)}\right) = 1$ for any $n \in \mathbb{N}$. It is, however, an open problem to determine this constant for any $d > 2$ and any $n \geq 2$. We will now show the inequality $\alpha_1\left(\mathcal{L}_d^{(n)}\right) \geq \frac{1}{2}$ for any $d \geq 2$ and $n \geq 1$, which is the best possible lower bound that is independent of the local dimension. Note that for $\sigma = \frac{\mathbb{1}_d}{d}$ inequality (3) for the channel $\left(\mathcal{T}_t{}^d\right)^{\otimes n}$ can be rewritten as the entropy production inequality,

$$S((T_t^d)^{\otimes n}(\rho)) \geq (1 - e^{-t})n\log(d) + e^{-t}S(\rho).$$

This inequality has been studied in Ref. 14 for the case where $d = 2$, for which, however, an incorrect proof was given. We will provide a proof of a more general statement, from which the claim $\alpha_1\left(\mathcal{L}_{\mathrm{dep}}^{(n)}\right) \geq \frac{1}{2}$ readily follows by the previous discussion.

Theorem 6.1. *For any $\sigma, \rho \in \mathcal{D}_d$ (not necessarily full rank), we have*

$$S((T_t^\sigma)^{\otimes n}(\rho)) \geq e^{-t}S(\rho) + (1 - e^{-t})S(\sigma^{\otimes n}).$$

For the proof, we will need a special case of the quantum Shearer's inequality. We will denote by $\rho \in \mathcal{D}\left(\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \otimes \cdots \otimes \mathbb{C}^{d_n}\right)$ a multipartite density matrix (where the d_i are the local dimensions of each tensor factor). Furthermore, we write $S(i_1, i_2, \ldots, i_k)_\rho$ for the entropy of the reduced density matrix ρ on the tensor factors specified by the indices $i_1, i_2, \ldots, i_k$. Similarly, we write

$$S(i_1, \ldots, i_k | j_1, \ldots, j_l)_\rho = S(i_1, \ldots, i_k, j_1, \ldots, j_l)_\rho - S(j_1, \ldots, j_l)_\rho$$

for a conditional entropy. The proof of the quantum version of Shearer's inequality is essentially the same as the proof given by Radhakrishnan and Llewellyn for the classical version (see Ref. 15). For convenience, we provide the full proof.

Lemma 6.1 (Quantum Shearer's inequality). Consider $t \in \mathbb{N}$ and a family $\mathcal{F} \subset 2^{\{1,\ldots,n\}}$ of subsets of $\{1, \ldots, n\}$ such that each $i \in \{1, \ldots, n\}$ is included in exactly t elements of $\mathcal{F}$. Then for any $\rho \in \mathcal{D}\left(\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \otimes \cdots \otimes \mathbb{C}^{d_n}\right)$, we have

$$S(1,2,\ldots,n)_\rho \leq \frac{1}{t} \sum_{F \in \mathcal{F}} S(F)_\rho\,. \tag{35}$$

Proof. For $F \subset \{1, \ldots, n\}$, denote its elements by $(i_1, \ldots, i_k)$, increasingly ordered. For any $\rho \in \mathcal{D}\left(\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \otimes \cdots \otimes \mathbb{C}^{d_n}\right)$ we have

$$\begin{aligned}\sum_{j=1}^{|F|} S(i_j | i_1, \ldots, i_{j-1})_\rho &= S(i_1)_\rho + S(i_2|i_1)_\rho + \cdots + S(i_{|F|}|i_1, i_2, \ldots, i_{|F|-1})_\rho \\ &= S(i_1, i_2, \ldots, i_{|F|})_\rho = S(F)_\rho,\end{aligned}$$

where we used a telescopic sum trick. By strong subadditivity[16] conditioning decreases the entropy. This implies

$$\sum_{j=1}^{|F|} S(i_j|1,2,\ldots,i_j - 1)_\rho \leq \sum_{j=1}^{|F|} S(i_j|i_1,\ldots,i_{j-1})_\rho = S(F)_\rho. \tag{36}$$

Now consider a family $\mathcal{F} \subset 2^{\{1,\ldots,n\}}$ with the properties stated in the assumptions. Using (36) for the first inequality gives

$$\sum_{F\in\mathcal{F}} S(F)_\rho \geq \sum_{F\in\mathcal{F}} \sum_{j=1}^{|F|} S(i_j|1,2,\ldots,i_j-1)_\rho \tag{37}$$

$$= t\sum_{i=1}^{n} S(i|1,2,\ldots,i-1)_\rho = tS(1,2,\ldots,n)_\rho. \tag{38}$$

Here we used the assumption that each $i \in \{1,\ldots,n\}$ is contained in exactly t elements of $\mathcal{F}$ and (36) in the special case of $F = \{1,\ldots,n\}$ for the final equality. □

Note that in the classical case, Shearer's inequality is true under the weaker assumption that any $i \in \{1,\ldots,d\}$ is contained in *at least* t elements of $\mathcal{F}$. However, as the quantum conditional entropy might be negative,[17] we have to use the stronger assumption to get the equality between (37) and (38) where an ≥ would be enough.

In the special case where $\mathcal{F} = \mathcal{F}_k := \{F \subseteq \{1,\ldots,n\} : |F| = k\}$ denotes the family of k-element subsets of $\{1,\ldots,n\}$ (i.e., every $i \in \{1,\ldots,d\}$ is contained in exactly $\binom{n-1}{k-1} = \frac{k}{n}\binom{n}{k}$ elements of $\mathcal{F}_k$) the quantum Shearer inequality gives

$$\frac{k}{n}S(1,\ldots,n) \leq \frac{1}{\binom{n}{k}} \sum_{F\in\mathcal{F}_k} S(F). \tag{39}$$

This inequality was also proved in Ref. 18, but in a more complicated way and without mentioning the more general quantum Shearer's inequality. It is also used as a lemma (with wrong proof) in Ref. 14, where the rest of the proof of their entropy production estimate is correct. The proof of Theorem 6.1 follows the same lines. For completeness, we will include the full proof here.

Proof of Theorem 6.1. In the following, we will abbreviate $p := e^{-t}$. For a subset $F \subset \{1,\ldots,n\}$, we denote by $\rho|_F$ the reduced density matrix on the tensor factors specified by F. Using this notation, we can write

$$(T_t^\sigma)^{\otimes n}(\rho) = \sum_{k=0}^{n} \sum_{F\in\mathcal{F}_k} (1-p)^k p^{n-k} \left(\bigotimes_{l\in F} \sigma \otimes \rho|_{F^c}\right),$$

where $F^c = \{1,\ldots,n\} \setminus F$. Concavity of the von Neumann entropy implies

$$\begin{aligned} S\left((T_t^\sigma)^{\otimes n}(\rho)\right) &\geq \sum_{k=0}^{n} \sum_{F\in\mathcal{F}_k} (1-p)^k p^{n-k} \left(kS(\sigma) + S(F^c)_\rho\right) \\ &\geq (1-p)nS(\sigma) + \sum_{k=0}^{n} \binom{n}{n-k} \frac{n-k}{n} (1-p)^k p^{n-k} S(\rho) \\ &= (1-p)S(\sigma^{\otimes n}) + pS(\rho). \end{aligned}$$

Here we used the elementary identity $\sum_{k=0}^{n}\binom{n}{k}(1-p)^k p^{n-k} k = (1-p)n$ and (39) for the $(n-k)$-element subsets F^c. □

ACKNOWLEDGMENTS

We thank Ion Nechita for pointing out Shearer's inequality to us and David Reeb for interesting and helpful discussions. D.S.F. acknowledges support from the graduate program TopMath of the Elite Network of Bavaria, the TopMath Graduate Center of TUM Graduate School at Technische Universität München and the Deutscher Akademischer Austauschdienst (DAAD). We also acknowledge financial support from the CHIST-ERA/BMBF project CQC (A.M.H. and M.M.W.).

APPENDIX A: QUASI-CONCAVITY OF A QUOTIENT OF RELATIVE ENTROPIES

In this appendix, we will prove the quasi-concavity of the function $y \mapsto q_y(x)$ for any $x \in (0,1)$. As defined in (8), the function $q_y : (0,1) \to \mathbb{R}$ denotes the continuous extension of $x \mapsto \frac{D_2(y\|x)}{D_2(x\|y)}$. In the following, we consider $f_x : [0,1] \to \mathbb{R}$ defined as $f_x(y) = q_y(x)$ for any $y \in (0,1)$ and with $f_x(0) = f_x(1) = 1$. It can be checked easily that f_x is continuous for any $x \in (0,1)$. We have the following Lemma:

Lemma A.1. For any $x \in (0,1)$ the function $f_x : [0,1] \to \mathbb{R}$ given by $f_x(y) = \frac{D_2(y\|x)}{D_2(x\|y)}$ for $y \notin \{0,x,1\}$ and extended continuously by $f_x(x) = 1$ and $f_x(0) = f_x(1) = 0$ is quasi-concave.

Proof. Note that without loss of generality we can assume $x \geq \frac{1}{2}$, as $f_x(y) = f_{1-x}(1-y)$. By continuity, it is clear that there exists an $m_f \in (0,1)$ (we can exclude the boundary points since $f_x(x) > f_x(0) = f_x(1)$) such that $f_x(m_f)$ is the global maximum. By Ref. 7 [p. 99] it is sufficient to show that f_x is unimodal, i.e., that f_x is monotonically increasing on $\left[0, m_f\right)$ and monotonically decreasing on $\left(m_f, 1\right]$. We will use the method of L'Hospital type rules for monotonicity developed in Refs. 19 and 20.

For any $x \in (0,1)$ and $y \in (0,1)$ with $x \neq y$, we compute

$$\partial_y D_2(y\|x) = \log\left(\frac{y(1-x)}{x(1-y)}\right), \qquad \partial_y D_2(x\|y) = \frac{y-x}{y(1-y)},$$
$$\partial_y \log\left(\frac{y(1-x)}{x(1-y)}\right) = \frac{1}{y(1-y)}, \qquad \partial_y \frac{y-x}{y(1-y)} = \frac{y^2+x-2yx}{(1-y)^2y^2}$$

and define

$$g_x(y) = \frac{\partial_y D_2(y\|x)}{\partial_y D_2(x\|y)} = \frac{\log\left(\frac{x(1-y)}{y(1-x)}\right) y(1-y)}{x-y}, \tag{A1}$$
$$h_x(y) = \frac{\partial_y \log\left(\frac{x(1-y)}{y(1-x)}\right)}{\partial_y \frac{x-y}{y(1-y)}} = \frac{y(1-y)}{y^2+x-2yx}, \tag{A2}$$

where again g_x is extended continuously by $g_x(0) = g_x(1) = 0$ and $g_x(x) = 1$. As $y \mapsto y^2 + x - 2yx$ has no real zeros for $x \in (0,1)$, the rational function h_x is continuously differentiable on $(0,1)$. A straightforward calculation reveals that for $x \geq \frac{1}{2}$ and on $(0,1)$, the derivative h_x' only vanishes in

$$m_h = \begin{cases} \dfrac{x - \sqrt{x(1-x)}}{2x-1} & \text{for } x > \dfrac{1}{2} \\ \dfrac{1}{2} & \text{for } x = \dfrac{1}{2} \end{cases},$$

which has to be a maximum as $h_x(0) = h_x(1) = 0$. By the lack of further points with vanishing derivative, we have $h_x'(y) < 0$ for any $y < m_h$ and also $h_x'(y) > 0$ for any $y > m_h$. Note that $m_h \leq x$ for any $x \geq \frac{1}{2}$.

Consider first the interval $(x,1) \subset (0,1)$. For $y \to x$ we have $\log\left(\frac{x(1-y)}{y(1-x)}\right) \to 0$ and $\frac{x-y}{y(1-y)} \to 0$. Also it is clear that $y \mapsto \frac{x-y}{y(1-y)}$ does not change sign on the interval $(x,1)$. Therefore and by (A2), we see that the pair g_x and h_x satisfy the assumptions of Ref. 19 [Proposition 1.1.] and as h_x is decreasing we have that $g_x'(y) < 0$ for any $y \in (x,1)$. We can use the same argument for the (possibly empty) interval (m_h, x) where h_x is decreasing as well and obtain $g_x'(y) < 0$ for any $y \in (m_h, x)$. By continuity of g_x in x, we see that g_x is decreasing on $(m_h, 1)$.

Note that in the case where $x = \frac{1}{2}$, we can directly apply[19] [Proposition 1.1.] to the remaining interval $(0, \frac{1}{2})$ where $h_{1/2}$ is increasing. This proves $g_{1/2}'(y) > 0$ for any $y \in (0, \frac{1}{2})$. By continuity, $m_g = \frac{1}{2}$ is the maximum point of $g_{1/2}$. For $x \neq \frac{1}{2}$, where the remaining interval is $(0, m_h)$, we apply the more general Proposition 2.1. in Ref. 20. It can be checked easily that the assumptions of this proposition are fulfilled for the pair g_x and h_x. As for $y \in (0, m_h)$, we have $\frac{y-x}{y(1-y)} \frac{y^2+x-2yx}{(1-y)^2y^2} < 0$ and as h_x is increasing the proposition shows that $g_x'(y) > 0$ for any $y \in (0, m_g)$ and $g_x'(y) < 0$ for any

$y \in (m_g, m_h)$. Here $m_g \in (0, m_h)$ denotes the maximum point of g_x (note that a maximum m_g has to exist due to continuity and $g_x(0) = g_x(1) = 0$).

The previous argument shows that for any $x \geq \frac{1}{2}$ there exists a point $m_g \in (0, m_h] \subset (0, x]$ (we have $m_g = m_h = \frac{1}{2}$ for $x = \frac{1}{2}$) such that $g'_x(y) > 0$ for $y \in (0, m_g)$ and $g'_x(y) < 0$ for $y \in (m_g, 1) \setminus \{x\}$. We can now repeat the above argument for the pair f_x and g_x. This gives the existence of a point $m_f \in (0, m_g]$ such that $f'_x(y) > 0$ for any $y \in (0, m_f)$ and $f'_x(y) < 0$ for any $y \in (m_f, 1) \setminus \{x\}$. By continuity in x this shows that the function f_x is unimodal and therefore quasi-concave. □

APPENDIX B: CONTINUOUS EXTENSION OF A QUOTIENT OF RELATIVE ENTROPIES

In this section, we show that the function $Q_\sigma : \mathcal{D}_d^+ \to \mathbb{R}$ as defined in (6) is indeed continuous. As Q_σ is clearly continuous in any point $\rho \neq \sigma$ we have to prove the following:

Lemma B.1. For $\sigma \in \mathcal{D}_d^+$ and $X \in \mathcal{M}_d$ with $X = X^\dagger$, $tr[X] = 0$ and $X \neq 0$, we have

$$\lim_{\epsilon \to 0} \frac{D(\sigma \| \sigma + \epsilon X)}{D(\sigma + \epsilon X \| \sigma)} = 1.$$

Proof. To show the claim, we will expand the relative entropy in terms of ϵ up to second order. Observe that for $\rho \in \mathcal{D}_d$ we have

$$D(\rho \| \sigma) = \int_0^\infty \operatorname{tr}\left[\rho\left((\rho + t)^{-1} - (\sigma + t)^{-1}\right)\right] dt. \tag{B1}$$

In the following, we assume $\epsilon > 0$ to be small enough such that $\sigma + \epsilon X \in \mathcal{D}_d^+$. To simplify the notation, we introduce $A(t) := (\sigma + t)^{-1}$ and $B(t) := (\sigma + \epsilon X + t)^{-1}$. Applying the recursive relation

$$B(t) = -\epsilon B(t) X A(t) + A(t),$$

twice leads to

$$\begin{aligned} B(t) - A(t) &= -\epsilon B(t) X A(t) = \epsilon^2 B(t) X A(t) X A(t) - \epsilon A(t) X A(t) \\ &= \epsilon^2 A(t) X A(t) X A(t) - \epsilon A(t) X A(t) + O\left(\epsilon^3\right). \end{aligned}$$

Inserting this into (B1) gives

$$D(\sigma \| \sigma + \epsilon X) = \int_0^\infty \operatorname{tr}\left[\epsilon \sigma A(t) X A(t) - \epsilon^2 \sigma A(t) X A(t) X A(t) + O(\epsilon^3)\right] dt \tag{B2}$$

and

$$D(\sigma + \epsilon X \| \sigma) = \int_0^\infty \operatorname{tr}\left[-\epsilon \sigma A(t) X A(t) + \epsilon^2 \sigma A(t) X A(t) X A(t) - \epsilon^2 X A(t) X A(t) + O(\epsilon^3)\right] dt. \tag{B3}$$

As $[A(t), \sigma] = 0$, we can diagonalize these operators in the same orthonormal basis $\{|i\rangle\} \subset \mathbb{C}^d$, which leads to

$$\int_0^\infty \operatorname{tr}[\sigma A(t) X A(t)]\, dt = \sum_{i=1}^d \langle i | X | i \rangle \int_0^\infty \frac{s_i}{(s_i + t)^2} dt = \sum_{i=1}^d \langle i | X | i \rangle = 0, \tag{B4}$$

where $\{s_i\}_{i=1}^d$ denotes the spectrum of σ. Note that again by diagonalizing σ and $A(t)$ in the same basis, we have

$$\int_0^\infty \operatorname{tr}\left[(2\sigma A(t) - \mathbb{1})\,(XA(t)XA(t))\right] dt \tag{B5}$$

$$= \sum_{i,j=1}^{d} |\langle i|X|j\rangle|^2 \int_0^\infty \frac{2s_i}{(s_i+t)^2(s_j+t)} - \frac{1}{(s_i+t)(s_j+t)} dt$$

$$= \sum_{i,j=1}^{d} \frac{|\langle i|X|j\rangle|^2}{(s_i-s_j)^2}\left(2(s_i-s_j) - (s_i+s_j)\log\left(\frac{s_i}{s_j}\right)\right). \tag{B6}$$

$$= 0.$$

The last equality follows from the fact that the expression in (B6) clearly changes its sign when s_i and s_j are exchanged. This is only possible if the value of the integral (B5) vanishes. Rearranging the integral (B5) gives

$$\int_0^\infty \operatorname{tr}\left[\sigma A(t)XA(t)XA(t)\right] dt = \int_0^\infty \operatorname{tr}\left[-\sigma A(t)XA(t)XA(t) + XA(t)XA(t)\right] dt. \tag{B7}$$

Finally applying (B4) and (B7) to the formulas for the relative entropies (B2) and (B3) gives

$$\frac{D\left(\sigma\|\sigma+\epsilon X\right)}{D\left(\sigma+\epsilon X\|\sigma\right)} = \frac{c+O(\epsilon)}{c+O(\epsilon)} \to 1$$

as $\epsilon \to 0$. Here, $c := \int_0^\infty \operatorname{tr}\left[\sigma A(t)XA(t)XA(t)\right] dt > 0$ as $\sigma, A(t) > 0$ for any $t \in [0,\infty)$ and $X \neq 0$ is Hermitian. □

[1] R. Olkiewicz and B. Zegarlinski, "Hypercontractivity in noncommutative L_p spaces," J. Funct. Anal. **161**(1), 246–285 (1999).

[2] M. J. Kastoryano and K. Temme, "Quantum logarithmic Sobolev inequalities and rapid mixing," J. Math. Phys. **54**(5), 052202 (2013).

[3] H. Spohn, "Entropy production for quantum dynamical semigroups," J. Math. Phys. **19**(5), 1227 (1978).

[4] I. Bengtsson, A. Ericsson, M. Kuś, W. Tadej, and K. Życzkowski, "Birkhoff's polytope and unistochastic matrices, N = 3 and N = 4," Commun. Math. Phys. **259**, 307–324 (2005).

[5] R. Bhatia, *Matrix Analysis* (Springer, 1997), Vol. 169.

[6] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*, Cambridge Series on Information and the Natural Sciences (Cambridge University Press, 2000).

[7] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).

[8] M. D. Sammer, "Aspects of mass transportation in discrete concentration inequalities," Ph.D. thesis, Georgia Institute of Technology, 2005.

[9] S. G. Bobkov and P. Tetali, "Modified logarithmic Sobolev inequalities in discrete settings," J. Theor. Probab. **19**(2), 289–336 (2006).

[10] I. H. Kim, "Modulus of convexity for operator convex functions," J. Math. Phys. **55**(8), 082201 (2014).

[11] K. M. Audenaert and J. Eisert, "Continuity bounds on the quantum relative entropy," J. Math. Phys. **46**(10), 102104 (2005).

[12] E. Ordentlich and M. Weinberger, "A distribution dependent refinement of Pinsker's inequality," IEEE Trans. Inf. Theory **51**(5), 1836–1840 (2005).

[13] W. Hoeffding, "Probability inequalities for sums of bounded random variables," J. Am. Stat. Assoc. **58**(301), 13–30 (1963).

[14] D. Aharonov, M. Ben-Or, R. Impagliazzo, and N. Nisan, "Limitations of noisy reversible computation," e-print arXiv:quant-ph/9611028 (1996).

[15] J. Radhakrishnan, "Entropy and counting," in *Computational Mathematics, Modelling and Algorithms*, edited by J. C. Misra (Narosa Publishing House, 2003), pp. 146–168.

[16] E. H. Lieb and M. B. Ruskai, "Proof of the strong subadditivity of quantummechanical entropy," J. Math. Phys. **14**(12), 1938 (1973).

[17] N. J. Cerf and C. Adami, "Negative entropy and information in quantum mechanics," Phys. Rev. Lett. **79**, 5194–5197 (1997).

[18] M. Junge and C. Palazuelos, "Cb-norm estimates for maps between noncommutative l_p-spaces and quantum channel theory," Int. Math. Res. Notices (published online); preprint arXiv:1407.7684 (2014).

[19] I. Pinelis, "L'Hôspital type rules for monotonicity, with applications," J. Ineq. Pure Appl. Math. **3**(1), article 5 (2002).

[20] I. Pinelis, "Non-strict L'Hôspital-type rules for monotonicity: Intervals of constancy," J. Ineq. Pure Appl. Math. **8**(1), article 14 (2007).

[21] A Liouvillian is primitive if, and only if, it has a unique full rank fixed point σ and for any $\rho \in \mathcal{D}_d^+$ we have $e^{t\mathcal{L}}(\rho) \to \sigma$ as $t \to \infty$.

Appendix B

Core Articles

B.1 Sandwiched Renyi Convergence for Quantum Evolutions

Sandwiched Rényi Convergence for Quantum Evolutions

D. Stilck França, A. Müller-Hermes

We study the convergence of quantum dynamical semigroups under sandwiched Rényi divergences. We derive expressions for the optimal convergence rates and relate these to other essential constants related to the convergence of semigroups, such as the spectral gap and the logarithmic Sobolev constant. These connections allow us to derive mixing time bounds from logarithmic Sobolev inequalities without any further assumptions and obtain bounds on the classical capacity of these semigroups as a function of time. Moreover, we combine these results with other results in the literature to obtain the first bounds on the capacity of stabilizer Hamiltonians under thermal noise.

B.1.1 Main Results

Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank stationary state $\sigma \in \mathcal{D}_d^+$. We will denote the $p-$sandwiched Rényi divergence [6, 7] by $D_p\left(\cdot\|\cdot\right)$, the $p-$Dirichlet form by $\mathcal{E}_p^{\mathcal{L}}$ and the noncommutative l_p-norm with respect to σ by $\|\cdot\|_{p,\sigma}$ for $p \in [1, \infty]$.

Theorem B.1.1 (Entropy Production for Sandwiched Rényi Divergences). *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be as above, $p \geq 1$ and $\beta_p\left(\mathcal{L}\right)$ be the largest constant β such that*

$$D_p\left(e^{t\mathcal{L}^*}\left(\rho\right)\|\sigma\right) \leq e^{-2\beta t} D_p\left(\rho\|\sigma\right) \tag{B.1}$$

holds for all $t \geq 0$ and $\rho \in \mathcal{D}_d$. Define $\kappa_p(X) = \frac{1}{p-1}\|X\|_{p,\sigma}^p \log\left(\frac{\|X\|_{p,\sigma}^p}{\|X\|_{1,\sigma}^p}\right)$ for $X \in \mathcal{M}_d^+$. Then we have

$$\beta_p\left(\mathcal{L}\right) \inf\left\{\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} : X \in \mathcal{M}_d, X > 0\right\}.$$

The theorem follows from reformulating Equation (B.1) as a differential inequality for some fixed state and then optimizing over all initial states. It is known [8, 36] that under technical conditions known as l_p-regularity, one can relate the convergence in the $1-$divergence to other convergence constants related to the semigroup, such as the logarithmic Sobolev constant and the spectral gap. The same holds here but without any extra technical assumptions.

Theorem B.1.2 (Comparison of Constants). *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be as before and denote its logarithmic Sobolev constant by $\alpha_2\left(\mathcal{L}\right)$ and its spectral gap by $\lambda\left(\mathcal{L}\right)$. Then*

$$\lambda\left(\mathcal{L}\right) \geq \beta_2\left(\mathcal{L}\right) \geq \frac{\alpha_2\left(\mathcal{L}\right)}{2}. \tag{B.2}$$

Both the upper and the lower bound in Equation (B.2) follow from showing relations between the functionals used in the definition of the involved constants. We show $\mathrm{Var}_{2,\sigma}(X) \geq \kappa_2(X) \geq 2^{-1}\mathrm{Ent}_{2,\sigma}(X)$, from which the claim follows.

B.1.2 Applications

One of the main applications of this framework is to derive mixing time bounds directly from LS inequalities, something which was only possible under extra technical assumptions before.

Theorem B.1.3 (Mixing time bounds). *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with stationary state $\sigma \in \mathcal{D}_d^+$. Then*

$$t_1(\epsilon) \leq \frac{1}{2\beta_2(\mathcal{L})} \log\left(\frac{2\log\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right)$$

and

$$t_1(\epsilon) \leq \frac{1}{\alpha_2(\mathcal{L})} \log\left(\frac{2\log\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right).$$

The theorem follows from applying Pinsker's inequality to the convergence bound in Equation (B.1) and combining this with the relations of the constants given in (B.2).

We can also explore the connection between sandwiched Rényi divergences and bounds on the classical capacity of quantum channels [7], tensorization results for logarithmic Sobolev inequalities [57] and estimates on spectral gaps of semigroups [10] to obtain the first estimates on the classical capacity of stabilizer Hamiltonians under thermal noise. To this end, we model thermal noise with Davies generators. Here we specialize in the case of the $2D$-toric code, a widely studied stabilizer code.

Theorem B.1.4. *Let H be the stabilizer Hamiltonian of the $2D$ toric code on a $N \times N$ lattice and $\mathcal{L}_\beta$ be its Davies generator at inverse temperature $\beta > 0$. Then the classical capacity $C(e^{t\mathcal{L}_\beta})$ is bounded by*

$$C(e^{t\mathcal{L}_\beta}) \leq \left(2N^2 + 2N + 1 + \log(2)4\beta N(N+1)\right) e^{-r(\beta,N)t}, \tag{B.3}$$

with

$$r(\beta, N) = \frac{e^{-8\beta}}{6\left((10N^2 + 10N + 5)\log(2) + 4\beta N(N+1)\right) + 66}.$$

Moreover, this is a bound in the strong converse sense.

B.1.3 Individual Contribution

I am the principal author of this article. The project's idea was motivated by discussions between Alexander Müller-Hermes and me. After our work on computing optimal convergence rates for depolarizing channels, we were looking for technical tools that could simplify such computations. Noticing many similarities in the expressions involved in the definitions of sandwiched Rényi divergences and logarithmic Sobolev inequalities, I proposed to further investigate these similarities and try to find deeper connections. I was responsible for making all the relevant computations of derivatives necessary to define the β_p constants, which are essentially the results of sections 3.1 and 3.2, and showing how they can be used to characterize the convergence of semigroups. I also did all the computations in section 3.3. I first proved Theorem B.1.2 for the case $p = 2$ and A. Müller-Hermes was responsible for proving the more general version of this statement found in the article (Theorem 4.1 and Theorem 4.3 in the article). My version of the proof that the spectral gap is an upper bound on β_2 is contained in Theorem 4.2. I was responsible for formulating the applications to mixing times and proving the discrete version of the convergence bounds in section 5, although their derivation from the results of the previous chapters was immediately clear to both of us. The converse bound relating the l_2 mixing time and β_2 in Theorem 5.3 was also formulated and proved by me. I was also responsible for proving the mixing time bounds for random Pauli channels in Corollary 5.3. The idea to use these convergence bounds to obtain bounds on classical capacities was inspired by our previous article also included in this dissertation, where we use similar techniques. I was responsible for showing the bounds on the classical capacity of stabilizer Hamiltonians. I proved the bounds on their smallest eigenvalue, such as in Lemma 7.1, and searched the literature for bounds on spectral gaps. Moreover, I made the first version of all plots, and A. Müller-Hermes then improved them. I wrote all sections of the paper, excluding Appendix A, Theorem 4.1 and Theorem 4.3.

Permission to include:

A. Müller-Hermes and D. Stilck França.
Sandwiched Renyi Convergence for Quantum Evolutions.
Quantum, 2, 55, 2018.

Gmail **Daniel Stilck França <dsfranca13@gmail.com>**

Permission to use Article in Thesis (doi:10.22331/q-2018-02-27-55)

2 messages

Daniel Stilck França <dsfranca13@gmail.com> Fri, Mar 9, 2018 at 1:38 AM
To: info@quantum-journal.org

Dear Quantum team,
I have recently published the article
Sandwiched Rényi Convergence for Quantum Evolutions
Alexander Müller-Hermes and Daniel Stilck Franca
(https://quantum-journal.org/papers/q-2018-02-27-55/)
in Quantum. I am planning to include it in my cumulative dissertation.
Therefore, I would like to ask for your permission to include it in my thesis.
Best,
Daniel

Mariana Munarriz <mmunarriz@quantum-journal.org> Fri, Mar 9, 2018 at 6:14 PM
To: Daniel Stilck França <dsfranca13@gmail.com>

Dear Daniel,

we are happy to inform you that you are free to include the paper in your dissertation and thesis, since authors retain copyright of any papers published in Quantum.
Nevertheless, we would highly appreciate a citation to the journal version of it, as I am sure you understand.
You can also refer to Quantum's Terms and conditions page for more information about our code of conduct:
https://quantum-journal.org/about/terms-and-conditions/

Please don't hesitate to contact us should you have any other queries.

Best regards,
Mariana

Management Assistant
Quantum
[Quoted text hidden]

Sandwiched Rényi Convergence for Quantum Evolutions

Alexander Müller-Hermes[1] and Daniel Stilck França[2]

[1]Department of Mathematical Sciences, University of Copenhagen, 2100 Copenhagen, Denmark
[2]Department of Mathematics, Technische Universität München, 85748 Garching, Germany

February 24, 2018

We study the speed of convergence of a primitive quantum time evolution towards its fixed point in the distance of sandwiched Rényi divergences. For each of these distance measures the convergence is typically exponentially fast and the best exponent is given by a constant (similar to a logarithmic Sobolev constant) depending only on the generator of the time evolution. We establish relations between these constants and the logarithmic Sobolev constants as well as the spectral gap. An important consequence of these relations is the derivation of mixing time bounds for time evolutions directly from logarithmic Sobolev inequalities without relying on notions like l_p-regularity. We also derive strong converse bounds for the classical capacity of a quantum time evolution and apply these to obtain bounds on the classical capacity of some examples, including stabilizer Hamiltonians under thermal noise.

Alexander Müller-Hermes: muellerh@posteo.net
Daniel Stilck França: dsfranca@mytum.de

Contents

1 Introduction 3

2 Notation and Preliminaries 4

2.1 Noncommutative l_p-spaces . . . 5

2.2 Quantum dynamical semigroups . . . 6

2.3 Logarithmic Sobolev inequalities and the spectral gap . . . 7

3 Convergence rates for sandwiched Rényi divergences 9

3.1 Rényi-entropy production . . . 10

3.2 Sandwiched Rényi convergence rates . . . 11

3.3 Computing β_p in simple cases . . . 12

4 Comparison with similar quantities 13

4.1 Comparison with spectral gap . . . 13

4.2 Comparison with logarithmic Sobolev constants . . . 15

5 Mixing times 16

5.1 Mixing in Continuous Time . . . 17

5.2 Mixing in Discrete Time . . . 19

6 Strong converse bounds for the classical capacity 21

7 Examples of bounds for the classical capacity of Semigroups 23

7.1 Depolarizing Channels . . . 23

7.2 Stabilizer Hamiltonians . . . 24

7.3 2D Toric Code . . . 26

7.4 Truncated harmonic oscillator . . . 26

8 Conclusion and open questions 28

Acknowledgments 29

A Taylor Expansion of the Dirichlet Form 29

B Interpolation Theorems and Proof of Theorem 5.3 31

References 33

1 Introduction

Consider a quantum system affected by Markovian noise modeled by a quantum dynamical semigroup T_t (with time parameter $t \in \mathbb{R}^+$) driving every initial state towards a unique full rank state σ. Using the framework of logarithmic Sobolev inequalities as introduced in [1, 2] the speed of the convergence towards the fixed point can be studied. Specifically, the α_1-logarithmic Sobolev constant (see [1, 2]) is the optimal exponent $\alpha \in \mathbb{R}^+$ such that the inequality

$$D(T_t(\rho)\|\sigma) \leq e^{-2\alpha t} D(\rho\|\sigma) \tag{1}$$

holds for the quantum Kullback-Leibler divergence, given by $D(\rho\|\sigma) = \operatorname{tr}[\rho(\ln(\rho) - \ln(\sigma))]$, for all $t \in \mathbb{R}^+$ and all states ρ.

The framework of logarithmic Sobolev constants is closely linked to properties of noncommutative l_p-norms, and specifically to hypercontractivity [1, 2]. Noncommutative l_p-norms also appeared recently in the definition of generalized Rényi divergences (so called "sandwiched Rényi divergences" [3, 4]). It is therefore natural to study the relationship between logarithmic Sobolev inequalities and noncommutative l_p-norms more closely. The approach used here is to define constants (which we call β_p for a parameter $p \in [1, \infty)$), which resemble the logarithmic Sobolev constants, but where the distance measure is a sandwiched Rényi divergence instead of the quantum Kullback-Leibler divergence. More specifically, the constants β_p will be the optimal exponents such that inequalities of the form (1) hold for the sandwiched Rényi divergences D_p, given by

$$D_p(\rho\|\sigma) = \begin{cases} \frac{1}{p-1} \ln\left(\operatorname{tr}\left[\left(\sigma^{\frac{1-p}{2p}} \rho \sigma^{\frac{1-p}{2p}}\right)^p\right]\right) & \text{if } \ker(\sigma) \subseteq \ker(\rho) \text{ or } p \in (0,1) \\ +\infty, & \text{otherwise,} \end{cases} \tag{2}$$

instead of the quantum Kullback-Leibler divergence D.

Our main results are two-fold:

- We derive inequalities between the new β_p and other quantities such as logarithmic Sobolev constants and the spectral gap of the generator of the time evolution. These inequalities not only reveal basic properties of the β_p, but can also be used as a technical tool to strengthen results involving logarithmic Sobolev constants.
- We apply our framework to derive bounds on the mixing time of quantum dynamical semigroups. Using the interplay between the β_p and the logarithmic Sobolev constants we show how to derive a mixing time bound with the same scaling as that of the one derived in [2] directly from a logarithmic Sobolev constant. Previously, this was only known under the additional assumption of l_p-regularity (see [2]) of the generator or for the α_1-logarithmic Sobolev constant. It is still an open question whether l_p-regularity holds for all primitive generators.

As an additional application of our methods we derive time-dependent strong converse bounds on the classical capacity of a quantum dynamical semigroup. We apply these to some examples of systems under thermal noise. These include stabilizer Hamiltonians, such as the $2D$ toric code, and a truncated harmonic oscillator. To the best of our knowledge, these are the first bounds available on the classical capacity of these channels. We also apply our bound to depolarizing channels, whose classical capacity is known [5], to benchmark our findings.

2 Notation and Preliminaries

Throughout this paper $\mathcal{M}_d$ will denote the space of $d \times d$ complex matrices. We will denote by $\mathcal{D}_d$ the set of d-dimensional quantum states, i.e. positive semi-definite matrices $\rho \in \mathcal{M}_d$ with trace 1. By $\mathcal{M}_d^+$ we denote the set of positive definite matrices and by $\mathcal{D}_d^+ = \mathcal{M}_d^+ \cap \mathcal{D}_d$ the set of full rank states.

In [3, 4] the following definition of sandwiched quantum Rényi divergences was proposed:

Definition 2.1 (Sandwiched p-Rényi divergence). *Let $\rho, \sigma \in \mathcal{D}_d$. For $p \in (0,1) \cup (1, \infty)$, the **sandwiched p-Rényi divergence** is defined as:*

$$D_p\left(\rho\|\sigma\right) = \begin{cases} \frac{1}{p-1} \ln\left(tr\left[\left(\sigma^{\frac{1-p}{2p}} \rho \sigma^{\frac{1-p}{2p}} \right)^p \right] \right) & \textit{if } ker(\sigma) \subseteq ker(\rho) \;\; \textit{or } p \in (0,1) \\ +\infty, & \textit{otherwise} \end{cases} \tag{3}$$

where $ker(\sigma)$ is the kernel of σ.

Note that we are using a different normalization than in [3, 4], which is more convenient for our purposes. The logarithm in our definition is in base e, while theirs is in base 2. When we write log in later sections we will mean the logarithm in base 2.

Taking the limit $p \to 1$ gives the usual quantum Kullback-Leibler divergence [6]

$$\lim_{p\to 1} D_p\left(\rho\|\sigma\right) = D\left(\rho\|\sigma\right) := \begin{cases} \mathrm{tr}[\rho\left(\ln(\rho) - \ln(\sigma)\right)] & \text{if } \ker\left(\sigma\right) \subseteq \ker\left(\rho\right) \\ +\infty, & \text{otherwise} \end{cases}.$$

Similarly by taking the limit $p \to \infty$ we obtain the max-relative entropy [3, Theorem 5]

$$\lim_{p\to\infty} D_p\left(\rho\|\sigma\right) = D_\infty\left(\rho\|\sigma\right) = \ln\left(\|\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}}\|_\infty \right).$$

The sandwiched Rényi divergences increase monotonically in the parameter $p \geq 1$ (see [7, Theorem 7]) and we have

$$D\left(\rho\|\sigma\right) = D_1\left(\rho\|\sigma\right) \leq D_p\left(\rho\|\sigma\right) \leq D_q\left(\rho\|\sigma\right) \leq D_\infty\left(\rho\|\sigma\right). \tag{4}$$

for any $q \geq p \geq 1$ and all $\rho, \sigma \in \mathcal{D}_d$. Next we state two simple consequences of this ordering, which will be useful later.

Lemma 2.1. *For $\sigma \in \mathcal{D}_d^+$ and $p \in [1, +\infty)$*

$$\sup_{\rho \in \mathcal{D}_d} D_p\left(\rho\|\sigma\right) = \ln\left(\|\sigma^{-1}\|_\infty \right). \tag{5}$$

Proof. Using (4) for $\rho \in \mathcal{D}_d$ we have

$$D_p(\rho\|\sigma) \leq D_\infty\left(\rho\|\sigma\right) = \ln\left(\|\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}}\|_\infty \right) \leq \ln\left(\|\sigma^{-1}\|_\infty \right).$$

Here we used that any quantum state $\rho \in \mathcal{D}_d$ fulfills $\rho \leq \mathbb{1}_d$. Clearly, choosing $\rho = |v_{\min}\rangle\langle v_{\min}|$ for an eigenvector $|v_{\min}\rangle \in \mathbb{C}^d$ corresponding to the eigenvalue $\|\sigma^{-1}\|_\infty$ of σ^{-1} achieves equality in the previous bound.

□

Using (4) together with the well-known Pinsker inequality [8, Theorem 3.1] for the quantum Kullback-Leibler divergence we have

$$\frac{1}{2}\|\sigma-\rho\|_1^2 \leq D\left(\rho\|\sigma\right) \leq D_p\left(\rho\|\sigma\right) \tag{6}$$

for any $p \geq 1$ and all $\rho, \sigma \in \mathcal{D}_d$. The constant $\frac{1}{2}$ has been shown to be optimal in the classical case (see [9]), i.e. restricting to ρ that commute with σ, and is therefore also optimal here.

2.1 Noncommutative l_p-spaces

In the following $\sigma \in \mathcal{D}_d^+$ will denote a full rank reference state. For $p \geq 1$ we define the **noncommutative p-norm** with respect to σ as

$$\|X\|_{p,\sigma} = \left(\operatorname{tr}\left[\left|\sigma^{\frac{1}{2p}} X \sigma^{\frac{1}{2p}}\right|^p\right]\right)^{\frac{1}{p}} \tag{7}$$

for any $X \in \mathcal{M}_d$. The space $(\mathcal{M}_d, \|\cdot\|_{p,\sigma})$ is called a (weighted) noncommutative l_p-space. For a linear map $\Phi : \mathcal{M}_d \to \mathcal{M}_d$ and $p, q \geq 1$ we define the **noncommutative $p \to q$-norm** with respect to σ as

$$\|\Phi\|_{p\to q,\sigma} = \sup_{Y \in \mathcal{M}_d} \frac{\|\Phi(Y)\|_{q,\sigma}}{\|Y\|_{p,\sigma}}.$$

We introduce the **weighting operator** $\Gamma_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ as

$$\Gamma_\sigma\left(X\right) = \sigma^{\frac{1}{2}} X \sigma^{\frac{1}{2}}.$$

For powers of the weighting operator we set

$$\Gamma_\sigma^p\left(X\right) = \sigma^{\frac{p}{2}} X \sigma^{\frac{p}{2}}$$

for $p \in \mathbb{R}$ and $X \in \mathcal{M}_d$. We define the so called **power operator** $I_{p,q} : \mathcal{M}_d \to \mathcal{M}_d$ as

$$I_{p,q}(X) = \Gamma_\sigma^{-\frac{1}{p}}\left(\left|\Gamma_\sigma^{\frac{1}{q}}\left(X\right)\right|^{\frac{q}{p}}\right) \tag{8}$$

for $X \in \mathcal{M}_d$. It can be verified that

$$\|I_{p,q}(X)\|_{p,\sigma}^p = \|X\|_{q,\sigma}^q$$

for any $X \in \mathcal{M}_d$. As in the commutative theory, the noncommutative l_2-space turns out to be a Hilbert space, where the **weighted scalar product** is given by

$$\langle X, Y\rangle_\sigma = \operatorname{tr}\left[\Gamma_\sigma\left(X^\dagger\right)Y\right] \tag{9}$$

for $X, Y \in \mathcal{M}_d$. With the above notions we can express the sandwiched p-Rényi divergence (3) for $p > 1$ in terms of a noncommutative l_p-norm as

$$D_p\left(\rho\|\sigma\right) = \frac{1}{p-1}\ln\left(\|\Gamma_\sigma^{-1}\left(\rho\right)\|_{p,\sigma}^p\right). \tag{10}$$

For a state $\rho \in \mathcal{D}_d$ the positive matrix $\Gamma_\sigma^{-1}\left(\rho\right) \in \mathcal{M}_d$ is called the **relative density** of ρ with respect to σ. Note that any $X \geq 0$ with $\|X\|_{1,\sigma} = 1$ can be written as $X = \Gamma_\sigma^{-1}\left(\rho\right)$ for some state $\rho \in \mathcal{D}_d$. We will simply call operators $X \geq 0$ that satisfy $\|X\|_{1,\sigma} = 1$ relative densities when the reference state is clear.

We refer to [1, 2] and references therein for proofs and more details about the concepts introduced in this section.

2.2 Quantum dynamical semigroups

A family of quantum channels, i.e. trace-preserving completely positive maps, $\{T_t\}_{t\in\mathbb{R}_0^+}$, $T_t : \mathcal{M}_d \to \mathcal{M}_d$, parametrized by a non-negative parameter $t \in \mathbb{R}_0^+$ is called a **quantum dynamical semigroup** if $T_0 = \mathrm{id}_d$ (the identity map in d dimensions), $T_{t+s} = T_t \circ T_s$ for any $s, t \in \mathbb{R}_0^+$ and T_t depends continuously on t. Any quantum dynamical semigroup can be written as $T_t = e^{t\mathcal{L}}$ (see [10, 11]) for a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ of the form

$$\mathcal{L}(X) = \mathcal{S}(X) - \kappa X - X\kappa^\dagger,$$

where $\kappa \in \mathcal{M}_d$ and $\mathcal{S} : \mathcal{M}_d \to \mathcal{M}_d$ is completely positive such that $\mathcal{S}^*(\mathbb{1}_d) = \kappa + \kappa^\dagger$, where $\mathcal{S}^*$ is the adjoint of $\mathcal{S}$ with respect to the Hilbert-Schmidt scalar product. We will also deal with tensor powers of semigroups. For a quantum dynamical semigroup $\{T_t\}_{t\in\mathbb{R}^+}$ with Liouvillian $\mathcal{L}$ we denote by $\mathcal{L}^{(n)}$ the Liouvillian of the quantum dynamical semigroup $\{T_t^{\otimes n}\}_{t\in\mathbb{R}^+}$.

In the following we will consider quantum dynamical semigroups having a full rank fixed point $\sigma \in \mathcal{D}_d^+$, i.e. the Liouvillian generating the semigroup fulfills $\mathcal{L}(\sigma) = 0$ (implying that $e^{t\mathcal{L}}(\sigma) = \sigma$ for any time $t \in \mathbb{R}_0^+$). We call a quantum dynamical semigroup (or the Liouvillian generator) **primitive** if it has a unique full rank fixed point σ. In this case for any initial state $\rho \in \mathcal{D}_d$ we have $\rho_t = e^{t\mathcal{L}}(\rho) \to \sigma$ as $t \to \infty$ (see [12, Theorem 14]).

The notion of primitivity can also be defined for discrete semigroups of quantum channels. For a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ we will sometimes consider the discrete semigroup $\{T^n\}_{n\in\mathbb{N}}$. Similar to the continuous case we will call this semigroup (or the channel T) primitive if there is a unique full rank state $\sigma \in \mathcal{D}_d^+$ with $\lim\limits_{n\to\infty} T^n(\rho) = \sigma$ for any $\rho \in \mathcal{D}_d$. We refer to [12] for other characterizations of primitive channels and sufficient conditions for primitivity.

To study the convergence of a primitive semigroup to its fixed point σ we introduce the time evolution of the relative density $X_t = \Gamma_\sigma^{-1}(\rho_t)$. For any Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$ define

$$\hat{\mathcal{L}} = \Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma \tag{11}$$

to be the generator of the time evolution of the relative density. Indeed it can be checked that

$$X_t = \Gamma_\sigma^{-1}\left(e^{t\mathcal{L}}(\rho)\right) = e^{t\hat{\mathcal{L}}}(X)$$

for any state $\rho \in \mathcal{D}_d$ and relative density $X = \Gamma_\sigma^{-1}(\rho)$. Note that $\|X_t\|_{1,\sigma} = \|X\|_{1,\sigma} = 1$ for all $t \in \mathbb{R}_0^+$. Clearly the semigroup generated by $\hat{\mathcal{L}}$ is completely positive and unital, but it is not trace-preserving in general. In the special case where

$$\Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma = \mathcal{L}^*, \tag{12}$$

the map $\hat{\mathcal{L}}$ generates the adjoint of the initial semigroup, i.e. the corresponding time evolution in Heisenberg picture. A semigroup fulfilling (12) is called **reversible** (or said to fulfill **detailed balance**), and in this case the Liouvillian $\hat{\mathcal{L}}$ is a Hermitian operator w.r.t. the σ-weighted scalar product. We again refer to [2, 1] for more details on these topics. For discrete semigroups we similarly set $\hat{T} = \Gamma_\sigma^{-1} \circ T \circ \Gamma_\sigma$.

One important class of semigroups are **Davies generators**, which describe a system weakly coupled to a thermal bath under an appropriate approximation [13]. Describing them in detail goes beyond the scope of this article and here we will only review their most basic properties. We refer to [14, 15, 16] for more details.

Suppose that we have a system of dimension d weakly coupled to a thermal bath of dimension d_B at inverse inverse temperature $\beta > 0$. Consider a Hamiltonian $H_{\text{tot}} \in \mathcal{M}_d \otimes \mathcal{M}_{d_B}$ of the system and the bath of the form

$$H_{\text{tot}} = H \otimes \mathbb{1}_B + \mathbb{1}_S \otimes H_B + H_I,$$

where $H \in \mathcal{M}_d$ is the Hamiltonian of the system, $H_B \in \mathcal{M}_{d_B}$ of the bath and

$$H_I = \sum_\alpha S^\alpha \otimes B^\alpha \in \mathcal{M}_d \otimes \mathcal{M}_{d_B} \tag{13}$$

describes the interaction between the system and the bath. Here the operators S^α and B^α are self-adjoint. Let $\{\lambda_k\}_{k\in[d]}$ be the spectrum of the Hamiltonian H. We then define the Bohr-frequencies $\omega_{i,j}$ to be given by the differences of eigenvalues of H, that is, $\omega_{i,j} = \lambda_i - \lambda_j$ for different values of λ. We will drop the indices on ω from now on to avoid cumbersome notation, as is usually done. Moreover, we introduce operators $S^\alpha(\omega)$ which are the Fourier components of the coupling operators S^α and satisfy

$$e^{iHt}S^\alpha e^{-iHt} = \sum_\omega S^\alpha(\omega)e^{i\omega t}.$$

The canonical form of the Davies generator at inverse temperature $\beta > 0$ in the Heisenberg picture, $\mathcal{L}^*_\beta$, is then given by

$$\mathcal{L}^*_\beta(X) = i[H, X] + \sum_{\omega,\alpha} \mathcal{L}^*_{\omega,\alpha}(X),$$

where

$$\mathcal{L}^*_{\omega,\alpha}(X) = G^\alpha(\omega)\left(S^\alpha(\omega)^\dagger X S^\alpha(\omega) - \frac{1}{2}\{S^\alpha(\omega)^\dagger S^\alpha(\omega), X\}\right).$$

Here $\{X, Y\} = XY + YX$ is the anticommutator and $G^\alpha : \mathbb{R} \to \mathbb{R}$ are the transition rate functions. Their form depends on the choice of the bath model [15]. For our purposes it will be enough to assume that these are functions that satisfy the KMS condition [17], that is, $G^\alpha(-\omega) = G^\alpha(\omega)e^{-\beta\omega}$. Although this presentation of the Davies generators is admittedly very short, for our purposes it will be enough to note that under some assumptions on the operators $S^\alpha(\omega)$ [18, 19] and on the transition rate functions, the semigroup generated by $\mathcal{L}_\beta$ converges to the thermal state $\frac{e^{-\beta H}}{\operatorname{tr}\left(e^{-\beta H}\right)}$ and is reversible [17]. In the examples considered here this will always be the case.

2.3 Logarithmic Sobolev inequalities and the spectral gap

To study hypercontractive properties and convergence times of primitive quantum dynamical semigroups the framework of logarithmic Sobolev inequalities has been developed in [1, 2]. Here we will briefly introduce this theory. For more details and proofs see [1, 2] and the references therein.

We define the **operator valued relative entropy** (for $p > 1$) of $X \in \mathcal{M}_d^+$ as

$$S_p(X) = -p\frac{d}{ds}I_{p+s,p}(X)\,|_{s=0}. \tag{14}$$

With this we can define the p-relative entropy:

Definition 2.2 (p-relative entropy)**.** *For any full rank $\sigma \in \mathcal{M}_d^+$ and $p > 1$ we define the* ***p-relative entropy*** *of $X \in \mathcal{M}_d^+$ as*

$$Ent_{p,\sigma}(X) = \langle I_{q,p}(X), S_p(X)\rangle_\sigma - \|X\|_{p,\sigma}^p \ln(\|X\|_{p,\sigma}), \tag{15}$$

where $\frac{1}{q} + \frac{1}{p} = 1$. For $p = 1$ we can consistently define

$$Ent_{1,\sigma}(X) = tr[\Gamma_\sigma(X)(\ln(\Gamma_\sigma(X)) - \ln(\sigma))].$$

by taking the limit $p \to 1$.

The p-relative entropy is not a divergence in the information-theoretic sense (e.g. it is not contractive under quantum channels). It was originally introduced to study hypercontractive properties of semigroups in [1], where they also show it is positive for positive operators. There is however a connection to the quantum relative entropy as

$$\mathrm{Ent}_{p,\sigma}\left(I_{p,1}\left(\Gamma_\sigma^{-1}(\rho)\right)\right) = \frac{1}{p} D(\rho\|\sigma).$$

As a special case of the last equation we have

$$\mathrm{Ent}_{1,\sigma}\left(\Gamma_\sigma^{-1}(\rho)\right) = D(\rho\|\sigma).$$

We may also use it to obtain an expression for $\mathrm{Ent}_{2,\sigma}$:

$$\mathrm{Ent}_{2,\sigma}(X) = \mathrm{tr}\left[\left(\Gamma_\sigma^{\frac{1}{2}}(X)\right)^2 \ln\left(\frac{\Gamma_\sigma^{\frac{1}{2}}(X)}{\|X\|_{2,\sigma}}\right)\right] - \frac{1}{2}\mathrm{tr}\left[\left(\Gamma_\sigma^{\frac{1}{2}}(X)\right)^2 \ln(\sigma)\right].$$

We also need Dirichlet forms to define logarithmic Sobolev inequalities:

Definition 2.3 (Dirichlet form)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For $p > 1$ we define the* ***p-Dirichlet form*** *of $X \in \mathcal{M}_d^+$ as*

$$\mathcal{E}_p^{\mathcal{L}}(X) = -\frac{p}{2(p-1)}\left\langle I_{q,p}(X), \hat{\mathcal{L}}(X)\right\rangle_\sigma$$

where $\frac{1}{p} + \frac{1}{q} = 1$ and $\hat{\mathcal{L}} = \Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma$ denotes the generator of the time evolution of the relative density (cf. (11)). For $p = 1$ we may take the limit $p \to 1$ and consistently define the 1-Dirichlet form by

$$\mathcal{E}_1^{\mathcal{L}}(X) = -\frac{1}{2} tr\left[\Gamma_\sigma\left(\hat{\mathcal{L}}(X)\right)(\ln(\Gamma_\sigma(X)) - \ln(\sigma))\right].$$

Formally, by making this choice we introduce the logarithmic Sobolev framework for $\hat{\mathcal{L}}$ (i.e. the generator of the time-evolution of the relative density) instead of $\mathcal{L}^*$. While this is a slightly different definition compared to [2], where the Heisenberg picture is used, they are the same for reversible Liouvillians.

In [1] the Dirichlet forms were introduced to study hypercontractive properties of semigroups. As we will see in Theorem 3.1, they appear naturally when we compute the entropy production of the Sandwiched Rényi divergences. From Corollary 3.1 we will be able to infer that the Dirichlet form is positive for positive operators, a fact already proved in [1]. Both the $\mathrm{Ent}_{p,\sigma}$ and the Dirichlet form are intimately related to hypercontractive properties of semigroups, as we have for a relative density X, some constant $\alpha > 0$ and $p(t) = 1 + e^{2\alpha t}$ that

$$\frac{d}{dt} \ln \left(\|X_t\|_{p(t),\sigma} \right) = \frac{\alpha e^{\alpha t}}{(1 + e^{\alpha t}) \|X_t\|_{p(t),\sigma}^{p(t)}} \left(\mathrm{Ent}_{p(t),\sigma}(X_t) - \frac{1}{\alpha} \mathcal{E}_{p(t)}(X_t) \right),$$

as shown in [1].

Notice that when working with $\mathcal{E}_2^{\mathcal{L}}$ we may always suppose the Liouvillian is reversible without loss of generality. This follows from the fact that

$$\mathcal{E}_2^{\mathcal{L}}(X) = -\left\langle X, \hat{\mathcal{L}}(X) \right\rangle_\sigma$$

is invariant under the additive symmetrization $\hat{\mathcal{L}} \mapsto \frac{1}{2}\left(\mathcal{L}^* + \Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma\right)$ for $X \geq 0$.

We can now introduce the logarithmic Sobolev constants:

Definition 2.4 (Logarithmic Sobolev constants)**.** *For a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and $p \geq 1$ the p-**logarithmic Sobolev constant** is defined as*

$$\alpha_p(\mathcal{L}) = \sup\{\alpha \in \mathbb{R}^+ : \alpha Ent_{p,\sigma}(X) \leq \mathcal{E}_p^{\mathcal{L}}(X) \text{ for all } X > 0\} \tag{16}$$

As $\mathrm{Ent}_{2,\sigma}$ does not depend on $\mathcal{L}$ and, as remarked before, $\mathcal{E}_2^{\mathcal{L}}$ is invariant under an additive symmetrization, we may always assume without loss of generality that the Liouvillian is reversible when working with α_2.

For any $X \in \mathcal{M}_d^+$ we can define its **variance** with respect to $\sigma \in \mathcal{D}_d^+$ as

$$\mathrm{Var}_\sigma(X) = \|X\|_{2,\sigma}^2 - \|X\|_{1,\sigma}^2. \tag{17}$$

This defines a distance measure to study the convergence of the semigroup. Given a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with fixed point $\sigma \in \mathcal{D}_d^+$ we define its **spectral gap** as

$$\lambda(\mathcal{L}) = \sup\left\{\lambda \in \mathbb{R}^+ : \lambda \mathrm{Var}_\sigma(X) \leq \mathcal{E}_2^{\mathcal{L}}(X) \text{ for all } X > 0\right\} \tag{18}$$

where $\hat{\mathcal{L}} : \mathcal{M}_d \to \mathcal{M}_d$ is given by (11). We can always assume the Liouvillian to be reversible when dealing with the spectral gap, as it again depends on $\mathcal{E}_2^{\mathcal{L}}$.

The spectral gap can be used to bound the convergence in the variance (see [20]), as for any $X \in \mathcal{M}_d^+$ we have

$$\frac{d}{dt}\mathrm{Var}_\sigma(X_t) = 2\left\langle \hat{\mathcal{L}}(X), X \right\rangle_\sigma \tag{19}$$

and so

$$\mathrm{Var}_\sigma(X_t) \leq e^{-2\lambda t} \mathrm{Var}_\sigma(X). \tag{20}$$

3 Convergence rates for sandwiched Rényi divergences

In this section we consider the sandwiched Rényi divergences of a state evolving under a primitive quantum dynamical semigroup and the fixed point of this semigroup. It is clear that these quantities converge to zero as the time-evolved state approaches the fixed point. To study the speed of this convergence we introduce a differential inequality, which can be seen as an analogue of the logarithmic Sobolev inequalities for sandwiched Rényi divergences.

3.1 Rényi-entropy production

In [18] the entropy production for the quantum Kullback-Leibler divergence of a Liouvillian was computed. We will now derive a similar expression for the entropy production for the p-Rényi divergences for $p > 1$.

Theorem 3.1 (Derivative of the sandwiched p-Rényi divergence)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For any $\rho \in \mathcal{D}_d$ and $p > 1$ we have*

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} = \frac{p}{p-1} \frac{tr\left[\left(\sigma^{\frac{1-p}{2p}} \rho \sigma^{\frac{1-p}{2p}}\right)^{p-1} \sigma^{\frac{1-p}{2p}} \mathcal{L}(\rho) \sigma^{\frac{1-p}{2p}}\right]}{tr\left[\left(\sigma^{\frac{1-p}{2p}} \rho \sigma^{\frac{1-p}{2p}}\right)^{p}\right]}. \tag{21}$$

Using the relative density $X = \Gamma_\sigma^{-1}(\rho)$ and (11) this expression can be written as:

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} = \frac{p}{p-1} \|X\|_{p,\sigma}^{-p} \left\langle I_{q,p}(X), \hat{\mathcal{L}}(X)\right\rangle_\sigma \tag{22}$$

with $\frac{1}{p} + \frac{1}{q} = 1$.

Proof. Rewriting the p-Rényi divergence in terms of the relative density $X = \Gamma_\sigma^{-1}(\rho)$ and the corresponding generator $\hat{\mathcal{L}} = \Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma$ (see (11)) we have

$$D_p(e^{t\mathcal{L}}(\rho)\|\sigma) = \frac{1}{p-1} \ln\left(\|e^{t\hat{\mathcal{L}}}(X)\|_{p,\sigma}^p\right). \tag{23}$$

By the chain rule

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}\rho\|\sigma)\Big|_{t=0} = \frac{1}{p-1} \|X\|_{p,\sigma}^{-p} \left(\frac{d}{dt}\|e^{t\hat{\mathcal{L}}}(X)\|_{p,\sigma}^p\right)\Big|_{t=0}.$$

Define the curve $\gamma : \mathbb{R}_0^+ \to \mathcal{M}_d$ as $\gamma(t) = \sigma^{\frac{1}{2p}} e^{t\hat{\mathcal{L}}}(X) \sigma^{\frac{1}{2p}}$ and observe that

$$\|e^{t\hat{\mathcal{L}}}(X)\|_{p,\sigma}^p = \mathrm{tr}[\gamma(t)^p].$$

As the differential of the function $X \mapsto X^p$ at $A \in \mathcal{M}_d^+$ is given by pA^{p-1}, another application of the chain rule yields

$$\frac{d}{dt}\|e^{t\hat{\mathcal{L}}}(X)\|_{p,\sigma}^p\Big|_{t=0} = p\left\langle \gamma(0)^{p-1}, \frac{d\gamma}{dt}(0)\right\rangle.$$

It is easy to check that $\frac{d\gamma}{dt}(0) = \sigma^{\frac{1}{2p}} \hat{\mathcal{L}}(X) \sigma^{\frac{1}{2p}}$. Inserting this in the above equations and writing it in terms of the power operator (8) we finally obtain

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} = \frac{p}{p-1} \|X\|_{p,\sigma}^{-p} \left\langle I_{q,p}(X), \hat{\mathcal{L}}(X)\right\rangle_\sigma$$

with $\frac{1}{p} + \frac{1}{q} = 1$. Expanding this formula gives (21).

□

By recognizing the p-Dirichlet form in the previous theorem we get:

Corollary 3.1. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For any $\rho \in \mathcal{D}_d$ and $p > 1$ we have*

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} = -2\|X\|_{p,\sigma}^{-p}\mathcal{E}_p^{\mathcal{L}}(X) \leq 0, \tag{24}$$

where we used the relative density $X = \Gamma_\sigma^{-1}(\rho)$.

As we remarked before, Corollary 3.1 implies that the Dirichlet form is always positive for relative densities. To see this, recall that the divergences contract under quantum channels [7] and therefore we have that $\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0} \leq 0$. As $\mathcal{E}_p^{\mathcal{L}}(\lambda X) = \lambda^p \mathcal{E}_p^{\mathcal{L}}(X)$ for $\lambda > 0$, this shows that it is positive for all positive operators by properly normalizing X.

3.2 Sandwiched Rényi convergence rates

For any $p > 1$ we introduce the functional $\kappa_p : \mathcal{M}_d^+ \to \mathbb{R}$ as

$$\kappa_p(X) = \frac{1}{p-1}\|X\|_{p,\sigma}^p \ln\left(\frac{\|X\|_{p,\sigma}^p}{\|X\|_{1,\sigma}^p}\right) \tag{25}$$

for $X \in \mathcal{M}_d^+$. For $p = 1$ we may again take the limit $p \to 1$ and obtain $\kappa_1(X) := \lim_{p\to 1} \kappa_p(X) = \mathrm{Ent}_{1,\sigma}(X)$. Note that κ_p is well-defined and non-negative as $\|X\|_{p,\sigma} \geq \|X\|_{1,\sigma}$ for $p \geq 1$. Strictly speaking the definition also depends on a reference state $\sigma \in \mathcal{D}_d^+$, which we usually omit as it is always the fixed point of the primitive Liouvillian under consideration.

Given a Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$ it is a simple consequence of Corollary 3.1 that for $\rho \neq \sigma$

$$\frac{\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma)\Big|_{t=0}}{D_p(\rho\|\sigma)} = -2\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)}, \tag{26}$$

where we used the relative density $X = \Gamma_\sigma^{-1}(\rho)$, which fulfills $\|X\|_{1,\sigma} = 1$. This motivates the following definition.

Definition 3.1 (Entropic convergence constant for p-Rényi divergence)**.** *For any primitive Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ and $p \geq 1$ we define*

$$\beta_p(\mathcal{L}) = \sup\{\beta \in \mathbb{R}^+ : \beta\kappa_p(X) \leq \mathcal{E}_p^{\mathcal{L}}(X) \textit{ for all } X > 0\}. \tag{27}$$

Note that as a special case we have $\alpha_1(\mathcal{L}) = \beta_1(\mathcal{L})$. It should be also emphasized that the supremum in the previous definition goes over any positive definite $X \in \mathcal{M}_d^+$ and not only over relative densities. However, it is easy to see that we can equivalently write

$$\beta_p(\mathcal{L}) = \inf\left\{\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} : X > 0\right\} = \inf\left\{\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} : X > 0, \|X\|_{1,\sigma} = 1\right\} \tag{28}$$

as replacing $X \mapsto X/\|X\|_{1,\sigma}$ does not change the value of the quotient $\mathcal{E}_p^{\mathcal{L}}(X)/\kappa_p(X)$. Therefore, to compute β_p it is enough to optimize over relative densities (i.e. $X > 0$ fulfilling $\|X\|_{1,\sigma} = 1$). By inserting β_p into (26) we have

$$\frac{d}{dt} D_p(e^{t\mathcal{L}}(\rho)\|\sigma) \leq -2\beta_p(\mathcal{L}) D_p\left(e^{t\mathcal{L}}(\rho)\|\sigma\right)$$

for any $\rho \in \mathcal{D}_d$ and Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$. By integrating this differential inequality we get

Theorem 3.2. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. For any $p \geq 1$ and $\rho \in \mathcal{D}_d$ we have*

$$D_p\left(e^{t\mathcal{L}}(\rho)\|\sigma\right) \leq e^{-2\beta_p(\mathcal{L})t} D_p\left(\rho\|\sigma\right) \tag{29}$$

where $\beta_p(\mathcal{L})$ is the constant defined in (27).

3.3 Computing β_p in simple cases

In general it is not clear how to compute β_p and it does not depend on spectral data of $\mathcal{L}$ alone. This is not surprising, as the computation of the usual logarithmic Sobolev constants α_2 or α_1 is also challenging and the exact values are only known for few Liouvillians [21, 2, 22]. In the following we compute β_2 for the depolarizing semigroups.

Theorem 3.3 (β_2 for the depolarizing Liouvillian)**.** *Let $\mathcal{L}_\sigma : \mathcal{M}_d \to \mathcal{M}_d$ denote the depolarizing Liouvillian given by $\mathcal{L}_\sigma(\rho) = tr(\rho)\,\sigma - \rho$ with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$\beta_2(\mathcal{L}_\sigma) = \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\ln\left(\|\sigma^{-1}\|_\infty\right)}. \tag{30}$$

Proof. Without loss of generality we can restrict to $X > 0$ with $\|X\|_{1,\sigma} = 1$ in the minimization (28). Observe that the generator of the time evolution of the relative density (see (11)) for the depolarizing Liouvillian is

$$\hat{\mathcal{L}}_\sigma(X) = \mathrm{tr}\left(\sigma^{\frac{1}{2}} X \sigma^{\frac{1}{2}}\right)\mathbb{1} - X.$$

An easy computation yields $\mathcal{E}_2^{\mathcal{L}_\sigma}(X) = \|X\|_{2,\sigma}^2 - 1$ and so

$$\frac{\mathcal{E}_2^{\mathcal{L}_\sigma}(X)}{\kappa_2(X)} = \frac{1 - \frac{1}{\|X\|_{2,\sigma}^2}}{\ln\left(\|X\|_{2,\sigma}^2\right)}.$$

As the function $x \mapsto \frac{1-\frac{1}{x}}{\ln(x)}$ is monotone decreasing for $x \geq 1$, we have

$$\inf_{X>0} \frac{\mathcal{E}_2^{\mathcal{L}_\sigma}(X)}{\kappa_2(X)} = \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\ln\left(\|\sigma^{-1}\|_\infty\right)}, \tag{31}$$

where we used

$$\sup_{X \geq 0, \|X\|_{1,\sigma}=1} \|X\|_{2,\sigma}^2 = \|\sigma^{-1}\|_\infty,$$

which easily follows from Lemma 2.1 by exponentiating both sides of Equation (5) and using the correspondence between relative densities and states. □

The exact value of $\alpha_2(\mathcal{L}_\sigma)$ is open to the best of our knowledge, but in the case of $\sigma = \frac{\mathbb{1}}{d}$ we have $\alpha_2\left(\mathcal{L}_{\frac{\mathbb{1}}{d}}\right) = \frac{2(1-2/d)}{\ln(d-1)}$ [2, Theorem 24], which is of the same order of magnitude as β_2 for these semigroups.

Computing β_p for $p \neq 2$ seems not to be straightforward even for depolarizing channels, but for the semigroup depolarizing to the maximally mixed state we can at least provide upper and lower bounds.

Theorem 3.4 (β_p for the Liouvillian depolarizing to the maximally mixed state)**.** *Let* $\mathcal{L}(\rho) = tr(\rho)\frac{\mathbb{1}}{d} - \rho$*. For* $p \geq 2$ *we have*

$$\frac{p}{2(p-1)}\frac{1}{\ln(d)} \geq \beta_p(\mathcal{L}) \geq \frac{p}{2(p-1)}\frac{d^{\frac{p-1}{p}}-1}{d^{\frac{p-1}{p}}\ln(d)}.$$

Proof. The Dirichlet Form of this Liouvillian for $X > 0$ with $\|X\|_{1,\frac{\mathbb{1}}{d}} = 1$ is given by

$$\mathcal{E}_p^{\mathcal{L}}(X) = \frac{p}{2(p-1)}(\|X\|^p_{p,\frac{\mathbb{1}}{d}} - \|X\|^{p-1}_{p-1,\frac{\mathbb{1}}{d}}).$$

Dividing this expression by $\kappa_p(X)$ we get

$$\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} = \frac{1 - \frac{\|X\|^{p-1}_{p-1,\frac{\mathbb{1}}{d}}}{\|X\|^p_{p,\frac{\mathbb{1}}{d}}}}{2\ln\left(\|X\|_{p,\frac{\mathbb{1}}{d}}\right)}. \tag{32}$$

By the monotonicity of the weighted norms, we have

$$\frac{\|X\|^{p-1}_{p-1,\frac{\mathbb{1}}{d}}}{\|X\|^p_{p,\frac{\mathbb{1}}{d}}} \leq \frac{1}{\|X\|_{p,\frac{\mathbb{1}}{d}}}$$

and so

$$\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} \geq \frac{\|X\|_{p,\frac{\mathbb{1}}{d}} - 1}{2\|X\|_{p,\frac{\mathbb{1}}{d}}\ln\left(\|X\|_{p,\frac{\mathbb{1}}{d}}\right)} \tag{33}$$

The expression on the right-hand side of (33) is monotone decreasing in $\|X\|_{p,\frac{\mathbb{1}}{d}}$ and so the infimum is attained at

$$\sup_{\|X\|_{1,\frac{\mathbb{1}}{d}}=1} \|X\|_{p,\frac{\mathbb{1}}{d}} = d^{\frac{p-1}{p}},$$

which again easily follows from Lemma 2.1. The upper bound follows from (32) as

$$\frac{\mathcal{E}_p^{\mathcal{L}}(X)}{\kappa_p(X)} \leq \frac{1}{2\ln\left(\|X\|_{p,\frac{\mathbb{1}}{d}}\right)}.$$

which is again monotone decreasing in $\|X\|_{p,\frac{\mathbb{1}}{d}}$. □

From the relations between LS constants [2, Proposition 13], it follows that for the LS constants of the depolarizing channels we have $\alpha_p\left(\mathcal{L}_{\frac{\mathbb{1}}{d}}\right) \geq \alpha_2\left(\mathcal{L}_{\frac{\mathbb{1}}{d}}\right) = \frac{2(1-2/d)}{\ln(d-1)}$ for $p \geq 1$. The constants β_p and α_p are therefore of the same order in this case for small $p \geq 2$.

4 Comparison with similar quantities

4.1 Comparison with spectral gap

Here we show how β_p, see (27), compares to the spectral gap (18) of a Liouvillian. This is motivated by similar results for logarithmic Sobolev constants, where it was shown [2, Theorem 16] that $\alpha_1(\mathcal{L}) \leq \lambda(\mathcal{L})$ for reversible semigroups, a result we recover and generalize here.

Theorem 4.1 (Upper bound spectral gap)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive and reversible Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and $p \geq 1$. Then*

$$\beta_p(\mathcal{L}) \leq \lambda(\mathcal{L}). \tag{34}$$

Proof. Let $(s_i)_{i=1}^d$ denote the spectrum of $\sigma^{1/p}$ and choose a unitary U such that

$$\sigma^{1/p} = U \mathrm{diag}(s_1, s_2, \ldots, s_d) U^\dagger.$$

As $\mathcal{L}$ is reversible, there is a self-adjoint eigenvector $X \in \mathcal{M}_d$ of $\hat{\mathcal{L}}$ corresponding to the spectral gap, i.e. $\hat{\mathcal{L}}(X) = -\lambda(\mathcal{L})X$. Let $\epsilon_0 > 0$ be small enough such that $Y_\epsilon = \mathbb{1}_d + \epsilon X$ is positive for any $|\epsilon| \leq \epsilon_0$. For $|\epsilon| \leq \epsilon_0$ we use Lemma A.1 of the appendix to show

$$\beta_p(\mathcal{L}) \leq \frac{\mathcal{E}_p^{\mathcal{L}}(Y_\epsilon)}{\kappa_p(Y_\epsilon)} = \frac{\lambda(\mathcal{L})\frac{p}{2(p-1)}\left(2\epsilon^2 \sum_{1\leq i\leq j\leq d} f_p(s_i, s_j) b_{ij} b_{ji} + O(\epsilon^3)\right)}{\frac{\epsilon^2}{p-1}\left(p\sum_{1\leq i\leq j\leq d} f_p(s_i, s_j) b_{ij} b_{ji}\right) + O(\epsilon^3)} \tag{35}$$

where $b_{ij} = (U^\dagger \sigma^{1/2p} X \sigma^{1/2p} U)_{ij}$ and

$$f_p(x, y) = \begin{cases} (p-1)x^{p-2} & \text{if } x = y \\ \frac{x^{p-1} - y^{p-1}}{x-y} & \text{else.} \end{cases} \tag{36}$$

Observe that $f_p(s_i, s_j) > 0$ for $s_i, s_j > 0$. Moreover, as $U^\dagger \sigma^{1/2p} X \sigma^{1/2p} U$ is non-zero and self-adjoint we have $b_{ij} b_{ji} \geq 0$ for all i, j and this inequality is strict for at least one choice of i, j. Therefore, the terms of second order in ϵ in the numerator and denominator of (35) are strictly positive, and we obtain $\lambda(\mathcal{L})$ as the limit of the quotient as $\epsilon \to 0$.

□

A similar argument as the one given in the previous proof shows that all real, nonzero elements of the spectrum of $\hat{\mathcal{L}}$ are upper bounds to β_p without invoking reversibility.

Note that in the case of $p = 2$ (see the discussion after (16)) we may assume that the Liouvillian is reversible without loss of generality and drop the requirement of reversibility in the previous theorem. Alternatively, we can obtain the same statement directly from a simple functional inequality. In this case we can also give a lower bound on β_2 in terms of the spectral gap.

Theorem 4.2 (Upper and lower bound for β_2)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$\lambda(\mathcal{L}) \frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\ln(\|\sigma^{-1}\|_\infty)} \leq \beta_2(\mathcal{L}) \leq \lambda(\mathcal{L}). \tag{37}$$

To prove Theorem 4.2 we need the following Lemma.

Lemma 4.1. *For any $X \in \mathcal{M}_d$ we have*

$$Var_\sigma(X) \leq \kappa_2(X).$$

Proof. For $X > 0$ dividing both sides of the inequality by $\|X\|^2_{1,\sigma}$ yields

$$\frac{\|X\|^2_{2,\sigma}}{\|X\|^2_{1,\sigma}} - 1 \leq \frac{\|X\|^2_{2,\sigma}}{\|X\|^2_{1,\sigma}} \ln\left(\frac{\|X\|^2_{2,\sigma}}{\|X\|^2_{1,\sigma}}\right).$$

This follows from the elementary inequality $x - 1 \leq x\ln(x)$ for $x \geq 1$, where we use the ordering $\|X\|_{2,\sigma} \geq \|X\|_{1,\sigma}$ for any $X \in \mathcal{M}_d$.

□

Proof of Theorem 4.2. Using the definition of β_2 (see (27)) and Lemma 4.1 yields

$$\beta_2 \mathrm{Var}_\sigma(X) \leq \beta_2 \kappa_2(X) \leq \mathcal{E}^{\mathcal{L}}_2(X).$$

Now the variational definition of $\lambda(\mathcal{L})$ (see (18)) implies the second inequality of (37).

To prove the first inequality of (37) consider the depolarizing Liouvillian

$$\mathcal{L}_\sigma(X) = \mathrm{tr}(X)\sigma - X.$$

By Theorem 3.3 we have

$$\frac{1 - \frac{1}{\|\sigma^{-1}\|_\infty}}{\ln\left(\|\sigma^{-1}\|_\infty\right)}\kappa_2(X) \leq \mathcal{E}^{\mathcal{L}_\sigma}_2(X)$$

As $\mathcal{E}^{\mathcal{L}_\sigma}_2(X) = \mathrm{Var}_\sigma(X)$, we have $\mathcal{E}^{\mathcal{L}_\sigma}_2(X) \leq \frac{1}{\lambda(\mathcal{L})}\mathcal{E}^{\mathcal{L}}_2(X)$ by the variational definition of $\lambda(\mathcal{L})$ (see (18)). Inserting this in the above inequality finishes the proof.

□

4.2 Comparison with logarithmic Sobolev constants

Here we show how β_p, see (27), compares to the logarithmic Sobolev constant α_p.

Theorem 4.3. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank fixed point $\sigma \in \mathcal{D}^+_d$. Then for any $p \geq 1$ we have*

$$\beta_p\left(\mathcal{L}\right) \geq \frac{\alpha_p\left(\mathcal{L}\right)}{p}. \tag{38}$$

We will need the following Lemma.

Lemma 4.2. *For any full rank state $\sigma \in \mathcal{D}^+_d$, any $p > 1$ and $X \in \mathcal{M}^+_d$ with $\|X\|_{1,\sigma} = 1$ we have*

$$Ent_{p,\sigma}(X) \geq \frac{\kappa_p(X)}{p}. \tag{39}$$

Proof. The function $p \mapsto D_p\left(\rho\|\sigma\right)$ is monotonically increasing [3, 7] and differentiable (as the noncommutative l_p-norm is differentiable in p [1, Theorem 2.7]). Thus, with $f : \mathbb{R}^+ \to \mathbb{R}$ given by $f(t) = t + p$ we have

$$0 \leq \|X\|^p_{p,\sigma}\frac{d}{dt}\left(D_{f(t)}\left(\rho\|\sigma\right)\right)\Big|_{t=0} = -\frac{1}{(p-1)^2}\|X\|^p_{p,\sigma}\ln\left(\|X\|^p_{p,\sigma}\right) + \frac{1}{p-1}\frac{d}{dt}\left(\|X\|^{f(t)}_{f(t),\sigma}\right)\Big|_{t=0}.$$

where we used the relative density $X = \Gamma_\sigma^{-1}(\rho)$. The remaining derivative in the above equation has been computed in [1, Theorem 2.7] and we have

$$\frac{d}{dt}\left(\|X\|_{f(t),\sigma}^{f(t)}\right)\Big|_{t=0} = \langle I_{q,p}(X), S_p(X)\rangle_\sigma$$

with the operator valued entropy S_p defined in (14) and $\frac{1}{p}+\frac{1}{q}=1$. Inserting this expression in the above equation we obtain

$$\frac{1}{p-1}\|X\|_{p,\sigma}^p \ln\left(\|X\|_{p,\sigma}^p\right) \leq \langle I_{q,p}(X), S_p(X)\rangle_\sigma \,. \tag{40}$$

for any $X \in \mathcal{M}_d^+$ with $\|X\|_{1,\sigma} = 1$, i.e. for any $X = \Gamma_\sigma^{-1}(\rho)$ for some state $\rho \in \mathcal{D}_d$. Now we get

$$\begin{aligned} p\mathrm{Ent}_{p,\sigma}(X) &= p\,\langle I_{q,p}(X), S_p(X)\rangle_\sigma - \|X\|_{p,\sigma}^p \ln(\|X\|_{p,\sigma}^p) \\ &\geq \frac{p}{p-1}\|X\|_{p,\sigma}^p \ln(\|X\|_{p,\sigma}^p) - \|X\|_{p,\sigma}^p \ln(\|X\|_{p,\sigma}^p) \\ &= \kappa_p(X) \end{aligned}$$

where we used (40).

□

Proof of Theorem 4.3. There is nothing to show for $p = 1$ as $\alpha_1(\mathcal{L}) = \beta_1(\mathcal{L})$ and we can assume $p > 1$. For $X \in \mathcal{M}_d^+$ with $\|X\|_{1,\sigma} = 1$ we can use Lemma 4.2 and the definition of $\alpha_p(\mathcal{L})$ to compute

$$\frac{\alpha_p(\mathcal{L})}{p}\kappa_p(X) \leq \alpha_p(\mathcal{L})\,\mathrm{Ent}_{p,\sigma}(X) \leq \mathcal{E}_p^{\mathcal{L}}(X).$$

By the variational definition (27) of β_p the claim follows. □

Theorem 4.3 will be applied in Section 5 to obtain bounds on the mixing time of a Liouvillian with a positive logarithmic Sobolev constant without invoking any form of l_p-regularity (see [2]). As usually a logarithmic Sobolev is implied by a hypercontractive inequality [1], we would like to remark that one can also make a similar statement as that of Theorem 4.3 from a hypercontractive inequality. One can easily show that

$$||e^{t\hat{\mathcal{L}}}||_{p(t)\to p,\sigma} \leq 1 \tag{41}$$

for $p(t) = (p-1)e^{-\alpha_p t} + 1$ implies that $\beta_p(\mathcal{L}) \geq \frac{\alpha_p}{p}$.

5 Mixing times

In this section we will introduce the quantities of interest and prove the building blocks to prove mixing times from the entropy production inequalities of the last sections, distinguishing between continuous and discrete time semigroups. We will mostly focus on β_2, as this seems to be the most relevant constant for mixing time applications. This is justified by the fact that the underlying Dirichlet form is a quadratic form and the entropy related to it stems from a Hilbert space norm. Moreover, as the same Dirichlet form is also involved in computations of the spectral gap, it could be easier to adapt existing techniques, such as the ones developed in [23, 19].

Definition 5.1 (Mixing times)**.** *For either $I = \mathbb{R}^+$ or $I = \mathbb{N}$ let $\{T_t\}_{t\in I}$ be a primitive semigroup of quantum channels with fixed point $\sigma \in \mathcal{D}_d^+$. We define the l_1* ***mixing time*** *for $\epsilon > 0$ as*

$$t_1(\epsilon) = \inf\{t \in I \ : \ \|T_t(\rho) - \sigma\|_1 \leq \epsilon \text{ for all } \rho \in \mathcal{D}_d\}.$$

Similarly we define the l_2 ***mixing time*** *for $\epsilon > 0$ as*

$$t_2(\epsilon) = \inf\{t \in I \ : \ Var_\sigma\left(\hat{T}_t\left(X\right)\right) \leq \epsilon \text{ for all } X \in \mathcal{M}_d^+ \text{ with } \|X\|_{1,\sigma} = 1\}.$$

In the continuous case $I = \mathbb{R}^+$ we will often speak of the mixing times of the Liouvillian generator of a quantum dynamical semigroup which we identify with the mixing times of the semigroup according to the above definition.

5.1 Mixing in Continuous Time

It is now straightforward to get mixing times from the previous results.

Theorem 5.1 (Mixing time from entropy production)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$t_1(\epsilon) \leq \frac{1}{2\beta_p(\mathcal{L})} \ln\left(\frac{2\ln\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right).$$

Proof. From (6) and Lemma 2.1 we have

$$\ln\left(\|\sigma^{-1}\|_\infty\right) e^{-2\beta_p(\mathcal{L})t} \geq \frac{1}{2}\|e^{t\mathcal{L}}(\rho) - \sigma\|_1^2. \tag{42}$$

for any $\rho \in \mathcal{D}_d$. The claim follows after rearranging the terms. □

Using Theorem 4.3 we can lower bound β_p in terms of the usual logarithmic Sobolev constant α_p. Combining this with Theorem 5.1 shows the following Corollary.

Corollary 5.1 (Mixing time bound from logarithmic Sobolev inequalities)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$t_1(\epsilon) \leq \frac{p}{2\alpha_p(\mathcal{L})} \ln\left(\frac{2\ln\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right). \tag{43}$$

By Corollary 5.1 a nonzero logarithmic Sobolev constant always implies a nontrivial mixing time bound. One should say that the same bound was showed in [2] for $p = 2$, however under additional assumptions (specifically l_p-regularity [2]) on the Liouvillian in question. While these assumptions have been shown for certain classes of Liouvillians (including important examples like Davies generators and doubly stochastic Liouvillians [2]) they have not been shown in general. Moreover, the bound in Theorem 5.1 clearly does not depend on p and one could in principle optimize over all β_p. However, as the computations in subsection 3.3 already indicate, it does not seem to be feasible to compute or bound β_p for $p \neq 2$ even in simple cases and one will probably only work with β_2 in applications.

The bound from Corollary 5.1 also has the right scaling properties needed in recent applications of rapid mixing, such as the results in [24, 25]. In particular, together with the results in [26], the last Corollary shows that the hypothesis of Theorem 4.2 in [24] is always satisfied for product evolutions and not only for the special classes considered in [2].

One may also use these techniques to get mixing times in the l_2 norms which are stronger than the ones obtained just by considering that β_2 is a lower bound to the spectral gap.

Theorem 5.2 (l_2-mixing time bound)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$t_2(\epsilon) \leq \frac{1}{2\beta_2(\mathcal{L})} \ln\left(\frac{\ln\left(\|\sigma^{-1}\|_\infty\right)}{\ln(1+\epsilon)}\right). \tag{44}$$

Proof. For $X > 0$ with $\|X\|_{1,\sigma} = 1$ we have $\mathrm{Var}_\sigma(X) = \|X - \mathbb{1}\|_{2,\sigma}^2 = \|X\|_{2,\sigma}^2 - 1$ and thus

$$\kappa_2(X) = (1 + \mathrm{Var}_\sigma(X)) \ln\left(1 + \mathrm{Var}_\sigma(X)\right).$$

In the following let $X_t = e^{t\hat{\mathcal{L}}}(X)$ denote the time evolution of the relative density X. Using (19) and the definition of $\beta_2(\mathcal{L})$ (see (27)) we obtain

$$\frac{d}{dt}\mathrm{Var}_\sigma(X_t) = -2\mathcal{E}_2^{\mathcal{L}}(X_t) \leq -2\beta_2(\mathcal{L})(1 + \mathrm{Var}_\sigma(X_t))\ln(1 + \mathrm{Var}_\sigma(X_t)).$$

Integrating this differential inequality we obtain

$$\begin{aligned}\ln\left(\frac{\ln(1+\mathrm{Var}_\sigma(X))}{\ln(1+\epsilon)}\right) &\leq \int_0^{t_2(\epsilon)} \frac{1}{(1+\mathrm{Var}_\sigma(X_t))\ln\left(1+\mathrm{Var}_\sigma(X_t)\right)}\left[\frac{d}{dt}\mathrm{Var}_\sigma(X_t)\right] dt \\ &\leq -2\beta_2 t_2(\epsilon).\end{aligned}$$

As $1 + \mathrm{Var}_\sigma(X) \leq \|\sigma^{-1}\|_\infty$, the claim follows after rearranging the terms. □

In the remaining part of the section we will discuss a converse to the previous mixing time bounds, i.e. a lower bound on the logarithmic Sobolev constant in terms of a mixing time. This excludes the possibility of a reversible semigroup with both small β_2 and short mixing time with respect to the l_2 distance. For this we generalize [21, Corollary 3.11] to the noncommutative setting.

Theorem 5.3 (LS inequality from l_2 mixing time)**.** *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive, reversible Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$\frac{1}{2} \leq \alpha_2(\mathcal{L})\, t_2\left(e^{-1}\right) \leq 2\beta_2(\mathcal{L})\, t_2\left(e^{-1}\right). \tag{45}$$

Moreover, this inequality is tight.

Proof. We refer to Appendix B for a proof. □

As remarked in [21], even the classical result does not hold anymore if we drop the reversibility assumption. Therefore, this assumption is also needed in the noncommutative setting. By considering a completely depolarizing channel it is also easy to see that no such bound can hold in discrete time.

Theorem 5.3 implies that for reversible Liouvillians β_2 and α_2 cannot differ by a large factor. More specifically we have the following corollary.

Corollary 5.2. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive, reversible Liouvillian with fixed point $\sigma \in \mathcal{D}_d^+$. Then*

$$2\beta_2(\mathcal{L}) \geq \alpha_2(\mathcal{L}) \geq \beta_2(\mathcal{L}) \ln\left(\frac{\ln\left(\|\sigma^{-1}\|_\infty\right)}{\ln(1+e^{-1})}\right). \tag{46}$$

Proof. We showed the first inequality in Theorem 4.3. The second inequality follows by combining (45) and (44). □

5.2 Mixing in Discrete Time

In this section we will obtain mixing time bounds and also entropic inequalities for discrete-time quantum channels $T : \mathcal{M}_d \to \mathcal{M}_d$. We will then use these techniques to derive mixing times for random local channels, which we will define next. These include channels that usually appear in quantum error correction scenarios, such as random Pauli errors on qubits [27, Chapter 10]. They will be based on the following quantity:

Definition 5.2. *For a primitive quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$, we define*

$$\beta_D(T) = \beta_2(T^*\hat{T} - id_d). \tag{47}$$

Here we used $\hat{T} = \Gamma_\sigma^{-1} \circ T \circ \Gamma_\sigma$.

The definition of $\beta_D(T)$ can be motivated by the following improved data-processing inequality for the 2-sandwiched Rényi divergence.

Theorem 5.4. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive quantum channel with full rank fixed point $\sigma \in \mathcal{D}_d^+$. Then for all $\rho \in \mathcal{D}_d$ we have*

$$D_2\left(T(\rho)\|\sigma\right) \leq (1 - \beta_D(T)) D_2\left(\rho\|\sigma\right). \tag{48}$$

Proof. Let $X = \sigma^{-\frac{1}{2}}\rho\sigma^{-\frac{1}{2}}$ denote the relative density of ρ with respect to σ. Observe that the 2-Dirichlet form (see Definition 2.3) of the semigroup $\mathcal{L} = T^*\hat{T} - \mathrm{id}_d$ can be written as

$$\mathcal{E}_2^{\mathcal{L}}(X) = \|X\|_{2,\sigma}^2 - \|\hat{T}(X)\|_{2,\sigma}^2.$$

From the definition of $\beta_2(\mathcal{L})$ (see (27)) it follows that

$$\|X\|_{2,\sigma}^2 - \|\hat{T}(X)\|_{2,\sigma}^2 \geq \beta_2\kappa_2(X),$$

which is equivalent to

$$\ln(\|\hat{T}(X)\|_{2,\sigma}^2) - \ln(\|X\|_{2,\sigma}^2) \leq \ln(1 - \beta_2 \ln(\|X\|_{2,\sigma}^2).$$

Using the elementary inequality $\ln(1-\beta_2 \ln(\|X\|_{2,\sigma}^2) \leq -\beta_2 \ln(\|X\|_{2,\sigma}^2)$, that $\ln(\|\hat{T}(X)\|_{2,\sigma}^2) = D_2(T(\rho)\|\sigma)$ and $\ln(\|X\|_{2,\sigma}^2) = D_2(\rho\|\sigma)$ hold, the statement of the theorem follows after rearranging the terms. □

One should note that, unlike in Theorem 3.2, the constant β_D is *not* optimal in (48). As an example take $T(\rho) = \mathrm{tr}[\rho]\frac{\mathbb{1}_d}{d}$ for which $\beta_D(T) = \frac{1-d^{-1}}{\ln(d)}$, but $D_2\left(T(\rho)\|\frac{\mathbb{1}_d}{d}\right) = 0$. Also, $\beta_D(T) > 0$ is not a necessary condition for primitivity, as there are primitive quantum channels that are not strict contractions with respect to D_2. To see this, consider the map $T : \mathcal{M}_2 \to \mathcal{M}_2$ which acts as follows on Pauli operators:

$$T(\mathbb{1}) = \mathbb{1}, \quad T(\sigma_x) = 0, \quad T(\sigma_y) = 0 \text{ and } \quad T(\sigma_z) = \sigma_x.$$

One can check that this is a a primitive quantum channel with $T^2(\rho) = \frac{\mathbb{1}}{2}$ for any state $\rho \in \mathcal{D}_d$. However, T maps the pure state $\frac{1}{2}(\mathbb{1} + \sigma_z)$ to the pure state $\frac{1}{2}(\mathbb{1} + \sigma_x)$, which implies that D_2 does not strictly contract under T. We can now prove the following bound on the discrete mixing time.

Theorem 5.5 (Discrete mixing time)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive quantum channel with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and $\beta_D(T) > 0$. Then*

$$t_1(\epsilon) \leq -\frac{1}{\ln(1 - \beta_D(T))} \ln\left(\frac{2\ln\left(\|\sigma^{-1}\|_\infty\right)}{\epsilon^2}\right).$$

Proof. By Theorem 5.4 we have

$$D_2(T^n(\rho)\|\sigma) \leq (1 - \beta_D(T))^n D_2(\rho\|\sigma).$$

for any $\rho \in \mathcal{D}_d$. The claim then follows from (6) and Lemma 2.1. □

Convergence results for primitive continuous-time semigroups can often be lifted to their tensor powers. In discrete time a similar result holds for the following class of channels:

Definition 5.3 (Random Local Channels)**.** *For a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ and probabilities $\mathbf{p} = (p_1, \ldots, p_n)$ with $p_i \geq 0$ and $\sum_i p_i = 1$ we define a **random local channel** $T_{\mathbf{p}}^{(n)} : \mathcal{M}_d^{\otimes n} \to \mathcal{M}_d^{\otimes n}$ by*

$$T_{\mathbf{p}}^{(n)} = \sum_{i=1}^{n} p_i \; id_d^{\otimes i-1} \otimes T \otimes id_d^{\otimes n-i}. \tag{49}$$

The previous definition can be generalized to the case where not all local channels are identical, i.e. if we have $T_i : \mathcal{M}_d \to \mathcal{M}_d$ acting on the ith system in the expression (49). As long as the local channels are all primitive our results also hold for this more general class of channels. However, for simplicity we will restrict here to the above definition.

Theorem 5.6. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive quantum channel with full rank fixed point $\sigma \in \mathcal{D}_d^+$ such that the Liouvillian $\hat{\mathcal{L}} = T^*\hat{T} - id_d$ fulfills $\beta_2\left(\mathcal{L}^{(n)}\right) \geq q$ for some $q > 0$ and all $n \in \mathbb{N}$. Then for any $n \in \mathbb{N}$ and probabilities $\mathbf{p} = (p_1, \ldots, p_n)$ with $p_i \geq 0$ and $\sum_i p_i = 1$ we have*

$$D_2(T_{\mathbf{p}}^{(n)}(\rho)\|\sigma^{\otimes n}) \leq (1 - qp_{\min}^2) D_2(\rho\|\sigma^{\otimes n}), \tag{50}$$

for any $\rho \in \mathcal{D}_{d^n}$ and where $p_{\min} = \min p_i$.

Proof. By Theorem 5.4 it is enough to show that $\beta_D(T_{\mathbf{p}}) \geq qp_{\min}^2$.

Observe that the Dirichlet form of $(T_{\mathbf{p}}^{(n)})^*\hat{T}_{\mathbf{p}}^{(n)} - \mathrm{id}_d$ is given by

$$\mathcal{E}_2^{\mathcal{L}}(X) = \sum_{i \neq j} p_i p_j \left\langle X - T_i^*\hat{T}_j(X), X\right\rangle_{\sigma^{\otimes n}} + \sum_i p_i^2 \left\langle X - T_i^*\hat{T}_i(X), X\right\rangle_{\sigma^{\otimes n}}$$

where the map $T_i^*\hat{T}_j$ acts as T^* on the i-th system, $\hat{T}$ on the j-th and as the identity elsewhere. As $T_i^*\hat{T}_j \leq \mathrm{id}_d$ with respect to $\langle\cdot,\cdot\rangle_{\sigma^{\otimes n}}$ we have

$$\mathcal{E}_2^{\mathcal{L}}(X) \geq \sum_i p_i^2 \left\langle X - T_i^*\hat{T}_i(X), X\right\rangle_{\sigma^{\otimes n}} \geq p_{\min}^2 \mathcal{E}_2^{\mathcal{L}^{(n)}}(X).$$

From the comparison inequality $\mathcal{E}_2^{\mathcal{L}} \geq p_{\min}^2 \mathcal{E}_2^{\mathcal{L}^{(n)}}$ and the assumption $\beta_2\left(\mathcal{L}^{(n)}\right) \geq q$ it then follows that $\beta_D(\Phi) \geq q p_{\min}^2$. □

As an application we can bound the entropy production and the mixing time in a system of n qubits affected (uniformly) by random Pauli errors. The time evolution of this system is given by the channel $T_n : \mathcal{M}_2^{\otimes n} \to \mathcal{M}_2^{\otimes n}$ given by

$$T_n = \frac{1}{n}\sum_{i=1}^{n} \mathrm{id}_2^{\otimes i-1} \otimes T \otimes \mathrm{id}_2^{\otimes n-i} \tag{51}$$

with $T(\rho) = \mathrm{tr}(\rho)\frac{\mathbb{1}}{2}$.

Theorem 5.7. *For T_n defined as in equation (51) we have*

$$D_2\left(T_n(\rho)\|\frac{\mathbb{1}_{2^n}}{2^n}\right) \leq \left(1 - \frac{1}{2n^2}\right) D_2\left(\rho\|\frac{\mathbb{1}_{2^n}}{2^n}\right).$$

for any $\rho \in \mathcal{D}_{2^n}$.

Proof. From [2] it is known that

$$\alpha_2\left(\mathcal{L}_{\frac{\mathbb{1}}{2}}^{(n)}\right) = 1.$$

Now combining Theorem 4.3 and Theorem 5.6 gives

$$\beta_D(T_n) \geq \frac{1}{2n^2}\alpha_2\left(\mathcal{L}_{\frac{\mathbb{1}}{2}}^{(n)}\right).$$

□

Corollary 5.3. *Let T_n be defined as in (51). Then we have*

$$t_1(\epsilon) \leq -\frac{1}{\ln\left(1 - \frac{1}{2n^2}\right)} \ln\left(\frac{n}{\epsilon^2}\right). \tag{52}$$

Proof. This follows directly from the previous theorem and Theorem 5.5. □

6 Strong converse bounds for the classical capacity

When classical information is sent via a quantum channel, the classical capacity is the supremum of transmission rates such that the probability for a decoding error vanishes in the limit of infinite channel uses. In general it is not possible to retrieve the information perfectly when it is sent over a finite number of uses of the channel, and the probability for successful decoding will be smaller than 1. Here we want to derive bounds on this probability for quantum dynamical semigroups. More specifically we are interested in strong converse bounds on the classical capacity. An upper bound on the capacity is called a strong converse bound if whenever a transmission rate exceeds the bound the probability of successful decoding goes to zero in the limit of infinite channel uses.

We refer to [27, Chapter 12] for the exact definition of the classical capacity and to [28, 29, 30, 4, 31] for more details on strong converses and strong converse bounds.

In [4] the following quantity was used to study strong converses.

Definition 6.1 (p-information radius)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel. The p-**information radius** T is defined as*[1]

$$K_p(T) = \frac{1}{\ln(2)} \min_{\sigma\in\mathcal{D}_d} \max_{\rho\in\mathcal{D}_d} D_p(T(\rho)\|\sigma).$$

We will often refer to a (m, n, p)-coding scheme for classical communication using a quantum channel T. By this we mean a coding-scheme for the transmission of m classical bits via n uses of the channel T for which the probability of successful decoding is p (see again [27, Chapter 12] for an exact definition). The following theorem shown in [4, Section 6] relates the information radius and the probability of successful decoding.

Theorem 6.1 (Bound on the success probability in terms of information radius)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, $n \in \mathbb{N}$ and $R \geq 0$. For any (nR, n, p_{succ})-coding scheme for classical communication via T we have*

$$p_{succ} \leq 2^{-n\left(\frac{p-1}{p}\right)\left(R-\frac{1}{n}K_p(T^{\otimes n})\right)}. \tag{53}$$

We will now apply the methods developed in the last sections to obtain strong converse bounds on the capacity of quantum dynamical semigroups.

Theorem 6.2. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$ such that for some $p \in (1, \infty)$ there exists $c > 0$ fulfilling $\beta_p(\mathcal{L}^{(n)}) \geq c$ for all $n \in \mathbb{N}$. Then for any (nR, n, p_{such})-coding scheme for classical communication via the quantum dynamical semigroup $T_t = e^{t\mathcal{L}}$ we have*

$$p_{succ} \leq 2^{-n\left(\frac{p-1}{p}\right)\left(R-e^{-2ct}\log(\|\sigma^{-1}\|_\infty)\right)}.$$

Proof. Using Theorem 3.2 and Lemma 2.1 we have

$$K_p(T_t^{\otimes n}) \leq \frac{1}{\ln(2)} \max_{\rho\in\mathcal{D}_{d^n}} D_p(T_t^{\otimes n}(\rho)\|\sigma) \leq ne^{-2\beta_p(\mathcal{L}^{(n)})t}\log(\|\sigma^{-1}\|_\infty).$$

Now Theorem 6.1 together with the assumption $\beta_p(\mathcal{L}^{(n)}) \geq c$ finishes the proof. □

Together with Theorem 4.3 the previous theorem shows that a quantum memory can only reliably store classical information for small times when it is subject to noise described by a quantum dynamical semigroup with "large" logarithmic Sobolev constant, as we will see more explicitly later in Section 7. Moreover, we can use the results from [26] to give a universal lower bound to the decay of the capacity in terms of the spectral gap and the fixed point.

[1] The ln(2) factor is due to our different choice of normalization for the divergences.

Corollary 6.1. *Let $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive Liouvillian with full rank fixed point $\sigma \in \mathcal{D}_d^+$ and spectral gap $\lambda(\mathcal{L})$. Then for any (nR, n, p_{such})-coding scheme for classical communication via the quantum dynamical semigroup $T_t = e^{t\mathcal{L}}$ we have*

$$p_{such} \leq 2^{-\frac{n}{2}\left(R - e^{-k(\lambda,\sigma)t}\log(\|\sigma^{-1}\|_\infty)\right)}$$

where $k(\lambda, \sigma) := \frac{\lambda}{2(\ln(d^4\|\sigma^{-1}\|_\infty)+11)}$.

Proof. It was shown in [26, Theorem 9] that $\alpha_2(\mathcal{L}^{(n)}) \geq 2k(\lambda(\mathcal{L}), \sigma)$ for all $n \in \mathbb{N}$. Using Theorem 4.3 we have $\beta_2(\mathcal{L}^{(n)}) \geq k(\lambda(\mathcal{L}), \sigma)$ for all $n \in \mathbb{N}$. Together with Theorem 6.2 this gives the claim. □

For unital semigroups, i.e. for $\sigma = \frac{\mathbb{1}_d}{d}$, one can improve the bound from the previous theorem slightly using (see [32, Theorem 3.3])

$$k\left(\lambda, \frac{\mathbb{1}_d}{d}\right) = \frac{\lambda(1-2d^{-2})}{2(\ln(3)\ln(d^2-1)+2(1-2d^{-2})} \tag{54}$$

For $d = 2$ we even have $k(\lambda, \frac{\mathbb{1}_d}{2}) = \frac{\lambda}{2}$ (see [2]).

7 Examples of bounds for the classical capacity of Semigroups

We will now apply the estimate on the capacity given by Corollary 6.1 to some examples of semigroups. Here $C(T)$ will denote the classical capacity of a quantum channel T.

7.1 Depolarizing Channels

In [5] it is shown that for $\mathcal{L}_{\frac{\mathbb{1}}{d}}(X) = \text{tr}(X)\frac{\mathbb{1}}{d} - X$ we have

$$C\left(e^{t\mathcal{L}_{\frac{\mathbb{1}}{d}}}\right) = \log(d) + \left(e^{-t} + c(t,d)\right)\log\left(e^{-t} + c(t,d)\right) + (d-1)c(t,d)\log\left(c(t,d)\right) \tag{55}$$

with $c(t, d) = (1 - e^{-t})d^{-1}$. In [30] the strong converse property was established. The semigroup generated by $\mathcal{L}_{\frac{\mathbb{1}}{d}}$ is therefore a natural candidate to evaluate the quality of our bounds, as determining its classical capacity can be considered a solved problem. As $\mathcal{L}_{\frac{\mathbb{1}}{d}}$ is just the difference of a projection and the identity, it is easy to see that the spectral gap of $\mathcal{L}_{\frac{\mathbb{1}}{d}}$ is 1, which gives us the upper bound

$$C\left(e^{t\mathcal{L}_{\frac{\mathbb{1}}{d}}}\right) \leq \log(d)e^{-\frac{(1-2d^{-2})}{2(\ln(3)\ln(d^2-1)+2(1-2d^{-2})}t} \tag{56}$$

for $d > 2$ and

$$C\left(e^{t\mathcal{L}_{\frac{\mathbb{1}}{2}}}\right) \leq e^{-\frac{t}{2}} \tag{57}$$

for $d = 2$.

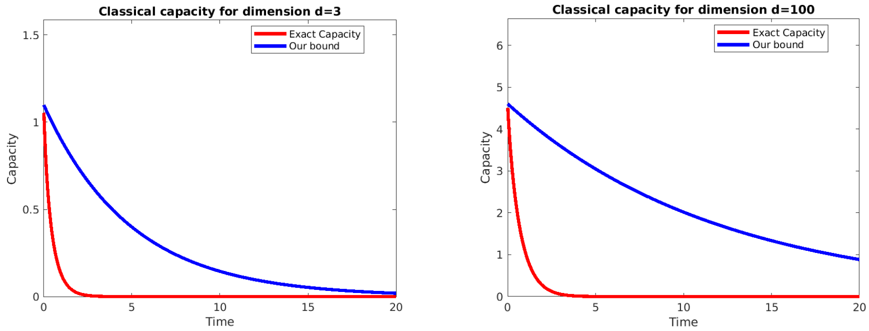

Figure 1: Comparison of the capacity of the depolarizing channel, given in Equation (55), and the bound obtained by our methods, given in Equation (56).

7.2 Stabilizer Hamiltonians

Estimates on the spectral gap of Davies generators of stabilizer Hamiltonians were obtained in [19]. In the following we will make the same assumptions as in [19] on the coupling of the system to the bath. That is, we assume that the operators S^α (see (13)) are given by single qubit Pauli operators σ_x, σ_y or σ_z. For the transition rates $G^\alpha(\omega)$ we only assume that they satisfy the KMS condition [17], that is, $G^\alpha(-\omega) = G^\alpha(\omega)e^{-\beta\omega}$. This condition implies that the semigroup is reversible. Recall that for Davies generators at inverse temperature $\beta > 0$, which we will denote by $\mathcal{L}_\beta$, the stationary state is always given by the thermal state $\frac{e^{-\beta H}}{\mathrm{tr}\left(e^{-\beta H}\right)}$.

We will not discuss stabilizer Hamiltonians and groups and their connection to error-correcting codes, but refer to [27, Section 10.5] for more details. Given some stabilizer group $S \subset \mathcal{P}_n$, where $\mathcal{P}_n$ is the group generated by the tensor product of n Pauli matrices, with commutative generators $S =< P_1, \ldots, P_k >$, we define the stabilizer Hamiltonian to be given by

$$H_S = -\sum_{i=1}^{k} P_i.$$

We then have:

Lemma 7.1. *Let H_S be the stabilizer Hamiltonian of the stabilizer group $S =< P_1, \ldots, P_k >$ on n qubits. Denote by $\sigma_\beta = \frac{e^{-\beta H_S}}{tr\left(e^{-\beta H_S}\right)}$ the corresponding thermal state at inverse inverse temperature $\beta > 0$. Then*

$$\|\sigma_\beta^{-1}\|_\infty \leq 2^n e^{2k\beta}.$$

Proof. The eigenvalues of each P_i are contained in $\{1, -1\}$, as they are just tensor products of Pauli matrices. From this we have

$$-k\mathbb{1} \leq H_S \leq k\mathbb{1}, \tag{58}$$

as H_S is just the sum of k terms such that $-\mathbb{1} \leq P_i \leq \mathbb{1}$. From (58) it follows that

$$\mathrm{tr}\left(e^{-\beta H_S}\right) \leq 2^n e^{\beta k}, \tag{59}$$

as we have 2^n eigenvalues, including multiplicities. Moreover, it also follows that

$$\|e^{\beta H_S}\|_\infty \leq e^{\beta k}. \tag{60}$$

As $\|\sigma_\beta^{-1}\|_\infty = \|e^{\beta H_S}\|_\infty \mathrm{tr}\left(e^{-\beta H_S}\right)$, the claim follows by putting (59) and (60) together. □

In [19, Theorem 15] they show

$$\lambda \geq \frac{h^*}{4\eta^*} e^{-2\beta\bar{\epsilon}}$$

for the spectral gap λ of the Davies generators of stabilizer Hamiltonians at inverse temperature $\beta > 0$. Here $\bar{\epsilon}$ is the generalized energy barrier, h^* is the smallest transition rate and η^* the longest path in Pauli space. We refer to [19] for the exact definition of these parameters. It is important to stress that in general η^* will scale with the number of qubits, so our estimate on the capacity will not be very good as the number of qubits increases.

However, in [19, Theorem 15] they also show the estimate

$$\lambda \geq \frac{h^*}{4} e^{-2\beta\bar{\epsilon}},$$

for the special case in which the generalized energy barrier can be evaluated with canonical paths Γ_1. We again refer to [19] for the exact definition. For these cases the gap does not scale with the dimension and our estimate is much better. Summing up we obtain:

Theorem 7.1. *Let H_S be the stabilizer Hamiltonian of the stabilizer group $S =< P_1, \ldots, P_k >$ on n qubits. Moreover, let $\mathcal{L}_\beta$ be its Davies generator at inverse temperature $\beta > 0$. Then the classical capacity $C(e^{t\mathcal{L}_\beta})$ is bounded by*

$$C(e^{t\mathcal{L}_\beta}) \leq (n + 2\beta k \log(e))\, e^{-r(\beta,n,k)t},$$

with

$$r(\beta, n, k) = e^{-2\beta\bar{\epsilon}} \frac{h^*}{8\eta^*\left(2k\beta + 5n\ln(2) + 11\right)}$$

and

$$r(\beta, n, k) = e^{-2\beta\bar{\epsilon}} \frac{h^*}{8\left(2k\beta + 5n\ln(2) + 11\right)}$$

in case the generalized energy barrier can be evaluated with canonical paths Γ_1. Moreover, this is a bound in the strong converse sense.

Proof. The claim follows immediately after inserting the bounds from Lemma 7.1 and [19, Theorem 15,16] into Corollary 6.1. □

In [19] one can find more explicit bounds for the parameters $\bar{\epsilon}$, η^* and h^* for some stabilizer groups. To the best of our knowledge this is the first bound available for the classical capacity of this class of quantum channels. To make the bound in Theorem 7.1 more concrete, we show what we obtain for the $2D$ toric code.

7.3 2D Toric Code

Here we consider the 2D toric code as originally introduced in [33], which is a stabilizer code. We consider only square lattices: We take an $N \times N$ lattice with N^2 vertical and $(N+1)^2$ horizontal edges; associating a qubit to each edge gives a total of $n = 2N^2+2N+1$ physical qubits. The stabilizer operators are $N(N+1)$ plaquette operators (including the "open" plaquettes along the rough boundary) and $N(N+1)$ vertex operators, all of which are independent. It goes beyond the scope of this article to explain the $2D$ toric code in detail and we refer to [34, Section 19.4] for a discussion. But from the previous observations we obtain that we have $k = 2N(N+1)$ generators for the stabilizer group of the $2D$ toric code on $n = 2N^2+2N+1$ qubits. We will make the same assumptions on the the Davies generators at inverse temperature $\beta > 0$ for the toric code as in [19]. These are discussed in the beginning of Subsection 7.2.

In [35] it was proved that the spectral gap for the Davies generators for the $2D$ toric code at inverse temperature β satisfies $\lambda \geq \frac{1}{3}e^{-8\beta}$, a result which was reproved in [19] using different techniques. We therefore obtain:

Corollary 7.1. *Let H be the stabilizer Hamiltonian of the $2D$ toric code on a $N \times N$ lattice and $\mathcal{L}_\beta$ be its Davies generator at inverse temperature $\beta > 0$. Then the classical capacity $C(e^{t\mathcal{L}_\beta})$ is bounded by*

$$C(e^{t\mathcal{L}_\beta}) \leq \left(2N^2 + 2N + 1 + \log(2)4\beta N(N+1)\right) e^{-r(\beta,L)t}, \tag{61}$$

with

$$r(\beta, N) = \frac{e^{-8\beta}}{6\left((10N^2 + 10N + 5)\ln(2) + 4\beta N(N+1)\right) + 66}.$$

Moreover, this is a bound in the strong converse sense.

Proof. The claim follows immediately from Lemma 7.1 and the spectral gap estimate of [35] for the toric code. □

From Figure 2 it becomes evident that we cannot retain information in the $2D$ toric for long times at small inverse temperatures and that we can get nontrivial estimates even for very high dimensions, as the size of the gap does not scale with the size of the lattice. It is conjectured that if the spectral gap of the Davies generators of a Hamiltonian with local, commuting terms satisfies a lower bound which is independent of the size of the lattice, then the logarithmic Sobolev 2 constant also satisfies such a bound [36]. As the Hamiltonian of the $2D$ toric code is of this form, proving this conjecture would lead to a bound similar to the one in Corollary 7.1, but with a rate $r(\beta, N)$ independent of the size of the lattice. This would of course lead to much better bounds for large lattice sizes.

7.4 Truncated harmonic oscillator

Consider the Hamiltonian of a truncated harmonic oscillator

$$H = \sum_{n=0}^{d} n|n\rangle\langle n| \in \mathcal{M}_{d+1}.$$

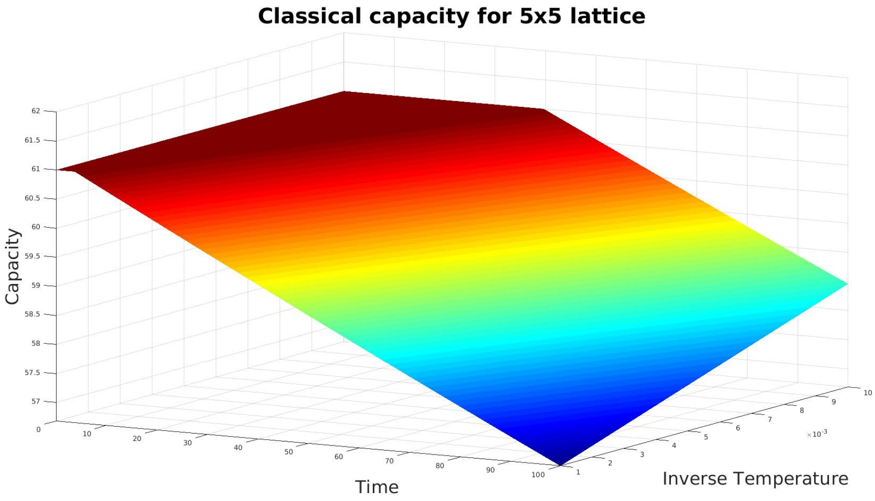

Figure 2: Plot for a 5×5 lattice of the minimum of the bound in Equation (61) and the trivial bound $C(e^{t\mathcal{L}_\beta}) \leq 2N^2 + 2N + 1 = 61$ as a function of the inverse temperature and time for the Davies generator of the $2D$ toric code.

Suppose that the systems couples to the bath via $S = (a + a^\dagger)$, with

$$a^\dagger = \sum_{n=1}^{d} \sqrt{n}\, |n\rangle \langle n-1| \tag{62}$$

and the transition rate function $G(x) = (1+e^{-x\beta})^{-1}$. Let $\sigma_\beta = \frac{e^{-\beta H}}{\operatorname{tr}(e^{-\beta H})}$. As the eigenvalues of $e^{-\beta H}$ are just a geometric sequence, we have

$$\|\sigma_\beta^{-1}\|_\infty = \frac{1 - e^{-\beta(d+1)}}{1 - e^{-\beta}} e^{\beta d}. \tag{63}$$

In [23, Section V, Example 1] they show

$$\lambda \geq \frac{1}{2} \min\{((1 + e^{-\beta})d)^{-1}, \left[(G(1)(\sqrt{d-1} - \sqrt{d})^2 + G(-1)(\sqrt{d-2} - \sqrt{d-1})^2)\right]\}, \tag{64}$$

for the spectral gap λ of the Davies generator $\mathcal{L}_\beta$ of the truncated harmonic oscillator at inverse temperature $\beta > 0$. We will denote the value of the lower bound in Equation (64) by $\mu(d, \beta)$. As we can compute $\|\sigma_\beta^{-1}\|$ exactly and have a bound on the spectral gap from we can apply Corollary 7.1 to these semigroups.

Note that in this case the bound scales with the dimension. Putting these inequalities together with the bound given in Corollary 6.1 for the capacity, we have for the classical capacity of this semigroup:

$$C(e^{t\mathcal{L}}) \leq \left(\log\left(\frac{1 - e^{-\beta(d+1)}}{(1 - e^{-\beta})}\right) + \beta d \log(e)\right) e^{-r(d,\beta)t}, \tag{65}$$

with

$$r(d, \beta) = \left(8 \ln{(d+1)} + 2 \ln\left(\frac{1 - e^{-\beta(d+1)}}{1 - e^{-\beta}}\right) + 2\beta d + 22\right)^{-1} \mu(d, \beta). \tag{66}$$

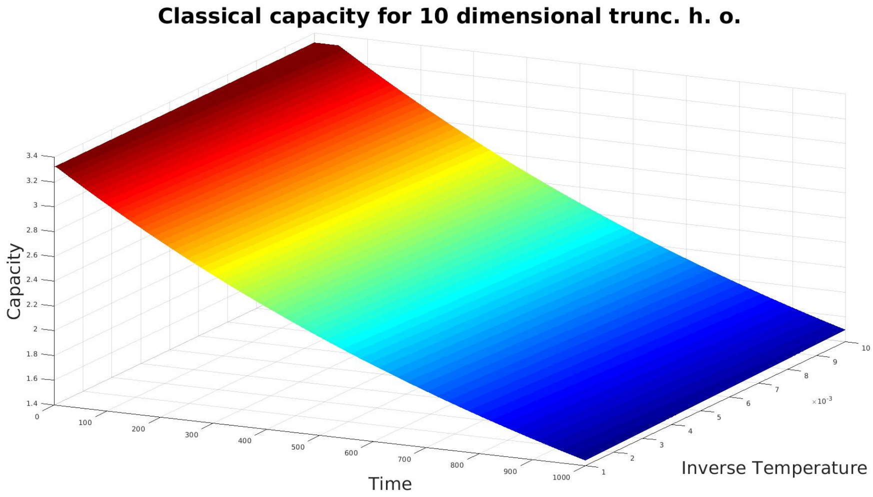

Figure 3: Plot of the minimum of the bound in Equation (65) and the trivial bound $C(e^{t\mathcal{L}_\beta}) \leq \log(10) \simeq 3.32$ as a function of the inverse temperature and time for the classical capacity of the Davies generator of the truncated harmonic oscillator with $d+1=10$.

In this example we see that, as the estimate available on the gap scales with the dimension, our estimates are not much better than the trivial $\log(d{+}1)$ for high dimensions unless we are looking at large times.

8 Conclusion and open questions

We have introduced a framework similar to logarithmic Sobolev inequalities to study the convergence of a primitive quantum dynamical semigroup towards its fixed point in the distance measure of sandwiched Rényi divergences. These techniques can be used to obtain mixing time bounds and strong converse bounds on the classical capacity of a quantum dynamical semigroup. Moreover, these results show that a logarithmic Sobolev inequality or hypercontractive inequality always implies a mixing time bound without the assumption of l_p-regularity (which is still not known to hold for general Liouvillians [2]). Although we have some structural results concerning the constants β_p, some questions remain open. For logarithmic Sobolev inequalities it is known that $\alpha_2 \leq \alpha_p$ for $p \geq 1$ under the assumption of l_p-regularity (see [2]). It would be interesting to investigate if a result of similar flavor also holds for the β_p. In all examples discussed here, β_2 and α_2 are of the same order and it would be interesting to know if this is always the case. The framework of logarithmic Sobolev inequalities has recently been extended to the nonprimitive case [37]. It should be possible to develop a similar theory for the sandwiched Rényi divergences to get rid of regularity assumptions present in their main results, as we did here for the usual logarithmic Sobolev constants.

We restricted our analysis to the sandwiched Rényi divergences, as they can be expressed in terms of relative densities and noncommutative l_p-norms. This allowed us to connect the convergence under the sandwiched divergences to the theory of hypercontractivity and to use tools from interpolation theory which were vital to prove estimates on capacities. There are however other noncommutative generalizations of the Rényi divergences that are known to contract under quantum channels, such as the one discussed in [38, p. 113]. It would be interesting to explore the entropy production and convergence under semigroups for this and other families of divergences in future work.

In a similar vein, it would be interesting to investigate the entropy production or convergence rate for the range $\frac{1}{2} < p < 1$, as the sandwiched Rényi divergences are known to contract under quantum channels for all $p > \frac{1}{2}$ [39]. However, looking closely at the proof of Theorem 3.1, we see that for $p < 1$ the sandwiched Rényi divergence is only differentiable at $t = 0$ if the initial state has full rank. The study of the convergence of these divergences for $p < 1$ therefore requires a different technical approach than that of this work. Finally, it would of course be relevant to obtain bounds on the β_p for more examples without relying on the estimate based on the spectral gap, such as Davies generators.

Acknowledgments

We thank David Reeb and Oleg Szehr for interesting and engaging discussions, which initiated this project. D.S.F acknowledges support from the graduate program TopMath of the Elite Network of Bavaria, the TopMath Graduate Center of TUM Graduate School at Technische Universität München, the Deutscher Akademischer Austauschdienst(DAAD) and by the Technische Universität München Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763.

A.M-H acknowledges financial support from the European Research Council (ERC Grant Agreement no 337603), the Danish Council for Independent Research (Sapere Aude), the Swiss National Science Foundation (project no PP00P2 150734), and the VILLUM FONDEN via the QMATH Centre of Excellence (Grant No. 10059).

This work was supported by the German Research Foundation (DFG) and the Technical University of Munich within the funding programme Open Access Publishing.

This work was initiated at the BIRS workshop 15w5098, Hypercontractivity and Log Sobolev Inequalities in Quantum Information Theory. We thank BIRS and the Banff Centre for their hospitality.

A Taylor Expansion of the Dirichlet Form

In order to compute the Taylor expansions of the Dirichlet forms and of the noncommutative l_p-norms we define $f_p : \mathbb{R}^2 \to \mathbb{R}$ and $g_p : \mathbb{R}^2 \to \mathbb{R}$ for $p > 1$ as

$$f_p(x, y) = \begin{cases} (p-1)x^{p-2} & \text{if } x = y \\ \frac{x^{p-1} - y^{p-1}}{x-y} & \text{else} \end{cases} \tag{67}$$

and

$$g_p(x, y) = \begin{cases} \frac{p(p-1)}{2} x^{p-2} & \text{if } x = y \\ \frac{(p-1)x^p - px^{p-1}y + y^p}{(x-y)^2} & \text{else.} \end{cases} \tag{68}$$

Note that the following identity holds

$$g_p(x, y) + g_p(y, x) = pf_p(x, y) \tag{69}$$

for any $x, y \in \mathbb{R}$.

Lemma A.1 (Taylor expansion)**.** *Consider a primitive, reversible Liouvillian $\mathcal{L} : \mathcal{M}_d \to \mathcal{M}_d$ with full rank fixed point $\sigma \in \mathcal{D}_d^+$. Let $X \in \mathcal{M}_d$ be an eigenvector of $\hat{\mathcal{L}} = \Gamma_\sigma^{-1} \circ \mathcal{L} \circ \Gamma_\sigma$ with corresponding eigenvalue $\lambda \in \mathbb{R}$ (i.e. $\hat{\mathcal{L}}(X) = \lambda X$) and $Y_\epsilon = \mathbb{1}_d + \epsilon X$. Then we have*

$$\mathcal{E}_p^{\mathcal{L}}(Y_\epsilon) = \frac{p}{2(p-1)} \left(2\epsilon^2 \sum_{1 \leq i \leq j \leq d} f_p(s_i, s_j) b_{ij} b_{ji} + O(\epsilon^3) \right). \tag{70}$$

and

$$\kappa_p(Y_\epsilon) = \frac{p\epsilon^2}{p-1} \sum_{1 \leq i \leq j \leq d} f_p(s_i, s_j) b_{ij} b_{ji} + O(\epsilon^3). \tag{71}$$

Where $\sigma^{1/p} = U\, diag\,(s_1, s_2, \ldots, s_d)\, U^\dagger$ and $b_{ij} = (U^\dagger \sigma^{1/2p} X \sigma^{1/2p} U)_{ij}$.

Proof. Using that $X \in \mathcal{M}_d$ is an eigenvector of $\hat{\mathcal{L}}$ a simple computation gives

$$\mathcal{E}_p^{\mathcal{L}}(Y_\epsilon) = \frac{p\epsilon\sigma}{2(p-1)} \text{tr}\left((A + \epsilon B)^{p-1} B \right)$$

for $A = \sigma^{1/p}$ and $B = \sigma^{1/2p} X \sigma^{1/2p}$. Note that

$$\frac{d^k}{d\epsilon^k} \text{tr}\left((A + \epsilon B)^{p-1} B \right) \Big|_{\epsilon=0} = \text{tr}\left(D^k F(A)(B, B, \ldots, B) B \right)$$

for the matrix power $F : \mathcal{M}_d \to \mathcal{M}_d$ given by $F(X) = X^{p-1}$. We apply the Daleckii-Krein formula (see [40] and [41, Theorem 2.3.1.] for the version used here) and obtain

$$\frac{d}{d\epsilon} \text{tr}\left((A + \epsilon B)^{p-1} B \right) \Big|_{\epsilon=0} = \sum_{i=1}^d \sum_{j=1}^d f_p(s_i, s_j) b_{ij} b_{ji}.$$

Using that

$$\text{tr}\left((A + \epsilon B)^{p-1} B \right) \Big|_{\epsilon=0} = \langle \mathbb{1}_d, X \rangle_\sigma = 0$$

by the orthogonality of eigenvectors, and that $f_p(x, y) = f_p(y, x)$ for any $x, y \in \mathbb{R}$ we obtain (70).

To obtain (71) we write

$$\|Y_\epsilon\|_{p,\sigma}^p = \text{tr}\left((A + \epsilon B)^p \right)$$

with $A = \sigma^{1/p}$ and $B = \sigma^{1/2p} X \sigma^{1/2p}$ as above. Again it is easy to see that

$$\frac{d^k}{d\epsilon^k} \text{tr}\left((A + \epsilon B)^p \right) \Big|_{\epsilon=0} = \text{tr}\left(D^k G(A)(B, B, \ldots, B) \right)$$

for the matrix power $G : \mathcal{M}_d \to \mathcal{M}_d$ given by $G(X) = X^p$. Using the Daleckii-Krein formulas we obtain the derivatives

$$\begin{aligned}
\frac{d}{d\epsilon} \text{tr}\left((A + \epsilon B)^p \right) \Big|_{\epsilon=0} &= p \langle \mathbb{1}_d, X \rangle_\sigma = 0 \\
\frac{d^2}{d\epsilon^2} \text{tr}\left((A + \epsilon B)^p \right) \Big|_{\epsilon=0} &= \sum_{i=1}^d \sum_{j=1}^d g_p(\sigma_i, \sigma_j) b_{ij} b_{ji} \\
&= p \sum_{1 \leq i \leq j \leq d} f_p(\sigma_i, \sigma_j) b_{ij} b_{ji}
\end{aligned}$$

where we used the identity (69) in the last step. The above shows that

$$\|Y_\epsilon\|_{p,\sigma}^p = 1 + \epsilon^2 p \sum_{1\leq i\leq j\leq d} f_p(\sigma_i,\sigma_j) b_{ij} b_{ji} + O(\epsilon^3). \tag{72}$$

With the well-known expansion $\ln(1+x) = x - \frac{x^2}{2} + O(x^3)$ we obtain

$$\kappa_p(Y_\epsilon) = \kappa_p(Y_\epsilon) = \frac{p\epsilon^2}{p-1} \sum_{1\leq i\leq j\leq d} f_p(\lambda_i,\lambda_j) b_{ij} b_{ji} + O(\epsilon^3).$$

which is (71).

□

B Interpolation Theorems and Proof of Theorem 5.3

In order to prove Theorem 5.3 we will need the following special case of the Stein-Weiss interpolation theorem [42, Theorem 1.1.1]. This classic result from interpolation spaces has been applied recently to solve problems from quantum information theory, such as in [7, Section III].

Theorem B.1 (Hadamard Three Line Theorem)**.** *Let $S = \{z \in \mathbb{C} : 0 \leq z \leq 1\}$ and $F : S \to \mathcal{B}(\mathcal{M}_d)$ be an operator-valued function holomorphic in the interior of S and uniformly bounded and continuous on the boundary. Let $\sigma \in \mathcal{D}_d^+$ and assume $1 \leq p \leq q \leq \infty$. For $0 < \theta < 1$ define $p_0 \leq p_\theta \leq p_1$ by*

$$\frac{1}{p_\theta} = \frac{1-\theta}{p_0} + \frac{\theta}{p_1}$$

Then for $0 \leq y \leq x \leq 1$ we have

$$\|F(y)\|_{2\to p_\theta,\sigma} \leq \sup_{a,b\in\mathbb{R}} \|F(ia)\|_{2\to p_0,\sigma}^{1-\theta} \|F(x+ib)\|_{2\to p_1,\sigma}^{\theta} \tag{73}$$

One important consequence of the Stein-Weiss interpolation theorem is the following interpolation result. We again refer to [42, Theorem 1.1.1] for a proof.

Theorem B.2 (Riesz-Thorin Interpolation Theorem)**.** *Let $L : \mathcal{M}_d \to \mathcal{M}_d$ be a linear map, $1 \leq p_0 \leq p_1 \leq +\infty$ and $1 \leq q_0 \leq q_1 \leq +\infty$. For $\theta \in [0,1]$ define p_θ to satisfy*

$$\frac{1}{p_\theta} = \frac{\theta}{p_0} + \frac{1-\theta}{p_1}$$

and q_θ analogously. Then for $\sigma \in \mathcal{D}_d^+$ we have:

$$\|L\|_{p_\theta\to q_\theta,\sigma} \leq \|L\|_{p_0\to q_0,\sigma}^{\theta} \|L\|_{p_1\to q_1,\sigma}^{1-\theta}$$

With these tools at hand we can finally prove Theorem 5.3:

Proof. Define $E : \mathcal{M}_d \to \mathcal{M}_d$ by $E(X) = \operatorname{tr}(\sigma X) \mathbb{1}_d$ and set $\tau = t_2(\epsilon)$ for some $\epsilon > 0$. In the following we use $T_z = e^{\tau z \mathcal{L}}$ for $z \in S = \{z \in \mathbb{C} : 0 \leq \mathrm{Re} z \leq 1\}$. We will show that for $s \in [0,1]$:

$$\|(T_s - E)(X)\|_{\frac{2}{1-s},\sigma} \leq \epsilon^s \|X\|_{2,\sigma}. \tag{74}$$

The family of operators $T_z - E$ clearly satisfies the assumptions of the Stein-Weiss interpolation theorem. We therefore have

$$\|T_s - E\|_{2\to\frac{2}{1-s},\sigma} \leq \sup_{a,b\in\mathbb{R}} \|T_{ia} - E\|^{1-s}_{2\to 2,\sigma}\|T_{1+ib} - E\|^{s}_{2\to\infty,\sigma}. \tag{75}$$

Observe that by reversibility of $\mathcal{L}$ the map T_{ia} is a unitary operator with respect to $\langle\cdot,\cdot\rangle_\sigma$. We also have $T_{ia}\circ E = E$, as $T_{ia}(\mathbb{1}_d) = \mathbb{1}_d$. This gives

$$\|\left(T_{ia} - E\right)(X)\|_{2,\sigma} = \|T_{ia}\left(X - E\left(X\right)\right)\|_{2,\sigma} = \|X - E\left(X\right)\|_{2,\sigma} \leq \|X\|_{2,\sigma},$$

where the last equality follows from $\|X - \operatorname{tr}\left(\sigma X\right)\mathbb{1}_d\|_{2,\sigma} = \min_{c\in\mathbb{R}} \|X - c\mathbb{1}_d\|_{2,\sigma}$. We therefore have

$$||T_{ia} - E||^{1-s}_{2\to 2,\sigma} \leq 1. \tag{76}$$

Furthermore, by the unitarity of T_{ib} we can compute

$$\|\left(T_{1+ib} - E\right)(X)\|_{\infty,\sigma} = \|T_{ib}\circ\left(T_1 - E\right)(X)\|_{\infty,\sigma} \leq \|T_1 - E\|_{2\to\infty,\sigma}\|X\|_{2,\sigma}$$

Using duality of the norms and that both T_1 and E are self-adjoint we have

$$\|T_{1+ib} - E\|_{2\to\infty,\sigma} \leq \|T_1 - E\|_{2\to\infty,\sigma} = \|T_1 - E\|_{1\to 2,\sigma} = \epsilon \tag{77}$$

using the definition of τ in the last equality. Inserting (76) and (77) into (75) we get

$$\epsilon^{-s}\|\left(T_s - E\right)(X)\|_{\frac{2}{1-s},\sigma} \leq \|X - E\left(X\right)\|_{2,\sigma}, \tag{78}$$

as $\|\left(T_s - E\right)(X)\|_{\frac{2}{1-s},\sigma} = \|\left(T_s - E\right)\left(X - E\left(X\right)\right)\|_{\frac{2}{1-s},\sigma}$.

Taking the derivative of (78) with respect to s on both sides at $s = 0$ we get

$$\frac{1}{2\|X - E\left(X\right)\|_{2,\sigma}}\left(-2\|X - E\left(X\right)\|^2_{2,\sigma}\ln\left(\epsilon\right) + \operatorname{Ent}_{2,\sigma}\left(|X - E\left(X\right)|\right) - 2\tau\mathcal{E}\left(X\right)\right) \leq 0. \tag{79}$$

Rearranging the terms in (79) we obtain

$$\operatorname{Ent}_{2,\sigma}\left(|X - E\left(X\right)|\right) \leq 2\tau\mathcal{E}\left(X\right) + 2\operatorname{Var}\left(X\right)\ln\left(\epsilon\right). \tag{80}$$

In [1, Theorem 4.2] the following inequality (known as Rothaus' inequality) was shown

$$\operatorname{Ent}_{2,\sigma}\left(X\right) \leq \operatorname{Ent}_{2,\sigma}\left(|X - E\left(X\right)|\right) + 2\operatorname{Var}\left(X\right). \tag{81}$$

Combining inequalities (81) with (80) and setting $\epsilon = \frac{1}{e}$ we get $t_2\left(\frac{1}{e}\right)\alpha_2\left(\mathcal{L}\right) \geq \frac{1}{2}$ by the definition of the LS constant.

To prove that the inequality is tight, consider the depolarizing Liouvillian $\mathcal{L}_\sigma(X) = \operatorname{tr}\left(X\right)\sigma - X$ for some full rank $\sigma \in \mathcal{D}_d^+$. It is easy to see that $\operatorname{Var}_\sigma\left(e^{t\hat{\mathcal{L}_\sigma}}X\right) = e^{-t}\operatorname{Var}_\sigma(X)$ and so $t_2\left(e^{-1}\right) = 1 + \ln\left(\|\sigma^{-1}\|_\infty - 1\right)$. Restricting to operators commuting with σ, it follows from [21, Theorem A.1] that

$$\alpha_2\left(\mathcal{L}_\sigma\right) \leq (1 - 2\|\sigma^{-1}\|^{-1}_\infty)\frac{1}{2\ln\left(\|\sigma^{-1}\|_\infty - 1\right)}.$$

Thus, for a sequence $\sigma_n \in \mathcal{D}_d^+$ converging to a state that is not full rank we have

$$\lim_{n\to\infty} t_2\left(e^{-1}\right)\alpha_2\left(\mathcal{L}_{\sigma_n}\right) = \frac{1}{2}$$

. □

References

[1] R. Olkiewicz and B. Zegarlinski. Hypercontractivity in noncommutative lp spaces. *J. Funct. Anal.*, 161(1):246 – 285, 1999. ISSN 0022-1236. DOI: 10.1006/jfan.1998.3342.

[2] M. J. Kastoryano and K. Temme. Quantum logarithmic Sobolev inequalities and rapid mixing. *J. Math. Phys.*, 54(5):052202, May 2013. DOI: 10.1063/1.4804995.

[3] M. Müller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel. On quantum Rényi entropies: A new generalization and some properties. *J. Math. Phys.*, 54(12): 122203, 2013. DOI: 10.1063/1.4838856.

[4] M. M. Wilde, A. Winter, and D. Yang. Strong Converse for the Classical Capacity of Entanglement-Breaking and Hadamard Channels via a Sandwiched Rényi Relative Entropy. *Commun. Math. Phys.*, 331:593–622, October 2014. DOI: 10.1007/s00220-014-2122-x.

[5] C. King. The capacity of the quantum depolarizing channel. *IEEE Trans. Inf. Theory.*, 49(1):221–229, 2003. DOI: 10.1109/TIT.2002.806153.

[6] H. Umegaki. Conditional expectation in an operator algebra. IV. Entropy and information. *Kodai Math. Sem. Rep.*, 14(2):59–85, 1962. DOI: 10.2996/kmj/1138844604.

[7] S. Beigi. Sandwiched Rényi divergence satisfies data processing inequality. *J. Math. Phys.*, 54(12):122202, December 2013. DOI: 10.1063/1.4838855.

[8] F. Hiai, M. Ohya, and M. Tsukada. Sufficiency, KMS condition and relative entropy in von Neumann algebras. *Pacific J. Math.*, 96(1):99–109. DOI: 10.1142/9789812794208˙0030.

[9] G. L. Gilardoni. On Pinsker's and Vajda's type inequalities for Csiszar's f -divergences. *IEEE Trans. Inf. Theory.*, 56(11):5377–5386, Nov 2010. ISSN 0018-9448. DOI: 10.1109/TIT.2010.2068710.

[10] G. Lindblad. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.*, 48(2):119–130. DOI: 10.1007/BF01608499.

[11] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan. Completely positive dynamical semigroups of N-level systems. *J. Math. Phys.*, 17:821–825, May 1976. DOI: 10.1063/1.522979.

[12] D. Burgarth, G. Chiribella, V. Giovannetti, P. Perinotti, and K. Yuasa. Ergodic and mixing quantum channels in finite dimensions. *New J. Phys.*, 15(7):073045, July 2013. DOI: 10.1088/1367-2630/15/7/073045.

[13] H. Spohn and J. L. Lebowitz. Irreversible thermodynamics for quantum systems weakly coupled to thermal reservoirs. *Adv. Chem. Phys*, 38:109–142. DOI: 10.1002/9780470142578.ch2.

[14] E.B. Davies. *Quantum Theory of Open Systems.* Academic Press, 1976. ISBN 9780122061509.

[15] H.P. Breuer and F. Petruccione. *The Theory of Open Quantum Systems.* OUP Oxford, 2007. ISBN 9780199213900. DOI: 10.1093/acprof:oso/9780199213900.001.0001.

[16] E.B. Davies. Generators of dynamical semigroups. *J. Funct. Anal.*, 34(3):421 – 432, 1979. ISSN 0022-1236. DOI: 10.1016/0022-1236(79)90085-5.

[17] A. Kossakowski, A. Frigerio, V. Gorini, and M. Verri. Quantum detailed balance and KMS condition. *Commun. Math. Phys.*, 57(2):97–110, Jun 1977. ISSN 1432-0916. DOI: 10.1007/BF01625769.

[18] H. Spohn. Entropy production for quantum dynamical semigroups. *J. Math. Phys.*, 19(5):1227–1230, 1978. DOI: 10.1063/1.523789.

[19] K. Temme. Thermalization time bounds for Pauli stabilizer hamiltonians. *Commun. Math. Phys.*, 350(2):603–637, Mar 2017. ISSN 1432-0916. DOI: 10.1007/s00220-016-2746-0.

[20] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete. The χ^2-divergence and mixing times of quantum Markov processes. *J. Math. Phys.*, 51 (12):122201, December 2010. DOI: 10.1063/1.3511335.

[21] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.*, 6(3):695–750, 08 1996. DOI: 10.1214/aoap/1034968224.

[22] A. Müller-Hermes, D. Stilck França, and M. M. Wolf. Relative entropy convergence for depolarizing channels. *J. Math. Phys.*, 57(2), 2016. DOI: 10.1063/1.4939560.

[23] K. Temme. Lower bounds to the spectral gap of Davies generators. *J. Math. Phys.*, 54(12):122110, December 2013. DOI: 10.1063/1.4850896.

[24] T. S. Cubitt, A. Lucia, S. Michalakis, and D. Perez-Garcia. Stability of local quantum dissipative systems. *Commun. Math. Phys.*, 337(3):1275–1315, 2015. ISSN 1432-0916. DOI: 10.1007/s00220-015-2355-3.

[25] F. G. S. L. Brandão, T. S. Cubitt, A. Lucia, S. Michalakis, and D. Perez-Garcia. Area law for fixed points of rapidly mixing dissipative quantum systems. *J. Math. Phys.*, 56(10):102202, October 2015. DOI: 10.1063/1.4932612.

[26] K. Temme, F. Pastawski, and M. J. Kastoryano. Hypercontractivity of quasi-free quantum semigroups. *J. Phys. A.*, 47(40):405303, 2014. DOI: 10.1088/1751-8113/47/40/405303.

[27] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000. ISBN 9780521635035. DOI: 10.1017/cbo9780511976667.

[28] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theory.*, 45(7):2481–2485, Nov 1999. ISSN 0018-9448. DOI: 10.1109/18.796385.

[29] T. Ogawa and H. Nagaoka. Strong converse to the quantum channel coding theorem. *IEEE Trans. Inf. Theory.*, 45(7):2486–2489, Nov 1999. ISSN 0018-9448. DOI: 10.1109/18.796386.

[30] R. König and S. Wehner. A Strong Converse for Classical Channel Coding Using Entangled Inputs. *Phys. Rev. Lett*, 103(7):070504, August 2009. DOI: 10.1103/PhysRevLett.103.070504.

[31] M. Tomamichel, M. M. Wilde, and A. Winter. Strong converse rates for quantum communication. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 2386–2390. IEEE, 2015. DOI: 10.1109/TIT.2016.2615847.

[32] A. Müller-Hermes, D. Stilck França, and M. M. Wolf. Entropy production of doubly stochastic quantum channels. *J. Math. Phys.*, 57(2):022203, 2016. DOI: 10.1063/1.4941136.

[33] S. B. Bravyi and A. Y. Kitaev. Quantum codes on a lattice with boundary. *arXiv preprint quant-ph/9811052*, 1998.

[34] D. A. Lidar and T. A. Brun. *Quantum error correction*. Cambridge University Press, 2013. DOI: 10.1017/CBO9781139034807.

[35] R. Alicki, M. Fannes, and M. Horodecki. On thermalization in Kitaev's 2D model. *J. Phys. A.*, 42(6):065303, 2009. DOI: 10.1088/1751-8113/42/6/065303.

[36] M. J. Kastoryano and F. G. S. L. Brandão. Quantum Gibbs samplers: The commuting case. *Commun. Math. Phys.*, 344(3):915–957, Jun 2016. ISSN 1432-0916. DOI: 10.1007/s00220-016-2641-8.

[37] I. Bardet. Estimating the decoherence time using non-commutative Functional Inequalities. *ArXiv preprint quant-ph/1710.01039*, October 2017.

[38] M. Ohya and D. Petz. *Quantum Entropy and Its Use*. Theoretical and Mathematical Physics. Springer, 2004. ISBN 9783540208068. DOI: 10.1007/978-3-642-57997-4.

[39] R. L. Frank and E. H. Lieb. Monotonicity of a relative Rényi entropy. *J. Math. Phys.*, 54(12):122201, 2013. DOI: 10.1063/1.4838835.

[40] J. L. Daletskii and S. G. Krein. Integration and differentiation of functions of hermitian operators and applications to the theory of perturbations. *AMS Translations (2)*, 47:1–30, 1965. DOI: 10.1090/trans2/047/01.

[41] F. Hiai. Matrix analysis: matrix monotone functions, matrix means, and majorization. *Interdisciplinary Information Sciences*, 16(2):139–246, 2010. DOI: 10.4036/iis.2010.139.

[42] J. Bergh and J. Lofstrom. *Interpolation spaces: an introduction*, volume 223. Springer, 2012. DOI: 10.1007/978-3-642-66451-9.

B.2 Perfect Sampling for quantum Gibbs States

Perfect Sampling for Quantum Gibbs States

D. Stilck França

We develop a quantum algorithm to obtain perfect samples from the distribution of any POVM measured on a quantum Gibbs states. The algorithm relies on the voter Coupling from the Past algorithm developed by Propp and Wilson [11] and requires us to be able to implement a lumpable quantum channel for the Gibbs state and phase estimation for the underlying Hamiltonian. The expected run-time of the algorithm depends on how degenerate the spectrum of the underlying Hamiltonian is, being more efficient the more degenerate the spectrum is. Moreover, we show that the algorithm is stable against noise in the implementation of the channel and faulty phase estimation.

B.2.1 The Algorithm

Consider the following algorithm for generating perfect samples of the Gibbs state at inverse temperature $\beta > 0$ of a Hamiltonian $H \in \mathcal{M}_d$ for a POVM $\{F_i\}_{i\in I}$. We assume H has d' distinct eigenvalues, we may implement a primitive lumpable quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ and a phase estimation routine for H. One example of such a channel is the one implemented in the quantum Metropolis algorithm discussed in [12]. For a lumpable channel, when we fix an eigenbasis of H the dynamics for states which are diagonal in that basis is a classical Markov chain. The equivalence relation that defines the lumped chain is then given by lumping together eigenvectors that have the same energy. The algorithm needs three registers ABC. A is of dimension d, and B and C are of dimension 2^m. Here m is large enough to tell different eigenvalues of H apart using phase estimation.

Theorem B.2.1. *Suppose we run the algorithm depicted in the flowchart in Figure* **??** *with H and T as described above. Moreover, assume that each eigenspace of H is of dimension at least $d/r(d)$, for a real function r. Denote the mixing time of the lumped chain by t_{mix}. Then the algorithm outputs a perfect sample after an expected number of*

$$\mathcal{O}(t_{mix} r(d)^2 \log(r(d)))$$

steps.

The fact that we obtain perfect samples follows from the fact that the classical algorithm by Propp and Wilson, voter coupling from the past [11], outputs perfect samples when the underlying Markov chain is primitive. The only difference here to the algorithm by Propp and Wilson is that we cannot choose the initial state of the chain deterministically. Therefore, we randomize the initial state. The expected number of steps follows from combining a bound for a generalized coupon collector problem, which comes from the fact that we randomize the initial state and the expected value for the classical algorithm by Propp and Wilson. This result shows that the algorithm is efficient for Hamiltonians with a uniformly highly degenerate spectrum.

B.2.2 Stability Results

There are two possible sources of noise for the algorithm. The first one comes from implementing the channel and the second comes from faulty phase estimation. The algorithm is stable against the two sources of error.

Theorem B.2.2. *Let T be as before and $T' : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel such that*

$$\|T - T'\|_{1\to 1} \leq \epsilon$$

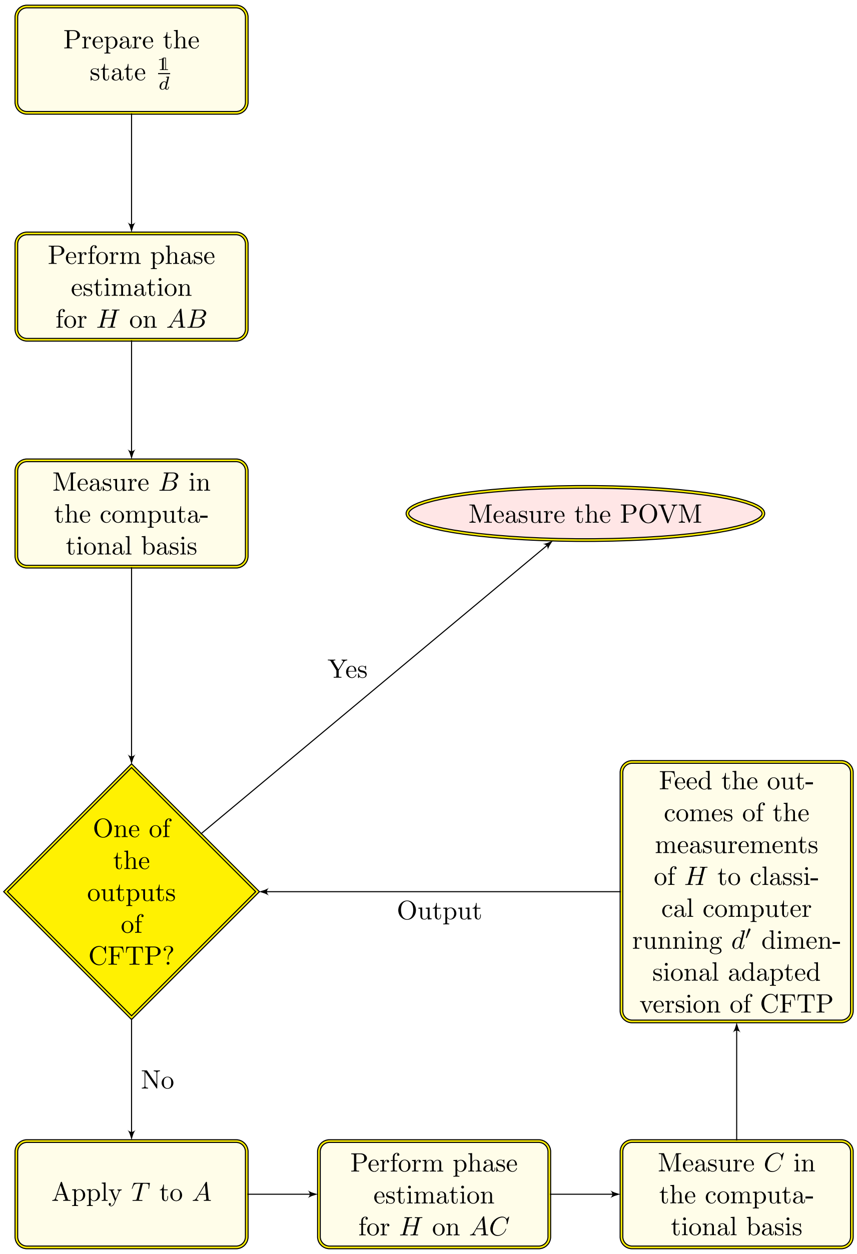
Prepare the state $\frac{\mathbb{1}}{d}$
Perform phase estimation for H on AB
Measure B in the computational basis
Measure the POVM
Yes
One of the outputs of CFTP?
Output
Feed the outcomes of the measurements of H to classical computer running d' dimensional adapted version of CFTP
No
Apply T to A
Perform phase estimation for H on AC
Measure C in the computational basis

Figure B.1: Flowchart of the Perfect Sampling Algorithm.

for some $\epsilon > 0$ and define

$$\kappa(T) = \sup_{X \in \mathcal{M}_d, \mathrm{Tr}(X)=0} \frac{\|(id - T + T_\infty)^{-1}(X)\|_1}{\|X\|_1}$$

with $T_\infty(X) = \mathrm{Tr}\,(X)\,\sigma$. Moreover, denote by p the probability distribution we obtain by measuring $\{F_i\}_{i\in I}$ on $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ and by p' the one we obtain by measuring the POVM on the output of the algorithm using T instead of T'. Then

$$\|p - p'\|_1 \leq (\kappa(T) + 2)\epsilon.$$

The algorithm is also stable against faulty phase estimation. The stability depends on how likely we are to obtain the right outcome when measuring system B and C after the phase estimation step. That is, given that the state of the system A is some eigenstate $|\psi_i\rangle$ of H, what is the probability that we correctly identify it after phase estimation. Denote the minimum probability over all eigenspaces of identifying the right eigenspace by ξ. The stability also depends on the probability of making the reverse error. That is, the probability of the system A being in the eigenstate $|\psi_i\rangle$ given that we observed this outcome. Denote the minimum over all possible outcomes by ξ'. Then we have

Theorem B.2.3. *Let ξ and ξ' be defined as above and T as before. Let p be the distribution we obtain at the output of the algorithm with perfect phase estimation and $p'(i)$ be the probability of observing F_i at the output of the algorithm with faulty phase estimation. Then*

$$\|p - p'\|_1 \leq 1 - \xi' + 2\,(\kappa(T) + 2)\,((1 - \xi\xi' + (1 - \xi)\xi' + \xi(1 - \xi') + (1 - \xi)(1 - \xi')).$$

The values of ξ and ξ' depend on the particular implementation of phase estimation, properties of the spectrum of H and the size of m. We also provide some more explicit estimates for their values for a standard, simple implementation of phase estimation.

B.2.3 Individual Contribution

This project was initiated after Michael M. Wolf inquired if it is possible to obtain perfect samples of quantum Gibbs states on a quantum computer. This question motivated this work and I am solely responsible for all the writing and results.

Permission to include:

D. Stilck França.
Perfect Sampling for Quantum Gibbs States.
To appear in *Quantum Information and Computation*, 18(5), 2018.

Daniel Stilck França <dsfranca13@gmail.com> Fri, Mar 9, 2018 at 1:45 AM
To: QIC Editorial <qic@rintonpress.com>
Dear Wei,
thank you very much!
I am currently preparing a cumulative dissertation. I would like to include the article
Perfect Sampling for Quantum Gibbs states
by Daniel Stilck Franca
and would like to ask for permission to include it in my dissertation.
Kind regards,
Daniel Stilck Franca
[Quoted text hidden]

QIC Editorial <qic@rintonpress.com> Fri, Mar 9, 2018 at 2:37 AM
To: Daniel Stilck França <dsfranca13@gmail.com>
Dear Daniel,
No problem at all. The QIC authors have the rights to use their published work anyway they wish. Good luck.
By the way, this work is most likely to appear in QIC Vol. 18 No.5, though all the papers for this issue are still yet to be
put together sometime later.
Regards, Wei

Quantum Information and Computation, Vol. 18, No. 5&6 (2018) 0361–0388
© Rinton Press

PERFECT SAMPLING FOR QUANTUM GIBBS STATES

DANIEL STILCK FRANÇA
Department of Mathematics, Technische Universität München, Boltzmannstrasse 3
Garching, 85748, Germany

Received April 28, 2017
Revised March 8, 2018

We show how to obtain perfect samples from a quantum Gibbs state on a quantum computer. To do so, we adapt one of the "Coupling from the Past"-algorithms proposed by Propp and Wilson. The algorithm has a probabilistic run-time and produces perfect samples without any previous knowledge of the mixing time of a quantum Markov chain. To implement it, we assume we are able to perform the phase estimation algorithm for the underlying Hamiltonian and implement a quantum Markov chain such that the transition probabilities between eigenstates only depend on their energy. We provide some examples of quantum Markov chains that satisfy these conditions and analyze the expected run-time of the algorithm, which depends strongly on the degeneracy of the underlying Hamiltonian. For Hamiltonians with highly degenerate spectrum, it is efficient, as it is polylogarithmic in the dimension and linear in the mixing time. For non-degenerate spectra, its runtime is essentially the same as its classical counterpart, which is linear in the mixing time and quadratic in the dimension, up to a logarithmic factor in the dimension. We analyze the circuit depth necessary to implement it, which is proportional to the sum of the depth necessary to implement one step of the quantum Markov chain and one phase estimation. This algorithm is stable under noise in the implementation of different steps. We also briefly discuss how to adapt different "Coupling from the Past"-algorithms to the quantum setting.

Keywords: quantum Gibbs states, perfect sampling, quantum algorithms

Communicated by: I Cirac & B Terhal

1 Introduction

Markov chain Monte Carlo methods are ubiquitous in science. They have a similar structure: the solution to a problem is encoded in the stationary distribution of a Markov chain that can be simulated. The chain is then simulated for a "long enough" time until the current state of the chain is "close enough" to a sample of the stationary distribution of interest.

It is expected that with the advent of quantum computers one could use similar methods to develop algorithms to simulate quantum many-body systems that do not suffer from the sign problem [1], and many quantum algorithms with this property were proposed [2, 3, 4, 5, 6, 7, 8, 9].

However, as for classical Monte Carlo methods, it is in general difficult to obtain rigorous bounds on how long is "long enough", as the huge literature dedicated to Markov chain mixing attests [10]. This prompted research on an algorithm that would "decide for itself" when the current state of the Markov chain is close to or is even a perfect sample of the stationary distribution, without any prior knowledge of the mixing properties of the chain.

One of the first algorithms to do so was developed in [11]. Later Propp and Wilson proposed the "Coupling from the Past" (CFTP)-algorithm [12] and showed how it can be applied to efficiently obtain perfect samples for the Ising model. They also showed how to sample perfectly from the stationary distribution of an unknown Markov chain for which we can only observe transitions in a subsequent paper [13] and many perfect sampling algorithms were developed[a] since. There are also proposals of quantum speedups to these algorithms [14].

In this article, we will generalize some of these algorithms to get perfect samples from a Gibbs state of a Hamiltonian on a quantum computer. By this, we mean that we are able to perform any measurement and observe the same statistics for the outcomes as if we were measuring the actual Gibbs state.

To implement them, we will need to be able to perform the phase estimation algorithm [15] for the underlying Hamiltonian. Furthermore, we assume we can implement a quantum Markov chain that drives the system to the desired Gibbs state and fulfills certain assumptions, such as reversibility of the chain, which we elaborate on below. We comment on which of the current proposals to prepare Gibbs states on quantum computers may be adapted for our purposes.

Like it is the case for the classical algorithms, our quantum algorithms do not require any previous knowledge about the mixing properties of the quantum Markov chain. They "decide for themselves" when the current state of the system corresponds to a perfect sample of the target Gibbs state and their run-time is probabilistic.

We will focus on adapting the "voter CFTP" algorithm. For the classical voter CFTP the expected run-time is quadratic in the dimension of the system and linear in the mixing time. The run-time of our version will depend highly on the number of distinct eigenvalues of the Hamiltonian and the dimension of the eigenspaces. In the worst case, which corresponds to Hamiltonians with a non-degenerate spectrum, our version of this algorithm will turn out to have the same expected run-time as its classical counterpart, up to a logarithmic factor in the dimension. However, for Hamiltonians with degenerate spectra, our algorithm can be more efficient than the classical CFTP. In the case of Hamiltonians with an extremely degenerate spectrum, our algorithm can even have a run-time which is proportional to the time necessary to obtain approximate samples. We discuss how to explore this fact to obtain certifiably good samples efficiently for Hamiltonians whose spectrum can be "lumped together" into a small number of intervals. We also briefly discuss how to generalize other variations of CFTP.

These algorithms are stable under noise and we give bounds on their stability with respect to the implementation of the quantum Markov chain and the phase estimation steps. A potential advantage of the "voter CFTP" in comparison to other methods proposed in the literature is that it only requires the implementation of a quantum circuit of low depth a (potentially prohibitive) number of times and significant classical post-processing to obtain

[a] The website `http://dimacs.rutgers.edu/~dbwilson/exact/`, maintained by Wilson, contains a comprehensive list of references concerning the topic and other related material.

a perfect sample without any prior knowledge of the mixing time. Other methods require the implementation of a circuit of (potentially prohibitive) length one time to obtain an approximate sample, but only under the previous knowledge of or assumptions on the mixing time. Therefore, these algorithms are qualitatively different from ours, as our algorithm provides a certificate that we are obtaining good samples.

Another motivation for this work is to explore how coupling techniques for classical Markov chains may be applied and generalized to the quantum setting. These are one of the most useful tools to derive mixing times [10, Section 4.2] and lie at the heart of many perfect sampling algorithms, as their name already suggests. The fact that we still lack a notion of a quantum analog of a coupling is, therefore, one of the main technical hurdles to generalize many results in the theory of classical Markov chains and is by itself an interesting open problem which hopefully this work can shed some light on.

2 Preliminaries

2.1 Notation and basic concepts

We begin by introducing some basic concepts we will need and fixing the notation. Throughout this paper, $\mathcal{M}_d$ will denote the space of $d \times d$ complex matrices and $[d] = \{1, \ldots, d\}$. We will denote by $\mathcal{D}_d$ the set of d-dimensional quantum states, i.e. positive semidefinite matrices $\rho \in \mathcal{M}_d$ with trace 1. We will call a Hermitian operator $H \in \mathcal{M}_d$ a Hamiltonian. We will always denote its spectral decomposition by $H = \sum_{i=1}^{d'} E_i P_i$, with P_i orthogonal projections. Here d' denotes the number of distinct eigenvalues of H. As we will see later, the expected run-time of the algorithm will depend more on this number than the dimension. The eigenspace corresponding to the energy level E_i of H will be denoted by S_i and we will denote its dimension by $|S_i|$. When we write $H = \sum_{i=1}^{d} E_i |\psi_i\rangle\langle\psi_i|$ we mean that $\{|\psi_i\rangle\}_{i=1}^{d}$ is an orthonormal eigenbasis of H. For an inverse temperature $\beta > 0$, we define $\mathcal{Z}_\beta = \text{tr}[e^{-\beta H}]$ to be its partition function and $e^{-\beta H}/\mathcal{Z}_\beta$ its Gibbs state. A quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is a trace preserving completely positive map. We will also refer to such a map as quantum Markov chain. A state $\sigma \in \mathcal{D}_d$ is a stationary state of T if we have $T(\sigma) = \sigma$. The channel is called primitive if we have $\forall \rho \in \mathcal{D}_d : \lim\limits_{n\to\infty} T^n(\rho) = \sigma$ and $\sigma > 0$. There is an equivalent spectral characterization of primitive quantum channels. A quantum channel is primitive if $\sigma > 0$ is the only eigenvector corresponding to eigenvalues of modulus 1 of the channel [16]. In particular, this implies that the property of being primitive is stable under small perturbations of the channel.

A collection of self-adjoint operators $\{F_i\}_{i\in I}$ is called a POVM (positive-operator valued measure) if the $F_i \in \mathcal{M}_d$ are all positive semidefinite and $\sum_{i\in I} F_i = \mathbb{1}$. Here $\mathbb{1} \in \mathcal{M}_d$ is the identity matrix. A state $\rho \in \mathcal{D}_d$ and a POVM $\{F_i\}_{i\in I}$ induce a probability distribution p through $p(i) = \text{tr}[F_i \rho]$. All the algorithms we will discuss have as their goal to produce exact samples of the distribution p generated by an arbitrary POVM in the case that ρ is a Gibbs state.

The following class of quantum channels will be one of the backbones of the algorithms we will present later.

Definition 2.1 (Eigenbasis preserving quantum channels)**.** *A quantum channel* $T : \mathcal{M}_d \to$

$\mathcal{M}_d$ *is called eigenbasis preserving for a Hamiltonian* $H = \sum_{i=1}^{d'} E_i P_i$ *and inverse temperature* $\beta > 0$ *if we have that for all* $i, j \in [d']$ *the commutator*

$$[T(P_i), P_j] = 0$$

vanishes and $T\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right) = \frac{e^{-\beta H}}{\mathcal{Z}_\beta}$.

By the commutator property, we can model the dynamics under T on states that commute with $e^{-\beta H}/\mathcal{Z}_\beta$ as a classical Markov chain. One should not take this condition to imply that the dynamics under T are classical, as will become clear in subsection 3.2, where we present some examples of eigenbasis preserving channels. We will first suppose that we can implement these channels exactly, but will later relax this condition and discuss the influence of noise in section 4.1.

We will need some distinguishability measures for quantum states and channels and convergence speed measures for primitive quantum channels. One of the main ones is through the Schatten 1−Norm $\|X\|_1 = \mathrm{tr}[|X|]$ for $X \in \mathcal{M}_d$. This is justified by the clear operational interpretation given by its variational expression [17, p. 404]. If we denote by $\mathcal{P}_d$ the set of orthogonal projections in $\mathcal{M}_d$, we have for $\rho, \sigma \in \mathcal{D}_d$

$$\frac{\|\rho - \sigma\|_1}{2} = \sup_{P \in \mathcal{P}_d} \mathrm{tr}[P(\rho - \sigma)]. \tag{1}$$

That is, $\|\rho - \sigma\|_1/2$ expresses the maximal probability of correctly distinguishing two states σ, ρ by a projective measurement. This norm induces the 1 → 1 norm on operators $T : \mathcal{M}_d \to \mathcal{M}_d$:

$$\|T\|_{1\to 1} = \sup_{X \in \mathcal{M}_d} \frac{\|T(X)\|_1}{\|X\|_1}. \tag{2}$$

As a measure of the convergence speed of a quantum channel, we define the l_1-mixing time threshold of a primitive quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ with unique stationary state σ, which is given by

$$t_{\mathrm{mix}} = \min\{n \in \mathbb{N} : \sup_{\rho \in D_d} \|T^n(\rho) - \sigma\|_1 \leq 2e^{-1}\}.$$

We will say that a channel $T : \mathcal{M}_d \to \mathcal{M}_d$ satisfies detailed balance or is reversible with respect to $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ if we have that

$$T\left(\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)^{\frac{1}{2}} X \left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)^{\frac{1}{2}}\right) = \left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)^{\frac{1}{2}} T^*(X) \left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)^{\frac{1}{2}}$$

holds for all $X \in \mathcal{M}_d$. Here T^* is the adjoint of the channel with respect to the Hilbert-Schmidt scalar product. Satisfying detailed balance with respect to $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ implies in particular that $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ is a stationary state of the channel.

A crucial ingredient for our sampling algorithm is the phase estimation algorithm, discovered originally in [15]. There are now many variations of it [18, 19, 20] and it is still the subject of active research. We will neither discuss in detail how to implement it nor its complexity and refer to [20] for that. For our purposes, we will just suppose that for a given Hamiltonian H acting on $\mathbb{C}^d$ we may implement a unitary U on $\mathbb{C}^d \otimes (\mathbb{C}^2)^{\otimes t}$ for some $t \in \mathbb{N}$ that acts as follows:

For $|\psi_i\rangle$ an eigenstate of a Hamiltonian H with $H|\psi_i\rangle = E_i|\psi_i\rangle$ we have $U|\psi_i\rangle \otimes |0\rangle = |\psi_i\rangle \otimes |E_i\rangle$, where $|E_i\rangle$ is the binary expansion of E_i in the computational basis of $(\mathbb{C}^2)^{\otimes t}$. We will first assume that we may implement U exactly, but later discuss how imperfections in the implementation of the phase estimation algorithm influence the output of the sampling algorithm in section 4.1.

We will now fix some notation and terminology for classical Markov chains. A sequence $X_0, X_1, X_2, \ldots$ of random variables taking values in a (finite) set S, referred to as the state space, is called a Markov chain if we have

$$P(X_{n+1} = j | X_n = i) = \pi(i,j)$$

for a $|S| \times |S|$ matrix π. π is called the transition matrix of the chain. We will always denote by π the transition matrix of a Markov chain that should be clear from context. Most of the times it will be the one induced by an eigenbasis-preserving channel through

$$\pi(i,j) = \operatorname{tr}\left(T(|\psi_i\rangle\langle\psi_i|)|\psi_j\rangle\langle\psi_j|\right).$$

A probability distribution μ on S is called stationary if we have that $\pi\mu = \mu$. A Markov chain is said to be irreducible if

$$\forall i,j \in S\ \exists n : \pi^n(i,j) > 0.$$

It is aperiodic if

$$\forall i \in S : \gcd\{n \in \mathbb{N} \setminus \{0\} : \pi^n(i,i) > 0\} = 1.$$

Analogously to the quantum case, we say that the transition matrix π satisfies detailed balance with respect to μ if

$$\pi(i,j)\mu(i) = \mu(j)\pi(j,i).$$

Satisfying detailed balance again implies that μ is stationary. It is a well-known fact that if a Markov chain is aperiodic and irreducible there exists a unique stationary distribution μ such that for any other distribution ν on S we have that $\lim\limits_{n\to\infty} \pi^n\nu = \mu$. We define the variational distance between probability distributions ν, μ as

$$\|\nu - \mu\|_1 = \sum_{i \in S} |\mu(i) - \nu(i)|. \tag{3}$$

With a slight abuse of notation, we will also denote the l_1-mixing time threshold in variation distance for a Markov chain by

$$t_{\text{mix}} = \min\{n \in \mathbb{N} : \sup_{\nu} \|\pi^n\nu - \mu\|_1 \leq 2e^{-1}\}. \tag{4}$$

Let

$$C = \min\{T | \forall i \in S \exists 1 \leq k \leq T \text{ such that } X_k = i\}.$$

We will denote by $E_i(C)$ the expected time it takes to observe all states starting from $X_0 = i$ and by $T_C = \max_{i \in S} E_i(C)$ the cover time of the chain. We refer to e.g. [10, Chapter 1] for a review of these basic concepts of Markov chains.

2.2 Lumpable Channels and Chains

Given an equivalence relation or equivalently a partition of the state space $S = \bigsqcup_{i=1}^{d'} S_i$ of a Markov chain $X_0, X_1, X_2, \ldots$, it is possible to define a new stochastic process in which the state space is given by the equivalence classes in the following way. Define the function $f : S \to [d']$ which maps a state to its equivalence class, that is

$$f(x) = i \Leftrightarrow x \in S_i.$$

This new stochastic process is then given by the random variables $f(X_0), (X_1), f(X_2), \ldots$. If this stochastic process is again a Markov chain for all possible initial probability distributions on S, the chain is said to be lumpable with respect to this equivalence relation. We refer to [21, Section 6] for more on lumpable chains. These are sometimes also called projective chains. The next Theorem gives necessary and sufficient conditions for lumpability.

Theorem 2.1 (Lumpable Chain)**.** *A necessary and sufficient condition for a Markov chain to be lumpable with respect to a partition $S = \bigsqcup_{i=1}^{d'} S_i$ is that for every pair S_l, S_k we have for all $l, l' \in S_l$*

$$\sum_{k \in S_k} \pi(l,k) = \sum_{k \in S_k} \pi(l',k) \tag{5}$$

Moreover, the transition probability between S_l and S_k in the lumpable chain is given by Eq. (5).

Proof: We refer to [21, Theorem 6.3.2] for a proof. □

In order to perfectly sample from Gibbs states with degenerate spectra, we will need the concept of a lumpable channel, which we introduce here.

Definition 2.2 (Lumpable channel)**.** *An eigenbasis preserving quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is called lumpable for a Hamiltonian $H = \sum_{i=1}^{d'} E_i P_i$ at inverse temperature $\beta > 0$ if it is reversible and there is a function $f : \mathbb{R} \times \mathbb{R} \to [0,1]$ such that*

$$tr(T(|\psi_i\rangle\langle\psi_i|)|\psi_j\rangle\langle\psi_j|) = f(E_i, E_j)$$

for any unit vectors $|\psi_i\rangle \in P_i(\mathbb{C}^d)$ and $|\psi_j\rangle \in P_j(\mathbb{C}^d)$.

Here $P_j(\mathbb{C}^d)$ denotes the image of the projection. That is, the transition probabilities for eigenstates of H depend only on their respective energies. Notice that as we demand that the quantum channel satisfies detailed balance, the classical transition matrix induced by the channel will also satisfy detailed balance. The definition of lumpable channels is motivated by the following lemma.

Lemma 2.1 (Lumping energy levels)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a lumpable quantum channel for a Hamiltonian $H = \sum_{i=1}^{d} E_i |\psi_i\rangle\langle\psi_i|$ and inverse temperature $\beta > 0$. Here $\{|\psi_i\rangle\}_{i=1}^{d}$ is an orthonormal eigenbasis of H. Define the classical Markov chain on $[d]$ with transition matrix* *given by*

$$\pi(i,j) = tr(T(|\psi_i\rangle\langle\psi_i|)|\psi_j\rangle\langle\psi_j|)$$

and partition the state space according to their energy into $\{S_1, \dots, S_{d'}\}$*, that is, the equivalence relation on the state space is given by*

$$i \sim j \Leftrightarrow tr(H|\psi_i\rangle\langle\psi_i|) = tr(H|\psi_j\rangle\langle\psi_j|). \tag{6}$$

Then the Markov chain is lumpable with respect to this partition and the lumped chain has transition matrix

$$\tilde{\pi}(S_l, S_k) = |S_k| f(E_l, E_k). \tag{7}$$

Moreover, the stationary distribution of the lumped chain is given by

$$\tilde{\mu}(S_i) = |S_i| \frac{e^{-\beta E_i}}{\mathcal{Z}_\beta}, \tag{8}$$

where $|S_i|$ *is the degeneracy of the energy level, and the chain satisfies detailed balance with respect to* $\tilde{\mu}$*.*

Proof: It follows from Theorem 2.1 that it is sufficient and necessary for the chain to be lumpable that for S_l, S_k we have for all $l, l' \in S_l$

$$\sum_{k \in S_k} \pi(l,k) = \sum_{k \in S_k} \pi(l',k). \tag{9}$$

Eq. (9) holds for lumpable quantum channels, as we have $\sum\limits_{k \in S_k} \pi(l,k) = |S_k| f(E_l, E_k)$, which clearly only depends on the equivalence class of l. Therefore, we may define a Markov chain with respect to this partition and from Theorem 2.1 it follows that the transition matrix of the lumpable chain is given by (7). We will now show that it satisfies detailed balance with respect to $\tilde{\mu}$ and therefore $\tilde{\mu}$ is the stationary distribution of the chain. We have

$$\tilde{\mu}(S_i)\tilde{\pi}(S_i, S_j) = |S_i| \frac{e^{-\beta E_i}}{\mathcal{Z}_\beta} |S_j| f(E_i, E_j). \tag{10}$$

Now, as the original chain satisfied detailed balance, it holds that

$$\frac{e^{-\beta E_i}}{\mathcal{Z}_\beta} f(E_i, E_j) = f(E_j, E_i) \frac{e^{-\beta E_j}}{\mathcal{Z}_\beta}. \tag{11}$$

Plugging Eq. (11) into (10) we see that the lumpable chain satisfies detailed balance with respect to $\tilde{\mu}$. This implies that $\tilde{\mu}$ is the stationary distribution of the lumped Markov chain. □

When working with lumpable channels, t_{mix} will always refer to the mixing time of the lumped chain. It is in general not clear how the mixing time of the lumpable chain relates to the mixing time of the original chain and this is a topic of current research. Surprisingly, the mixing time may even increase under lumping, as was shown recently in [22]. However, as is shown e.g. in [10, Lemma 12.8] and remarked in [22], important parameters that describe the convergence of a Markov chain, such as the spectral gap or Cheeger constant [10, Chapter 12], can only increase when lumping chains. This implies that the mixing time cannot increase significantly by lumping. For example, in the counterexample found in [22], lumping increases the mixing time by a factor of $\Theta(\log(d))$.

2.3 (Classical) Voter CFTP

We will briefly describe a perfect sampling algorithms based on CFTP for Markov chains introduced in [13], called "voter CFTP". We mostly stick to their terminology and notation. The goal of this algorithm is to produce perfect samples of the stationary distribution μ of some Markov chain. One of the main advantages of this algorithm is that we only need to be able to observe valid transitions of this Markov chain to obtain perfect samples of the target distribution.

One should note that this is in general not the most efficient algorithm for perfect sampling [13], but arguably the simplest to understand. Besides the pedagogical motivation to present it, it turns out that this version is of interest in the quantum case, as we will see later. For this algorithm we suppose we have access to a randomized procedure **RandomSuccessor** : $S \to S$ such that $P(\textbf{RandomSuccessor}(i) = j) = \pi(i, j)$, where π is a transition matrix having μ as a stationary measure. Let G be a vertex-labeled graph with vertices $-\mathbb{N}_0 \times S$ and labels S. We will define the labels and edges as the algorithm runs and denote by $G(k, i)$ the label of the vertex (k, i). Pseudocode for the algorithm is provided below in algorithm 1 .

One does not need to add the edge in step 7. This only helps to visualize the process. The expected run-time of this algorithm and its complexity, of course, depend on properties of **RandomSuccessor** : $S \to S$. We will discuss these when we analyze the same questions for our algorithm in the quantum case. We now provide a proof that algorithm 1 indeed produces a perfect sample if it terminates almost surely.

Theorem 2.2. *Suppose algorithm 1 terminates with probability* 1 *and denote the output by* Y. *Then* $P(Y = i) = \mu(i)$.

Proof: Let $\epsilon > 0$. As the algorithm terminates with probability 1, there is a N_ϵ such that

$$P(\text{algorithm terminates after at most } N_\epsilon \text{ steps}) \geq 1 - \epsilon.$$

Denote by A_ϵ the event that the algorithm terminates after at most N_ϵ steps. Define a Markov chain $X_{-M}, X_{-M+1}, X_{-M+2}, \ldots, X_0$ for some $M \in \mathbb{N}$ and choose X_{-M} according to μ, i.e. $P(X_{-M} = i) = \mu(i)$. The transitions are defined by the graph, which we suppose was labeled for all (k, i) with $k > -M$. Given $X_k = j$ we set $X_{k+1} = i$, where $\{(k, j), (k+1, i)\}$ is an edge of the graph G. As we chose X_{-M} according to μ and **RandomSuccessor** has μ as a stationary distribution, $P(X_k = i) = \mu(i)$ for all $-M \leq k \leq 0$. We have

$$P(X_0 \neq Y) = P(X_0 \neq Y | A_\epsilon) P(A_\epsilon) + P(X_0 \neq Y | A_\epsilon^C) P(A_\epsilon^C).$$

One can check that the label on the graph at $(-M, i)$ is nothing but the value of X_0. Thus, if we assume that the algorithm has terminated, the value of X_0 does not depend on the initial value and will always be equal to Y. Therefore $P(X_0 \neq Y | A_\epsilon) = 0$ if $-M \leq N_\epsilon$. Also,

```
1: procedure VOTER CFTP
2:     Set $G(0, i) = i$ and $k = 0$.
3:     while $\exists j, i \in S$ s.t. $G(k, i) \neq G(k, j)$ do
4:         for $i \in S$ do
5:             Let $j =$ RandomSuccessor($i$).
6:             Set $G(k - 1, i) = G(k, j)$.
7:             Add the edge $\{(k - 1, i), (k, j)\}$
8:         end for
9:         Set $k \to k - 1$
10:    end while
11: return $G(k, i_0)$ for some $i_0 \in S$
12: end procedure
```

Fig. 1. Voter CFTP [13]

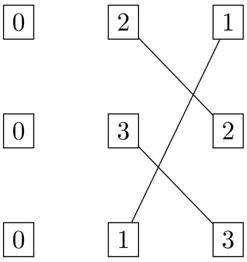

Fig. 2. Possible first two columns of the graph after running the for-loop in the fourth step one time for $d = 3$. Notice that the third column has still not been labeled.

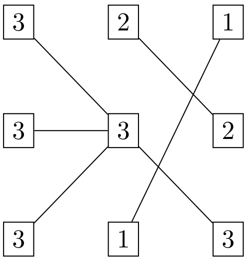

Fig. 3. Possible graph after running the for-loop one more time. Notice the algorithm has terminated and outputs the sample 3.

by construction, $P(X_0 \neq Y|A_\epsilon^C)P(A_\epsilon^C) \leq \epsilon$. We then conclude $P(X_0 \neq Y) \leq \epsilon$ and so the value of Y and X_0 coincide, as ϵ was arbitrary. As X_0 is distributed according to μ, so is Y. □

3 CFTP for quantum Gibbs states

3.1 Voter CFTP

We will now show how to adapt the voter CFTP algorithm to quantum Gibbs state. We start by focusing on Hamiltonians that have a non-degenerate spectrum, as we need less assumptions in this case and the proof is simpler. We later generalize to arbitrary Hamiltonians. Given a Hamiltonian $H \in \mathcal{M}_d$ with non-degenerate spectrum, an eigenbasis preserving quantum channel T for some inverse temperature $\beta > 0$ and a POVM $\mathcal{F} = \{F_i\}_{i\in I}$, the following algorithm allows us to obtain perfect samples from the distribution $p(i) = \operatorname{tr}\left[F_i \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right]$. The algorithm uses three registers corresponding to the tensor factors $\mathbb{C}^d \otimes \left(\mathbb{C}^2\right)^{\otimes t} \otimes \left(\mathbb{C}^2\right)^{\otimes t}$, where t is large enough to perform phase estimation for H and tell apart the different eigenvalues of H. We will discuss how to choose t in section 4.1. The first one will encode the current state of our system, while the other two will be used to record the output of two phase estimation steps. Define a labeled graph G with vertices $V = -\mathbb{N}_0 \times \{1, \ldots, d\}$ and labels given by $\{0, \ldots, d\}$. We assume that G has no edges at the beginning of the algorithm and the vertices are labeled as

$$G(k, j) = \begin{cases} j & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases}. \tag{12}$$

We assume we can prepare the maximally mixed state $\frac{\mathbb{1}}{d}$. This can be done by picking a uniformly distributed integer between 1 and d and preparing the corresponding state of the computational basis, for example. We will assume that the Hamiltonian has a spectral decomposition given by $H = \sum_{i=1}^d E_i |\psi_i\rangle\langle\psi_i|$. The number n denotes how many samples we wish to obtain in total and c will denote a counter for the number of samples we still wish to obtain. The pseudocode for the algorithm is below in algorithm 4.

We now prove it indeed outputs perfect samples.

Theorem 3.1. *Let T be a primitive, eigenbasis preserving quantum channel for a Hamiltonian H and inverse temperature $\beta > 0$. Then algorithm 4 terminates with probability 1 and generates n perfect samples of the distribution p defined above.*

Proof: We will first show that with probability 1 there is a $k \in -\mathbb{N}$ and $l \in [d]$ such that $\forall i \in [d]$ $G(k, i) = l$. The probability that we observe an eigenstate $|\psi_i\rangle$ at step 6 is $\frac{1}{d}$, so with probability 1 we will observe it if we run the loop at step 3 often enough. This implies that we will assign a label different to 0 to arbitrary vertices of the graph G if we run the while-loop at step 3 for long enough. Observe that as T is an eigenbasis preserving quantum channel, the dynamics on the eigenbasis of H under T is just a classical Markov chain. As T is primitive and the stationary state has full rank, this Markov chain is aperiodic and irreducible [16]. Because of that, the probability that we will obtain a k such that $G(k, i) = l$ $\forall i \in [d]$ is 1, using the same argument as the one given in [13] for the classical case. By the same argument, the probability that this label is l is given by $\frac{e^{-\beta E_l}}{\mathcal{Z}_\beta}$, as this is the stationary distribution of

```
procedure QUANTUM VOTER CFTP (NON-DEGENERATE CASE)
    Set $R = \varnothing$ and $c = n$.
    while $c \neq 0$ do
        Prepare the state $\frac{\mathbb{1}}{d} \otimes |0\rangle\langle 0| \otimes |0\rangle\langle 0|$
        Run phase estimation on the first and second register.
        Measure the second register in the computational basis.
        if $i \in R$ then
            Measure $\mathcal{F}$ on the first register.
            Update $c$ to $c - 1$.
        else
            Apply $T \otimes \mathrm{id}_{2^t} \otimes \mathrm{id}_{2^t}$ to the system.
            Run phase estimation on the first register and third register.
            Measure the third register in the computational basis. Let the result be $j$.
            For the largest $k$ s.t. $G(k,i) = 0$ we add the edge $\{(k,i),(k+1,j)\}$.
            if $G(k+1,j) \neq 0$ then
                Change the labels on all the vertices $(k',i')$ with $k' < k$ for which there is a path to $(k,i)$ from 0 to $G(k+1,j)$.
            end if
            if There is $k_0 \in -\mathbb{N}$ and $l \in [d]$ s.t. $\forall i \in [d]$ $G(k_0,i) = l$ then
                Append $l$ to $R$.
                Erase all edges to the vertices $(k_0,i)$ and set the labels to $G(k_0,i) = i$ and $G(k,i) = 0$ for $k < k_0$.
            end if
        end if
    end while
end procedure
```

Fig. 4. Voter CFTP for quantum Gibbs states

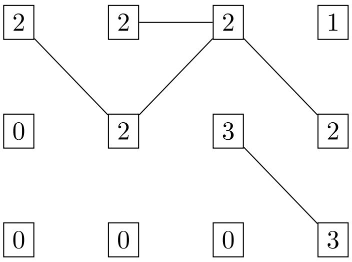

Fig. 5. Possible first four columns of the graph after running the while-loop in step 3 five times for $d = 3$.

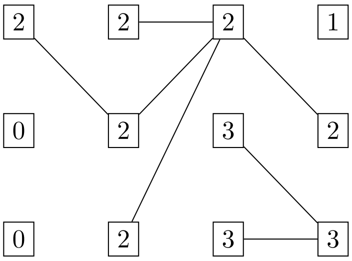

Fig. 6. Possible graph after running the while-loop two more times. Notice the algorithm has terminated and outputs the sample 2.

the underlying classical Markov chain. As before, we will observe $|\psi_l\rangle\,|E_l\rangle\,|0\rangle$ at step 6 with probability 1 if we run the while-loop at step 3 often enough and this will then be a perfect sample by the previous discussion. □

Note that we could check if the measurement outcome we observe at step 13 is one of the desired outcomes to increase our probability of observing it.

Given how many distinct eigenvalues the Hamiltonian has and that we are able to implement a lumpable channel, we may run a modified version of algorithm 4 and obtain perfect samples. The steps of the algorithm are exactly the same and we do not write them out in detail. The only difference is the graph we feed the transitions to and what we feed. Let d' be again the number of distinct eigenvalues of H. In the case of degenerate Hamiltonians, we define a labeled graph G with vertices $V = -\mathbb{N}_0 \times \{1, \ldots, d'\}$ and labels given by $\{0, \ldots, d'\}$. We assume that G has no edges at the beginning of the algorithm and the vertices are labeled as

$$G(k, j) = \begin{cases} j & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases}. \tag{13}$$

That is, the graph is essentially the same as in the non-degenerate case but with d' instead of d labels and vertices. At step 14 we then label the graph according to the energy levels we measured before at steps 6 and 13, as we can only tell apart states with different energies using phase estimation. We then have:

Theorem 3.2. *Let T be a primitive, lumpable quantum channel for a Hamiltonian H and inverse temperature $\beta > 0$. Suppose further that H has d' distinct eigenvalues. Then, if we run algorithm 4 with a graph modified as explained above, it terminates with probability 1 and generates n perfect samples of the distribution $p(i) = tr\left(F_i \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)$.*

Proof: It should be clear that in this case the classical CFTP algorithm we are running based on the measurement outcomes will generate perfect samples from the stationary distribution of the lumped chain defined in Lemma 2.1. The convergence is guaranteed by the same argument as in the proof of Theorem 3.1. From Lemma 2.1 it follows that we will obtain the sample S_j with probability

$$|S_j| \frac{e^{-\beta E_j}}{\mathcal{Z}_\beta}. \tag{14}$$

Now, given that we have observed the label associated to S_j after the first phase estimation step, we know that the state of the first register is given by

$$\rho_j = \frac{1}{|S_j|} P_j. \tag{15}$$

Measuring F_i on the outputs of the algorithm, therefore, gives perfect samples from the distribution p. □

3.2 Examples of eigenbasis preserving and lumpable channels

In order to run algorithm 4, we need to be able to implement a primitive eigenbasis preserving quantum channel for the Gibbs state we want to sample from in the case of non-degenerate spectrum and further that it is lumpable for the general case. In recent years many algorithms have been proposed to approximately prepare quantum Gibbs states on a quantum computer [2, 3, 4, 5, 6, 7, 8, 9]. We will here briefly discuss how some of them provide us with eigenbasis preserving or lumpable quantum channels for Gibbs states.

One class of eigenbasis preserving channels in the non-degenerate case are quantum dynamical semigroups with Davies generators. These are Markovian approximations for a quantum system weakly coupled to a thermal reservoir. A detailed description of the derivation and structure of Davies generators is beyond the scope of this article and can be found in [23, 24]. Under some conditions on the Hamiltonian and the coupling of the system to the bath, the Davies semigroup is primitive. The exact speed of this convergence is the subject of current research. We refer to [25] for a discussion of the conditions under which the Davies generators are primitive and some bounds on the convergence speed. In [9] Davies generators are proposed as a way of preparing thermal states on a quantum computer.

For our purposes, their main relevant property is that if the underlying Hamiltonian has a non-degenerate spectrum, the dynamics in the eigenbasis of the Hamiltonian does not couple diagonal terms to off-diagonal terms. They are therefore eigenbasis preserving. This was observed by many authors since the beginning of their study [24, 26, 27] and we refer to those for a proof of this claim.

More generally, it can be shown that dynamical semigroups that satisfy a quantum version of the detailed balance condition and whose stationary state has a non-degenerate spectrum are always eigenbasis preserving [28]. Our two previous examples fall into that category. This gives us a simple sufficient criterion to check whether a given implementation is eigenbasis preserving.

Note, however, that it is not a priori clear that a quantum dynamical semigroup can be implemented efficiently or by only using local operations. We refer to [9, 29] for a discussion of these topics.

An example of a lumpable channel is given by the implementation of the quantum Metropolis algorithm proposed in [7], as the quantum channel implemented at each step maps eigenstates of H to eigenstates of H and the transition probabilities are a function of their energy difference. However, it can be simplified for our purposes. As in the usual Metropolis algorithm, at each step we have to accept or reject a move that was made. One of the main difficulties to implement the quantum algorithm is reversing the evolution of the system if we reject the move. This is because, by the No-Cloning Theorem [30], we can't make a copy of the previous state of the system. But the information that we rejected the move is enough for our algorithm, as we may simply copy the previous label when labeling the vertices. Therefore, we may skip the procedure of reversing the move.

3.3 Lumping Eigenstates together to obtain good samples

Until now we assumed we are able to implement phase estimation exactly and know the number of distinct eigenvalues of H. We may loosen this assumption and lump different eigenvalues together.

Definition 3.1 (ϵ-Spectral Covering). *Let $H \in \mathcal{M}_d$ be a Hamiltonian and $\epsilon > 0$ be given. We call $\{e_1, \ldots, e_{d'}\} \subset \mathbb{R}$ a ϵ-Spectral Partition for H if*

$$\sigma(H) \subset \bigsqcup_{i=1}^{d'} (e_i - \epsilon, e_i + \epsilon)$$

and for all $i \in [d'] : \sigma(H) \cap (e_i - \epsilon, e_i + \epsilon) \neq \varnothing$. Here $\sigma(H)$ denotes the spectrum of H. We will refer to d' as the size of the covering.

It should be clear that ϵ-spectral coverings are not unique and may have different sizes for fixed ϵ. Although an ϵ-spectral covering will not be readily available in most cases, there are some methods to obtain them. One can use e.g. the Gershgorin circle Theorem [31, Section VIII] to obtain a covering. If we can decompose H into local commuting terms it is also possible to obtain an ϵ-spectral covering by considering that the spectrum of H must consist of sums of the eigenvalues of the local terms. Spectral coverings will be useful later to quantify the stability of algorithm 4 with respect to measuring the wrong energy with phase estimation. Here we will focus on showing how we may still obtain good samples based on an ϵ-spectral covering and that the algorithm is stable w.r.t. introducing degeneracies into the spectrum because we can only obtain an estimate of it to a finite precision. Given an ϵ-spectral covering for the Hamiltonian, we may run algorithm 4 with the number of labels being given by the size of the covering. If we use the Metropolis algorithm from [7] with the e_i as the possible energies to define the transition probabilities we obtain:

Theorem 3.3. *Let $H \in \mathcal{M}_d$ be a Hamiltonian and $\{e_1, \dots, e_{d'}\} \subset \mathbb{R}$ be an ϵ-spectral covering. Suppose we run algorithm 4 with this ϵ-spectral covering as described above. Then the probability distribution $\tilde{p}$ of samples obtained from outputs of algorithm 4 satisfies*

$$\|\tilde{p} - p\|_1 \leq \sqrt{4\epsilon\beta} \tag{16}$$

Proof: The stationary distribution of the lumped chain will be

$$\tilde{\mu}(i) = |\sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)| \frac{e^{-\beta e_i}}{\tilde{\mathcal{Z}}_\beta},$$

with $\tilde{\mathcal{Z}}_\beta = \sum_{i=1}^{d'} |\sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)| e^{-\beta e_i}$. Let $\tilde{P}_i$ be the projection onto the subspace spanned by the eigenvectors of H corresponding to eigenvalues in $\sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)$. From the proof of Theorem 3.2, it follows that algorithm 4 will output the state

$$\tilde{\rho} = \frac{1}{\tilde{\mathcal{Z}}_\beta} \sum_{i=1}^{d'} \frac{e^{-\beta e_i}}{|\sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)|} \tilde{P}_i. \tag{17}$$

We will now show $\left\|\tilde{\rho} - \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right\|_1 \leq \sqrt{4\epsilon\beta}$, from which the claim again follows from the variational definition of the trace norm. From Pinsker's inequality [32, Theorem 3.1], it follows that

$$\left\|\tilde{\rho} - \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right\|_1 \leq \sqrt{2D\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta} || \tilde{\rho}\right)}, \tag{18}$$

where $D\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta} || \tilde{\rho}\right) = \operatorname{tr}\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}(\log(\frac{e^{-\beta H}}{\mathcal{Z}_\beta}) - \log(\tilde{\rho})\right)$ is the relative entropy. As $\tilde{\rho}$ and

$$\frac{e^{-\beta H}}{\mathcal{Z}_\beta} = \frac{1}{\operatorname{tr}\left(e^{-\beta H}\right)} \sum_{E_j \in \sigma(H)} e^{-\beta E_j} P_j$$

commute, we have

$$D\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta} || \tilde{\rho}\right) = \sum_{i=1}^{d'} \sum_{E_j \in \sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)} \frac{e^{-\beta E_j}}{\mathcal{Z}_\beta} \left(\log\left(\frac{\tilde{\mathcal{Z}}_\beta}{\mathcal{Z}_\beta}\right) + \beta\left(E_j - e_i\right)\right). \tag{19}$$

As we have an ϵ-spectral covering, we have $E_j - e_i \leq \epsilon$ for $E_j \in \sigma(H) \cap (e_i - \epsilon, e_i + \epsilon)$ and $\frac{\tilde{\mathcal{Z}}_\beta}{\mathcal{Z}_\beta} \leq e^{\beta\epsilon}$. From this we obtain

$$D\left(\frac{e^{-\beta H}}{\mathcal{Z}_\beta} || \tilde{\rho}\right) \leq 2\beta\epsilon. \tag{20}$$

Plugging Eq. (20) into (18) we obtain the claim. □

This result may be interpreted as a first stability result. This shows that if we lump eigenvalues that are very close together, the Gibbs state does not change a lot. That is, if we introduce artificial degeneracies by not being able to tell apart eigenvalues that are very close through phase estimation this will not change the output of the algorithm significantly. As observed in [7], one could argue that a similar effect could in principle also affect classical Markov chain methods, as we are only able to compute the transition probabilities up to a finite precision. This does not seem to affect them in practice. Moreover, if we want samples that are certifiably at most δ apart in total variation distance at inverse temperature $\beta > 0$, we may lump together eigenvalues that are at most $\delta^2/4\beta$ apart. As we will see later, high levels of degeneracy can reduce the run-time of the algorithm and this can be used to obtain good samples more efficiently.

4 Stability of the Algorithm

We will now address two possible sources of noise for algorithm 4 and show it is stable under these two. First, in the implementation of the channel and second in the phase estimation steps.

4.1 Stability in the implementation of the Channel

As shown in [33], one may quantify the stability of primitive quantum Markov chains with the following constant:

Definition 4.1. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive quantum channel with stationary state $\sigma \in \mathcal{D}_d$. We define*

$$\kappa(T) = \sup_{X \in \mathcal{M}_d, tr(X)=0} \frac{\|(id - T + T_\infty)^{-1}(X)\|_1}{\|X\|_1}$$

with $T_\infty(X) = tr(X)\sigma$.

We refer to [33] for bounds on it and how it can be used to quantify the stability of a quantum Markov chain with respect to different perturbations. Note that due to the spectral characterization of primitive quantum channels [34], the set of primitive quantum channels is relatively open in the convex set of quantum channels.

Theorem 4.1. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a primitive eigenbasis preserving channel for a Hamiltonian H and inverse temperature $\beta > 0$ and $T' : \mathcal{M}_d \to \mathcal{M}_d$ a quantum channel satisfying*

$$\|T - T'\|_{1\to 1} \leq \epsilon \tag{21}$$

for some $\epsilon > 0$ small enough for T' to be primitive too. For a POVM $\{F_i\}_{i\in I}$, let p and p' be probability distributions we obtain by measuring $\{F_i\}_{i\in I}$ on the output of algorithm 4 using T and T' respectively. Then

$$\|p - p'\|_1 \leq (\kappa(T) + 2)\epsilon. \tag{22}$$

Proof: Let $\{P_i\}_{1\leq i\leq d'}$ be the eigenprojections of H and define $Q : \mathcal{M}_d \to \mathcal{M}_d$ to be the quantum channel given by

$$Q(X) = \sum_{i=1}^{d'} \mathrm{tr}(P_i X)\frac{P_i}{|S_i|}. \tag{23}$$

Note that as T is an eigenbasis preserving channel, QTQ is an eigenbasis preserving channel with stationary state $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$. As T' is assumed to be primitive, $QT'Q$ is primitive too, as $\|QT'Q - QTQ\|_{1\to1} \leq \|T' - T\|_{1\to1}$. Denote by ρ the stationary state of the channel $QT'Q$. By the variational expression for the trace distance, we have that

$$\|p - p'\|_1 \leq \left\| \frac{e^{-\beta H}}{\mathcal{Z}_\beta} - \rho \right\|_1 . \tag{24}$$

From theorem 1 in [33] it follows that

$$\left\| \frac{e^{-\beta H}}{\mathcal{Z}_\beta} - \rho \right\|_1 \leq \kappa(QTQ)\|Q(T - T')Q\|_{1\to1}. \tag{25}$$

As Q is a quantum channel, it follows that $\|Q\|_{1\to1} \leq 1$ and so

$$\|Q(T - T')Q\|_{1\to1} \leq \|T - T'\|_{1\to1}. \tag{26}$$

Eq. (22) would then follow from $\kappa(QTQ) \leq 2 + \kappa(T)$. Note that as T is primitive, we have that $\|T - T_\infty\| < 1$, where we use the operator norm. Also, $QT_\infty Q = T_\infty$ and Q is a projection. Thus

$$\|QTQ - T_\infty\| \leq \|T - T_\infty\| < 1.$$

As T is an eigenbasis preserving channel, we have that

$$Q(T - T_\infty)Q = (T - T_\infty)Q$$

and so

$$(Q(T - T_\infty)Q)^n = Q(T - T_\infty)^n Q.$$

We therefore have

$$(\mathrm{id} - (Q(T - T_\infty)Q))^{-1} = \sum_{n=1}^{\infty} \frac{(Q(T - T_\infty)Q)^n}{n!} = \tag{27}$$

$$\mathrm{id} - Q + Q\left(\sum_{n=0}^{\infty} \frac{(T - T_\infty)^n}{n!}\right) Q = \mathrm{id} - Q + Q(\mathrm{id} - (T - T_\infty))^{-1}Q. \tag{28}$$

As $\|Q\|_{1\to1}, \|\mathrm{id}\|_{1\to1} \leq 1$ and from (27) we obtain

$$\kappa(QTQ) = \sup_{X\in\mathcal{M}_d, \mathrm{tr}(X)=0} \frac{\|[Q(\mathrm{id} - (T - T_\infty))^{-1}Q + \mathrm{id} - Q](X)\|_1}{\|X\|_1} \leq \tag{29}$$

$$\sup_{X\in\mathcal{M}_d, \mathrm{tr}(X)=0} \frac{\|Q(\mathrm{id} - (T - T_\infty))^{-1}Q(X)\|_1 + \|(\mathrm{id} - Q)(X)\|_1}{\|X\|_1} \leq 2 + \kappa(T),$$

which completes the proof. □

Theorem 4.1 shows that the algorithm is stable under perturbations of the eigenbasis preserving channel. The stability for lumpable channels follows by observing that every lumpable channel is in particular eigenbasis preserving.

4.2 Faulty Phase Estimation

We will now analyze the errors stemming from faulty phase estimation. It is important to differentiate two different types of error that are caused by the phase estimation procedure. The first type of error comes from the fact that we are only able to obtain an estimate of the energy up to t bits from the phase estimation procedure. This leads to round-off errors and may introduce degeneracies. As discussed in section 3.3 in theorem 3.3, algorithm 4 is stable against this sort of error. Moreover, as we will see later, this can even lead to the algorithm being more efficient.

The second kind of error comes from the fact that the phase estimation procedure only gives the correct energy of the state with high probability. This can cause some transitions we record to be corrupted. We will now show that algorithm 4 is stable against this kind of error.

We note that the exact distribution of the outcomes of the phase estimation procedure depends on which version is being used and this is still a topic of active research [20].

However, it is reasonable to assume that for any phase estimation routine the distribution of the outcomes will concentrate on the correct output. We will give bounds in terms of how large this peak is and explicit bounds for the implementation discussed in [17, Section 5.2]. We now assume we have some rule to assign the labels based on the outcome of the phase estimation step. In case we have an ϵ-spectral covering, this might just be a function which assigns the label based on which interval of the covering the outcome of the measurement belongs to. Let $X_1 \in [d']$ be the random variable which describes which label we assign to the graph after the measurement and $Y_1 \in \{E_1, \dots, E_{d'}\}$ the random variable which describes in which eigenspace the system finds itself after the first measurement at step 6. Let analogously Y_2 be the random variable which is distributed according to the probability of each eigenspace at step 13 and X_2 the second label which we assign. We will now assume that the errors stemming from the phase estimation steps are independent and have the same distribution. That is, given that the system is in a given eigenstate, the probability distribution of the measurement outcomes is the same in the two steps. Let the stochastic matrix $\Xi \in \mathcal{M}_{d'}$ be given by

$$\Xi(i,j) = P(Y_1 = E_j | X_1 = i).$$

Then, given that we have assigned the label i to the graph after step 16 in algorithm 4, the state of the system is described by

$$\rho = \sum_{j=1}^{d'} \Xi(i,j) \frac{P_j}{|S_j|},$$

where P_j is the projection onto the eigenspace corresponding to E_j. After we apply an eigenbasis preserving channel T, the state of the system is described by the state

$$T(\rho) = \sum_{k=1}^{d'} \sum_{j=1}^{d'} \pi(j,k) \Xi(i,j) \frac{P_k}{|S_k|}. \tag{30}$$

Furthermore, denote by $\Xi' \in \mathcal{M}_{d'}$ the stochastic matrix

$$\Xi'(i,j) = P(X_2 = j | Y_2 = E_i).$$

From Eq. (30) it then follows that the probability that the second label is l given that the first label was i is

$$P(X_2 = l|X_1 = i) = \sum_{k,j,l=1}^{d'} \Xi'(k,l)\pi(j,k)\Xi(i,j). \tag{31}$$

From Eq. (31) it is clear that the transition matrix for the labels is given by

$$\pi' = \Xi'\pi\Xi \tag{32}$$

when we have faulty phase estimation.

As mentioned before, we expect $P(X_1 = i|Y_1 = E_j) \simeq \delta_{i,j}$, that is, that the distribution peaks around the right outcome. To quantify this we define

$$\xi = \min_{i\in[d']} \Xi_{i,i} \tag{33}$$

and ξ' analogously. We then have

Lemma 4.1. *Let ξ and ξ' be defined as above and $\{F_i\}_{i\in I}$ a POVM. For a primitive lumpable channel $T : \mathcal{M}_d \to \mathcal{M}_d$ for a Hamiltonian H and inverse temperature $\beta > 0$, let $p(i) = tr\left(F_i \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right)$ and $p'(i)$ be the probability of observing F_i at the output of algorithm 4 with faulty phase estimation. Then*

$$\|p - p'\|_1 \leq 1 - \xi' + 2\left(\kappa(T) + 2\right)\left((1 - \xi\xi' + (1-\xi)\xi' + \xi(1-\xi') + (1-\xi)(1-\xi')\right).$$

Proof: As discussed in Eq. (32), the transition matrix for the observed energy labels is given by $\pi' = \Xi'\pi\Xi$. It easily follows from the definition of ξ and ξ' that

$$\Xi = \xi\mathbb{1} + (1-\xi)\tilde{\Xi},$$
$$\Xi' = \xi'\mathbb{1} + (1-\xi')\tilde{\Xi}',$$

where $\tilde{\Xi}$ and $\tilde{\Xi}'$ are again stochastic matrices. We may therefore write

$$\pi' = \xi\xi'\pi + \xi(1-\xi')\tilde{\Xi}'\pi + (1-\xi)\xi'\pi\tilde{\Xi} + (1-\xi)(1-\xi')\tilde{\Xi}'\pi\tilde{\Xi}.$$

This transition matrix will still be primitive for ξ and ξ' sufficiently large. Let μ be the stationary distribution of π and μ' the one of π'. Observe that, as $\kappa(\pi) \leq \kappa(QTQ)$, we may use the bound $\kappa(QTQ) \leq 2 + \kappa(T)$ from the proof of theorem 4.1 and obtain

$$\|\mu - \mu'\|_1 \leq$$
$$(\kappa(T)+2)\,\|\xi\xi'\pi + \xi(1-\xi')\tilde{\Xi}'\pi + (1-\xi)\xi'\pi\tilde{\Xi} + (1-\xi)(1-\xi')\tilde{\Xi}'\pi\tilde{\Xi} - \pi\|_{1\to1}$$
$$\leq 2\left(\kappa(T)+2\right)\left((1-\xi\xi' + (1-\xi)\xi' + \xi(1-\xi') + (1-\xi)(1-\xi')\right).$$

Here we have used that $\|\pi\|_{1\to 1} \leq 1$ for a stochastic matrix π. At the output of the algorithm, we would be measuring the POVM on the state $\rho' = \sum_{i=1}^{d'} \mu'(i)\frac{P_i}{|S_i|}$ if no error occurs at step 6 of algorithm 4. But as an error might occur when we try to identify a given eigenstate, we will be measuring the state $\rho_{EM} = \sum_{i=1}^{d'}(\Xi'\mu')(i)\frac{P_i}{|S_i|}$. By the definition of ξ', we have $\Xi'\mu' = \xi'\mu' + (1-\mu')\tilde{\Xi}'\mu'$. We will measure the POVM on the state

$$\rho_{EM} = \xi'\rho' + (1-\xi')\rho''. \tag{34}$$

Here ρ'' is some density matrix. It then follows that

$$\left\|\rho_{EM} - \frac{e^{-\beta H}}{\mathcal{Z}_\beta}\right\|_1 \leq 1-\xi' + 2\left(\kappa(T)+2\right)\left((1-\xi\xi' + (1-\xi)\xi' + \xi(1-\xi') + (1-\xi)(1-\xi')\right).$$

The claim then follows from the variational expression for the trace distance as in the proof of theorem 4.1. □

Using Bayes' rule it is possible to express the entries of the matrix Ξ in terms of those of Ξ', which are more readily accessible. We have

$$\Xi(i,j) = P(Y_1 = E_j)\frac{P(X_1 = i|Y_1 = E_j)}{\sum\limits_{l=1}^{d'} P(X_1 = i|Y_1 = E_l)P(Y_1 = E_l)}. \tag{35}$$

As the initial state is the maximally mixed one, we have that $P(Y_1 = E_j) = |S_j|d^{-1}$. From this discussion it follows that:

Theorem 4.2. *Let ξ be defined as in Eq.* (33)*. Then*

$$\xi = \min_{j\in[d']} \frac{P(X_1 = j|Y_1 = E_j)}{|S_j|^{-1}\sum\limits_{l=1}^{d'} P(X_1 = j|Y_1 = E_l)|S_l|}.$$

Proof: See the discussion above. □

This shows that the algorithm is stable if we do not have eigenvalues that we can misidentify with considerable probability and s.t. the degeneracy levels are of different order.

We now give estimates of ξ and ξ' for the implementation of phase estimation considered in [17, Section 5.2] in case we have an ϵ-spectral covering of the Hamiltonian or know that different eigenvalues are ϵ apart. In [17] it is shown that if we use

$$t \geq n + \log\left(2 + (2\delta)^{-1}\right)$$

qubits to perform phase estimation, then we obtain E_i accurate to n bits with probability at least $1-\delta$. This implies that $\xi' \geq 1-\delta$. To estimate ξ we need to control the terms of the form $P(X_1 = i|Y_1 = E_j)$ for $i \neq j$. To this end we define

$$\Delta(i,j) = \inf\{|2^t x - 2^t y \mod 2^t||(x,y) \in A(i,j)\} \tag{36}$$

with

$$A(i,j) = \{x \in \sigma(H) \cap (E_i - \epsilon, E_i + \epsilon)\} \times \{y \in (E_j - \epsilon, E_j + \epsilon)\}.$$

Lemma 4.2. *Let $H \in \mathcal{M}_d$ be a Hamiltonian and $\{e_1, \dots, e_{d'}\}$ be an ϵ-spectral covering of it. Suppose we implement phase estimation for H using t qubits. Then*

$$P(X_1 = j | Y_1 = e_i) \leq \frac{2^{t+1}\epsilon + 1}{\Delta(i,j)^2}$$

for $j \neq i$ and $\Delta_{i,j}$ defined as in Eq. (36).

Proof: In [17, Section 5.2] it is shown that given that the eigenstate of the system is E_i, we have that the probability that the observed outcome is E is bounded by

$$|2^t(E_i - E) \mod 2^t|^{-2}.$$

For any point of the spectrum of H in $(e_i - \epsilon, e_i + \epsilon)$ and for a point in $E \in (e_j - \epsilon, e_j + \epsilon)$ we have that $|2^t(e_i - E) \mod 2^t|^{-2} \geq \Delta(i,j)$. There are at most $2\epsilon 2^t + 1$ possible outcomes that lie in the interval $(e_j - \epsilon, e_j + \epsilon)$. We therefore have

$$P(X_1 = j | Y_1 = e_i) \leq \frac{2^{t+1}\epsilon + 1}{\Delta(i,j)^2}. \quad \square$$

Note that we have $2^t \leq \Delta(i,j)$, so the probability of misidentifying the labels goes to zero exponentially fast with the number of qubits for fixed ϵ. We then obtain for ξ and ξ':

Corollary 4.1. *Let $H \in \mathcal{M}_d$ be a Hamiltonian and $\{e_1, \dots, e_{d'}\}$ be an ϵ-spectral covering of it with $\epsilon \geq 2^{-n}$. Suppose we implement phase estimation for H using $t \geq n+1+\log\left(2 + (2\delta)^{-1}\right)$ qubits. Then $\xi' \geq 1 - \delta$ and*

$$\xi \geq \min_{j \in [d']} \frac{1-\delta}{1 - \delta + (2^{t+1}\epsilon + 1)|S_j|^{-1} \sum\limits_{l \neq j} \Delta(j,l)^{-2}|S_l|}. \tag{37}$$

Proof: As $\epsilon \geq 2^{-n}$ and with probability at least $1-\delta$ we will obtain an output which is accurate up to $n+1$ bits, with probability at least $1-\delta$ we will correctly identify in which element of the covering we are from the output. From this, it follows that $\xi' \geq 1-\delta$. As the function $(x,y) \mapsto \frac{x}{x+y}$ is monotone increasing in x and decreasing in y for $x, y > 0$, we obtain Eq. (37) by inserting the bound on ξ' and the result of Lemma 4.2 into the expression we derived for ξ in theorem 4.2. □

Corollary 4.1 clarifies that the algorithm requires a larger number of qubits to be reasonably stable if we have close E_j and $E_{j'}$ s.t. $|S_j| \ll |S_{j'}|$. The converse is also true; if $|S_j| \backsimeq |S_{j'}|$ for all j, j' the algorithm is already stable with a small precision.

5 Expected run-time, Memory Requirements and Circuit Depth

We will now address the expected run-time of algorithm 4. To this end, we will only consider the number of calls of the phase estimation and eigenbasis preserving or lumpable channel and not the necessary classical post-processing, as we consider the quantum routines the more expensive resources.

In [13, Theorem 5], it was shown that the expected time to obtain a sample using algorithm 1 is $\mathcal{O}(t_{\text{mix}}|S|^2)$ steps, where again t_{mix} is the time such that the chain is e^{-1} close to stationarity and $|S|$ the size of our state space. In the case of Hamiltonians with degenerate spectrum t_{mix} will denote the mixing time of the classical lumped chain induced by the lumpable channel (see Eq. (7)).

Recall that d' denotes the number of distinct eigenvalues of the Hamiltonian or the size of the ϵ-spectral covering being used. That is, d' is just the number of different labels of the graph. We will say that a column indexed by $k \in -\mathbb{N}$ of G is complete if $\forall i \in [d']$ $G(k,i) \neq 0$. In [13], it is shown that we need to complete on average $\mathcal{O}(t_{\text{mix}}d')$ columns of the graph G before the labels on a column become constant. As each step to complete a column needs $\mathcal{O}(d')$ calls of **RandomSuccessor**, this leads to a total of $\mathcal{O}(t_{\text{mix}}{d'}^2)$ calls of **RandomSuccessor**. The dynamics in the eigenbasis of H is classical, so we may use the exact same reasoning to conclude that we will need an expected $\mathcal{O}(t_{\text{mix}}d')$ number of complete columns until we obtain one perfect sample.

But in our case, we may need more uses of the channel and phase estimation, as we may not prepare an arbitrary eigenstate of $H \in \mathcal{M}_d$ which might be necessary to complete a column deterministically. We will denote the expected number of measurements necessary to complete a column by $\phi(H)$ and in theorem A.1 in the Appendix A we give an explicit expression for this quantity.

In Appendix A we prove bounds on $\phi(H)$ for various cases of interest and remark that in the worst case, namely Hamiltonians with a non-degenerate spectrum, $\phi(H) = \mathcal{O}\left(d\log(d)\right)$. Preparing the initial states probabilistically does not significantly change the overall efficiency of the algorithm, as illustrated by the next theorem.

Theorem 5.1. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a lumpable quantum channel for a Hamiltonian H at inverse temperature $\beta > 0$ with mixing time t_{mix}. Then the expected number of steps until algorithm 4 returns a perfect sample is $\mathcal{O}(t_{mix}d'\phi(H))$.*

Proof: We will need an average of $\phi(H)$ measurements to complete a column. From the result [13, Theorem 5] we know that we will need an expected number of $\mathcal{O}(t_{\text{mix}}d')$ number of complete columns to obtain a sample. As the number of measurements needed to complete a column and complete columns to obtain a sample are independent, we have an expected $\mathcal{O}(t_{\text{mix}}d'\phi(H))$ number of steps to obtain a sample. □

It should be clear from theorem 5.1 that algorithm 4 is considerably less efficient than other algorithms such as quantum Metropolis [7] if we are willing to settle for an approximate sample for Hamiltonians with a non-degenerate spectrum. In this case we have $d' = d$ and $\phi(H) = \mathcal{O}\left(d\log(d)\right)$, giving a total complexity of $\mathcal{O}(t_{\text{mix}}d^2\log(d))$ in the worst case. After all, to obtain a sample that is e^{-1} close in trace distance to the Gibbs state, one only needs t_{mix} steps of the Metropolis algorithm instead of the $\mathcal{O}(t_{\text{mix}}d^2\log(d))$ needed for CFTP. Therefore, it is important to stress again that these algorithms are very different in nature. Algorithm 4 provides us with perfect, not approximate samples, and it is the first algorithm of this form for quantum Gibbs states to the best of our knowledge. It provides a certificate that we are indeed sampling from the right distribution when it terminates, while most other algorithms require some mixing time bounds to obtain a sample that can be considered close to the target distribution. Moreover, it only requires us to be able to implement one step of the chain.

However, for the case of Hamiltonians with a highly degenerate spectrum, our algorithm is efficient, as is illustrated by the next theorem:

Theorem 5.2. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a lumpable quantum channel for a Hamiltonian $H \in \mathcal{M}_d$ at inverse temperature $\beta > 0$ with mixing time t_{mix}. Moreover, assume that we have $d \leq |S_i|r(d)$ for some function $r : \mathbb{R} \to \mathbb{R}$ and all eigenspaces S_i. Then the expected number of steps until algorithm 4 returns a perfect sample is $\mathcal{O}(t_{mix}r(d)^2\log(r(d)))$.*

Proof: From $d \leq |S_i|r(d)$ it follows that $d' \leq r(d)$. It follows from theorem A.2 that we will need an average of $\mathcal{O}\left(r(d)\log(r(d))\right)$ measurements to complete a column. From the result [13, Theorem 5] we know that we will need an expected number of $\mathcal{O}(t_{\text{mix}}d')$ number of complete columns to obtain a sample. As the number of measurements needed to complete a column and complete columns to obtain a sample are independent, we have an expected $\mathcal{O}(t_{\text{mix}}r(d)^2\log(r(d)))$ number of steps to obtain a sample. □

In particular, for the cases $r(d) = c$ for some $c \in \mathbb{R}$, which corresponds to having eigenspaces with a degeneracy proportional to the dimension, we have that we only need $\mathcal{O}\left(t_{\text{mix}}\right)$ steps to obtain a perfect sample. That is, the time necessary to obtain perfect samples with our algorithm and approximate ones are the same up to a constant factor. Slightly more generally, for $r(d) = c\log(d)^m$ our algorithm still has a polylogarithmic runtime and is efficient. Admittedly such level of degeneracy is not usual for Hamiltonians of physical relevance. One could use the strategy discussed in section 3.3 and still obtain certifiably good samples by lumping together eigenvalues that are close. Moreover, as we will only need to run a $r(d)$ dimensional version of classical CFTP, the classical part of the algorithm will be efficient.

Although the worst case $\mathcal{O}(t_{\text{mix}}d^2\log(d))$ scaling is prohibitive for applications, this is still more efficient than explicitly diagonalizing H as long as $t_{\text{mix}}\log(d) = \mathcal{O}(d^{\omega-2})$. Here $2 < \omega < 2.373$ is the optimal exponent of matrix multiplication, which has the same complexity as diagonalization [35]. That is, as long as approximate sampling is efficient, obtaining perfect samples is faster than diagonalizing even in the worst case.

We now analyze the circuit depth and memory requirements to obtain a sample.

Theorem 5.3. *Let C_{PT} and C_T be the circuit depth needed to implement the phase estimation for H and the eigenbasis preserving channel, respectively. Then one needs to implement a quantum circuit of depth $\mathcal{O}(C_{PT}+C_T)$ to obtain a sample and moreover an expected $\mathcal{O}(\phi(H))$ classical memory.*

Proof: The circuit length part follows easily from just going through the steps of algorithm 4, as to label the new vertex we need to implement two phase estimation steps and apply the eigenbasis preserving channel once.

To see the that we only need $\mathcal{O}(\phi(H))$ classical memory, notice that we only need to store the information contained in the last complete column to perform the later steps. This is because it contains all possible labels for future columns. By corollary A.2, we have that the expected number of labels we obtain before completing a column is $\mathcal{O}(\phi(H))$, and so we need a total classical memory of size $\mathcal{O}(\phi(H))$. □

The quantum part of algorithm 4 can be easily parallelized, as we could use different quantum computers feeding a classical computer with valid transitions. Note that the classical resources necessary to run the algorithm are also not very large in the cases in which we have a highly degenerate spectrum, as discussed before.

6 Adapting other Variations of CFTP

In [13] the authors discuss other variations of CFTP that can be more efficient, such as the cover time CFTP algorithm. We will not discuss in detail how to adapt these other proposals, but it should be straightforward to do so from the results in the last sections. In this section, we will just mention the main ideas. Note that the only thing necessary to implement all these variations is a valid **RandomSuccessor** function and the outputs of the measurements in steps 6 and 13 of algorithm 4 do exactly that. This information could then be fed to a classical computer running a variation of CFTP. The only difference to the classical case is that we may not choose arbitrary initial states, but do so probabilistically. However, by waiting until each initial state is observed, we may circumvent this and do not have a significant overhead by the result of corollary A.2.

For some variations of CFTP, like again the cover time CFTP, one needs to iterate **RandomSuccessor**. This is also straightforward. If we want to obtain a given number of iterations of **RandomSuccessor**, we just apply an eigenbasis preserving or lumpable channel T to the first register, repeated by a phase estimation step and a measurement in the computational basis. We then repeat this procedure to obtain the iterations.

One could then repeat the analysis done in this section and see that the run-time is again of the same order of magnitude as the classical version of the CFTP algorithm and obtain a perfect sampling algorithm with a run-time proportional to the cover time of the lumped chain.

7 Conclusion and Open Problems

We have shown how to adapt perfect sampling algorithms for classical Markov chains to obtain perfect samples of quantum Gibbs states on a quantum computer. These algorithms have an average run-time which in the worst case is similar to their classical counterparts. For highly degenerate Hamiltonians this algorithm gives an efficient sampling scheme and in the extreme case of having degeneracies proportional to the dimension, the time required to sample perfectly is even proportional to the time necessary to obtain an approximate sample. In these cases, the classical post-processing required can be done efficiently. We showed how to increase the efficiency of the sampling scheme and still obtain certifiably good samples by lumping close eigenstates together. We argue that one of its main advantages is its short circuit depth. We show that the algorithm is stable under noise in different steps of the implementation. It would be interesting to find sampling applications or models that satisfy the conditions under which our algorithms are efficient. Moreover, it would be worthwhile to investigate if there is a class of models to which we can tailor the perfect sampling algorithms to be efficient, as was done with success for attractive spin systems [12].

Acknowledgments

D.S.F. would like to thank Michael M. Wolf for raising the question of how to perfectly sample from a quantum Gibbs state and the constant support and Andreas Bluhm for helpful discussions. D.S.F acknowledges support from the graduate program TopMath of the Elite Network of Bavaria, the TopMath Graduate Center of TUM Graduate School at Technische Universität München, the Deutscher Akademischer Austauschdienst(DAAD) and by the Technische Universität München Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763.

References

1. M. Suzuki. *Quantum Monte Carlo Methods in Equilibrium and Nonequilibrium Systems*, volume 74. Springer, 2012.
2. B. M. Terhal and D. P. DiVincenzo. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A*, 61:022301, Jan 2000.
3. D. Poulin and P. Wocjan. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.*, 103(22):220502, 2009.
4. Y. Ge, A. Molnár, and J. I. Cirac. Rapid adiabatic preparation of injective projected entangled pair states and Gibbs states. *Phys. Rev. Lett.*, 116:080503, Feb 2016.
5. M. Yung and A. Aspuru-Guzik. A quantum–quantum Metropolis algorithm. *Proc. Natl. Acad. Sci. U.S.A*, 109(3):754–759, 2012.
6. A. Riera, C. Gogolin, and J. Eisert. Thermalization in nature and on a quantum computer. *Phys. Rev. Lett.*, 108(8):080402, 2012.
7. K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete. Quantum Metropolis sampling. *Nature*, 471(7336):87–90, 2011.
8. E. Bilgin and S. Boixo. Preparing thermal states of quantum systems by dimension reduction. *Phys. Rev. Lett.*, 105(17):170405, 2010.
9. M. J. Kastoryano and F. G. S. L. Brandão. Quantum Gibbs samplers: The commuting case. *Comm. Math. Phys.*, pages 1–43, 2016.
10. D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times.* Providence, R.I. American Mathematical Society, 2009.
11. L. Lovász and P. Winkler. Exact mixing in an unknown Markov chain. *Electron. J. Combin.*, 2, 1995.
12. J. G. Propp and D. B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random structures and Algorithms*, 9(1-2):223–252, 1996.
13. J. G. Propp and D. B. Wilson. How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *Journal of Algorithms*, 27(2):170 – 217, 1998.
14. N. Destainville, B. Georgeot, and O. Giraud. Quantum algorithm for exact Monte Carlo sampling. *Phys. Rev. Lett.*, 104:250502, Jun 2010.
15. A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 2007.
16. D. Burgarth, G. Chiribella, V. Giovannetti, P. Perinotti, and K. Yuasa. Ergodic and mixing quantum channels in finite dimensions. *New J. Phys.*, 15(7):073045, July 2013.
17. M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information.* Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000.
18. D. S. Abrams and S. Lloyd. Quantum algorithm providing exponential speed increase for finding

eigenvalues and eigenvectors. *Phys. Rev. Lett.*, 83(24):5162, 1999.

19. R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. In *Proc. Royal Soc. A*, volume 454, pages 339–354. The Royal Society, 1998.
20. K. M. Svore, M. B. Hastings, and M. Freedman. Faster phase estimation. *Quantum Info. Comput.*, 14(3-4):306–328, March 2014.
21. J.G. Keméy and J.L. Snell. *Finite Markov chains.* University series in undergraduate mathematics. Van Nostrand, 1960.
22. J. Hermon and Y. Peres. On sensitivity of mixing times and cutoff. *arXiv preprint arXiv:1610.04357*, 2016.
23. R. Dümcke and H. Spohn. The proper form of the generator in the weak coupling limit. *Z. Phys. B*, 34(4):419–422, 1979.
24. E. B. Davies. *Quantum theory of open systems.* IMA, 1976.
25. K. Temme. Lower bounds to the spectral gap of Davies generators. *J. Math. Phys.*, 54(12):122110, 2013.
26. E. B. Davies. Generators of dynamical semigroups. *J. Funct. Anal.*, 34(3):421–432, 1979.
27. W. Roga, M. Fannes, and K. Życzkowski. Davies maps for qubits and qutrits. *Rep. Math. Phys.*, 66:311–329, 2010.
28. R. Alicki. On the detailed balance condition for non-Hamiltonian systems. *Rep. Math. Phys.*, 10(2):249 – 258, 1976.
29. M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano, and J. Eisert. Dissipative quantum Church-Turing theorem. *Phys. Rev. Lett.*, 107(12):120501, 2011.
30. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
31. R. Bhatia. *Matrix analysis*, volume 169. Springer, 1997.
32. F. Hiai, M. Ohya, and M. Tsukada. Sufficiency, KMS condition and relative entropy in von Neumann algebras. *Pacific J. Math.*, 96(1):99–109, 1981.
33. O. Szehr and M. M. Wolf. Perturbation bounds for quantum Markov processes and their fixed points. *J. Math. Phys.*, 54(3):032203–032203, March 2013.
34. O. Szehr, D. Reeb, and M. M. Wolf. Spectral convergence bounds for classical and quantum Markov processes. *Comm. Math. Phys.*, pages 1–31, 2014.
35. J. Demmel, I. Dumitriu, and O. Holtz. Fast linear algebra is stable. *Numer. Math.*, 108(1):59–91, 2007.
36. P. Neal. The generalised coupon collector problem. *J. Appl. Probab.*, 45(3):621629, 2008.
37. P. Flajolet, D. Gardy, and L. Thimonier. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Appl. Math.*, 39(3):207–229, 1992.

Appendix A Expected number of observations to sample all possible outputs

To estimate the expected run-time of algorithm 4, we need to determine how often, on average, we must measure projections $\{P_i\}_{1\leq i\leq d'}$ on the state $\frac{\mathbb{1}}{d}$ until we observe all possible outcomes i. Here the P_i correspond to projections onto different eigenspaces of the Hamiltonian Gibbs state $\frac{e^{-\beta H}}{\mathcal{Z}_\beta}$ we are trying to sample from. The number of measurements corresponds to the time necessary to complete a column in algorithm 4 and will be denoted by $\phi(H)$.

One can see that this corresponds to the classical problem of determining how many coupons one must collect to obtain at least one of each, the coupon collector problem [36]. But in our case, we have unequal probabilities for the coupons or outcomes, as the probability $q(i)$ of observing i is given by

$$q(i) = \operatorname{tr}\left(P_i\frac{\mathbb{1}}{d}\right) = \frac{|S_i|}{d}. \tag{38}$$

Theorem A.1 (Coupon collector with unequal probabilities)**.** *Let $q \in \mathbb{R}^{d'}$ be the probability distribution of the measurement outcomes as in Eq.* (38)*. Let Y be the random variable given by the number of measurements necessary to observe all possible outputs. Then*

$$\phi(H) = E(Y) = \sum_{j=0}^{d'-1} (-1)^{d'-1-j} \sum_{|J|=j} (1 - Q_J)^{-1}, \tag{39}$$

where Q_J is defined as $Q_J = \sum_{i \in J} q(i)$ for $J \subset [d']$ and the second sum is over all subsets of size j.

Proof: We refer to [37, Corollary 4.2] for a proof. □

Although the expression in Eq. (39) is exact, its asymptotic scaling is not clear from it. Therefore, we show the following bound which is more directly accessible.

Theorem A.2. *Let $q \in \mathbb{R}^{d'}$ be the probability distribution of the measurement outcomes as in Eq.* (38)*. Let Y be the random variable given by the number of measurements necessary to observe all possible outputs. Moreover, assume that $d \leq |S_i| r(d)$ for some $r(d) : \mathbb{R} \to \mathbb{R}$. Then*

$$\phi(H) = E(Y) \leq r(d)\varphi_{d'} = \mathcal{O}\left(r(d) \log\left(r(d)\right)\right),$$

where $\varphi_{d'} = \sum\limits_{l=1}^{d'} \frac{1}{l}$.

Proof: We mimic the proof of the classical result for the coupon collector problem with uniform probability distribution. Denote by t_l the expected time to collect a new coupon after $l-1$ have been collected. We then have

$$E(Y) = \sum_{l=1}^{d'} E(t_l).$$

We clearly have $E(t_1) = 1$. Define $k(d) = \frac{d}{r(d)}$. For $l \geq 2$, note that as we have that we get each coupon with probability at least $\frac{k(d)}{d}$, the probability of getting a new coupon after having collected $l-1$ is at least $(d' - l + 1)\frac{k(d)}{d}$. From this, it follows that $E(t_l) \leq \frac{d}{k(d)} \left(d' - l + 1\right)^{-1}$ and so

$$E(Y) \leq 1 + \frac{d}{k(d)} \sum_{l=2}^{d'} \frac{1}{d' - l + 1} \leq \frac{d}{k(d)} \sum_{l=1}^{d'} \frac{1}{l}.$$

The claim follows from observing that $\varphi_d = \mathcal{O}(\log(d))$ and that $d' \leq r(d)$. □

From this, it is easier to get estimates for cases that might be of interest. Here we collect the bounds for the extreme cases of highly degenerate spectra, that is, with each eigenspace having dimension $\Omega\left(d \log(d)^{-m}\right)$ for $m \in \mathbb{N}$ and the non-degenerate case.

Corollary A.1. *Let $q \in \mathbb{R}^{d'}$ be the probability distribution of the measurement outcomes as in Eq.* (38)*. Let Y be the random variable given by the number of measurements necessary to observe all possible outputs. Moreover, assume that $|S_i| \geq c\frac{d}{\log(d)^m}$ for some $c \in \mathbb{R}$ and all $i \in [d']$. Then*

$$\phi(H) = E(Y) \leq \frac{\log(d)^m}{c}\varphi_{c^{-1}\log(d)^m} = \mathcal{O}\left(\log\left(d\right)^m \log\left[m\log(d)\right]\right).$$

Proof: Just take $r(d) = c^{-1}\log(d)^m$ in theorem A.2. □

Corollary A.2. *Let $q \in \mathbb{R}^{d'}$ be the probability distribution of the measurement outcomes as in Eq.* (38)*. Let Y be the random variable given by the number of measurements necessary to observe all possible outputs. Moreover, assume that $q(i) = \frac{1}{d}$. Then*

$$\phi(H) = E(Y) = \mathcal{O}\left(d\log\left(d\right)\right)$$

Proof: Just take $r(d) = d$ in theorem A.2. □

That is, we might go from a constant number of samples necessary to complete a column in the case of degeneracies proportional to the dimension to a scaling like $d\log(d)$. One should note that applying the bound in corollary A.2 to analyze the runtime of algorithm 4 probably leads to bounds that are too pessimistic for spectra that are not very degenerate. To see why this is the case, note that in algorithm 4 we do not discard measurements outcome we have already observed, but rather use them to complete other columns, which we do not take into account in this analysis.

Appendix C

Further articles as principal author under review

C.1 Dimensionality reduction of SDPs through sketching

Dimensionality reduction of SDPs through sketching

D. Stilck França, A. Bluhm

Semidefinite programs (SDPs) are a natural framework to formulate and solve many optimization problems encountered in quantum information theory. One of the main bottlenecks for their more widespread application in practice is the prohibitive amount of memory required to solve them. In this work, we investigate if one can apply concepts from quantum information theory to solve SDPs using less memory. In particular, we show how to construct positive maps whose output dimension is much smaller than its input and that approximately preserve the Hilbert Schmidt scalar product between matrices based on Johnson-Lindenstrauss transforms. We call such maps positive linear sketches. We then use them to approximately solve SDPs with inequality constraints using substantially less memory, as such map approximately preserve feasibility sets of SDPs. The effectiveness of our methods depends on the Schatten $1-$norm of the matrices that define the constraints. We show how to apply similar ideas to linear matrix inequality feasibility problems. Moreover, we show some no-go results that clarify the limitations of this approach to the solution of SDPs and the limitations of positive, linear sketches, results which might be of independent interest to the quantum information community.

C.1.1 SDPs and Quantum Information

Many optimization problems in quantum information theory can be naturally cast as SDPs. Examples related to the problems discussed here before are e.g. upper bounds on the classical capacity of quantum channels [58]. Other interesting problems related to the convergence of (classical) semigroups can be cast as SDPs, such as the fastest mixing Markov chain on a certain graph [59]. Although SDPs are solvable in polynomial time, the large amount of memory required to solve them limits their application in practice to problems of moderate dimension. It is therefore of great importance to develop methods that can (approximately) solve larger instances of SDPs using fewer resources. Although the previous works in this dissertation are concerned with *how noisy* a certain semigroup of quantum channels is, here it will be central to see *how much we can compress* matrices while approximately preserving positivity and the geometry induced by the Hilbert-Schmidt scalar product. We will focus on real matrices now, but note that all results can be generalized to complex matrices. Given a matrix $A \in \mathcal{M}_d$, we will denote its transpose by A^T. Moreover, we denote the set of symmetric matrices, i.e. A such that $A = A^T$, by $\mathcal{M}_D^{\mathrm{sym}}$.

Given that in SDPs one optimizes a linear functional over positive matrices that satisfy linear constraints, one natural approach to "compress" SDPs and reduce their dimension is to apply positive linear maps that approximately projects the feasible set to a set of smaller dimension without changing the target value too much. To find such positive maps we drew inspiration from the field of sketching [60]. Sketching is an idea that is currently studied intensively in theoretical computer science and mathematics.

C.1.2 Sketching the HS scalar product

The main idea in sketching is to (probabilistically) compress the input to a problem and to give a (probabilistic) algorithm that yields an approximate solution to the original problem with high probability. A handy set of tools in this setting are Johnson-Lindenstrauss transforms, which we also explore in this work.

Definition C.1.1 (Johnson-Lindenstrauss transforms)**.** *A random matrix $S \in \mathcal{M}_{d,D}$ is a JLT with parameters (ϵ, δ, k) if with probability at least $1-\delta$ for any k-element subset $V \subset \mathbb{R}^D$ and for all v, $w \in V$ it holds that*

$$|\langle Sv, Sw\rangle - \langle v, w\rangle| \leq \epsilon \|v\|_2 \|w\|_2.$$

It is known that there are (ϵ, δ, k)-JLT with $d = \mathcal{O}(\epsilon^{-2} \log(\delta k))$. Those can also be chosen sparse.

As mentioned before, we want to construct random maps that approximately preserve the HS scalar product, and the following construction delivers exactly that.

Lemma C.1.2. *Let $B_1, \ldots, B_m \in \mathcal{M}_D^{\mathrm{sym}}$ and $S \in \mathcal{M}_{d,D}$ be an (ϵ, δ, k)-JLT with $\epsilon \leq 1$ and k such that*

$$k \geq \sum_{i=1}^{m} \operatorname{rank}(B_i).$$

Then

$$\mathbb{P}\left[\forall i, j \in [m] : |\operatorname{Tr}\left(SB_iS^TSB_jS^T\right) - \operatorname{Tr}\left(B_iB_j\right)| \leq 3\epsilon\|B_i\|_1\|B_j\|_1\right] \geq 1 - \delta. \tag{C.1}$$

This claim follows immediately if we diagonalize the matrices B_i and apply the JLT property to their eigenvectors combined with the Cauchy-Schwarz inequality. The usual JLT theorem gives a scaling of the error with the Schatten $2-$norm, which is smaller than the $1-$norm, and thus the scaling in Equation (C.1) is worse than expected. Given that this proof is admittedly crude, it is natural to ask if it is possible to obtain better bounds using more sophisticated techniques or a better family of positive maps. Some no-go theorems in this direction were already proved in [14], but we also prove the following theorem that shows that our results cannot be improved significantly:

Theorem C.1.3 (No-go for Sketching HS). *Let $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ be a random positive map such that with positive probability for any $Y_1, \ldots, Y_{D+1} \in \mathcal{M}_D$ and $0 < \epsilon < \frac{1}{4}$ we have*

$$|\operatorname{Tr}\left(\Phi(Y_i)^T\Phi(Y_j)\right) - \operatorname{Tr}\left(Y_i^TY_j\right)| \leq \epsilon\|Y_i\|_2\|Y_j\|_2. \tag{C.2}$$

Then $d = \Omega(D)$.

The proof of this statement uses the fact that for any set $\{P_i\}_{i\in I}$ of $|I| \geq d$ positive semidefinite matrices in $\mathcal{M}_d$ such that $\operatorname{Tr}\left(P_i^2\right) = 1$ we have that

$$\sum_{i\neq j} \operatorname{Tr}\left(P_iP_j\right)^2 \geq \frac{(|I| - d)^2|I|}{(|I| - 1)d^2}.$$

We then use this to show that if d is not $\Omega(D)$ we could violate this bound.

C.1.3 Sketching SDPs to approximate the optimum value

The kind of SDPs we consider are given in terms of inequality constraints, i.e. they have the form

$$\begin{array}{lll} \text{maximize} & \operatorname{Tr}(AX) & \\ \text{subject to} & \operatorname{Tr}(B_iX) \leq \gamma_i, & i \in \{1, \ldots, m\} \\ & X \geq 0 & \end{array}$$

where A, $B_i \in \mathcal{M}_D^{\mathrm{sym}}$ and $\gamma_i \in \mathbb{R}$ for all $i \in \{1, \ldots, m\}$. We will always assume that our SDPs satisfy Slater's condition. The sketched SDP for this problem will be:

Definition C.1.4 (Sketched SDP). *Let A, $B_i \in \mathcal{M}_D^{\mathrm{sym}}$ and η, $\gamma_i \in \mathbb{R}$ for all $i \in \{1, \ldots, m\}$ and $\epsilon > 0$. Given the existence of an optimal solution $X^* \geq 0$ satisfying $\operatorname{Tr}(X^*) \leq \eta$ and given further an (ϵ, δ, k)-JLT $S \in \mathcal{M}_{d,D}$, we call the d-dimensional optimization problem*

$$\begin{array}{lll} \textit{maximize} & \operatorname{Tr}\left(SAS^TY\right) & \\ \textit{subject to} & \operatorname{Tr}\left(SB_iS^TY\right) \leq \gamma_i + 3\epsilon\eta\|B_i\|_1, & i \in \{1, \ldots, m\} \\ & Y \geq 0 & \end{array} \tag{C.3}$$

the sketched SDP. Here, $k \geq \operatorname{rank}(X^*) + \operatorname{rank}(A) + \sum_{i=1}^{m} \operatorname{rank}(B_i)$.

The sketched SDP gives an approximation of the value of the original problem with high probability.

Theorem C.1.5 (Bounds on sketched value)**.** *Let α_S be the value of the sketched SDP defined by A, B_i and S. Then*

$$\alpha_S + 3\epsilon\eta\|A\|_1 \geq \alpha \tag{C.4}$$

with probability at least $1-\delta$. Moreover, we also have

$$\alpha \geq \alpha_S - \epsilon C. \tag{C.5}$$

Here C scales linearly with $\|B_i\|_1, \eta$, $(\alpha - \mathrm{Tr}\,(X_0 A))$ and $\max_i(\gamma_i - \mathrm{Tr}\,(X_0 B_i))^{-1}$. X_0 is a strictly feasible solution to the original problem.

The upper bound in Equation (C.4) follows from noting that by our result on sketching the HS norm and our choice of the JLT S we have that $\mathrm{Tr}\left(SB_iS^TSX^*S^T\right) \leq \gamma_i + 3\epsilon\eta\|B_i\|_1$ for all $1 \leq i \leq m$. This implies that SX^*S^T, which is again positive semidefinite, will be a feasible point of the sketched SDP with value at least $\alpha - 3\epsilon\eta\|A\|_1$, from which the claim follows. The lower bound follows from observing that, by the cyclicity of the trace, $S^TX_S^*S$, where X_S^* is an optimal point of the sketched SDP, will be a feasible point of a perturbed version of the sketchable SDP. Using continuity estimates for SDPs we obtain the claim.

These results suggest the following simple algorithm to solve SDPs approximately. Sample a JLT S with the desired parameters, compute the sketched matrices SB_iS^T and SAS^T and solve the sketched SDP. The previous Theorem then guarantees that the error will not be significant.

We also note that we can obtain approximately optimal points of the original SDP in the case that all $\gamma_i \geq 0$.

C.1.4 Linear Matrix Inequality feasibility problems

Using similar ideas, we can also handle linear matrix inequality feasibility problem.

Theorem C.1.6. *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{sym}\backslash\{0\}$ such that*

$$\sum_{i=1}^m c_iB_i - A \not\geq 0$$

for all $c \in \mathbb{R}_+^m$. Suppose further that

$$\Lambda = \mathrm{cone}\{B_1, \ldots, B_m\}$$

is pointed and $\Lambda \cap S_D^+ = \{0\}$. Moreover, let $\rho \in \mathcal{S}_D^+$ be such that for all $i \in [m]$

$$\mathrm{Tr}\,(\rho B_i) < 0, \quad \mathrm{Tr}\,(-A\rho) < 0 \quad \text{and} \quad \mathrm{Tr}\,(\rho) = 1.$$

Set

$$\epsilon = \frac{1}{2}\min\left\{\left|\frac{\mathrm{Tr}\,(\rho B_1)}{\|B_1\|_1}\right|, \ldots, \left|\frac{\mathrm{Tr}\,(\rho B_m)}{\|B_m\|_1}\right|, \left|\frac{\mathrm{Tr}\,(\rho A)}{\|A\|_1}\right|\right\}$$

and take $S \in \mathcal{M}_{d,D}$ to be an (ϵ, δ, k)-JLT. Here,

$$k \geq \mathrm{rank}(A) + \mathrm{rank}(\rho) + \sum_{i=1}^m \mathrm{rank}(B_i).$$

Then

$$\sum_{i=1}^m c_iSB_iS^T - SAS^T \not\geq 0 \tag{C.6}$$

for all $c \in \mathbb{R}_+^m$, with probability at least $1-\delta$.

The proof relies on the fact that under these assumptions such a ρ always exists and defines a hyperplane that separates the LMI and the cone of positive semidefinite matrices. It follows from our result on positive sketches of the HS scalar product that $S\rho S^T$ will also define a separating hyperplane with high probability. This results suggest an algorithm to certify that a certain LMI is not feasible. We first check if the LMI

$$\sum_{i=1}^{m} c_i S B_i S^T - SAS^T \not\geq 0$$

is feasible. If it is not, then the original problem was not feasible as well, as any feasible point remains feasible when we conjugate with S. The results of the last theorem assure that we will be able to detect that the original problem was not feasible with a high enough probability if we choose ϵ small enough.

C.1.5 Complexity and Memory Considerations

The main bottleneck of the algorithm is to compute the matrices SB_iS^T. For a fixed error ϵ, probability of success δ, $\|B_i\|_1 = \mathcal{O}(1) = \|A\|_1$, $m = \text{poly}(D)$ and solving SDPs to a fixed accuracy we have:

	General Solver	Sketched	Ellipsoid method
Complexity	$\text{SDP}(m, D)$	$\mathcal{O}(D^2 m \log(mD) + \text{SDP}(m, \log(mD))$	$\mathcal{O}(\max\{m, D^2\}D^6)$
Memory	$\mathcal{O}(D^2 m)$	$\mathcal{O}(m \log(mD))$	$\mathcal{O}(D^2 m)$

We obtain further improvements in the complexity and memory if we suppose that the B_i are sparse. If the SDP can be sketched, doing so gives a speedup as long as $\text{SDP}(m, D) = \Omega(mD^{2+\mu})$, $\mu > 0$.

C.1.6 Individual Contribution

The project's idea was motivated by discussions between Andreas Bluhm and me. I am the principal author of this article. I had the idea that one could use conjugations with JLTs to reduce the dimensionality of SDPs while approximately preserving their feasible set. I proved and formulated Lemma C.1.2 (Lemma 3.1 in the article) and Theorem C.1.3 (Theorem 3.2 in the article). A further no-go result not included in this summary, Theorem 5.1 in the article, was also formulated and proved by me. In the case of Theorem C.1.5, I was responsible for proving and formulating the upper bound on the value, which corresponds to Theorem 5.3 in the article, while A. Bluhm formulated and proved the lower bound, which corresponds to Theorem 5.5, although the idea of obtaining the bound through continuity bounds on the relaxed SDP was mine. He then observed that duality of SDPs would deliver the desired bound. The idea of relaxing the original SDP as done in the sketched SDP also goes back to me. The observation that we can also obtain a point which is close to optimal from SDP-packing problems, formulated in Theorem 5.6, goes back to me and I formulated and proved the theorem. Regarding Theorem C.1.6 (Lemma 4.1 in the article), I was responsible for formulating and proving a slightly less general version of it and A. Bluhm for arguing that the same proof would apply to the current more general setting. A. Bluhm was responsible for the discussion on complexity and memory considerations and I was responsible for the numerics and applications section. I wrote all the code necessary for our examples and thought of the applications. I wrote sections 1, 2, 3, 4, 5, 7 and Appendix B. A. Bluhm was responsible for the appendices on complexifying JLTs and we contributed equally to the section on random feasibility problems. After the completion of the first draft, we discussed together how we could improve the presentation of the article, leading to its current form.

Dimensionality reduction of SDPs through sketching

Andreas Bluhm* [1] and Daniel Stilck França† [1]

[1] *Department of Mathematics, Technical University of Munich, 85748 Garching, Germany*

July 31, 2017

We show how to sketch semidefinite programs (SDPs) using positive maps in order to reduce their dimension. More precisely, we use Johnson-Lindenstrauss transforms to produce a smaller SDP whose solution preserves feasibility or approximates the value of the original problem with high probability. These techniques allow to improve both complexity and storage space requirements. They apply to problems in which the Schatten 1-norm of the matrices specifying the SDP and of a solution to the problem is constant in the problem size. Furthermore, we provide some no-go results which clarify the limitations of positive, linear sketches in this setting. Finally, we discuss numerical examples to benchmark our methods.

Contents

1. Introduction **2**

2. Preliminaries **3**

3. Sketching the Hilbert-Schmidt product with positive maps **6**

4. Sketching linear matrix inequality feasibility problems **8**

5. Approximating the value of semidefinite programs through sketching **10**

6. Complexity and memory gains **16**

*andreas.bluhm@ma.tum.de
†dsfranca@mytum.de

7. Applications and numerical examples 17
7.1. Estimating the value of a semidefinite packing problem 17
7.2. Linear matrix inequality feasibility problems 18

8. Conclusion 19

A. Complexifying Johnson-Lindenstrauss transforms 23

B. Proof of Theorem 3.2 24

C. Lower bound on the value of SDPs through sketching 26

D. Random feasibility problems 29

1. Introduction

Semidefinite programs (SDPs) are a prominent class of optimization problems [LA16]. They have applications across different areas of science and mathematics, such as discrete optimization [WA02] or control theory [BEFB94].

However, although there are many different algorithms that solve an SDP up to an error ϵ in a time that scales polynomially with the dimension and logarithmically with ϵ^{-1} [Bub15], solving large instances of SDPs still remains a challenge. This is not only due to the fact that the number and cost of the iterations scale superquadratically with the dimension for most algorithms to solve SDPs, but also due to the fact that the memory required to solve large instances is beyond current capabilities. This has therefore motivated research on algorithms that can solve SDPs, or at least obtain an approximate solution, with less memory requirements. One such example is the recent [YUAC17], where ideas similar to ours were applied to achieve optimal storage requirements necessary to solve a certain class of SDPs. While their work proposes a new way to solve an SDP using linear sketches, our approach relies on standard convex optimization methods.

In this work, we develop algorithms to estimate the value of an SDP with linear inequality constraints and to determine if a given linear matrix inequality (LMI) is feasible or not. These algorithms convert the original problem to one of the same type, but of smaller dimension, which we call the sketched problem. Subsequently, this new problem can be solved with the same techniques as the original one, but potentially using less memory and achieving a smaller runtime. Therefore, we call this a black box algorithm. With high probability an optimal solution to the sketched problem allows us to infer something about the original problem.

In the case of LMIs, if the sketched problem is infeasibile, we obtain a certificate that the original problem is also infeasibile. If the sketched problem is feasible, we are able to infer that the original problem is either "close to feasible" or feasible with high probability, under some technical assumptions.

In the case of estimating the value of SDPs, we are able to give an upper bound that holds with high probability and a lower bound on the value of the SDP from the value of the sketched problem, again under some technical assumptions. For a certain class of SDPs, which includes the so-called semidefinite packing problems [IPS05], we are able to find a feasible point of the original problem which is close to the optimal point and most technical aspects simplify significantly. For this class it can be checked whether this feasible point is indeed optimal.

Our algorithms work by conjugating the matrices that define the constraints of the SDP with Johnson-Lindenstrauss transforms [Woo14], thereby preserving the structure of the problem. Similar ideas have been proposed to reduce the memory usage and complexity of solving linear programs [VPL15]. While those techniques aim to reduce the number of constraints, our goal is to reduce the dimension of the matrices involved.

Unfortunately, the dimension of the sketch needed to have a fixed error with high probability scales with the Schatten 1-norm of the constraints and that of an optimal solution to the SDP, which significantly restricts the class of problems to which these methods can be applied. We are able to show that one cannot significantly improve this scaling and that one cannot sketch general SDPs using linear maps.

This paper is organized as follows: in Section 2, we fix our notation and recall some basic notions from matrix analysis, Johnson-Lindenstrauss transforms, semidefinite programs and convex analysis which we will need throughout the paper. We then proceed to show how to sketch the Hilbert-Schmidt scalar product with positive maps in Section 3. We apply these techniques in Section 4 to show how to certify that certain LMIs are infeasible by showing the infeasibility of an LMI of smaller dimension. In Section 5, we apply similar ideas to estimate the value of an SDP with linear inequality constraints by solving an SDP of lower dimension. This is followed by a discussion of the possible gains in the complexity of solving these problems and for the memory requirements in Section 6. Furthermore, we make some numerical simulations in Section 7 to benchmark our findings by applying our techniques to a problem from the field of optimal designs of experiments and to a random LMI with matrices sampled from the Gaussian unitary ensemble.

2. Preliminaries

We begin by fixing our notation. For brevity, we will write the set $\{1,\ldots,d\}$ as $[d]$. The set of $d\times D$ matrices over some field $\mathbb{K}\in\{\mathbb{R},\mathbb{C}\}$ will be written as $\mathcal{M}_{d,D}(\mathbb{K})$ and just $\mathcal{M}_d(\mathbb{K})$ if $d=D$. We will often omit the underlying field if it is not relevant for the statement. We will denote by $\mathcal{M}_d^{sym}$ the set of symmetric $d\times d$ matrices in the real case and the set of Hermitian matrices in the complex case. For $A\in\mathcal{M}_d$, A^T will denote the transpose of A in the real case and the Hermitian conjugate in the complex case. To avoid cumbersome notation and redundant theorems, we will prove most of the statements only for real matrices. However, note that all statements translate to the complex case in a straightforward fashion. We will state most of the definitions just for real matrices, but it should be clear how to generalize them to the complex case. For $A\in\mathcal{M}_d^{sym}$ we will write $A\geq 0$ if A is positive semidefinite. We will denote the cone of $d\times d$ positive semidefinite matrices by $\mathcal{S}_d^+$ and its interior, the positive definite matrices, by $\mathcal{S}_d^{++}$.

Definition 2.1 (Schatten p-norm and Hilbert-Schmidt scalar product). *Let $A \in \mathcal{M}_d$ and $s_1, \ldots, s_d$ be its singular values. We define the Schatten p-norm of A for $p \geq 1$, denoted by $\|A\|_p$, to be given by*

$$\|A\|_p^p = \sum_{i=1}^{d} s_i^p$$

and for $p = \infty$ by $\|A\|_\infty = \max_{1 \leq i \leq d} s_i$ The Schatten $2-norm$ is induced by the Hilbert-Schmidt scalar product, which is given by

$$\langle A, B \rangle_{HS} = \operatorname{Tr}\left(A^T B\right)$$

for $A, B \in \mathcal{M}_d$.

We will sometimes refer to the case $p = 2$ as the Hilbert-Schmidt (HS) norm and $p = \infty$ as the operator norm.

Definition 2.2 (Positive Map). *A linear map $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ is called positive if $\Phi(\mathcal{S}_D^+) \subseteq \mathcal{S}_d^+$.*

The structure of the set of positive maps is still not very well understood [Stø13]. For our purposes, however, we will only need maps of the form $\Phi(X) = SXS^T$ with $S \in \mathcal{M}_{d,D}$, which are easily seen to be positive.

We will adopt the standard Big $\mathcal{O}$ notation for the asymptotic behavior of functions. That is, for two functions $f, g : \mathbb{R} \to \mathbb{R}$, we will write $g = \mathcal{O}(f)$ if there exists a constant $M > 0$ such that for all $x > x_0$ we have $|g(x)| \leq M|f(x)|$. Analogously, we write $g = \Omega(f)$ if there exists a constant $M > 0$ such that for all $x > x_0$ we have $|g(x)| \geq M|f(x)|$.

The following families of matrices will play a crucial role for our purposes:

Definition 2.3 (Johnson-Lindenstrauss transform). *A random matrix $S \in \mathcal{M}_{d,D}(\mathbb{K})$ is a Johnson-Lindenstrauss transform (JLT) with parameters (ϵ, δ, k) if with probability at least $1 - \delta$, for any k-element subset $V \subseteq \mathbb{K}^D$, for all $v, w \in V$ it holds that*

$$|\langle Sv, Sw \rangle - \langle v, w \rangle| \leq \epsilon \|v\|_2 \|w\|_2.$$

Note that one usually only demands that the norm of the vectors involved is distorted by at most ϵ in the definition of JLTs, but this is equivalent to the definition we chose by the polarization identity. There are many different examples of JLTs in the literature and we refer to [Woo14] and references therein for more details. Most of the constructions of JLTs focus on real matrices, but in Section A we show how to lift some of these results to cover complex matrices. The most prominent JLT is probably the following:

Theorem 2.4. *Let $0 < \epsilon, \delta < 1$ and $S = \frac{1}{\sqrt{d}} R \in \mathcal{M}_{d,D}(\mathbb{R})$, where the entries of R are i.i.d. standard Gaussian random variables. If $d = \Omega(\epsilon^{-2} \log(k\delta^{-1}))$, then S is an (ϵ, δ, k)-JLT.*

Proof. We refer to [Woo14, Lemma 2.12] for a proof. □

The main drawback of using Gaussian JLTs is that these are dense matrices. We will denote by $\text{nnz}(X)$ the number of nonzero elements of a matrix $X \in \mathcal{M}_D$. As we later want to compute products of the form SXS^T, it will be advantageous to have a sparse S to speed up the computation of this product. The computational cost of forming this product will most often be the bottleneck of our algorithms. Fortunately, there has been a lot of recent work on sparse JLTs. In particular, we have the following almost optimal result.

Theorem 2.5 (Sparse JLT [KN14, Section 1.1])**.** *There is an (ϵ, δ, k)-JLT $S \in \mathcal{M}_{d,D}$ with $d = \mathcal{O}\left(\epsilon^{-2}\log\left(k\delta^{-1}\right)\right)$ and $s = \mathcal{O}(\epsilon^{-1}\log\left(k\delta^{-1}\right))$ nonzero entries per column.*

Proof. We refer to [KN14, Section 1.1] for a proof and remark that the proof is constructive. □

Given some JLT $S \in \mathcal{M}_{d,D}$, the positive map $\Phi : \mathcal{M}_D \to \mathcal{M}_d$, $X \mapsto SXS^T$ will be called the sketching map and d the sketching dimension.

We will now fix our notation for semidefinite programs. Semidefinite programs are a class of optimization problems in which a linear functional is optimized under linear constraints over the set of positive semidefinite matrices. We refer to [LA16] for an introduction to the topic. There are many equivalent ways of formulating SDPs. In this work, we will assume w.l.o.g. that the SDPs are given in the following form:

Definition 2.6 (Sketchable SDP)**.** *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{\text{sym}}$ and $\gamma_1, \ldots, \gamma_m \in \mathbb{R}$. We will call the constrained optimization problem*

$$\begin{array}{lll} \text{maximize} & \operatorname{Tr}(AX) & \\ \text{subject to} & \operatorname{Tr}(B_iX) \leq \gamma_i, & i \in [m] \\ & X \geq 0, & \end{array} \tag{1}$$

a sketchable SDP.

Sometimes we will also refer to a sketchable SDP as the original problem. We will see later how to approximate the value of these SDPs. SDPs have a rich duality theory, in which, instead of optimizing over positive semidefinite matrices that satisfy certain constraints, one optimizes over the points that satisfy a linear matrix inequality (LMI). The dual problem of a sketchable SDP is given by the following:

$$\begin{array}{ll} \text{minimize} & \langle c, \gamma \rangle \\ \text{subject to} & \sum_{i=1}^{m} c_i B_i - A \geq 0 \\ & c \in \mathbb{R}_+^m, \end{array} \tag{2}$$

where $\gamma \in \mathbb{R}^m$ is the vector with coefficients γ_i. Here, $\mathbb{R}_+^m = \{\, x \in \mathbb{R}^m : x_i \geq 0 \,\}$. SDPs and LMIs will be called feasible if there is at least one point satisfying all the constraints, otherwise we will call it infeasible. A sketchable SDP will be called strictly feasible if there is a point $X \geq 0$ such that all the constraints in (1) are satisfied with strict inequality.

Under some conditions the primal problem (1) and the dual problem (2) have the same value. One widely used sufficient condition is Slater's condition [LA16], which asserts that if we have a strictly feasible point of full rank for the primal problem and if the dual problem is feasible, then the primal and the dual have the same value and there is an optimal solution to the dual problem.

We will need some standard concepts from convex analysis. Given $a_1, \ldots, a_n \in V$ for a vector space V, we denote by $\operatorname{conv}\{a_1, \ldots, a_n\}$ the convex hull of the points. By $\operatorname{cone}\{a_1, \ldots, a_n\}$ we will denote the cone generated by these elements and a convex cone C will be called pointed if $C \cap -C = \{0\}$.

3. Sketching the Hilbert-Schmidt product with positive maps

One of our main ingredients to sketch an SDP or LMI will be a random positive map $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ that preserves the Hilbert-Schmidt scalar product with high probability. We demand positivity to assure that the structure of the SDP or LMI is preserved. Below, we first consider the example $\Phi(X) = SXS^T$ with S a JLT.

Lemma 3.1. *Let $B_1, \ldots, B_m \in \mathcal{M}_D^{\mathrm{sym}}$ and $S \in \mathcal{M}_{d,D}$ be an (ϵ, δ, k)-JLT with $\epsilon \leq 1$ and k such that*

$$k \geq \sum_{i=1}^{m} \operatorname{rank} B_i.$$

Then

$$\mathbb{P}\left[\forall i, j \in [m] : |\operatorname{Tr}\left(SB_iS^TSB_jS^T\right) - \operatorname{Tr}\left(B_iB_j\right)| \leq 3\epsilon \|B_i\|_1 \|B_j\|_1\right] \geq 1 - \delta. \quad (3)$$

Proof. Observe that the eigenvectors of the B_i corresponding to nonzero eigenvalues of the B_i form a subset of cardinality at most k of $\mathbb{K}^D$. Let $A, B \in \{B_1, \ldots, B_m\}$. As S is an (ϵ, δ, k)-JLT, with probability at least $1 - \delta$ we have for all normalized eigenvectors a_i of A and b_j of B that

$$\left||\langle Sa_i, Sb_j\rangle| - |\langle a_i, b_j\rangle|\right| \leq \epsilon$$

by the reverse triangle inequality. We also have that for any a_i, b_j

$$\|Sa_i\|_2; \|Sb_j\|_2 \leq \sqrt{1+\epsilon},$$

again by the fact that S is a JLT. As $\epsilon \leq 1$ and by the Cauchy-Schwarz inequality, it follows that

$$|\langle Sa_i, Sb_j\rangle| + |\langle a_i, b_j\rangle| \leq 3$$

and hence

$$\left||\langle Sa_i, Sb_j\rangle|^2 - |\langle a_i, b_j\rangle|^2\right| \leq 3\epsilon. \quad (4)$$

Now let λ_i and μ_j be the eigenvalues of A and B, respectively. We have:

$$\begin{aligned} \left|\mathrm{Tr}\left(SAS^TSBS^T\right) - \mathrm{Tr}\left(AB\right)\right| &= \left|\sum_{i,j=1}^{D} \lambda_i\mu_j(|\langle Sa_i, Sb_j\rangle|^2 - |\langle a_i, b_j\rangle|^2)\right| \\ &\leq 3\epsilon \sum_{i,j=1}^{D} |\lambda_i||\mu_j| = 3\epsilon\|A\|_1\|B\|_1 \end{aligned}$$

with probability at least $1-\delta$. As A, B were arbitrary, the claim follows.

□

The scaling of the error with the Schatten 1-norm of the matrices involved in Lemma 3.1 is highly undesirable and the estimates used to prove it are admittedly crude. We note that a similar estimate was proved in [SH15]. Moreover, just using the fact that $\mathcal{M}_D^{\mathrm{sym}}(\mathbb{R}) \simeq \mathbb{R}^{\frac{D(D+1)}{2}}$ as a Hilbert space, we could use an (ϵ, δ, k)-JLT L for $\mathbb{R}^{\frac{D(D+1)}{2}}$ and isometrically embed the resulting vector into a symmetric matrix. Denoting this transformation by $\tilde{L}$, we obtain

$$\mathbb{P}\left[\forall i, j \in [n] : \left|\mathrm{Tr}\left(\tilde{L}(B_i)\tilde{L}(B_j)\right) - \mathrm{Tr}\left(B_iB_j\right)\right| \leq 3\epsilon\|B_i\|_2\|B_j\|_2\right] \geq 1-\delta.$$

That is, if only demand the sketching map to map symmetric matrices to symmetric matrices, we clearly obtain a better scaling of the error with this procedure. Note, however, that the map L may not be positive, one of the requirements to later sketch SDPs. The next theorem shows that a scaling of the error with the Schatten $2-$norm of the matrices involved is not possible with positive maps if we want to achieve a non-trivial compression.

Theorem 3.2. *Let $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ be a random positive map such that with strictly positive probability for any $Y_1, \dots Y_{D+1} \in \mathcal{M}_D$ and $0 < \epsilon < \frac{1}{4}$ we have*

$$\left|\mathrm{Tr}\left(\Phi(Y_i)^T\Phi(Y_j)\right) - \mathrm{Tr}\left(Y_i^TY_j\right)\right| \leq \epsilon\|Y_i\|_2\|Y_j\|_2. \tag{5}$$

Then $d = \Omega(D)$.

Proof. We refer to Appendix B for a proof. □

One could hope to achieve a better bound for low rank matrices, but we note that this does not significantly improve our bound, as for $A \in \mathcal{M}_D$ of rank r we have $\|A\|_1 \leq \sqrt{r}\|A\|_2$. That is, by choosing an $\left(\frac{\epsilon}{r}, \delta, k\right)$-JLT, we may ensure that inequality (3) holds with the HS norm if the rank of the matrices involved is bounded by $r \ll d$. This just increases the dimension of the involved JLT matrices by a factor of r^2 if we have the usual ϵ^{-2} dependence on the dimension for the JLTs. It remains open if one could achieve a better compression for a sublinear number of matrices.

4. Sketching linear matrix inequality feasibility problems

In this section we will show how to use JLTs to certify that certain linear matrix inequalities (LMI) are infeasible by showing that an LMI of smaller dimension is infeasible. We will consider inequalities like the ones in the following lemma:

Lemma 4.1. *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{sym} \setminus \{ 0 \}$ such that*

$$\sum_{i=1}^{m} c_i B_i - A \ngeq 0 \tag{6}$$

for all $c \in \mathbb{R}_+^m$. Suppose further that

$$\Lambda = \operatorname{cone}\{B_1, \ldots, B_m\}$$

is pointed and $\Lambda \cap S_D^+ = \{0\}$. Then there exists a $\rho \in \mathcal{S}_D^+$ such that for all $i \in [m]$

$$\operatorname{Tr}(\rho B_i) < 0, \quad \operatorname{Tr}(-A\rho) < 0 \quad \textit{and} \quad \operatorname{Tr}(\rho) = 1.$$

Proof. Let $E = \operatorname{conv}\{-A, B_1, \ldots, B_m\}$. We will show that $S_D^+ \cap E = \emptyset$. Suppose there exists an $X = -p_0 A + \sum\limits_{i=1}^{m} p_i B_i \in S_D^+ \cap E$ with $p \in [0,1]^{m+1}$. If $p_0 > 0$, we could rescale X by p_0^{-1} and obtain a feasible point for (6), a contradiction. If $p_0 = 0$ and $X \neq 0$, this would in turn contradict $\Lambda \cap S_D^+ = \{0\}$. And if $X = 0$, the cone Λ would not be pointed. From these arguments it follows that $0 \notin E$. The set E is therefore closed, convex, compact and disjoint from the convex and closed set $\mathcal{S}_D^+$. We may thus find a hyperplane that strictly separates $\mathcal{S}_D^+$ from E. That is, a $\rho \in \mathcal{M}_D^{\mathrm{sym}}$ such that w.l.o.g. $\operatorname{Tr}(\rho X) \geq 0$ for all $X \in \mathcal{S}_D^+$, as $0 \in \mathcal{S}_D^+$, and $\operatorname{Tr}(Y\rho) < 0$ for all $Y \in E$. As $\operatorname{Tr}(\rho X) \geq 0$ for all $X \geq 0$, it follows that ρ is positive semidefinite and it is clear that by normalizing ρ we may choose ρ with $\operatorname{Tr}(\rho) = 1$. □

The main idea is now to show that under these conditions we may sketch the hyperplane in a way that it still separates the set of positive semidefinite matrices and the sketched version of the set $\{\sum_{i=1}^m \gamma_i B_i - A | \gamma_i \geq 0\}$.

Theorem 4.2. *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{sym} \setminus \{ 0 \}$ such that*

$$\sum_{i=1}^{m} c_i B_i - A \ngeq 0$$

for all $c \in \mathbb{R}_+^m$. Suppose further that

$$\Lambda = \operatorname{cone}\{B_1, \ldots, B_m\}$$

is pointed and $\Lambda \cap S_D^+ = \{0\}$. Moreover, let $\rho \in \mathcal{S}_D^+$ be such that for all $i \in [m]$

$$\operatorname{Tr}(\rho B_i) < 0, \quad \operatorname{Tr}(-A\rho) < 0 \quad \textit{and} \quad \operatorname{Tr}(\rho) = 1.$$

Set

$$\epsilon = \frac{1}{2}\min\left\{ \left|\frac{\operatorname{Tr}(\rho B_1)}{\|B_1\|_1}\right|, \dots, \left|\frac{\operatorname{Tr}(\rho B_m)}{\|B_m\|_1}\right|, \left|\frac{\operatorname{Tr}(\rho A)}{\|A\|_1}\right| \right\}$$

and take $S \in \mathcal{M}_{d,D}$ to be an (ϵ, δ, k)-JLT. Here,

$$k \geq \operatorname{rank} A + \operatorname{rank}\rho + \sum_{i=1}^{m} \operatorname{rank} B_i.$$

Then

$$\sum_{i=1}^{m} c_i S B_i S^T - SAS^T \not\geq 0 \tag{7}$$

for all $c \in \mathbb{R}_+^m$, with probability at least $1 - \delta$.

Proof. It should first be noted that ρ exists and $\epsilon > 0$ by Lemma 4.1. The matrix ρ defines a hyperplane that strictly separates the set

$$E = \left\{ \sum_{i=1}^{m} c_i B_i - A \,\middle|\, c \in \mathbb{R}_+^m \right\}$$

and S_D^+. We will now show that $S\rho S^T$ strictly separates the sets

$$E_S = \left\{ \sum_{i=1}^{m} c_i S B_i S^T - SAS^T \,\middle|\, c \in \mathbb{R}_+^m \right\}$$

and S_d^+ with probability at least $1 - \delta$, from which the claim follows. Note that by our choice of ρ and ϵ, it follows from Lemma 3.1 that we have

$$\operatorname{Tr}\left(S\rho S^T S B_i S^T\right) \leq \operatorname{Tr}(\rho B_i) + \epsilon \|B_i\|_1 < 0$$

with probability at least $1 - \delta$ and similarly for $-A$ instead of B_i. Therefore, it follows that $\operatorname{Tr}(Z S\rho S^T) < 0$ for all $Z \in E_S$. As $S\rho S^T$ is a positive semidefinite matrix, it follows that $\operatorname{Tr}\left(Y S\rho S^T\right) \geq 0$ for all $Y \in \mathcal{S}_d^+$. We have therefore found a strictly separating hyperplane for E_S and S_D^+ and the LMI (7) is infeasible. □

Theorem 4.2 suggests a way of sketching feasibility problems of the form

$$\sum_{i=1}^{m} c_i B_i - A \geq 0, \qquad c \in \mathbb{R}_+^m. \tag{8}$$

If we are interested in the case in which it is infeasible, we investigate if the LMI

$$\sum_{i=1}^{m} c_i S B_i S^T - SAS^T \geq 0, \qquad c \in \mathbb{R}_+^m \tag{9}$$

is feasible, for S a suitably chosen JLT as in Theorem 4.2. If Equation (9) is infeasible, we know that Equation (8) is infeasible, as any choice of a feasible c leads to a feasible c for Equation (9). Moreover, it follows from Theorem 4.2 that if the cone spanned by the B_i is well-behaved enough and the JLT is suitably chosen, it only happens with very low probability that Equation (9) is feasible and Equation (8) is not. To obtain more concrete bounds on the probability that the original problem is feasible, one would need to know the parameter ϵ, which is not possible in most applications. We emphasize that this is a black box algorithm. That is, we can decide whether a large instance of an LMI is infeasible by showing that a smaller instance of an LMI is infeasible. In Section 6 we will discuss the implications for the complexity and memory usage of the last theorems.

5. Approximating the value of semidefinite programs through sketching

We will now show how to approximate with high probability the value of a sketchable SDP by conjugating the target matrix and the matrices that describe the constraints with JLTs and subsequently solving a smaller instance of an SDP. The next theorem shows that in general it is not possible to approximate the value of a sketchable SDP using linear sketches with high probability and that we need to make further assumptions on the problem to achieve a non-trivial compression using linear maps.

Theorem 5.1. *Let $\Phi : \mathcal{M}_{2D} \to \mathbb{R}^d$ be a random linear map such that for all sketchable SDPs there exists an algorithm which allows us to estimate the value of an SDP up to a constant factor $1 \leq \tau < \frac{2}{\sqrt{3}}$ given the sketch $\{\Phi(A), \Phi(B_1), \ldots, \Phi(B_m)\}$ with probability at least $9/10$. Then $d = \Omega(D^2)$.*

Proof. By the duality relations for Schatten p-norms, it is easy to see that the value of the SDP

$$\begin{array}{ll} \text{maximize} & \operatorname{Tr}(AX) \\ \text{subject to} & \operatorname{Tr}(X) \leq 1 \\ & X \geq 0 \end{array} \tag{10}$$

with

$$A = \begin{pmatrix} 0 & G \\ G^T & 0 \end{pmatrix}$$

is twice the operator norm of the matrix $G \in \mathcal{M}_D$. In [Woo14, Theorem 6.5] it was shown that any algorithm that estimates the operator norm of a matrix from a linear sketch with probability larger than $9/10$ must have sketch dimension $\Omega(D^2)$. As the sketch $\{ \Phi(A), \Phi(\mathbb{1}) \}$ would thus allow to sketch the operator norm, the assertion follows. □

The above result remains true even if we restrict to SDPs that have optimal points with small Schatten 1-norm and low rank. This follows from the fact that the SDP given in Equation (10) has an optimal solution with rank 1 and trace 1.

We may even restrict to SDPs whose value scales sublinearly. To see that, notice that to show that the operator norm cannot be sketched, [Woo14] constructs two families of random matrices whose operator norm is of order $\sqrt{D}$ (cf. [Woo14, Lemma 6.3]). As the class of SDPs considered here covers this problem, we obtain the claim.

As we will see soon, the main hurdle to sketch SDPs and thus overcome the last no-go theorem is that we also need to suppose that the matrices that define the constraints and the target function have a small Schatten 1-norm, not only one optimal solution. To this end, we define:

Definition 5.2 (Sketched SDP)**.** *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{\text{sym}}$, $\eta, \gamma_1, \ldots, \gamma_m \in \mathbb{R}$ and $\epsilon > 0$. Given that an optimal point $X^* \in \mathcal{S}_D^+$ of the sketchable SDP defined through these matrices satisfies $\operatorname{Tr}(X^*) \leq \eta$ and given a random matrix $S \in \mathcal{M}_{d,D}$, we call the optimization problem*

$$\begin{array}{lll} \text{maximize} & \operatorname{Tr}\left(SAS^TY\right) & \\ \text{subject to} & \operatorname{Tr}\left(SB_iS^TY\right) \leq \gamma_i + \mu\|B_i\|_1, & i \in [m] \\ & Y \geq 0 & \end{array} \tag{11}$$

with $\mu = 3\epsilon\eta$ the sketched SDP.

The motivation for defining the sketched SDP is given by the following theorem.

Theorem 5.3. *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{\text{sym}}$, $\eta, \gamma_1, \ldots, \gamma_m \in \mathbb{R}$ and $\epsilon > 0$. Denote by α the value of the sketchable SDP and assume it is attained at an optimal point X^* which satisfies $\operatorname{Tr}(X^*) \leq \eta$. Moreover, let $S \in \mathcal{M}_{d,D}$ be an (ϵ, δ, k)-JLT, with*

$$k \geq \operatorname{rank} X^* + \operatorname{rank} A + \sum_{i=1}^{m} \operatorname{rank} B_i.$$

Let α_S be the value of the sketched SDP defined by A, B_i and S. Then

$$\alpha_S + 3\epsilon\eta\|A\|_1 \geq \alpha$$

with probability at least $1 - \delta$.

Proof. It follows from Lemma 3.1 that SX^*S^T is a feasible point of the sketched SDP with probability at least $1 - \delta$. Again by Lemma 3.1, we have that

$$\operatorname{Tr}\left(SAS^TSX^*S^T\right) \geq \operatorname{Tr}(AX^*) - 3\epsilon\eta\|A\|_1.$$

It then follows that $\alpha_S + 3\epsilon\eta\|A\|_1 \geq \alpha$. □

In Section 6, we will discuss the implications for memory usage and complexity of approximating the value of an SDP.

Note that the optimal value of an SDP is not necessarily attained. We could also demand X^* to be only close to optimality which would slightly increase the error made by the sketch. Since this makes notation more cumbersome, we assume the existence of such an optimal point.

Although it is not customary to assume a bound on the Schatten 1-norm of an optimal solution to an SDP, it is common to assume that for example the solution lies in a given ellipsoid when using the ellipsoid method to solve SDPs [LA16]. From such assumptions it is straightforward to derive bounds on the HS norm of the solution. If we are also given that an optimal solution to the SDP is low rank, the HS norm gives a good upper bound on the Schatten 1-norm, as remarked after the proof of Lemma 3.1. Solutions of low rank of SDPs have been extensively studied over the past years and there are many results available in the literature which guarantee that the optimal solution to an SDP has low rank. In general, it has been shown [Bar95, Pat98] that if we have m constraints and the SDP is feasible, there is an optimal solution with rank at most $r = \lfloor \frac{\sqrt{8m+1}-1}{2} \rfloor$.

Notice that Theorem 5.3 does not rule out the possibility that the value of the sketched problem is much larger than that of the sketchable SDP. To investigate this issue, we introduce the following:

Definition 5.4 (Relaxed SDP). *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{\text{sym}}$, $\eta, \gamma_1, \ldots, \gamma_m \in \mathbb{R}$ and $\epsilon > 0$. Given that an optimal point X^* of the sketchable SDP defined through these matrices satisfies $\operatorname{Tr}(X^*) \leq \eta$, we call the optimization problem*

$$\begin{array}{lll} \text{maximize} & \operatorname{Tr}(AX) & \\ \text{subject to} & \operatorname{Tr}(B_iX) \leq \gamma_i + \mu\|B_i\|_1, & i \in [m] \\ & X \geq 0 & \end{array} \tag{12}$$

with $\mu = 3\epsilon\eta$ the relaxed SDP.

Notice that, given a feasible point Y to the sketched SDP, S^TYS is a feasible point for the relaxed problem by the cyclicity of the trace. It follows that if the values of the original and the relaxed are close, the values of the original and the sketched problem are close as well. We formalize this intuition and prove the following bound in Appendix C.

Theorem 5.5. *We are in the setting of Definition 2.6. Assume that there exists an $X_0 > 0$ such that all the constraints of the sketchable SDP are strictly satisfied and that the dual problem is feasible. Then the value of the sketched SDP α_S is bounded by*

$$\alpha_S \leq \alpha(0) + \epsilon C_1 \left(\alpha(0) - \operatorname{Tr}(AX_0)\right) / C_2.$$

Here

$$C_1 = \max\left\{\, 3\eta\|B_i\|_1 \mid i \in [m] \,\right\},$$
$$C_2 = \min\left\{\, (\gamma_i - \mathrm{Tr}\,(B_i X_0)) \mid i \in [m] \,\right\},$$

where $\eta \geq \mathrm{Tr}\,(X^*)$ for an optimal point X^* of the sketchable SDP.

Proof. We refer to Appendix C for the proof. □

Note that this statement is not probabilistic and holds regardless of the choice of the sketching map. It could also be the case that $\mathrm{Tr}\,(AX_0)$ itself gives a better lower bound on the value than the one in Theorem 5.5. One can therefore say that Theorem 5.5 guarantees that in general the value of the sketched SDP cannot differ significantly from the value of the original one if we have feasible points which are not too close to the boundary. Combining Theorem 5.3 and Theorem 5.5 it is possible to pick ϵ small enough to have an arbitrarily small additive error under some structural assumptions on the SDP. That is, we need bounds on the Schatten 1-norms of A and B_i, be given a strictly feasible point for the relaxed SDP and a bound on the Schatten 1-norm of an optimal solution to the sketchable SDP.

In the case that all the $\gamma_i > 0$ for a sketchable SDP we may obtain a bound on the value and an approximate solution to it in a much simpler way. This class includes the so-called semidefinite packing problems [IPS05], where we have in addition that $B_i \geq 0$. Note that we may set all $\gamma_i = 1$ w.l.o.g. by dividing B_i by γ_i. We then obtain:

Theorem 5.6. *For a sketchable SDP with $\gamma_i = 1$ and $\kappa = \max\limits_{i\in[m]} \|B_i\|_1$, we have that*

$$\frac{\alpha_S}{1+\nu} \leq \alpha, \tag{13}$$

where $\nu = 3\epsilon\eta\kappa$. Moreover, denoting by X_S^* an optimal point of the sketched SDP, we have that $\frac{1}{1+\nu} S^T X_S^* S$ is a feasible point of the sketchable SDP that attains this lower bound. Furthermore, if $\|B_i\|_1 = \kappa$ for all $i \in [m]$ it can be checked if this lower bound is the optimal value.

Proof. The lower bound in Equation (13) follows immediately from the cyclicity of the trace, as $\frac{1}{1+\nu} S^T X_S^* S$ is a feasible point of the sketchable SDP. Given an optimal solution $\{c_i\}_{i\in[m]}$ to the dual of the sketched SDP and that $\|B_i\|_1 = \kappa$ for all $i \in [m]$, it is possible to check if the lower bound given by the sketched SDP is indeed optimal as follows. Slater's condition holds for the sketched and sketchable SDP, as $\vartheta\mathbb{1}$ is a strictly feasible point for $\vartheta > 0$ small enough. If we have

$$\sum_{i=1}^{m} c_i B_i - A \geq 0,$$

then the obtained feasible point is indeed optimal for the sketchable SDP by strong duality. □

It is possible to relax the condition that all the Schatten 1-norms of the matrices that define the constraints are the same and still obtain a lower bound for which it can be checked whether it is indeed optimal. To achieve this, it is necessary to modify all the constraints of the sketched SDP, as in Equation (11), to $1+\mu\kappa$ instead of $1+\mu\|B_i\|_1$. In the primal picture, we will find an optimal point of the sketchable SDP through the sketched SDP whenever there is a $Y \in \mathcal{S}_D^+$ such that $S^T Y S = X^*$, for X^* an optimal point of the sketchable SDP. Note that for semidefinite packing problems it is possible to derive a bound on the Schatten 1-norm of an optimal solution in a straightforward way.

Lemma 5.7. *Let $B_1,\ldots,B_m \in \mathcal{M}_D$ be positive semidefinite matrices such that the smallest strictly positive eigenvalue of $\sum_{i=1}^m B_i$ is given by λ. Then for a sketchable SDP with constraints B_i, $\gamma_1,\ldots,\gamma_m \in \mathbb{R}_+$, $A \geq 0$ and finite value there exists an optimal point X^* such that*

$$\operatorname{Tr}(X^*) \leq \frac{1}{\lambda}\sum_{i=1}^m \gamma_i. \tag{14}$$

Proof. As we assume that the SDP has a finite value, we may restrict to solutions whose support is contained in the support of $\sum_{i=1}^m B_i$. Denote by P the projection onto the support of $\sum_{i=1}^m B_i$. We then have

$$\lambda P \leq \sum_{i=1}^m B_i. \tag{15}$$

Conjugating both sides with $(X^*)^{\frac{1}{2}}$ and taking the trace we obtain

$$\lambda \operatorname{Tr}(X^*) \leq \sum_{i=1}^m \operatorname{Tr}(X^* B_i),$$

as we supposed w.l.o.g. that the support of X^* is contained in the support of $\sum_{i=1}^m B_i$. As X^* is a feasible point, we have $\operatorname{Tr}(X^* B_i) \leq \gamma_i$ and we obtain the claim. □

Of all the assumptions we needed for the results of Theorem 5.3, the bound on the Schatten 1-norm of an optimal solution to the SDP is arguably the most difficult to show, as bounds of this form are not readily available in the literature. Moreover, some SDPs have $\operatorname{Tr}(X) \leq \eta$ as constraint and would be natural candidates to apply these methods to. As $\operatorname{Tr}(\mathbb{1}) = D$ we will not be able to obtain any non-trivial compression with the scheme discussed so far. However, if one is only interested in obtaining an upper bound on the value of the SDP, it is still possible to have $\mathbb{1}$ as a constraint or as the target matrix and achieve a non-trivial compression.

Theorem 5.8. *Let $A, B_1,\ldots,B_{m-1} \in \mathcal{M}_D^{\text{sym}}$, $B_m = \mathbb{1}$. Further, let $\gamma_1,\ldots,\gamma_{m-1} \in \mathbb{R}$, $\gamma_m = \eta$ and $\epsilon > 0$. Denote the value of the sketchable SDP by α and assume it is attained* *at an optimal point X^*. Moreover, let $S \in \mathcal{M}_{d,D}$ be an (ϵ, δ, k)-JLT, with*

$$k \geq \operatorname{rank} X^* + \operatorname{rank} A + \sum_{i=1}^{m} \operatorname{rank} B_i.$$

Let α'_S be the value of the modifed sketched SDP defined by A, B_i and S, given by

$$\begin{array}{ll} \text{maximize} & \operatorname{Tr}\left(SAS^T X\right) \\ \text{subject to} & \operatorname{Tr}\left(SB_i S^T X\right) \leq \gamma_i + \mu \|B_i\|_1, \quad i \in [m-1] \\ & \operatorname{Tr}(X) \leq (1+\epsilon)\eta, \\ & X \geq 0 \end{array} \tag{16}$$

with $\mu = 3\epsilon\eta$. Then

$$\alpha'_S + 3\epsilon\eta \|A\|_1 \geq \alpha$$

with probability at least $1 - \delta$.

Proof. The proof is essentially the same as the one of Theorem 5.3, as we have that $|\operatorname{Tr}\left(SX^*S^T\right) - \operatorname{Tr}(X^*)| \leq 3\epsilon \operatorname{Tr}(X^*)$ and so SX^*S^T is a feasible point with probability at least $1 - \delta$. □

The main difference between the modified sketched SDP and the sketched SDP is that here we do not conjugate the identity with S, only the other constraints. With this, we do not have that $S^T X^* S$ is a feasible point of the relaxed SDP, but we do not need the assumption $\operatorname{Tr}(X^*) \leq \eta$ to obtain an upper bound. It should be clear from Theorem 5.8 that we may also optimize over the trace, i.e. $A = \mathbb{1}$, without conjugating with S and still have an upper bound and non-trivial compression.

We may therefore summarize the results of this section as follows. If we want to obtain an upper bound on the value of the sketchable SDP with our techniques, it is necessary to have upper bounds on the Schatten 1-norm of all the matrices that define the constraints, the target matrix and of an optimal solution. We may then choose a JLT of suitable dimension to solve the sketched SDP, whose value will allow us to infer an upper bound to the original problem with high probability. If we are addtionaly given a strictly feasible point of the sketchable SDP or if we are solving a semidefinite packing problem, we also obtain a lower bound on the value of the sketchable SDP in terms of the sketched one. In the case of semidefinite packing problems, we even obtain a feasible point of the sketchable SDP whose value is close to the sketched value. If we are not given a bound on the Schatten 1-norm of an optimal solution, we may impose it as a constraint as in Theorem 5.8 and obtain an upper bound on the value of the sketchable SDP constrained to points which have their Schatten 1-norm bounded by η. Although we are able to drop the assumption on the Schatten 1-norm of an optimal solution, we are not able to prove that this upper bound cannot differ significantly from the true value in this case.

6. Complexity and memory gains

In this section, we will discuss how much we gain by considering the sketched SDP instead of the sketchable SDP. We focus on the results of Section 5, but the discussion carries over to the results of Section 4. Throughout this section we will assume that we are guaranteed that the Schatten $1-$norm of an optimal solution to our SDP and of the matrices that define the constraints is $\mathcal{O}(1)$. It is therefore Theorem 5.3 for which we need a sketch of appropriate size. The theorem states that we need an (ϵ, δ, k)-JLT for the upper bound on α to hold with probability at least $1-\delta$. As stated in Theorem 5.3, we can choose a sketching matrix $S \in \mathcal{M}_{d,D}$ with $d = \mathcal{O}(\epsilon^{-2}\log\bigl(k\delta^{-1}\bigr))$ and $s = \mathcal{O}(\epsilon^{-1}\log\bigl(k\delta^{-1}\bigr))$ nonzero entries per column. Here k is as in Theorem 5.3. The cost of generating S is at most $\mathcal{O}(dD)$, which will be of smaller order than the necessary matrix multiplications. We will therefore not take this cost into account for the rest of the analysis.

One could argue that one needs to know the Schatten 1-norm of the different matrices that define our constraints for estimating the value or obtaining more concrete bounds for the feasibility problems. We will, however, suppose that an upper bound on the Schatten 1-norm of an optimal solution and the constraints is given or that this can be computed in a time which is $\mathcal{O}(D^2)$. This is the case if for example we have a semidefinite packing problem, where the matrices are positive semidefinite and we can compute their Schatten 1-norm in $\mathcal{O}(D)$ time.

To generate the sketched SDP, we need to compute $m+1$ matrices of the form SBS^T, where $B \in \mathcal{M}_D$. Each of this computations needs $\mathcal{O}(\max\{\, \mathrm{nnz}(B), Dd \,\}\,\epsilon^{-1}\log\bigl(k\delta^{-1}\bigr))$ operations. In the worst case, when all matrices $\{\, A, B_1, \dots, B_m \,\}$ are dense and have full rank, this becomes $\mathcal{O}(mD^2\log(mD))$ operations to generate the sketched SDP for fixed ϵ and δ.

Let us collect these considerations in a proposition:

Proposition 6.1. *Let $A, B_1, \dots, B_m \in \mathcal{M}_D^{\mathrm{sym}}$, $\gamma_1, \dots, \gamma_m \in \mathbb{R}$ of a sketchable SDP be given. Furthermore, let $z := \max\{\, \mathrm{nnz}(A), \mathrm{nnz}(B_1), \dots, \mathrm{nnz}(B_m) \,\}$ and $\mathrm{SDP}(m,d,\zeta)$ be the complexity of solving a sketchable SDP (up to accuracy ζ) of dimension d. Then a number of*

$$\mathcal{O}(\max\bigl\{\, z, D\epsilon^{-2}\log\bigl(k\delta^{-1}\bigr) \,\bigr\}\,\epsilon^{-1}m\log\bigl(k\delta^{-1}\bigr) + \mathrm{SDP}(m, \epsilon^{-2}\log\bigl(k\delta^{-1}\bigr), \zeta))$$

operations is needed to generate and solve the sketched SDP, where k is defined as in Theorem 5.3.

It is easy to see that we can parallelize computing the matrices SB_iS^T. Typically, the costs of forming the sketched matrices SB_iS^T dominates the overall complexity. For example, using the ellipsoid method [GS88, Chapter 3], the complexity of solving an SDP becomes $\mathrm{SDP}(m,D,\zeta) = \mathcal{O}(\max\bigl\{\, m, D^2 \,\bigr\}\, D^6\log(1/\zeta))$ (cf. [Bub15, p.250]). Assuming that ϵ, δ and ζ are fixed, we need $\mathcal{O}(\max\bigl\{\, m, D^2 \,\bigr\}\, D^6)$ operations to solve the sketchable SDP, compared to $\mathcal{O}(mD^2\log(mD))$ operations to obtain an approximate solution via first forming the sketched problem and then solving it. Admittedly, the ellipsoid method is not used in practice, but using interior point methods, we still need $\mathrm{SDP}(m, D, \zeta) = \mathcal{O}(\max\{\, m^3, D^2m^2, mD^\omega \,\} D^{0.5} \log(D/\zeta))$ operations [dK02, Chapter 5], where ω is the exponent of matrix multiplication. The best known algorithms achieve $\omega \approx 2.37$ [LG14]. If the SDP can be sketched, doing so gives a speedup as long as the complexity of solving the SDP directly is $\Omega(mD^{2+\mu})$, where $\mu > 0$.

A great advantage is that for the sketched problem, we only need to store $m + 1$ matrices of size $d \times d$ instead of $D \times D$. We collect this in a proposition.

Proposition 6.2. *Let $A, B_1, \ldots, B_m \in \mathcal{M}_D^{\mathrm{sym}}$, $\gamma_1, \ldots, \gamma_m \in \mathbb{R}$ be a sketchable SDP. Then we need only store $\mathcal{O}(m\epsilon^{-4} \log(mk/\delta)^2)$ entries for the sketched problem, where k is defined as in Theorem 5.3.*

7. Applications and numerical examples

7.1. Estimating the value of a semidefinite packing problem

Inspired by [Sag11] we will test our techniques on an SDP stemming from the field of optimal design of experiments. The problem is the following: an experimenter wishes to estimate the quantity $\langle c, \theta \rangle$, where $\theta \in \mathbb{R}^D$ is an unknown D-dimensional parameter and $c \in \mathbb{R}^D$ is given. To this end, one is given linear measurements of the parameter $y_i = A_i\theta$, up to a (centered) measurement noise for $A_i \in \mathcal{M}_D$. We refer to [Puk06] for more details on the topic. To find the amount of effort to spend on the $i-$th experiment to minimize the variance is given by the SDP

$$\begin{aligned} \text{maximize} \quad & \operatorname{Tr}\left(cc^T X\right) \\ \text{subject to} \quad & \operatorname{Tr}\left(M_i X\right) \leq 1, \qquad i \in [m] \\ & X \geq 0 \end{aligned} \tag{17}$$

with $M_i = A_i^T A_i$. This problem always admits optimal solutions of rank 1 [Sag11]. We generated random instances of this SDP in the following way:

1. We sampled four matrices A_i distributed as follows: the first three rows of A_i are sampled independently from the uniform distribution on the unit sphere in $\mathbb{R}^D$. The other $D - 3$ rows of A_i are set to 0.

2. Given the A_i, we generate c by getting four samples k_1, k_2, k_3, k_4 from the uniform distribution on $\{\, 1, 2, 3 \,\}$ and four samples x_1, x_2, x_3, x_4 from the standard normal distribution. c is then given by $\sum\limits_{i=1}^{4} x_i (A_i)_{k_i}$, where $(A_i)_{k_i}$ is the k_i-th row of A_i.

This gives matrices M_i of rank 3 almost surely and Schatten 1-norm equal to 3. The fact that c is a linear combination of the rows of A_i ensures that the problem is bounded, as can be easily seen by looking at the dual problem [Sag11]. Note that this is a semidefinite packing problem, so we are able to use the results of Theorem 5.6 to obtain a lower bound. There exists an optimal solution whose Schatten 1-norm is bounded by 8 with very high

D	d	Value	Error L.B.	Error U.B.	M.R.T. Sketchable [s]	M.R.T Sketch [s]
100	10	2.52	0.0880	0.156	1.27	0.324
100	20	2.50	0.00	0.250	1.17	0.305
200	20	2.69	0.00	0.269	6.50	0.299
200	40	2.53	0.00	0.00102	6.82	0.375
500	50	2.55	0.00	0.255	98.0	0.453
500	100	2.66	0.00	0.266	97.6	1.23
700	70	2.57	0.00	0.257	557	1.38
700	140	2.49	0.00	0.249	548	3.53

Table 1: For each combination of the sketchable dimension (D) and dimension of the sketch (d) we have generated 40 instances of the SDP in Equation (17). Here "M.R.T." stands for mean running time, "L.B." stands for lower bound and "U.B." for upper bound and each column shows the mean of the sample. The column "Value" stands for the optimal value of the sketchable SDP.

probability. This easily follows from Theorem 5.7 and the fact that i.i.d. unit vectors are almost orthogonal. To obtain the upper bound we have used the results of Theorem 5.3 and for the lower bound Theorem 5.6. We have chosen $\eta = 0.1$ to generate these results. We used sparse JLTs with sparsity parameter $s = 1$ [KN14] to obtain faster matrix multiplications to form the sketches. We define the error of the lower bound to be given by $\alpha - \frac{1}{1+\eta}\alpha_S$ and of the upper to be $\alpha_S - \alpha$. To solve the SDP given in Equation (17) we used cvx, a package for specifying and solving convex programs [GB14, GB08].

As we can see, the results of Table 1 show that, excluding the case where the sketching dimension was 10, we were able to find feasible points which were numerically indistinguishable from being optimal by using our sketching methods. Moreover, the time needed to find an optimal solution was smaller by 1 or 2 orders of magnitude.

7.2. Linear matrix inequality feasibility problems

We will now apply our techniques to an LMI feasibility problem. Let $G_i \in \mathcal{M}_{d'}$, $i \in [m]$, be random matrices sampled independently from the Gaussian unitary ensemble (GUE). See for example [AGZ10, Section 2.2] or Section D for their definition. Consider the rescaled and shifted matrices $\tilde{G}_i(\alpha) = \frac{1}{\sqrt{d'}}G_i + \alpha\mathbb{1}$ for some $\alpha \in \mathbb{R}$. For $V : \mathcal{M}_{d'} \to \mathcal{M}_D$ a random isometry, we will test our techniques on the feasibility of the LMI

$$\sum_{i=1}^{m} t_i V\tilde{G}_i(\alpha)V^* - VV^* \geq 0, \quad t \in \mathbb{R}_+^m \tag{18}$$

depending on α. As V is an isometry, Equation (18) is clearly feasible if and only if the LMI

$$\sum_{i=1}^{m} t_i\tilde{G}_i(\alpha) - \mathbb{1} \geq 0, \quad t \in \mathbb{R}_+^m \tag{19}$$

is feasible. Using standard techniques from random matrix theory, we show that Equation (19), and so Equation (18), is feasible with high probability for $\alpha > \frac{2}{\sqrt{m}}$ and infeasible with high probability for $\alpha < \frac{2}{\sqrt{m}}$ in case $m \ll d'$.

We refer to Appendix Section D for a proof of this claim. This therefore allows us to quantify "how close to feasible" the LMI inequality is in terms of how close α is to $\frac{2}{\sqrt{m}}$ and to know whether the LMI was feasible or not. We will choose $d' \ll D$, as this way we avoid having a Schatten 1-norm of the matrices that define the LMI which is of the same order as the dimension. The technique we used to solve Equation (18) is the same as discussed in Section 4. That is, we will check for the feasibility of

$$\sum_{i=1}^{m} t_i S V \tilde{G}_i(\alpha) V^* S^T - S V V^* S^T \geq 0, \quad t \in \mathbb{R}_+^m \tag{20}$$

for a complex Gaussian JLT S. That is, $S = \frac{1}{\sqrt{2}}(S_1 + iS_2)$ with S_1 and S_2 independent Gaussian JLTs. We refer to Section A for a proof that this choice of random matrices indeed gives a JLT with the same scaling of the parameters as the real one. We refer to Theorem D.4 for a proof that the LMI defined in Equation (20) satisfies the assumptions of Theorem 4.2 with high probability. To solve the SDP given in Equation (20) we used cvx, a package for specifying and solving convex programs [GB14, GB08]. The results are summarized in Table 2. We can observe that by using our methods we were able to show that the LMI is infeasibile in a much smaller running time or even show that certain LMI were infeasible when a direct computation was not possible due to memory constraints in most choices of the parameters. In some cases it was, however, necessary to increase the sketch dimension to show that the inequality was infeasibile.

8. Conclusion

We have shown how to obtain sketches of the HS product using positive maps obtained from JLTs, how to apply these to show that certain LMI are infeasible and to obtain approximations of the value of certain SDPs. In some cases, these techniques can lead to significant improvements in the runtime necessary to solve the instances of the SDPs and significant gains in the memory needed to solve them. However, the class of problems to which these techniques can be applied is significantly restricted by the fact that the matrices that define the constraints of the problems and a solution must have Schatten 1-norms which do not scale with the dimension for them to be advantageous. Moreover, the no-go theorems proved here show that one cannot significantly improve our results using positive linear maps to sketch the HS norm or to approximate the value of SDPs.

Acknowledgements

We would like to thank Ion Nechita for helpful discussions. A.B. acknowledges support from the ISAM Graduate Center at Technical University of Munich. D.S.F. acknowl-

D	d	d'	α	M.R.T. Original [s]	M.R.T. Sketch [s]	Error Rate
200	50	100	0.3	452	4.65	0
200	50	100	0.4	407	4.31	0
200	50	100	0.5	383	5.71	0
200	50	100	0.6	407	6.46	0
200	100	100	0.3	449	66.3	0
200	100	100	0.4	344	86.6	0
200	100	100	0.5	393	102	0
200	100	100	0.6	457	91.4	0
400	50	200	0.3	-	4.02	0
400	50	200	0.4	-	7.04	0
400	50	200	0.5	-	3.39	0.975
400	50	200	0.6	-	2.35	1.0
400	100	200	0.3	-	114	0
400	100	200	0.4	-	122	0
400	100	200	0.5	-	118	0
400	100	200	0.6	-	115	0

Table 2: For each combination of the dimension of the image of the random isometry (D), dimension of the domain of the random isometry (d'), dimension of the sketch (d) and α we have generated 40 instances of the random LMI in Equation (18) with $m = 9$. Here "M.R.T." stands for mean running time. The error rate gives the ratio of infeasible problems that were not detected to be infeasible by sketching. A dash in the running time means that we were not able to solve the LMI because we ran out of memory.

edges support from the graduate program TopMath of the Elite Network of Bavaria, the TopMath Graduate Center of TUM Graduate School at Technical University of Munich. D.S.F. is supported by the Technical University of Munich – Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763.

References

[AGZ10] G. W. Anderson, A. Guionnet, and O. Zeitouni. *An Introduction to Random Matrices*. Cambridge University Press, 2010.

[Bar95] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(2):189–202, 1995.

[BEFB94] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, June 1994.

[Bha97] R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer, 1997.

[Bub15] S. Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3-4):231–357, November 2015.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[dK02] E. de Klerk. *Aspects of Semidefinite Programming: Interior Point Algorithms and Selected Applications*. Applied Optimization. Springer US, 2002.

[GB08] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer, 2008. `http://stanford.edu/~boyd/graph_dcp.html`.

[GB14] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014.

[GS88] Lovász L. Grötschel, M. and A. Schrijver. *Geometric algorithms and combinatorial optimization*. Springer, 1988.

[IPS05] G. Iyengar, D. J. Phillips, and C. Stein. *Approximation Algorithms for Semidefinite Packing Problems with Applications to Maxcut and Graph Coloring*, pages 152–166. Springer, 2005.

[KN14] D. M. Kane and J. Nelson. Sparser Johnson-Lindenstrauss transforms. *J. ACM*, 61(1):4:1–4:23, January 2014.

[LA16] J.B. Lasserre and M.F. Anjos. *Handbook on Semidefinite, Conic and Polynomial Optimization.* International Series in Operations Research & Management Science Series. Springer, 2016.

[LG14] F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC '14, pages 296–303. ACM, 2014.

[LR10] M. Ledoux and B. Rider. Small deviations for beta ensembles. *Electronic Journal of Probability*, 15(41):1319 – 1343, 2010.

[Pat98] G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of Operations Research*, 23(2):339–358, 1998.

[Puk06] F. Pukelsheim. *Optimal Design of Experiments.* Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 2006.

[RV13] M. Rudelson and R. Vershynin. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Journal of Probability*, 18(82):1–9, 2013.

[Sag11] G. Sagnol. A class of semidefinite programs with rank-one solutions. *Linear Algebra and its Applications*, 435(6):1446 – 1463, 2011.

[SH15] C. J. Stark and A. W. Harrow. Compressibility of positive semidefinite factorizations and quantum models. In *ISIT*, pages 2777–2781. IEEE, 2015.

[Stø13] E. Størmer. *Positive Linear Maps of Operator Algebras.* Springer Monographs in Mathematics. Springer, 2013.

[Ver12] R. Vershynin. *Compressed sensing*, chapter Introduction to the non-asymptotic analysis of random matrices, pages 210–268. Cambridge University Press, 2012.

[VPL15] K. Vu, P.-L. Poirion, and L. Liberti. Using the Johnson-Lindenstrauss lemma in linear and integer programming. *ArXiv 1507.00990*, July 2015.

[WA02] H. Wolkowicz and M. F. Anjos. Semidefinite programming for discrete optimization and matrix completion problems. *Discrete Appl. Math.*, 123(1-3):513–577, November 2002.

[Wol12] M. M. Wolf. Quantum channels and operations: Guided tour. Lecture notes available at `http://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf`, 2012.

[Woo14] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1–2):1–157, 2014.

[YUAC17] A. Yurtsever, M. Udell, Tropp J. A., and V. Cevher. Sketchy Decisions: Convex Low-Rank Matrix Optimization with Optimal Storage. *ArXiv 1702.06838*, February 2017.

A. Complexifying Johnson-Lindenstrauss transforms

In this Appendix we will generalize some of the results we need concerning JLTs to complex vector spaces. We will see that, up to a constant, most of the statements that hold in the real case also hold in the complex case. We will consider matrices of the form $\frac{1}{\sqrt{2d}}(S+iT)$, with S, T independent and with i.i.d. entries that are sub-gaussian and show that they give us complex JLTs. Note that these constructions clearly give sparse JLTs if the real JLTs used are sparse.

Definition A.1 (Sub-Gaussian Distribution). *The probability distribution of a random variable X is called sub-gaussian if there exist $C, v > 0$ such that $\forall t > 0$*

$$\mathbb{P}(|X| > t) \leq Ce^{-vt^2}.$$

A random variable is sub-gaussian if and only if for $p \geq 1$, $E(|X|^p) = \mathcal{O}(p)^{\frac{p}{2}}$ and the sub-gaussian norm of X is defined as:

$$\|X\|_{\psi_2} = \sup_{p \geq 1} p^{-\frac{1}{2}} (E(|X|^p))^{\frac{1}{p}}. \tag{21}$$

See for example [Ver12, Section 5.2.3] for more details on this. The main ingredient to show how to generalize JLTs to the complex case will be the following theorem.

Theorem A.2 ([RV13, Theorem 2.1]). *Let $A \in \mathcal{M}_{d,D}(\mathbb{R})$ be fixed. Consider a random vector $X = (X_1, \ldots, X_D)$, where X_i are independent random variables satisfying $\mathbb{E}[X_i] = 0$, $\mathbb{E}[X_i^2] = 1$ and $\|X_i\|_{\psi_2} \leq K$. Then for any $t \geq 0$, we have*

$$\mathbb{P}\left[\left| \|AX\|_2 - \|A\|_{HS} \right| > t \right] \leq 2 \exp\left[-\frac{ct^2}{K^4 \|A\|_\infty^2} \right].$$

Using this, we have a different way of proving the Johnson-Lindenstrauss lemma which generalizes to the complex case. We follow the proof of [RV13, Theorem 3.1].

Lemma A.3. *Let $S, T \in \mathcal{M}_{d,D}(\mathbb{R})$ have independent sub-gaussian entries with $\mathbb{E}[X_{ij}] = 0$, $\mathbb{E}[X_{ij}^2] = 1$ and $\|X_{ij}\|_{\psi_2} \leq K$, for all $X \in \{ S, T \}$. Then for $d = \mathcal{O}\left(\epsilon^{-2} \ln(2/\delta)\right)$, we have*

$$\mathbb{P}\left[\left\| \frac{S+iT}{\sqrt{2d}} x \right\|_2 \in (1 \pm \epsilon) \|x\|_2 \right] \geq 1 - \delta$$

for any fixed $x \in \mathbb{C}^D$.

Proof. Define the linear operator $\Phi : \mathcal{M}_{d,D}(\mathbb{C}) \to \mathbb{C}^d$, $G \mapsto Gx$, where $x \in \mathbb{C}^d$ is a fixed vector. We use the standard isomorphisms $\mathcal{M}_{d,D}(\mathbb{C}) \simeq \mathbb{R}^{2dD}$ and $\mathbb{C}^d \simeq \mathbb{R}^{2d}$ and denote by $\tilde{\Phi}$ the map Φ composed with these isomorphisms, which is now a linear map from a real vector space to another real vector space. Moreover, observe that a matrix of the form $X + iY \in \mathcal{M}_{d,D}(\mathbb{C})$ with $X, Y \in \mathcal{M}_{d,D}(\mathbb{R})$ is mapped to (X, Y) under the isomorphism. The map $\tilde{\Phi}$ will play the role of A in the statement of Theorem A.2. It is straightforward to compute the norms involved in the statement, as explained in [RV13, Section 3.1]. We have

$$\left\|\tilde{\Phi}\right\|_2^2 = 2d\|x\|_2^2, \quad \left\|\tilde{\Phi}\right\|_\infty^2 = \|x\|_2^2, \quad \left\|\tilde{\Phi}((S,T)^T)\right\|_2^2 = 2\|(S+iT)x\|_2^2.$$

As S, T have sub-gaussian entries, the vector (S, T) satisfies the assumptions of Theorem A.2 and the statement follows. □

Unfortunately, the sparse JLTs discussed in Theorem 2.5 are not of this form. The entries of these JLTs are not independent from each other, one of the assumptions of Lemma A.3.

B. Proof of Theorem 3.2

Theorem 3.2. *Let $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ be a random positive map such that with positive probability for any $Y_1, \ldots, Y_{D+1} \in \mathcal{M}_D$ and $0 < \epsilon < \frac{1}{4}$ we have*

$$|\mathrm{Tr}\left(\Phi(Y_i)^T\Phi(Y_j)\right) - \mathrm{Tr}\left(Y_i^T Y_j\right)| \leq \epsilon\|Y_i\|_2\|Y_j\|_2. \tag{22}$$

Then $d = \Omega(D)$.

Proof. Let $\{e_i\}_{1\leq i\leq D}$ be an orthonormal basis of $\mathbb{C}^D$ and define $X_i = e_i e_i^T$. As Equation (22) is satisfied with positive probability, there must exist a positive map $\Phi : \mathcal{M}_D \to \mathcal{M}_d$ such that Equation (22) is satisfied for $Y_i = X_i$, $i \in [D]$, and $Y_{D+1} = \mathbb{1}$. As the X_i are orthonormal w.r.t. the Hilbert-Schmidt scalar product and by the positivity of Φ we have for $i, j \in [D]$

$$\mathrm{Tr}\left(\Phi(X_i)\Phi(X_j)\right) \in \begin{cases} [0, \epsilon], & \text{for } i \neq j \\ [1-\epsilon, 1+\epsilon], & \text{for } i = j. \end{cases} \tag{23}$$

Define the matrix $A \in \mathcal{M}_D$ with $(A)_{ij} = \mathrm{Tr}\left(\Phi(X_i)\Phi(X_j)\right)$ for $i, j \in [D]$. It is clear that A is Hermitian and that its entries are positive. We have

$$\sum_{i,j\in[D]} A_{ij} = \mathrm{Tr}\left(\Phi(\mathbb{1})\Phi(\mathbb{1})\right) \in \left[(1-\epsilon)D, (1+\epsilon)D\right].$$

As $A_{ii} \geq (1-\epsilon)$, it follows that

$$\sum_{i\neq j} A_{ij} \leq 2\epsilon D. \tag{24}$$

Let

$$J = \{ \, (i,j) \in [D] \times [D] \mid i \neq j, A_{ij} \leq \frac{1}{D} \, \} \, .$$

It follows from Equation (24) that $| \{ \, (i,j) \in [D] \times [D] \mid i \neq j, (i,j) \notin J \, \} | \leq 2D^2\epsilon$ and so

$$|J| \geq \big((1-2\epsilon)D^2 - D\big) \, .$$

Since for $(i,j) \in J$ also $(j,i) \in J$, we can write $J = (I \times I) \backslash \{(i,i) | i \in I\}$ for $I \subseteq [D]$. Thus, we infer for $D \geq 2$

$$|J| = |I|(|I|-1) \geq ((1-2\epsilon)D^2 - D) \geq \left(\frac{1}{2} - 2\epsilon\right) D^2.$$

From this it follows that

$$|I|^2 \geq |I|(|I|-1) \geq (\frac{1}{2} - 2\epsilon)D^2,$$

and we finally obtain

$$|I| \geq \sqrt{1/2 - 2\epsilon} D. \tag{25}$$

Notice that it follows from Equation (23) that we may rescale all the X_i to X_i' such that $\mathrm{Tr}\left(\Phi(X_i')^2\right) = 1$ and the pairwise scalar product still satisfies $\mathrm{Tr}\left(\Phi(X_i')\Phi(X_j')\right) \leq \frac{1}{D(1-\epsilon)}$ for $(i,j) \in J$. If there is an $N \in \mathbb{N}$ such that $d > \sqrt{1/2 - 2\epsilon} D$ for all $D \geq N$, the claim follows. We therefore now suppose that $d \leq \sqrt{1/2 - 2\epsilon} D$. Hence, $d \leq |I|$ by Equation (25). By the positivity of Φ and the fact that the X_i' are positive semidefinite, we have that $\Phi(X_i')$ is positive semidefinite. In [Wol12, Proposition 2.7] it is shown that for any set $\{P_i\}_{i \in I}$ of $|I| \geq d$ positive semidefinite matrices in $\mathcal{M}_d$ such that $\mathrm{Tr}\left(P_i^2\right) = 1$ we have that

$$\sum_{i \neq j} \mathrm{Tr}\left(P_i P_j\right)^2 \geq \frac{(|I| - d)^2 |I|}{(|I|-1)d^2}.$$

By the definition of the set J, we have that

$$\sum_{(i,j) \in J} \mathrm{Tr}\left(X_i' X_j'\right)^2 \leq \frac{|J|}{(1-\epsilon)^2 D^2} \leq \frac{1}{(1-\epsilon)^2},$$

as $|J| \leq D^2$. From Equation (25) it follows that

$$\frac{1}{(1-\epsilon)^2} \geq \left(\frac{\sqrt{1/2 - 2\epsilon} D}{d} - 1\right)^2$$

and after some elementary computations we finally obtain

$$d \geq \frac{(1-\epsilon)\sqrt{1/2 - 2\epsilon}}{2 - \epsilon} D.$$

□

C. Lower bound on the value of SDPs through sketching

We will obtain lower bounds on the value of the sketchable SDP in terms of the value of the sketched SDP through continuity bounds on the relaxed SDP. As the continuity bound we use is for SDPs given in equality form, we begin by giving an equivalent formulation of a sketchable SDP with equality constraints. The method of using duality to derive perturbation bounds on a convex optimization problem used here is standard and we refer to [BV04, Section 5.6] for a similar derivation. Given a sketchable SDP, define the maps $\Phi : \mathcal{M}_D \to \mathcal{M}_m$

$$\Phi(X) = \sum_{j=1}^{m} \operatorname{Tr}(B_j X)\, e_j e_j^T$$

for $\{ e_j \}_{j=1}^m$ an orthonormal basis of $\mathbb{R}^m$ and $\Psi : \mathcal{M}_{D+m} \to \mathcal{M}_m$

$$\Psi\left(\begin{bmatrix} X & * \\ * & Z \end{bmatrix}\right) = \Phi(X) + [Z_{jj}]_j.$$

With the help of the matrix

$$G = \sum_{j=1}^{m} \gamma_j e_j e_j^T, \tag{26}$$

the sketchable SDP can be written in equality form as

$$\begin{aligned} \text{maximize} \quad & \operatorname{Tr}\left(\begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} X & * \\ * & Z \end{bmatrix}\right) \\ \text{subject to} \quad & \Psi\left(\begin{bmatrix} X & * \\ * & Z \end{bmatrix}\right) = G \\ & \begin{bmatrix} X & * \\ * & Z \end{bmatrix} \geq 0, \end{aligned} \tag{27}$$

where $*$ are submatrices which are not relevant for our discussion. The dual problem of Equation (27) may be written as

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^{m} \gamma_j Y_{jj} \\ \text{subject to} \quad & \sum_{j=1}^{m} B_j Y_{jj} \geq A \\ & Y_{jj} \geq 0, \qquad \forall j \in [m]. \end{aligned} \tag{28}$$

Lemma C.1. *If there is an $X > 0$ such that the sketchable SDP is satisfied with strict inequality and the dual problem is feasible, then both the primal problem given in Equation* (27) *and dual in Equation* (28) *are feasible, there is no duality gap and there is a dual solution which attains the optimal value. Furthermore, this condition is equivalent to Slater's condition.*

Proof. If we can show equivalence with Slater's condition, the first statement follows automatically. Slater's condition for the primal problem is the following. If there is a $\begin{bmatrix} X & * \\ * & Z \end{bmatrix} > 0$ which satisfies the constraints, then the duality gap is zero and there is a Y which attains the optimal value. For such a block matrix, we also need $X > 0$, $Z > 0$. Hence, we can formulate Slater's condition with off-diagonal entries $* = 0$. We observe

$$\Psi\left(\begin{bmatrix} X & * \\ * & Z \end{bmatrix}\right) = G \Leftrightarrow \begin{pmatrix} \gamma_1 - \operatorname{Tr}(B_1 X) & & \\ & \ddots & \\ & & \gamma_m - \operatorname{Tr}(B_m X) \end{pmatrix} = Z > 0.$$

Hence $X > 0$ satisfies the constraints with strict inequality. The converse is clear. □

Now we can bound the optimal solution to Equation (11). We denote by $\mathcal{A}(\epsilon)$ the feasible set of the relaxed SDP as in Definition 5.4 for some $\epsilon > 0$. With this notation, $\mathcal{A}(0)$ is the feasible set for the primal problem. Analogously, we denote by $\alpha(\epsilon)$ the optimal value of the relaxed problem, by $\alpha(0)$ the optimal value of the sketchable SDP.

Lemma C.2. *Assume that there exists $X_0 \in \mathcal{A}(0)$ such that $X_0 > 0$ and the constraints are strictly satisfied. Then*

$$\alpha(0) \leq \alpha(\epsilon) \leq \alpha(0) + \langle \tilde{\epsilon}, y^* \rangle,$$

where y^* is an optimal solution to the dual problem to the sketchable SDP and $\tilde{\epsilon} \in \mathbb{R}^m$ with $\tilde{\epsilon}_i = 3\epsilon\eta\|B_i\|$.

Proof. The first inequality is obvious, since any $X \in \mathcal{A}(0)$ is also in $\mathcal{A}(\epsilon)$. By Lemma C.1, strong duality holds and there is a $y^* \geq 0$ which achieves the optimal value. Hence

$$\begin{aligned} \alpha(0) &= \sum_{j=1}^{m} y_j^* \gamma_j \\ &\geq \sum_{j=1}^{m} y_j^* \gamma_j - \operatorname{Tr}\left(\left[\sum_{i=1}^{m} y_i^* B_i - A\right] X\right), \qquad X \geq 0 \\ &= \operatorname{Tr}(AX) - \sum_{i=1}^{m} y_i^* \left[\operatorname{Tr}(B_i X) - \gamma_i\right]. \end{aligned} \tag{29}$$

The first line holds by duality. If we take the supremum over $X \in \mathcal{A}(\epsilon)$, we infer

$$\alpha(0) \geq \alpha(\epsilon) - \sum_{i=1}^{m} y_i^* \tilde{\epsilon}_i,$$

since $y_i^* \geq 0$. □

Corollary C.3. *Assume that there exists $X_0 \in \mathcal{A}(0)$ such that $X_0 > 0$ and the constraints are strictly satisfied. Then*

$$\alpha(\epsilon) \leq \alpha(0) + \left[\max_{i\in[m]} \tilde{\epsilon}_i\right] (\alpha(0) - \operatorname{Tr}(AX_0)) \,/ \left(\min_k (\gamma_k - \operatorname{Tr}(B_k X_0))\right),$$

where $\tilde{\epsilon} \in \mathbb{R}^m$ is defined as in Lemma C.2.

Proof. By Equation (29), we have that

$$\alpha(0) \geq \operatorname{Tr}(AX_0) - \sum_{i=1}^{m} y_i^* \left[\operatorname{Tr}(B_i X_0) - \gamma_i\right].$$

Since $[\operatorname{Tr}(B_i X_0) - \gamma_i] < 0$, it follows that

$$\sum_{i=1}^{m} y_i^* \leq (\alpha(0) - \operatorname{Tr}(AX_0)) / \min_{i\in[m]} \left[\gamma_i - \operatorname{Tr}(B_i X_0)\right].$$

With $\langle \tilde{\epsilon}, y^* \rangle \leq [\max_{i\in[m]} \epsilon_i] \sum_{i=1}^m y_i^*$ for $y^* \geq 0$, the Corollary follows. □

Theorem C.4. *Assume that there exists $X_0 \in \mathcal{A}(0)$ such that $X_0 > 0$ and the constraints are strictly satisfied. Then the value of the sketched SDP α_S is bounded by*

$$\alpha_S \leq \alpha(0) + \epsilon C_1 \left(\alpha(0) - \operatorname{Tr}(AX_0)\right) / C_2.$$

Here

$$\begin{aligned} C_1 &= \max\left\{ 3\eta \|B_i\|_1 \mid i \in [m] \right\}, \\ C_2 &= \min\left\{ (\gamma_i - \operatorname{Tr}(B_i X_0)) \mid i \in [m] \right\}, \end{aligned}$$

where $\eta = \operatorname{Tr}(X^*)$ for an optimal point X^* of the sketchable SDP.

Proof. The key step is to recognize that $\operatorname{Tr}\left(SCS^T X\right)$ is equal to $\operatorname{Tr}\left(CS^T XS\right)$ by the cyclicity of the trace. Thus, the relaxed SDP gives an upper bound for the sketched SDP. The theorem then follows by Corollary C.3. □

Note that this result is not probabilistic and holds regardless of the sketching matrix S used.

D. Random feasibility problems

In this section, we investigate under which conditions the convex hull of m random GUE [AGZ10, Section 2.2] matrices shifted by a multiple of the identity both contains a positive semidefinite matrix and the cone they define is pointed. This is used in Section 7.2.

Let $G_i \in \mathcal{M}_{d'}$, $i \in [m]$, be random matrices sampled independently from the GUE. This means that G_i is Hermitian and that

$$(G_i)_{kl} = \begin{cases} (G_i)_{kl} = Y_k & \text{for } k = l \\ (G_i)_{kl} = Z_{kl} & \text{for } k > l \end{cases},$$

where Y_k is a real normal random variable with mean 0 and variance 1 and Z_{kl} is a complex normal random variable with mean 0 and variance 1. Since we will need similar matrices with different variance, we will call this distribution $\mathrm{GUE}(0,1)$. Consider the rescaled and shifted matrices $\tilde{G}_i(\alpha) = \frac{1}{\sqrt{d'}} G_i + \alpha \mathbb{1}$ for some $\alpha \in \mathbb{R}$. We call the convex hull of these matrices

$$\mathcal{X}_\alpha := \operatorname{conv}\left(\left\{ \tilde{G}_i(\alpha) \,\middle|\, i \in [m] \right\} \right).$$

Lemma D.1. *Let $\alpha \geq \frac{2}{\sqrt{m}}(1+\epsilon)$. Then*

$$\mathbb{P}\left[\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^{++} \neq \emptyset \right] \geq 1 - Ce^{-2d'\epsilon^{3/2}/C},$$

where C is a numerical constant independent of m, d', ϵ.

Proof. Let $t \in \mathbb{R}_+^m$ such that $\sum_{i=1}^m t_i = 1$. By the definition of the GUE,

$$G(t) := \sum_{i=1}^m t_i G_i$$

is again in $\mathrm{GUE}(0, \sum_{i=1}^n t_i^2)$. By [LR10, Theorem 1] and the fact that the GUE is invariant under unitary transformations, it holds that

$$\mathbb{P}\left[\lambda_{\min}(G(t)/\sqrt{d'}) \leq -2\|t\|_2(1+\epsilon) \right] \leq Ce^{-2d'\epsilon^{3/2}/C}.$$

The expression $-2\|t\|_2(1+\epsilon)$ is maximized by $t_{\min} = (1/m, \ldots, 1/m)$, for which we obtain the value $-2/\sqrt{m}$ (see [Bha97, Remark II.3.7]). Hence, for $\alpha \geq 2(1+\epsilon)/\sqrt{m}$, we infer that

$$\begin{aligned} \mathbb{P}\left[\lambda_{\min}(G(t_{\min})/\sqrt{d'} + \alpha \mathbb{1}) \leq 0 \right] &= \mathbb{P}\left[\lambda_{\min}(G(t_{\min})/\sqrt{d'}) \leq -\alpha \right] \\ &\leq \mathbb{P}\left[\lambda_{\min}(G(t_{\min})/\sqrt{d'}) \leq -\frac{2}{\sqrt{m}}(1+\epsilon) \right] \\ &\leq Ce^{-2d'\epsilon^{3/2}/C}. \end{aligned}$$

As $\lambda_{\min}(G(t_{\min})/\sqrt{d'} + \alpha\mathbb{1}) > 0$ implies that $G(t_{\min})/\sqrt{d'} + \alpha\mathbb{1}$ is positive definite, the assertions follows. □

Lemma D.2. *Let $\alpha \leq \frac{2}{\sqrt{m}}(1-\epsilon)$ with $\epsilon \in \left(0, \frac{1}{2}\right)$. Then*

$$\mathbb{P}\left[\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^+ = \emptyset\right] \geq 1 - m\left(1 + \frac{8\sqrt{m}}{\epsilon}\right)^m C^4 e^{-\frac{1}{4C}(8+d'\epsilon^{3/2})\epsilon^{3/2}d'},$$

where C is a numerical constant independent of m, d', ϵ.

Proof. Take an $\epsilon/(4\sqrt{m})$-net $\mathcal{N}$ for the ℓ_1-sphere S_1^{m-1} in $\mathbb{R}^m$. This means that for all $t \in S_1^{m-1}$, there is an $s \in \mathcal{N}$ such that $\|t-s\|_1 < \epsilon/(4\sqrt{m})$. It can be shown as in [Ver12, Section 5.2.2] that we can choose $\mathcal{N}$ such that $|\mathcal{N}| \leq (1 + 8\sqrt{m}/\epsilon)^m$. Now assume that

$$\lambda_{\max}(G_i/\sqrt{d'}) \leq 2(1+\epsilon) \wedge \lambda_{\min}(G_i/\sqrt{d'}) \geq -2(1+\epsilon) \qquad \forall i \in [m].$$

This implies that $\left\|G_i/\sqrt{d'}\right\|_\infty \leq 2(1+\epsilon)$. By Weyl's perturbation theorem [Bha97, Theorem II.2.6], it follows that for $t \in S_1^{m-1} \cap \mathbb{R}_+^m$.

$$\begin{aligned}|\lambda_{\min}(G(t)/\sqrt{d'}) - \lambda_{\min}(G(s)/\sqrt{d'})| &\leq \left\|G(t)/\sqrt{d'} - G(s)/\sqrt{d'}\right\|_\infty \\ &\leq \sum_{i=1}^m |t_i - s_i| \left\|G_i/\sqrt{d'}\right\|_\infty \\ &\leq 2(1+\epsilon)\|t-s\|_1 \leq \frac{\epsilon}{\sqrt{m}}.\end{aligned}$$

Now assume further that $\lambda_{\min}(G(s)/\sqrt{d'} + \alpha\mathbb{1}) \leq -\epsilon/\sqrt{m}$ for all $s \in \mathcal{N}$. Then clearly $\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^+ = \emptyset$ by the above. We thus have to estimate the probability with which our assumptions are met. Using a union bound, we obtain

$$\begin{aligned}&\mathbb{P}\left[\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^+ = \emptyset\right] \\ \geq&\mathbb{P}\left[\left\{\lambda_{\min}\left(\alpha\mathbb{1} + \frac{G_i}{\sqrt{d'}}\right) \leq -\frac{\epsilon}{\sqrt{m}}\right\}\forall s \in \mathcal{N} \wedge \left\{\left\|\frac{G_i}{\sqrt{d'}}\right\|_\infty \leq 2(1+\epsilon)\right\}\forall i \in [m]\right] \\ \geq& 1 - \prod_{s\in\mathcal{N}} \mathbb{P}\left[\lambda_{\min}\left(\frac{G(s)}{\sqrt{d'}}\right) \geq -\frac{2}{\sqrt{m}}\left(1 - \frac{\epsilon}{2}\right)\right] \prod_{i\in[m]} \mathbb{P}\left[\lambda_{\min}\left(\frac{G_i}{\sqrt{d'}}\right) \leq -2(1+\epsilon)\right] \\ &\mathbb{P}\left[\lambda_{\max}\left(\frac{G_i}{\sqrt{d'}}\right) \geq 2(1+\epsilon)\right].\end{aligned}$$

Using again [LR10, Theorem 1] we infer that for all $i \in [m]$

$$\begin{aligned}
\mathbb{P}\left[\lambda_{\min}\left(\frac{G_i}{\sqrt{d'}}\right) \leq -2(1+\epsilon)\right] &\leq Ce^{-2d'\epsilon^{3/2}/C}, \\
\mathbb{P}\left[\lambda_{\max}\left(\frac{G_i}{\sqrt{d'}}\right) \geq 2(1+\epsilon)\right] &\leq Ce^{-2d'\epsilon^{3/2}/C}, \\
\mathbb{P}\left[\lambda_{\min}\left(\frac{G(s)}{\sqrt{d'}}\right) \geq -\frac{2}{\sqrt{m}}\left(1-\frac{\epsilon}{2}\right)\right] &\leq \mathbb{P}\left[\lambda_{\min}\left(\frac{G(s)}{\sqrt{d'}}\right) \geq -2\|s\|_2\left(1-\frac{\epsilon}{2}\right)\right] \\
&\leq C^2 e^{-2d'^2(\epsilon/2)^3/C}.
\end{aligned}$$

For the last estimate, we have used that $G(s)$ is again a GUE element with different variance (cf. proof of Lemma D.1). Combining this with the estimates concerning $|\mathcal{N}|$, the assertion follows. □

We obtain as a corollary that the cone generated by the $\tilde{G}(\alpha)$ is pointed with high probability if $m << d'$.

Corollary D.3. *Let* $\alpha \leq \frac{2}{\sqrt{m}}(1-\epsilon)$ *for* $\epsilon \in \left(0, \frac{1}{2}\right)$ *and denote by* $C_\alpha = \operatorname{cone}\{\tilde{G}_1(\alpha), \ldots, \tilde{G}_m(\alpha)\}$. *Then*

$$\mathbb{P}\left[C_\alpha \cap -C_\alpha = \{0\}\right] \geq 1 - m\left(1 + \frac{8\sqrt{m}}{\epsilon}\right)^m C^4 e^{-\frac{1}{4C}(8+d'\epsilon^{3/2})\epsilon^{3/2}d'},$$

where C *is a numerical constant independent of* m, d', ϵ.

Proof. We know from Lemma D.2 that we have the same lower bound for the probability of the event $A = \{\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^+ = \emptyset\}$. But note that the event A implies that the cone is pointed. That is because if the cone was not pointed, there would exist $\gamma, \mu \in \mathbb{R}_+^m \setminus \{0\}$ such that

$$\sum_{i=1}^m \gamma_i \tilde{G}_i(\alpha) = -\sum_{i=1}^m \mu_i \tilde{G}_i(\alpha)$$

and so

$$\frac{1}{\sum\limits_{i=1}^m (\gamma_i + \mu_i)} \sum_{i=1}^m (\mu_i + \gamma_i) G_i(\alpha) = 0,$$

which implies that $\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^+ \neq \emptyset$. □

Theorem D.4. *Consider the LMI*

$$\sum_{i=1}^m t_i \tilde{G}_i(\alpha) - \mathbb{1}.$$

Let $\epsilon \in \left(0, \frac{1}{2}\right)$. *Then for* $\alpha \geq \frac{2}{\sqrt{m}}(1+\epsilon)$, *this LMI is feasible with probability at least*

$$1 - Ce^{-2d'\epsilon^{3/2}/C}. \tag{30}$$

Moreover, for $\alpha \leq \frac{2}{\sqrt{m}}(1+\epsilon)$, this LMI is infeasible, the cone generated by the $\tilde{G}(\alpha)$ is pointed and we have $\mathcal{X}_\alpha \cap \mathcal{S}_{d'}^{+} = \emptyset$ with probability at least

$$1 - m\left(1 + \frac{8\sqrt{m}}{\epsilon}\right)^m C^4 e^{-\frac{1}{4C}(8+d'\epsilon^{3/2})\epsilon^{3/2}d'}. \tag{31}$$

Proof. For the first assertion, by Lemma D.1 $\mathcal{X}_\alpha$ contains a positive definite element $\sum_{i=1}^m r_i \tilde{G}_i(\alpha)$ with probability lower bounded by the expression in Equation (30). Then $G(\mu r)$ is feasible for $\mu \in \mathbb{R}_+$ large enough. For the second assertion, we note that the LMI being infeasible is equivalent to $\mathcal{X}_\alpha$ not containing a positive definite element. From Lemma D.2 the lower bound in (31) for the probability that the LMI is infeasible. Moreover, the fact that the same lower bound holds for the probability that the cone is pointed follows from Corollary D.3. □

C.2 Approximate Randomized Benchmarking for Finite Groups

Approximate Randomized Benchmarking for Finite Groups

D. Stilck França, A.K. Hashagen

We extend and generalize the randomized benchmarking protocol in three directions. First, we show how to adapt the protocol to estimate the average gate fidelity of a set of gates forming a representation of a finite group. With a few exceptions, the previous literature only focused on implementing Clifford gates. The usual randomized benchmarking protocol works under the assumption that we may sample from the Haar measure of the group. We relax this assumption and show that the protocol also works using approximate Haar samples, which allows us to use Markov chain Monte Carlo techniques to obtain the samples. Inspired by these results, we show how to implement the protocol only implementing gates from a set of generators of the group that is closed under inverses and one arbitrary element of the group. Proceeding this way significantly reduces the experimental cost of implementing the protocol and is a more natural framework for the error model. This last version of the protocol works under the assumption that the quantum channel that describes the noise is close to being covariant. Moreover, we discuss some sets of quantum gates that might be relevant for applications and fit into our framework and perform numerical tests of our results.

C.2.1 Main Results

Randomized benchmarking [15] is a protocol to efficiently estimate the average fidelity of the implementation of a set of quantum gates U_g that is a representation of a finite group G. The protocol works under the assumption that the quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ that describes the noise in the implementation is independent of the gates. It works by applying random sequences of gates $\mathcal{U}_k$ from the group to the system initially in a state ρ and measuring a POVM element E on it afterward. The decay of the probability of observing the particular outcome with the sequence length m is called the survival probability and we denote it by $F(m, E, \rho)$.

Lemma C.2.1. *Let G be a finite group and $\{U_g\}_{g\in G} \subset U(d)$ be a representation such that $U_g \otimes \bar{U}_g \in \oplus_{\alpha\in\hat{G}} \mathbb{C}_{d_\alpha} \otimes \mathbb{1}_{m_\alpha}$ is the decomposition of $U_g \otimes \bar{U}_g$ into irreps. Then there exist $\lambda_1, \ldots, \lambda_k \in \overline{B_1(0)}$ and $a_0, a_1, \ldots, a_k$ with $k \leq \sum_{\alpha\in\hat{G}} m_\alpha$ such that for all $m \geq \max m_\alpha$*

$$F(m, E, \rho) = a_0 + \sum_{l=1}^{k} a_l e^{\lambda_l m}. \tag{C.7}$$

The proof of this Lemma follows from the fact that channels $T : \mathcal{M}_d \to \mathcal{M}_d$ that are covariant with respect to this representation can be diagonalized for $m \geq \max_{\alpha\in\hat{G}} m_\alpha$ with a block structure which is determined by the decomposition of the representation $U_g \otimes \bar{U}_g$. After diagonalizing the channel, Equation (C.7) follows immediately. Given the λ_l, it is then possible to obtain estimates on the average fidelity of the quantum channel T.

The usual randomized benchmarking protocol assumes one has access to Haar distributed samples of the group. We relax this to samples that are close to Haar.

Theorem C.2.2. *Let μ be the Haar measure on G and $\nu_1, \ldots, \nu_m$ probability measures on G s.t.*

$$\|\mu - \nu_k\|_1 \leq \epsilon_k. \tag{C.8}$$

for all $1 \leq k \leq m$ and $\epsilon_k \geq 0$. Suppose we pick the $\mathcal{U}_k$ independently from ν_k and denote the resulting expectation value for the survival probability by $\tilde{F}(m, E, \rho)$. Then

$$|\tilde{F}(m, E, \rho) - F(m, E, \rho)| \leq 4\sqrt{\frac{\log(|G|)}{1 - |G|^{-1}} \sum_{k=1}^{m} \epsilon_k}.$$

The proof of this statements estimates the 1 → 1 norm of the channel $\mathcal{T}(T)^m$ and the expected channel if we perform randomized benchmarking with the gates chosen approximately at random. This is done by exploring the invariance of the total variation distance under permutations and a tight reverse Pinsker inequality. From the estimate on the 1 → 1 norm, it is then straightforward to obtain the claim using Hölder's inequality.

Moreover, for quantum channels that are close to being covariant, we show that one can also perform randomized benchmarking by just implementing a set of generators A of the group closed under taking inverses and one arbitrary element. The channels are close to being covariant in the sense that there is a $\delta > 0$ such that $T = (1-\delta)T_c + \delta T_n$ with T_c, T_n quantum channels and T_c covariant, which we call δ-covariance. The protocol requires that we first apply a random sequence of gates from A of length b, with b being proportional to the time it takes for the random walk generated by A on G to mix. We can then show:

Theorem C.2.3. *Let T be $\delta-$covariant w.r.t. a representation U_g of a group G, A a subset of G that generates G and is closed under taking inverses and $\delta > 0$. Suppose we run the protocol above with $b = t_{mix}(m^{-1}\epsilon)$ for some $\epsilon > 0$. denote the resulting expectation value for the survival probability for the protocol by $F_{gen}(m, E, \rho)$. Then*

$$|F_{gen}(m, E, \rho) - F(m, E, \rho)| \leq \epsilon + \mathcal{O}\left(\delta^2 m\right).$$

We prove this Theorem by expressing powers of $\mathcal{T}(T)$ in terms of T_c and T_n and noting that by choosing b this way all terms of first order in δ vanish.

C.2.2 Applications

We applied our generalized benchmarking protocol to the group of unitaries generated by products of $d-$dimensional permutation matrices and diagonal unitaries with entries that are $n-$th roots of unity. This group contains the T gate for $n \geq 8$ and thus these gates together with Cliffords allow for universal quantum computation. We show that one only needs to estimate two parameters when performing randomized benchmarking with these gates and one can multiply and invert them with complexity $\mathcal{O}(d)$, which makes it feasible to perform the protocol for a moderate number of qubits. Additionally, we test our results for approximate samples and generator based benchmarking numerically using the Clifford group and random quantum channels that are approximately covariant. We take the set of generators to be given by the Hadamard H, the $\pi-$gate and its inverse and the $CNOT$ between different qubits. The numerical results for the approximate sampling benchmarking protocol and the usual protocol were indistinguishable. For the range $\delta < 0.3$ the generator randomized benchmarking performance was comparable to the usual protocol.

C.2.3 Individual Contribution

This project started after Anna-Lena K. Hashagen asked me to review a paper of hers on a similar topic. I realized that many techniques could be generalized and some assumptions of the usual randomized benchmarking protocol were too strong and should be relaxed. Therefore, I proposed to pursue these ideas. I am the principal author of this article. We formulated the protocol, which can be found in Section 4, for general groups together after discussions and she derived the expression in Lemma 11 of the article for the case of irreducibly covariant representations. I then generalized the expression to arbitrary representations of finite groups in the version that can now be found in Lemma C.2.1 of this summary or Lemma 11 of the article. I was responsible for proving that the protocol also works using approximate samples of the Haar measure. That is, I proved and formulated all results of Section 5, and the idea of generalizing the randomized benchmarking in this direction was also mine. Furthermore, I formulated and proved the version of the protocol using generators of the group discussed in Section 6. The idea of implementing the protocol only using generators goes back to me. A.K. Hashagen initially suggested the group of monomial unitary matrices as a possible application of our protocol. However, as this is not a finite group, it did not fit our framework properly.

I was then responsible for showing that the protocol would also work if we considered the finite subgroup of monomial matrices with entries that are $n-$th roots of unity. Applying our techniques to the Clifford group was a natural choice, and we both had this in mind. I was responsible for all the numerics in the article and wrote all the code necessary for it. I was also responsible for the contents of Appendix A and B, as the statistical and numerical considerations included there were inspired by some observations I made during the process of writing this article. A.K. Hashagen wrote section 2 on Notation and Preliminaries, and I wrote all other sections but the Introduction, which we wrote together. After the first draft of the paper was ready, we had many discussions on how to improve the presentation until we arrived at the present form of the article.

Approximate randomized benchmarking for finite groups

D S França and A K Hashagen

E-mail: dsfranca@mytum.de, hashagen@ma.tum.de
Department of Mathematics, Technical University of Munich, Germany

March 2018

Abstract. We investigate randomized benchmarking in a general setting with quantum gates that form a representation, not necessarily an irreducible one, of a finite group. We derive an estimate for the average fidelity, to which experimental data may then be calibrated. Furthermore, we establish that randomized benchmarking can be achieved by the sole implementation of quantum gates that generate the group as well as one additional arbitrary group element. In this case, we need to assume that the noise is close to being covariant. This yields a more practical approach to randomized benchmarking. Moreover, we show that randomized benchmarking is stable with respect to approximate Haar sampling for the sequences of gates. This opens up the possibility of using Markov chain Monte Carlo methods to obtain the random sequences of gates more efficiently. We demonstrate these results numerically using the well-studied example of the Clifford group as well as the group of monomial unitary matrices. For the latter, we focus on the subgroup with nonzero entries consisting of n-th roots of unity, which contains T gates.

Keywords: Randomized benchmarking, quantum gates, Clifford gates, monomial unitary, random walks on groups, fidelity estimation.

Submitted to: *J. Phys. A: Math. Theor.*

1. Introduction

One of the main obstacles to build reliable quantum computers is the need to implement quantum gates with high fidelity. Therefore, it is key to develop techniques to estimate the quality of quantum gates and thus certify the quality of a quantum computer. To this end, one could perform tomography for the underlying noise in the implementation and in principle obtain a complete description of it [1, 2]. However, in general, the number of measurements necessary to estimate for a complete tomography of the noise scales exponentially with the system size and is not a practical solution to the problem. Thus, it is vital to develop techniques to estimate the level of noise in systems more efficiently, even if we only obtain partial information.

Randomized benchmarking (RB) is a protocol to estimate the average fidelity of a set of quantum gates forming a representation of a group [3, 4, 5, 6]. The very important case of Clifford gates has already been widely studied and some rigorous results that show its efficiency under some noise scenarios are available [7, 8], such as when the noise is independent of the gate and time. Besides its efficiency, another highlight of the protocol is that it is robust against state preparation and measurement errors. This makes it very attractive from an experimental point of view and its applicability was demonstrated successfully [9, 10, 11, 12, 13, 14, 15, 16, 17].

In this work, we show how to extend these protocols to gates that are representations of a finite group‡; these must not necessarily be irreducible or form a 2-design. Although other works, such as [18, 19, 20, 21], already extended the protocol to other specific groups of interest, we focus on showing how to estimate the average fidelity based on properties of the particular representation at hand for arbitrary finite groups. To this end, we investigate the structure of quantum channels that are covariant under a unitary representation of a group and derive formulas for their average fidelity in terms of their spectra. We then show that one can use RB to estimate the average fidelity of these gates under the assumption that they are subject to time and gate independent noise.

In order for this procedure to be efficient, it is necessary that we may multiply, invert and sample uniformly distributed elements of the group efficiently and that the given representation does not decompose into too many irreducible unitary representations, as we will discuss in more detail later. This is the case for the well-studied case of Cliffords.

The usual RB protocol assumes that we can implement sequences of gates that are sampled from the Haar distribution of the group [3, 4, 5, 6]. We further generalize the RB protocol by showing that it is possible to implement sequence gates that are approximately Haar distributed instead. Therefore, it is possible to use Markov chain Monte Carlo methods to obtain the samples, potentially more efficiently. This result is of independent interest to the RB literature, as it shows that the protocol is stable

‡ Most of the results in this work can easily be extended to compact groups. However, as it is not clear that implementing the RB protocol for compact groups is relevant for applications and given that this would make some proofs less accessible, we restrict to finite groups here.

against small errors in the sampling.

Moreover, we show how one can perform RB by just implementing gates that generate the group and one arbitrary element of the group. This might provide a more natural framework to the protocol, as often one is only able to implement a certain number of gates that generate the group and must, therefore, decompose the gates into generators. However, this protocol works under the assumption that the noise affecting the gates is already close to being covariant w.r.t. the group and not for arbitrary quantum channels, as in the usual setting.

To illustrate our techniques, we apply them to subgroups of the monomial unitary matrices, i.e. products of $d-$dimensional permutation and diagonal unitary matrices. These can be seen as a generalization of stabilizer groups [22]. We focus on the subgroup of monomial unitary matrices whose nonzero entries are roots of unity. We show that we only need to estimate two parameters and multiplying and inverting elements of it can be done in time $\mathrm{O}(d)$. Moreover, they include the T gate, which is known to form a universal set for quantum computation together with the Clifford gates [23]. Therefore, one can use the protocol described here to estimate the noise from T gates more efficiently. We make numerical simulations for our protocol and these subgroups and show that it is able to reliably estimate the average gate fidelity. Moreover, we numerically compare our techniques based on approximate Haar samples and implementation of generators to the usual protocol for Cliffords and show that the three yield indistinguishable results in the high fidelity regime.

This paper is structured as follows: we start by fixing our notation and reviewing basic results on Markov chains and covariant quantum channels; needed in section 2. In section 3 we derive the average fidelity of quantum channels in terms of their spectra and we give basic results on the decay of the probability of measurement outcomes under covariant quantum channels. These form the basis for the RB protocol for general groups, which we discuss and analyze in section 4. In section 5 we prove that it is also possible to implement the protocol using approximate samples. We then discuss the generalized RB protocol based on implementing random sequences of gates that generate the group in section 6. In this section, we also discuss the conditions under which this protocol applies. Finally, in section 7, we apply our techniques to the subgroup of monomial unitary matrices and perform numerical experiments for it. In the same section, we also compare numerically the RB protocols developed here with the usual one in the case of the Clifford group.

2. Notation and Preliminaries

We will be interested in finite dimensional quantum systems. Denote by $\mathcal{M}_d$ the space of $d \times d$ complex matrices. We will denote by $\mathcal{D}_d$ the set of d-dimensional quantum states, i.e., positive semi-definite matrices $\rho \in \mathcal{M}_d$ with trace 1. We will call a linear map $T : \mathcal{M}_d \to \mathcal{M}_{d'}$ a quantum channel if it is trace preserving and completely positive. We will call a collection of positive semidefinite matrices $\{E_i\}_{i=1}^{l}$ a positive operator valued measure (POVM) if the POVM elements E_i, called effect operators, sum up to the identity. Throughout this paper, we will use the channel-state duality that provides a one-to-one correspondence between a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ and its Choi-Jamiolkowski state $\tau_T \in \mathcal{M}_{d^2}$ obtained by letting T act on half of a maximally entangled state, i.e.,

$$\tau_T := (T \otimes \mathrm{id}_d)\left(|\Omega\rangle\langle\Omega|\right), \tag{1}$$

where $|\Omega\rangle\langle\Omega| \in \mathcal{M}_{d^2}$ is a maximally entangled state, that is,

$$|\Omega\rangle\langle\Omega| = \frac{1}{d}\sum_{i,j=1}^{d} |ii\rangle\langle jj|, \tag{2}$$

where $\{|i\rangle\}_{i=1}^{d}$ is an orthonormal basis in $\mathbb{C}^d$. Please refer to [24] for more on these concepts. To measure the distance between two states we will use the Schatten 1−norm for $A \in \mathcal{M}_d$, denoted by $\|\cdot\|_1$ and given by

$$\|A\|_1 := \mathrm{Tr}\left((A^*A)^{\frac{1}{2}}\right), \tag{3}$$

where $*$ denotes the adjoint. Then, given two states $\rho, \sigma \in \mathcal{D}_d$, their trace distance is given by $\|\rho - \sigma\|_1/2$. This norm on $\mathcal{M}_d$ induces a norm on linear operators $\Phi : \mathcal{M}_d \to \mathcal{M}_d$ through

$$\|\Phi\|_{1\to 1} := \sup_{X \in \mathcal{M}_d, X \neq 0} \frac{\|\Phi(X)\|_1}{\|X\|_1}. \tag{4}$$

Given a random quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$, we will denote its expectation value by $\mathbb{E}(T)$.

We will also need some basic facts from the representation theory of finite groups. We refer to e.g. [25] for more on this and the proofs of the statements we use here. We will be particularly interested in the commutant of the algebra generated by the group. To this end we introduce:

Definition 1 (Commutant)**.** Let $\mathcal{A}$ be an algebra of operators on a Hilbert space $\mathcal{H}$. Then the commutant $\mathcal{A}'$ of $\mathcal{A}$ is defined by

$$\mathcal{A}' := \{B | BA = AB \text{ for all } A \in \mathcal{A}\}. \tag{5}$$

Let $U : G \to \mathcal{M}_d$ be a unitary representation of a finite group G on a finite-dimensional Hilbert space $\mathcal{H} \simeq \mathbb{C}^d$. We will denote the unitary corresponding to g by U_g. From basic results of representation theory, we know that there exists distinct $\alpha_1, \ldots, \alpha_k \in \hat{G}$, where $\hat{G}$ denotes the set of equivalence classes of irreducible unitary representations (irreps), such that the unitary representation can be written as a direct sum of irreps, i.e. $U \cong \oplus U^{\alpha_i} \otimes \mathbb{I}_{m_\alpha}$ with $m_\alpha > 0$ denoting the degeneracy of the α_i-th irrep. The structure of the commutant is then described in the following theorem.

Theorem 2 ([25, Theorem IX.11.2])**.** *Let U be a unitary representation of a finite group G on $\mathcal{H}$. Write $\mathcal{H} = \oplus_{\alpha\in\hat{G}}\left(\mathbb{C}^{d_\alpha}\otimes\mathbb{C}^{m_\alpha}\right)$ so that $U_g = \oplus_{i=1}^{k} U_g^{\alpha_i}\otimes\mathbb{I}_{m_\alpha}$ with $\{\alpha_i\}_{i=1}^{k}$ distinct elements in $\hat{G}$. Let $\mathcal{A}(U)$ be the algebra of operators generated by the $\{U_g\}_{g\in G}$, and $\mathcal{A}(U)'$ its commutant. Then*

$$\mathcal{A}(U) = \left\{\oplus_{i=1}^{k} A_i\otimes\mathbb{I}_{m_\alpha}\middle| A_i\in\mathcal{M}_{d_{\alpha_i}}\right\}, \tag{6a}$$

$$\mathcal{A}(U)' = \left\{\oplus_{i=1}^{k}\mathbb{I}_{d_{\alpha_i}}\otimes B_i\middle| B_i\in\mathcal{M}_{m_\alpha}\right\}. \tag{6b}$$

Given a finite group G, we will call the uniform probability distribution on it its Haar measure. For a proof of its existence and basic properties, we refer to [25, Section VII.3]. We will denote the character of a unitary representation $\alpha\in\hat{G}$ by χ^α and remark that one can find the decomposition in theorem 2 through characters [25, Section III.2].

2.1. Covariant Quantum Channels and Twirls

The definition of covariance of quantum channels is central to the study of their symmetries and will be one of the building blocks of the generalized RB protocol:

Definition 3 (Covariant quantum channel [26])**.** A quantum channel $T:\mathcal{M}_d\to\mathcal{M}_d$ is covariant with respect to a unitary representation $U: G\to\mathcal{M}_d$ of a finite group G, if for all $g\in G$

$$T\left(U_g\cdot U_g^*\right) = U_g T\left(\cdot\right)U_g^*. \tag{7}$$

In general, one allows different unitary representations of the group in the input and output of the channel in the definition of covariance, but here we will restrict to the case when we have the same unitary representation. There are many different and equivalent characterizations of covariance. Here we mention that covariance is equivalent to the Choi-Jamiolkowski state τ_T commuting with $U_g\otimes\bar{U}_g$ for all $g\in G$. To see this, note that given a unitary representation U of G we may define its adjoint representation $\mathcal{U}: G\to\mathrm{End}(\mathcal{M}_d)$ through its action on any $X\in\mathcal{M}_d$ by conjugation,

$$\mathcal{U}_g(X) = U_g X U_g^*. \tag{8}$$

Through the Choi-Jamiolkowski isomorphism, it is easy to see that the adjoint representation is equivalent to the unitary representation $U_g\otimes\bar{U}_g\in\mathcal{M}_{d^2}$. As we can rephrase (7) as T commuting with the adjoint representation, this translates to the Choi-Jamiolkowski state commuting with $U_g\otimes\bar{U}_g$. This means in particular that we may use structural theorems, like theorem 2, to investigate covariant channels, as covariance implies that the channel is in the commutant of the adjoint representation.

Theorem 4. *Let $T:\mathcal{M}_d\to\mathcal{M}_d$ be a quantum channel that is covariant w.r.t. a unitary representation U of a finite group G and let $\oplus_{\alpha\in\hat{G}}\left(\mathbb{C}^{d_\alpha}\otimes\mathbb{C}^{m_\alpha}\right)$ be the decomposition of the underlying Hilbert space into irreps α of G with multiplicity m_α for the unitary representation $U\otimes\bar{U}$. Then:*

$$T = \oplus_{\alpha\in\hat{G}}\mathbb{I}_{d_\alpha}\otimes B_\alpha \tag{9}$$

with $B_\alpha\in\mathcal{M}_{m_\alpha}$.

Proof. As T is covariant, it must be an element of the commutant of the adjoint representation, i.e. $T \in \mathcal{A}(\mathcal{U})'$. The decomposition then follows from theorem 2. □

This decomposition further simplifies when no multiplicities in the decomposition of the unitary representations into its irreducible components are present. We call such channels irreducibly covariant. Here we briefly mention some of the results of [27], where the structure of such channels is investigated.

Theorem 5 ([27, Theorem 40]). *A quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is irreducibly covariant with respect to an irrep $U : G \to \mathcal{M}_d$ of a finite group G if and only if it has a decomposition of the following form:*

$$T = l_{\mathrm{id}} P^{\mathrm{id}} + \sum_{\alpha \in \hat{G}, \alpha \neq \mathrm{id}} l_\alpha P^\alpha, \tag{10}$$

with $l_{\mathrm{id}} = 1$, $l_\alpha \in \mathbb{C}$ and where $P^{\mathrm{id}}, P^\alpha : \mathcal{M}_d \to \mathcal{M}_d$ are projectors defined as

$$P^\alpha(\cdot) = \frac{\chi^\alpha(e)}{|G|} \sum_{g \in G} \chi^\alpha\left(g^{-1}\right) U_g \cdot U_g^*, \tag{11}$$

with $\alpha \in \hat{G}$ and $e \in G$ the identity of the group. They have the following properties:

$$P^\alpha P^\beta = \delta_{\alpha\beta} P^\alpha, \qquad (P^\alpha)^* = P^\alpha \quad \textit{and} \quad \sum_{\alpha \in \hat{G}} P^\alpha = \mathrm{id}_d, \tag{12}$$

where $\mathrm{id}_d : \mathcal{M}_d \to \mathcal{M}_d$ is the identity map and the coefficients l_α are the eigenvalues of the quantum channel T.

That is, in the case of an irreducibly covariant channel we can also write down the projections onto different eigenspaces and diagonalize the channel.

One of the most important concepts in this paper is that of the twirl of a channel.

Definition 6 (Twirl). Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, G a finite group with Haar measure μ and $U : U \to \mathcal{M}_d$ a unitary representation of G. We define the twirl of T w.r.t. G, denoted by $\mathcal{T}(T) : \mathcal{M}_d \to \mathcal{M}_d$, as

$$\mathcal{T}(T)(\cdot) = \int_G \mathcal{U}_g^* \circ T \circ \mathcal{U}_g(\cdot) \mathrm{d}\mu. \tag{13}$$

Strictly speaking the twirled channel, of course, depends on the particular group and unitary representation at hand. However, we will omit this in the notation, as the group in question should always be clear from context. It is then easy to show that $\mathcal{T}(T)$ is a quantum channel that is covariant w.r.t. this representation.

2.2. Random Walks on Groups

We will need some basic tools from the field of random walks on groups to motivate and explain our protocol to perform RB with generators or with approximate samples.

Therefore, we review these basic concepts here and refer to e.g. [28, Chapter 2.6] for more details and proofs. Given a finite group G and a probability measure μ on G, we denote the set of probability measures on G by $\mathcal{P}(G)$. If X, Y are two independent random variables on G with distributions $\mu, \nu \in \mathcal{P}(G)$, respectively, we denote their joint distribution on $G \times G$ by $\mu \otimes \nu$. Analogously, we will denote the joint distribution of $Y_1, \ldots, Y_n$ i.i.d. variables with distribution ν by $\nu^{\otimes n}$ and the m-fold Cartesian product of G with itself by G^m. The random walk on G with increment distribution ν is defined as follows: it is a Markov chain with state space G. Given that the current state X_n of the chain is g, the next state X_{n+1} is given by multiplying the current state on the left by a random element of G selected according to ν. That is, we have

$$P(X_{n+1} = g_2 | X_n = g_1) = \nu\left(g_2 g_1^{-1}\right). \tag{14}$$

Another way of tracking the transition probabilities for these chains is through the transition matrix of the chain, π. For $g_1, g_2 \in G$, this matrix is defined as

$$\pi(g_1, g_2) = \nu\left(g_2 g_1^{-1}\right). \tag{15}$$

If X_0 is distributed according to $\mu \in \mathcal{P}(G)$, we have that the distribution of X_n is given by $\pi^n \mu$, where we just expressed μ as a vector in $\mathbb{R}^{|G|}$. We recall the following fundamental result about random walks on groups:

Theorem 7. *Let G be a finite group and A be a set of generators of G that is closed under inversion. Moreover, let ν be the uniform distribution on A and $X_1, X_2, \ldots$ be a random walk with increment distribution ν. Then the distribution of X_n converges to the Haar distribution on G as $n \to \infty$.*

Proof. We refer to e.g. [28, Section 2.6.1] for a proof and more details on this. □

Given a generating subset A of G that is closed under inverses and ν the uniform distribution on A, we will refer to the random walk with increment ν as the random walk generated by A. This result provides us with an easy way of obtaining samples which are approximately Haar distributed if we have a set of generators by simulating this random walk for long enough. The speed of this convergence is usually quantified in the total variation distance. Given two probability measure μ, ν on G, we define their total variation distance to be given by:

$$\|\mu - \nu\|_1 := \frac{1}{2} \sum_{g \in G} |\mu(g) - \nu(g)|. \tag{16}$$

We then define the mixing time of the random walk as follows:

Definition 8 (Mixing Time of Random Walk)**.** Let G be a finite group and A a set of generators closed under inverses and μ be the Haar measure on the group. For $\epsilon > 0$, the mixing time of the chain generated by A, $t_1(\epsilon)$, is defined as

$$t_1(\epsilon) := \inf\{n \in \mathbb{N} | \forall \nu \in \mathcal{P}(G) : \|\pi^n \nu - \mu\|_1 \leq \epsilon\}. \tag{17}$$

We set t_{mix} to be given by $t_1(4^{-1})$. One can then show that $t_1(\epsilon) \leq \lceil \log_2(\epsilon^{-1}) \rceil t_{\text{mix}}$ (see [28, Section 4.5] for a proof). There is a huge literature devoted to determining the mixing time of random walks on groups and we refer to [29] and references therein for more details. For our purposes it will be enough to note that in most cases we have that $t_1(\epsilon)$ scales logarithmically with ϵ^{-1} and $|G|$. Another distance measure which is quite useful in the study of convergence of random variables is the relative entropy D. For two probabilities measures μ, ν on $\{1, \ldots, d\}$ we define their relative entropy to be

$$D(\mu||\nu) := \begin{cases} \sum\limits_{i=1}^{d} \mu(i) \log\left(\frac{\mu(i)}{\nu(i)}\right), & \text{if } \mu(i) = 0 \text{ for all } i \text{ s.t. } \nu(i) = 0, \\ +\infty, \quad \text{else.} \end{cases} \tag{18}$$

One of its main properties is that for $\mu, \nu \in \mathcal{P}(G)$ we have

$$D(\mu^{\otimes n}||\nu^{\otimes n}) = nD(\mu||\nu). \tag{19}$$

3. Fidelities

Given a quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ and a unitary channel $\mathcal{U} : \mathcal{M}_d \to \mathcal{M}_d$, the average fidelity between them is defined as

$$F(T,\mathcal{U}) = \int \operatorname{Tr}\left(T(|\psi\rangle\langle\psi|)\mathcal{U}(|\psi\rangle\langle\psi|)\right) \mathrm{d}\psi, \tag{20}$$

where we are integrating over the Haar measure on quantum states. In case $\mathcal{U}$ is just the identity, we refer to this quantity as being the average fidelity of the channel and denote it by $F(T)$. As shown in [30], the average fidelity of a channel is a simple function of its entanglement fidelity, given by

$$F_e(T) = \operatorname{Tr}\left(T \otimes \operatorname{id}\left(|\Omega\rangle\langle\Omega|\right)|\Omega\rangle\langle\Omega|\right), \tag{21}$$

with $|\Omega\rangle\langle\Omega|$ the maximally entangled state. One can then show that

$$F(T) = \frac{dF_e(T) + 1}{d + 1}. \tag{22}$$

Thus, we focus on estimating the entanglement fidelity instead of estimating the average fidelity. This can be seen to be just a function of the trace of the channel and the dimension, as we now show.

Lemma 9. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel. Then $F_e(T) = d^{-2}\operatorname{Tr}(T)$. Here we mean the trace of T as a linear operator between the vector spaces $\mathcal{M}_d$.*

Proof. The entanglement fidelity is

$$\begin{aligned} F_e(T) &= \operatorname{Tr}\left(T \otimes \operatorname{id}\left(|\Omega\rangle\langle\Omega|\right)|\Omega\rangle\langle\Omega|\right) \\ &= \frac{1}{d^2}\sum_{i,j,k,l=1}^{d} \operatorname{Tr}\left(\left[T\left(|i\rangle\langle j|\right) \otimes |i\rangle\langle j|\right]|l\rangle\langle k| \otimes |l\rangle\langle k|\right) \\ &= \frac{1}{d^2}\sum_{i,j=1}^{d} \operatorname{Tr}\left(T\left(|i\rangle\langle j|\right)\left(|i\rangle\langle j|\right)^*\right). \end{aligned}$$

Note that $\{|i\rangle\langle j|\}_{i,j=1}^{d}$ is an orthonormal basis of $\mathcal{M}_d$ and $\operatorname{Tr}\left(T\left(|i\rangle\langle j|\right)\left(|i\rangle\langle j|\right)^*\right)$ corresponds to the Hilbert-Schmidt scalar product between $T\left(|i\rangle\langle j|\right)$ and $|i\rangle\langle j|$. Therefore, we have that

$$\sum_{i,j=1}^{d}\operatorname{Tr}\left(T\left(|i\rangle\langle j|\right)\left(|i\rangle\langle j|\right)^*\right)=\operatorname{Tr}\left(T\right),$$

where again $\operatorname{Tr}(T)$ is again meant as the trace of T as a linear operator. □

That is, if we know the eigenvalues or the diagonal elements of T w.r.t. some basis, we may determine its entanglement and average fidelity. The RB protocol explores the fact that twirling a channel does not change its trace and that the trace of covariant channels has a much simpler structure, as made clear in the next corollary.

Corollary 10. *Let $T:\mathcal{M}_d\to\mathcal{M}_d$ be a quantum channel that is covariant w.r.t. a unitary representation $U:G\to\mathbb{C}^d$ of a finite group G and let $\oplus_{\alpha\in\hat{G}}\left(\mathbb{C}^{d_\alpha}\otimes\mathbb{C}^{m_\alpha}\right)$ be the decomposition of $\mathbb{C}^d\otimes\mathbb{C}^d$ into irreps α of G with multiplicity m_α for the unitary representation $U\otimes\bar{U}$. Choose a basis s.t.*

$$T=\oplus_{\alpha\in\hat{G}}\mathbb{I}_{d_\alpha}\otimes B_{m_\alpha} \tag{23}$$

with $B_\alpha\in\mathcal{M}_{m_\alpha}$. Then

$$F_e(T)=d^{-2}\sum_{\alpha\in\hat{G}}d_\alpha\operatorname{Tr}\left(B_\alpha\right). \tag{24}$$

Proof. The claim follows immediately after we combine theorem 4 and lemma 9. □

This shows that the spectrum of quantum channels that are covariant with respect to a unitary representation of a finite group has much more structure and is simpler than that of general quantum channels. In particular, if the unitary representation $U\otimes\bar{U}$ is such that $\sum_\alpha m_\alpha\ll d^2$, then we know that the spectrum of the quantum channel is highly degenerate and we only need to know a few points of it to estimate the trace. We will explore this fact later in the implementation of the RB protocol.

We will now show in lemma 11 the probability of measurement outcomes has a very simple form for covariant channels and their powers.

Lemma 11. *Let $T:\mathcal{M}_d\to\mathcal{M}_d$ be a quantum channel that is covariant w.r.t. a unitary representation $U:G\to\mathcal{M}_d$ of a finite group G and let $\oplus_{\alpha\in\hat{G}}\left(\mathbb{C}^{d_\alpha}\otimes\mathbb{C}^{m_\alpha}\right)$ be the decomposition of $\mathbb{C}^d\otimes\mathbb{C}^d$ into irreps α of G with multiplicity m_α for the unitary representation $U\otimes\bar{U}$. Moreover, let $\rho\in\mathcal{D}_d$, $E\in\mathcal{M}_d$ be a POVM element and $m\geq\max m_\alpha$. Then there exist $\lambda_1,\ldots,\lambda_k\in\overline{B_1(0)}$, the unit ball in the complex plane, and $a_0,a_1,\ldots,a_k\in\mathbb{C}$ s.t.*

$$\operatorname{Tr}\left(T^m(\rho)E\right)=a_0+\sum_{i=1}^{k}a_k\lambda_i^m. \tag{25}$$

Moreover,

$$k \leq \sum_{\alpha \in \hat{G}} m_\alpha - 1 \tag{26}$$

corresponds to the number of distinct eigenvalues of T *and* λ_i *are its eigenvalues.*

Proof. As T is a linear map from $\mathcal{M}_d$ to $\mathcal{M}_d$ it has a Jordan decomposition [31]. That is, there exists an invertible linear operator $X : \mathcal{M}_d \to \mathcal{M}_d$ such that

$$X^{-1} \circ T \circ X = D + N, \quad [D, N] = 0.$$

Here $D : \mathcal{M}_d \to \mathcal{M}_d$ is diagonal in the standard basis $\{|i\rangle\langle j|\}_{i,j=1}^d$ of $\mathcal{M}_d$ with diagonal entries given by the eigenvalues of T and $N : \mathcal{M}_d \to \mathcal{M}_d$ nilpotent. As we have that T is covariant, it follows from the decomposition in theorem 4 that the eigenvalues can be at most $\max m_\alpha = m_0-$fold degenerate and $N^{m_0} = 0$. Thus, it follows that T^m is diagonalizable, as $m \geq \max m_\alpha$. We then have

$$X^{-1} \circ T^m \circ X = D^m.$$

We can then rewrite the scalar product

$$\mathrm{Tr}\left(T^m(\rho)E\right) = \mathrm{Tr}\left(X \circ D^m \circ X^{-1}(\rho)E\right) = \mathrm{Tr}\left(D^m(X^{-1}(\rho))X^*(E)\right).$$

Let $b_{i,j}$ and $c_{i,j}$ be the matrix coefficient of $X^*(E)$ and $X^{-1}(\rho)$, respectively, in the standard basis. That is

$$X^*(E) = \sum_{i,j=1}^d b_{i,j} \, |i\rangle\langle j| \,, \qquad X^{-1}(\rho) = \sum_{i,j=1}^d c_{i,j} \, |i\rangle\langle j| \,.$$

Exploring the fact that D is diagonal in this basis we obtain

$$\mathrm{Tr}\left(T^m(\rho)E\right) = \sum_{i,j=1}^d b_{i,j} c_{i,j} d_{i,j}^m,$$

where $d_{i,j}$ are just the eigenvalues of T, including multiplicities. To arrive at the curve in (25), we group together all terms corresponding to the same eigenvalue λ_i. Moreover, note that quantum channels always have 1 in their spectrum, which gives the a_0 term that does not depend on m. The fact that $\lambda_i \in \overline{B_1(0)}$ follows from the fact that they are given by the eigenvalues of the channel and these are always contained in the unit circle of the complex plane [32]. □

Finally, we show that twirling does not change the entanglement fidelity and thus does not change the average fidelity, as observed in [30] and elsewhere in the literature. Thus, when we want to estimate the average fidelity of a channel $T : \mathcal{M}_d \to \mathcal{M}_d$ we may instead work with the twirled channel $\mathcal{T}(T)$ and explore its rich structure.

Theorem 12. *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, G be a finite group and $U : G \to \mathcal{M}_d$ be a unitary representation. Then*

$$F_e(T) = F_e(\mathcal{T}(T)). \tag{27}$$

Proof. We present a slightly different proof of this fact here. Note that $\mathcal{U}_g^* \circ T \circ \mathcal{U}_g$ is just a similarity transformation of T and thus $\operatorname{Tr}\left(\mathcal{U}_g^* \circ T \circ \mathcal{U}_g\right) = \operatorname{Tr}(T)$, where again we mean the trace of these channels as linear operators. Thus, integrating over all $\mathcal{U}_g$ does not change the entanglement fidelity, as $F_e(T) = d^{-2}\operatorname{Tr}(T)$. □

4. Randomized benchmarking protocol

The RB protocol, as discussed in [3, 4, 5, 6, 8, 33, 34, 35, 36, 37, 38] is a protocol to estimate the average fidelity of the implementation of gates coming from some group G. Its usual setting is the Clifford group, but we discuss it for general groups here. Other papers have investigated the protocol for gates beyond Cliffords, such as [18, 19, 21]. But all of these have restricted their analysis to some other specific group. As we will see later, we can analyse the protocol for arbitrary groups by just investigating properties of the given unitary representation. We mostly follow the notation of [37]. We assume that the error quantum channel is gate and time independent. That is, whenever we want to implement a certain gate $\mathcal{U}_g$, where $\mathcal{U}_g(\cdot) = U_g \cdot U_g^*$ with $U_g \in U(d)$, we actually implement $\mathcal{U}_g \circ T$ for some quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$. We assume that we are able to multiply and invert elements of G and draw samples from the Haar measure on G efficiently to implement this protocol, but will later relax this sampling condition. The protocol is as follows:

Step 1 Fix a positive integer $m \in \mathbb{N}$ that varies with every loop.

Step 2 Generate a sequence of $m+1$ quantum gates. The first m quantum gates $\mathcal{U}_{g_1}, \ldots, \mathcal{U}_{g_m}$ are independent and Haar distributed. The final quantum gate, $\mathcal{U}_{g_{m+1}}$ is chosen such that in the absence of errors the net sequence is just the identity operation,

$$\mathcal{U}_{g_{m+1}} \circ \mathcal{U}_{g_m} \circ \ldots \circ \mathcal{U}_{g_2} \circ \mathcal{U}_{g_1} = \mathrm{id}, \tag{28}$$

where $\circ$ represents composition. Thus, the whole quantum gate sequence is

$$\mathcal{S}_m = \bigcirc_{j=1}^{m+1} \mathcal{U}_{g_j} \circ T, \tag{29}$$

where T is the associated error quantum channel.

Step 3 For every sequence, measure the sequence fidelity

$$\operatorname{Tr}\left(\mathcal{S}_m(\rho) E\right), \tag{30}$$

where ρ is the initial quantum state, including preparation errors, and E is an effect operator of some POVM including measurement errors.

Step 4 Repeat steps 2-3 and average over M random realizations of the sequence of length m to find the averaged sequence fidelity given by

$$\bar{F}(m,E,\rho)=\frac{1}{M}\sum_{m}\operatorname{Tr}\left(\mathcal{S}_m(\rho)E\right). \tag{31}$$

Step 5 Repeat steps 1-4 for different values of m and obtain an estimate of the expected value of the sequence fidelity

$$F(m,E,\rho)=\operatorname{Tr}\left(\mathbb{E}(\mathcal{S}_m)(\rho)E\right). \tag{32}$$

4.1. Analysis of the Protocol

We will now show how we can estimate the average fidelity from the data produced by the protocol, that is, an estimate on the curve $F(m,E,\rho)=\operatorname{Tr}\left(\mathbb{E}(\mathcal{S}_m)(\rho)E\right)$.

Theorem 13. *Let $T:\mathcal{M}_d\to\mathcal{M}_d$ be a quantum channel and G a group with a unitary representation $U:G\to\mathcal{M}_d$. If we perform the RB protocol for G we have*

$$\mathbb{E}(\mathcal{S}_m)=\mathcal{T}(T)^m. \tag{33}$$

Proof. Although the proof is identical to the case in which G is given by the Clifford group, we will cover it here for completeness. Given some sequence $\{\mathcal{U}_{g_1},\ldots,\mathcal{U}_{g_{m+1}}\}$ of unitary gates from G, define the unitary operators

$$\mathcal{D}_i=\bigcirc_{j=1}^{i}\mathcal{U}_{g_i}.$$

Note that we have

$$\begin{aligned}\mathcal{S}_m =&\mathcal{U}_{g_{m+1}}\circ T\circ\mathcal{U}_{g_m}\circ T\circ\ldots\circ\mathcal{U}_{g_2}\circ T\circ\mathcal{U}_{g_1}\\ =&\overbrace{\mathcal{U}_{g_{m+1}}\circ(\mathcal{U}_{g_m}\circ\ldots\circ\mathcal{U}_{g_1}}^{=\mathbb{I}}\circ\overbrace{\mathcal{U}^*_{g_1}\circ\ldots\circ\mathcal{U}^*_{g_m})}^{=\mathcal{D}^*_m}\circ T\circ\mathcal{U}_{g_m}\circ T\circ\\ &\ldots T\circ\underbrace{\mathcal{U}_{g_3}\circ(\mathcal{U}_{g_2}\circ\mathcal{U}_{g_1}}_{=\mathcal{D}_3}\circ\underbrace{\mathcal{U}^*_{g_1}\circ\mathcal{U}^*_{g_2})}_{=\mathcal{D}^*_2}\circ T\circ\underbrace{\mathcal{U}_{g_2}\circ(\mathcal{U}_{g_1}}_{=\mathcal{D}_2}\circ\underbrace{\mathcal{U}^*_{g_1}}_{=\mathcal{D}^*_1})\circ T\circ\underbrace{\mathcal{U}_{g_1}}_{=\mathcal{D}_1}\\ =&\mathcal{D}^*_m\circ T\circ\mathcal{D}_m\circ\ldots\circ\mathcal{D}^*_2\circ T\circ\mathcal{D}_2\circ\mathcal{D}^*_1\circ T\circ\mathcal{D}_1\\ =&\bigcirc_{j=1}^{m}\left(\mathcal{D}^*_j\circ T\circ\mathcal{D}_j\right).\end{aligned} \tag{34}$$

As we have that each of the $\mathcal{U}_{g_i}$ is independent and Haar-distributed, it follows that the $\mathcal{D}_i$ are independent and Haar distributed as well. It then follows from (34) that

$$\mathbb{E}\left(\mathcal{S}_m\right)=\mathbb{E}\left(\bigcirc_{j=1}^{m}\left(\mathcal{D}^*_j\circ T\circ\mathcal{D}_j\right)\right)=\bigcirc_{j=1}^{m}\mathbb{E}\left(\mathcal{D}^*_j\circ T\circ\mathcal{D}_j\right)=\mathcal{T}(T)^m. \qquad \square$$

We can then use our structural results on covariant quantum channels to obtain a more explicit form for the curve $F(m,E,\rho)$.

Corollary 14. *Suppose we perform RB for a unitary representation $U : G \to \mathcal{M}_d$ of a finite group G s.t. $U \otimes \bar{U} = \oplus_{\alpha \in \hat{G}} \left(\mathbb{C}^{d_\alpha} \otimes \mathbb{C}^{m_\alpha}\right)$ and a channel T. Then there exist $\lambda_1, \ldots, \lambda_k \in \overline{B_1(0)}$ and $a_0, a_1, \ldots, a_k \in \mathbb{C}$ s.t.*

$$F(m, E, \rho) = a_0 + \sum_{i=1}^{k} a_k \lambda_i^m. \tag{35}$$

for $m \geq \max m_\alpha$. Moreover, $k \leq \sum_\alpha m_\alpha$ corresponds to the number of distinct eigenvalues of $\mathcal{T}(T)$ and λ_i are its eigenvalues.

Proof. The claim follows immediately from theorem 13 and lemma 11. □

That is, by fitting the curve to experimental data we may obtain estimates on the λ_i and thus on the spectrum of $\mathcal{T}(T)$. If we know the multiplicity of each eigenvalue, then we can estimate the trace as well and thus the average fidelity. However, in the case in which we have more than one parameter to estimate, it is not clear which eigenvalue corresponds to which irrep and we therefore cannot simply apply the formula in corollary 10. Thus, given an estimate $\{\hat{\lambda}_1, \ldots, \hat{\lambda}_k\}$ of the parameters, we define the minimal fidelity, $F_{\min}$, to be given by

$$F_{\min} = \min \sum d_\alpha \hat{\lambda}_i \tag{36}$$

and the maximum fidelity, $F_{\max}$, to be given by

$$F_{\max} = \max \sum d_\alpha \hat{\lambda}_i. \tag{37}$$

That is, we look at the pairings of d_α and $\hat{\lambda}_i$ that produces the largest and the smallest estimate for the fidelity. These then give the most pessimistic and most optimistic estimate, respectively. The fact that we cannot associate a λ_i to each irrep causes some problems in this approach from the numerical point of view and we comment on them in appendix A. Moreover, we also offer some preliminary ideas of how to overcome these issues.

5. Approximate Twirls

In the description of our RB protocol, we assume that we are able to obtain samples from the Haar measure of the group G. It is not possible or efficient to obtain samples of the Haar measure for most groups, but a lot of research has been done on how to obtain approximate samples efficiently using Markov chain Monte Carlo methods, as discussed in section 2.2. Here we discuss how to use samples which are approximately Haar distributed for RB. Note that these results may also be interpreted as a stability result w.r.t. not sampling exactly from the Haar measure of G. We will assume we are able to pick the $\mathcal{U}_{g_k}$ independently and that they are distributed according to a measure ν_k s.t.

$$\|\nu_k - \mu\|_1 \leq \epsilon_k, \tag{38}$$

for $\epsilon_k \geq 0$. Our goal is to show that under these assumptions we may still implement the RB protocol discussed before and obtain measurement statistics that are close to the ones obtained using Haar samples.

Motivated by this, we define the $\tilde{\nu}$-twirl of a channel.

Definition 15 ($\tilde{\nu}$-twirl to the power m)**.** Let $\tilde{\nu}$ be a probability measure on G^m, $T : \mathcal{M}_d \to \mathcal{M}_d$ a quantum channel and $U : G \to \mathcal{M}_d$ a $d-$dimensional unitary representation of G. We define the $\tilde{\nu}$-twirl to the power m to be given by

$$\mathcal{T}_{\tilde{\nu},m}(T) = \sum_{i_1,\ldots,i_m=1}^{|G|} \tilde{\nu}\left(g_{i_1},\ldots,g_{i_m}\right) \bigcirc_{k=1}^{m} \mathcal{U}_{g_{i_k}} \circ T \circ \mathcal{U}^*_{g_{i_k}}. \tag{39}$$

This definition boils down to the regular twirl for $\tilde{\nu} = \mu^{\otimes m}$, μ the Haar measure on G. We will now show that by sampling $\mathcal{U}_{g_k}$ close to Haar we have that the $\tilde{\nu}$-twirl of a channel is also close to the usual twirl.

Lemma 16 (Approximate Twirl)**.** *Let $T : \mathcal{M}_d \to \mathcal{M}_d$ be a quantum channel, G a finite group with a $d-$dimensional unitary representation $U : G \to \mathcal{M}_d$ and $\tilde{\nu}$ a probability measure on G^m. Let $\mathcal{T}_{\tilde{\nu},m}(T)$ be the $\tilde{\nu}$-twirl to the power m and $\mathcal{T}(T)$ be the twirl w.r.t. the Haar measure on G given by μ. Moreover, let $\|\cdot\|$ be a norm s.t. $\|T\| \leq 1$ for all quantum channels. Then*

$$\|\mathcal{T}_{\tilde{\nu},m}(T) - \mathcal{T}(T)^m\| \leq 2\|\tilde{\nu} - \mu^{\otimes m}\|_1. \tag{40}$$

Proof. Observe that we may write

$$\|\mathcal{T}_{\tilde{\nu},m}(T) - \mathcal{T}(T)^m\| = \left\| \sum_{i_1,\ldots,i_m=1}^{|G|} \left(\tilde{\nu}\left(g_{i_1},\ldots,g_{i_m}\right) - \frac{1}{|G|^m}\right) \bigcirc_{k=1}^{m} \mathcal{U}_{g_{i_k}} \circ T \circ \mathcal{U}^*_{g_{i_k}} \right\|.$$

The claim then follows from the triangle inequality and the fact that $\| \bigcirc_{k=1}^{m} \mathcal{U}_{g_{i_k}} \circ T \circ \mathcal{U}^*_{g_{i_k}} \| \leq 1$. □

Thus, in order to bound $\|\mathcal{T}_{\tilde{\nu},m}(T) - \mathcal{T}(T)^m\|$ in any norm in which quantum channels are contractions, it suffices to bound $\|\tilde{\nu} - \mu^{\otimes m}\|_1$. Examples of such norms are the $1 \to 1$ norm and the diamond norm [39, Theorem 2.1]. We remark that other notions of approximate twirling were considered in the literature [38, 40], but these works were mostly concerned with the case of the unitary group and not arbitrary finite groups. Although it would be straightforward to adapt their definitions to arbitrary finite groups, it is not clear at first sight that their notions of approximate twirls behave well when taking powers of channels that have been twirled approximately. This is key for RB. Given random unitaries $\{U_i\}_{i=1}^m$ from G, let $\mathcal{D}_k = \bigcirc_{i=1}^{k} \mathcal{U}_i$, as before.

Theorem 17. *Let μ be the Haar measure on G and $\nu_1, \ldots, \nu_m$ probability measures on G s.t.*

$$\|\mu - \nu_k\|_1 \leq \epsilon_k, \tag{41}$$

for all $1 \leq k \leq m$ *and* $\epsilon_k \geq 0$. *Denote by* $\tilde{\nu}$ *the distribution of* $(\mathcal{D}_1, \ldots, \mathcal{D}_m)$ *if we pick the* $\mathcal{U}_k$ *independently from* ν_k . *Then*

$$\|\mathcal{T}_{\tilde{\nu},m}(T) - \mathcal{T}(T)^m\|_{1\to 1} \leq 4\sqrt{\frac{\log(|G|)}{1-|G|^{-1}}\sum_{k=1}^{m}\epsilon_k}. \tag{42}$$

Proof. We refer to appendix C for a proof. □

Note that the same result holds for any norm that contracts under quantum channels, such as the diamond norm.

Corollary 18. *Let* μ *be the Haar measure on* G *and* $\nu_1, \ldots, \nu_m$ *probability measures on* G *s.t.*

$$\|\mu - \nu_k\|_1 \leq \epsilon_k, \tag{43}$$

for all $1 \leq k \leq m$ *and* $\epsilon_k \geq 0$. *Denote by* $\tilde{\nu}$ *the distribution of* $(\mathcal{D}_1, \ldots, \mathcal{D}_m)$ *if we pick the* $\mathcal{U}_k$ *independently from* ν_k. *Then*

$$|\operatorname{Tr}(\mathcal{T}_{\tilde{\nu},m}(T)(\rho)E) - F(m,E,\rho)| \leq 4\sqrt{\frac{\log(|G|)}{1-|G|^{-1}}\sum_{k=1}^{m}\epsilon_k}. \tag{44}$$

Proof. It follows from Hölder's inequality that

$$\begin{aligned}|\operatorname{Tr}(\mathcal{T}_{\tilde{\nu},m}(T)(\rho)E) - F(m,E,\rho)| =&|\operatorname{Tr}\left(E\left(\mathcal{T}_{\tilde{\nu},m}(T)(\rho) - \mathcal{T}\left(T\right)^m(\rho)\right)\right)| \\ \leq&\|E\|_\infty\|\mathcal{T}_{\tilde{\nu},m}(T) - \mathcal{T}(T)^m\|_{1\to 1}.\end{aligned}$$

As E is the element of a POVM, we have $\|E\|_\infty \leq 1$ and the claim then follows from theorem 17. □

This shows that we may use approximate twirls instead of exact ones and obtain expectation values that are close to the perfect twirl. Given that we want to assure that the statistics we obtain for some $m \in \mathbb{N}$ are $\delta > 0$ close to our target distribution, we would have to sample the U_{g_k} such that

$$\|\mu - \nu_k\|_1 \leq \frac{\delta^2(1-|G|^{-1})}{16\log(|G|)m}, \tag{45}$$

as can be seen by plugging in this bound in the result of corollary 18. If we use a random walk on a group to sample from the Haar distribution we have to run each chain for $t_1\left(\frac{\delta^2(1-|G|^{-1})}{16\log(|G|)m}\right)$ steps, which gives a total runtime of $\mathrm{O}\left(t_{\mathrm{mix}}\log\left(\frac{16\log(|G|)m}{\delta^2(1-|G|^{-1})}\right)\right)$. For a fixed δ, this will be efficient if the chain mixes rapidly, that is, t_{mix} is small, and we choose m to be at most of the order of the dimension.

6. Randomized benchmarking with generators

One of the downsides of the usual RB protocol [3, 4, 5, 6, 33, 34, 35, 36, 37, 38] is that we assume that we may implement any gate of the group. Usually, gates have to be broken down into generators, as discussed in [41, Section 1.2.3 and Chapter 8]. Therefore, it would be desirable both from the point of view of justifying the noise model and the implementation level of the protocol to mostly need to implement gates from a set of generators. We describe here a protocol to perform RB by just implementing gates from a set of generators closed under inversion and one arbitrary gate. We also make the additional assumption that the quantum channel that describes the noise is already approximately covariant in a sense we will make precise soon. This protocol is inspired by results of the last section that suggest a way of performing RB by just implementing gates coming from a set A that generates the group G and is closed under inversion and one additional arbitrary gate from G at each round of the protocol. From the basic results of random walks discussed in section 2.2, we know that if we pick gates $U_{g_1}, U_{g_2}, \ldots$ uniformly at random from A, it follows that $U_{g_b}U_{g_{b-1}}\ldots U_{g_1}$ will be approximately distributed like the Haar measure on G for $b \simeq t_{\text{mix}}$. However, one should note that in this setting the $\mathcal{D}_i$ will not be independent of each other. To see this, note that given $\mathcal{D}_i = \mathcal{U}_g$, we know that the distribution of the D_{i+1} is restricted to elements $h \in G$ of the form $h = ag$ with $a \in A$, which clearly show that they are not independent in general. However, if we look at $\mathcal{D}_{i+l}$ for $l \sim t_{\text{mix}}$, then their joint distribution will be close to Haar. That is, looking at $\mathcal{D}_i$ and $\mathcal{D}_j$ which are far enough apart from each other, we may again assume that they are both almost Haar distributed and if we look at each $\mathcal{D}_i$ individually we may assume that they are almost Haar distributed. One way to explore this observation for RB protocols only having to implement the generators is to look at the following class of quantum channels:

Definition 19 (δ-covariant quantum channel)**.** A quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ is called δ-covariant with respect to a unitary representation $U : G \to \mathcal{M}_d$ of a group G, if there exist quantum channels $T_c, T_n : \mathcal{M}_d \to \mathcal{M}_d$ such that

$$T = (1-\delta)T_c + \delta T_n, \tag{46}$$

and T_c is covariant with respect to U.

That is, T is almost covariant with respect to the group. Similar notions of approximate covariance were also introduced in [42]. The standard example for our purposes are quantum channels that are close to the identity channel, i.e., we have δ small and T_c the identity channel.

We will need to fix some notation before we describe the protocol. For a given sequence of unitaries $s_i = (U_{g_1}, U_{g_2}, \ldots)$ we let $\mathcal{S}_{s_i,c,d} = \bigcirc_{j=c}^{d} \mathcal{U}_{g_j} \circ T$ for $c, d \in \mathbb{N}$ and the gates chosen according to the sequence.

Thus, if we apply random generators b times as an initialization procedure and only start fitting the curve after this initialization procedure we may also estimate the average fidelity.

This yields the following protocol.

Step 1 Fix a positive integer $m \in \mathbb{N}$ that varies with every loop and another integer $b \in \mathbb{N}$.

Step 2 Generate a sequence of $b+m+1$ quantum gates, s_i. The first $b+m$ quantum gates $\mathcal{U}_{g_1}, \ldots, \mathcal{U}_{g_{b+m}}$ are chosen independently and uniformly at random from A. The final quantum gate, $\mathcal{U}_{g_{b+m+1}}$ is chosen as

$$\mathcal{U}_{g_{b+m+1}} = (\mathcal{U}_{g_{b+m}} \circ \ldots \circ \mathcal{U}_{g_2} \circ \mathcal{U}_{g_1})^{-1}. \tag{47}$$

Step 3 For each sequence s_i, measure the sequence fidelity

$$\operatorname{Tr}\left(\mathcal{S}_{s_i,b+1,b+m+1}(\mathcal{S}_{s_i,1,b}(\rho))E\right), \tag{48}$$

where ρ is the initial quantum state and E is an effect operator of a POVM.

Step 4 Repeat steps 2-3 and average over M random realizations of the sequence of length m to find the averaged sequence fidelity

$$\bar{F}(m, E, \rho) = \frac{1}{M}\sum_{i=1}^{M} \operatorname{Tr}\left(\mathcal{S}_{s_i,b+1,b+m+1}(\mathcal{S}_{s_i,1,b}(\rho))E\right). \tag{49}$$

Step 5 Repeat steps 1-4 for different values of m to obtain an estimate of the expected value of the average survival probability

$$F(m, E, \rho) = \mathbb{E}\left(\operatorname{Tr}\left(\mathcal{S}_{s_i,b+1,b+m+1}(\mathcal{S}_{s_i,1,b}(\rho))E\right)\right). \tag{50}$$

We will now prove that this procedure gives rise to the same statistics as if we were using samples from the Haar distribution up to $\mathrm{O}(\delta^2)$.

Theorem 20. *Let T be $\delta-$covariant w.r.t. a unitary representation $U : G \to \mathcal{M}_d$ of a finite group G, A a subset of G that generates G and is closed under inversion and $\delta > 0$. Suppose we run the protocol above with $b = t_{mix}(m^{-1}\epsilon)$ for some ϵ. Furthermore, let π be the doubly-stochastic matrix associated with the random walk induced by A and suppose that*

$$\|\pi(\nu) - \mu\|_1 \leq \lambda \tag{51}$$

for all $\nu \in \mathcal{P}(G)$ and some $\lambda \in [0, 1)$. Then

$$\|\mathcal{T}(T)^m - \mathbb{E}(\mathcal{S}_{b,b+m+1})\|_{1\to 1} \leq \epsilon + \mathrm{O}\left(\delta^2 \frac{\lambda}{1-\lambda} m\right). \tag{52}$$

Proof. We refer to appendix D for a proof. □

Using standard methods from Markov chains it is possible to show that one can always choose $\lambda = (|G| - |A|)|G|^{-1}$.

Corollary 21. *Let $\mathcal{S}_{b,m+b+1}$ and λ be as in theorem 20. Then for any POVM element E and state $\rho \in \mathcal{M}_d$:*

$$|\operatorname{Tr}\left(\mathbb{E}\left(S_{b,m+1}\right)(\rho)E\right) - F(m,E,\rho)| \leq \epsilon + \mathrm{O}\left(\delta^2 \frac{\lambda}{1-\lambda} m\right). \tag{53}$$

Proof. The proof is essentially the same as that of corollary 18. □

This shows that performing RB by only implementing the generators is feasible as long as we have a $\delta-$covariant channel with δ small and a rapidly mixing set of generators.

7. Numerics and Examples

Here we show how to apply our methods to groups that might be of special interest and discuss some numerical examples. Many relevant questions for the practical application of our work are still left open and have two different flavors: the numerical and statistical side. From the numerical point of view, it is not clear at first how to fit the data gathered by a RB protocol to an exponential curve if we have several parameters. We refer to appendix A for a discussion of these issues and some proposals of how to overcome them. From a statistical point of view, it is not clear how to derive confidence intervals for the parameters and how large we should choose the different parameters of the protocol, such as m and M. We refer to appendix B for a discussion of these issues and preliminary results in this direction.

7.1. Monomial Unitary Matrices

We consider how to apply our methods of generalized RB to some subgroups of the monomial unitary matrices $MU(d)$.

Definition 22. Let $\{|i\rangle\}_{i=1}^d$ be an orthonormal basis of $\mathbb{C}^d$. We define the group of monomial unitary matrices, $MU(d)$ to be given by $U \in U(d)$ of the form $U = DP$ with $D, P \in U(d)$ and D diagonal w.r.t. $\{|i\rangle\}_{i=1}^d$ and P a permutation matrix.

Subgroups of this group can be used to describe many-body states in a formalism that is broader than the stabilizer formalism of Paulis and have other applications to quantum computation (see [22]). As the group above is not finite and it is unreasonable to assume that we may implement diagonal gates with phases of an arbitrary precision, we focus on the following subgroups:

Definition 23. We define $MU(d,n)$ to be the subgroup of the monomial unitary matrices of dimension d whose nonzero entries consist only of $n-$th roots of unity.

Another motivation to consider these subgroups is that they contain the T-gate [23],

$$T = |0\rangle\langle 0| + e^{\mathrm{i}\frac{\pi}{4}} |1\rangle\langle 1| \tag{54}$$

in case $n \geq 8$. Thus these gates, together with Cliffords, constitute a universal set of quantum gates [23]. We now show that we have to estimate two parameters for them.

Lemma 24 (Structure of channels covariant w.r.t. monomial unitaries)**.** *Let $MU(d,n)$ be such that $n \geq 3$ and $T : \mathcal{M}_d \to \mathcal{M}_d$ a quantum channel. Then the following are equivalent:*

(i) $T(\rho) = UT\big(U^*\rho U\big)U^* \quad \forall U \in MU(d,n), \rho \in \mathcal{S}_d$.

(ii) There are $\alpha, \beta \in \mathbb{R}$ so that

$$T(\cdot) = \mathrm{Tr}\left(\cdot\right)\frac{\mathbb{I}}{d} + \alpha\left(\mathrm{id} - \sum_{i=1}^{d} |i\rangle\langle i|\,\langle i|\cdot|i\rangle\right) + \beta\left(\sum_{i=1}^{d} |i\rangle\langle i|\,\langle i|\cdot|i\rangle - \mathrm{Tr}\left(\cdot\right)\frac{\mathbb{I}}{d}\right). \quad (55)$$

Moreover, the terms in the r.h.s. of (55) *are projections of rank* 1*,* $d^2 - d$ *and* $d - 1$*, respectively.*

Proof. We refer to appendix E for a proof. □

This result shows that we only need to estimate two parameters when performing RB with these subgroups. They are therefore a natural candidate to apply our methods to and we investigate this possibility further. We begin by analyzing the complexity of multiplying and inverting elements of $MU(d,n)$. We show this more generally for $MU(d)$, as it clearly gives an upper bound for its subgroups as well. We may multiply and invert elements of $MU(d)$ in time $\mathrm{O}(d)$. To multiply elements in $MU(d)$ we need to multiply two permutations of d elements, which can be done in time $\mathrm{O}(d)$, multiply a vector $u \in \mathbb{C}^d$ with a permutation matrix, which can be done in time $\mathrm{O}(d)$, and multiply d elements of $U(1)$ with each other, which again can be done in time $\mathrm{O}(d)$. This shows that multiplying elements of this group takes $\mathrm{O}(d)$ operations. To invert an element of $MU(d)$ we need to invert a permutation, which again takes $\mathrm{O}(d)$, invert d elements of $U(1)$ and apply a permutation to the resulting vector. This also takes $\mathrm{O}(d)$ operations. Moreover, one can generate a random permutation and an element of $U(1)^d$ in time $\mathrm{O}(d)$, giving $\mathrm{O}(Mmd)$ complexity for the classical part of the RB procedure. Although this scaling is not efficient in the number of qubits as in the case of Clifford gates [3], the fact that it is linear in the dimension and not superquadratic as in the general case still allows for our method to be applied to high dimensions.

To exemplify our methods, we simulate our algorithm for some dimensions and number of samples M. We run the simulations for $MU(d,8)$, as it is the smallest one that contains the T gate. We consider the case of a quantum channel T that depolarizes to a random state $\sigma \in \mathcal{D}_d$ with probability $(1-p)$, that is

$$T(\rho) = p\rho + (1-p)\sigma, \quad (56)$$

where $\sigma \in \mathcal{D}_d$ is chosen uniformly at random from the set of states. It is not difficult to see that in this case the entanglement fidelity $F_e(T) = (p(d^2-1)+1)/d^2$ and we, therefore, measure our error in terms of the parameter p. The results are summarized in table 1. These numerical results clearly show that we may estimate the fidelity to a good degree with our procedure.

Table 1. Error analysis of the RB protocol described in section 4 to the group $MU(d, 8)$. We take the initial state to be $|0\rangle\langle 0|$, the POVM element to be $|0\rangle\langle 0|$, $p = 0.9$ and we always choose $m = 40$. Moreover, we generate 100 different channels for each combination of dimension and number of samples. The table shows the resulting mean and median error as well as the standard deviation for different values of d and M.

d	M	Mean Error ($\times 10^{-3}$)	Median Error ($\times 10^{-3}$)	Standard Deviation ($\times 10^{-3}$)
64	1000	9.17	2.14	3.93
128	100	6.08	1.48	2.14
128	1000	5.17	1.01	1.13
1024	100	9.17	2.14	3.93
1024	1000	4.55	1.13	1.77

7.2. Clifford Group

As mentioned before, the Clifford group is the usual setup of RB, as we only have to estimate one parameter and it is one of the main building blocks of quantum computing [23]. Thus, we apply our protocols based on approximate samples of the Haar distribution and generator based protocols to Clifford gates. It is known that the Clifford group on n qubits, $\mathcal{C}(n)$, is generated by the Hadamard H, the $\pi-$gate and the $CNOT$ gate between different qubits. We refer to e.g. [43, Section 5.8] for a proof of this claim. We need a set of generators that is closed under taking inverses for our purposes. All but the $\pi-$gate are their own inverse, so we add the inverse of the $\pi-$gate to our set of generators to assure that the random walk converges to the Haar measure on the Clifford group. That is, we will consider the set A of generators of the Clifford group $\mathcal{C}(n)$ consisting of Hadamard gates, $\pi-$gates and its inverse on each individual qubit and $CNOT$ between any two qubits,

$$A = \{\pi_i, \pi_i^{-1}, H_i, CNOT_{i,j}\}. \tag{57}$$

To the best of our knowledge, there is no rigorous estimate available for the mixing time of the random walk generated by A and it would certainly be interesting to investigate this question further. However, based on our numerical results and the results of [40], we conjecture that it is rapidly mixing, i.e. $t_{\text{mix}} = \mathrm{O}(n^2 \log(n))$. This would be more efficient than the algorithm proposed in [44], which takes $\mathrm{O}(n^3)$ operations. To again test our methods we perform similar numerics as in the case of the monomial unitaries.

We simulate the following noise model: We first pick a random isometry $V : (\mathbb{C}^2)^{\otimes n} \to (\mathbb{C}^2)^{\otimes n} \otimes (\mathbb{C}^2)^{\otimes n}$ and generate the quantum channel $T(\rho) = p\rho + (1 - p)\mathrm{tr}_2\left(V \rho V^*\right)$, where tr_2 denotes the partial trace over the second tensor factor. That is, T is just the convex combination of the identity and a random channel and is δ-covariant w.r.t. to a group with $\delta = p$. From the discussion in section 6 we expect this to work best for p close to 1. The results for p close to 1 are summarized in table 2. The average

Table 2. For each combination of p, M and b we generate 20 different random quantum channels and perform generator RB for the Clifford group on 5 qubits. In all these cases we pick $m = 20$. The average error is defined as the average of the absolute value between the exact fidelity and the one estimated using our protocol. The table shows the average error and its standard deviation in terms of different choices of b, M and p.

p	b	M	Average Error ($\times 10^{-3}$)	Standard Deviation of Error ($\times 10^{-4}$)
0.98	10	10	5.49	1.38
0.95	10	100	1.44	3.92
0.95	5	100	1.52	7.94
0.95	5	20	1.56	7.44
0.90	10	20	3.20	1.58
0.80	10	50	8.63	6.01

error increases as the channel becomes noisier, but generally speaking we are able to obtain an estimate which is 10^{-3} close to the true value with M around 20 and $m = 20$.

We also performed some numerical experiments for p significantly away from 1, which are summarized in table 3.

Table 3. For each combination of p, M and b we generate 20 different random quantum channels and perform generator RB for the Clifford group on 5 qubits. In all these cases we pick $m = 20$. The average error is defined as the average of the absolute value between the exact fidelity and the one estimated using our protocol. The table shows the average error and its standard deviation in terms of different choices of b, M and p.

p	b	M	Average Error ($\times 10^{-2}$)	Standard Deviation of Error ($\times 10^{-3}$)
0.7	5	100	2.07	1.15
0.65	5	100	2.29	1.95
0.60	5	100	27.1	52.30
0.55	5	100	44.5	67.30

These results show that these methods are effective to estimate the average fidelity under less restrictive assumptions on the gates we may implement using RB if we have a high fidelity, as indicated in table 2. However, in case we do not have a high fidelity, these methods are not reliable, as can be seen in table 3. This should not severely restrict the applicability of these methods, as one is usually interested in the high fidelity regime when performing RB.

Finally, in figure 1 we compare the three different RB protocols discussed in this paper. We compare the usual RB protocol, which we call the perfect sampling protocol, to the one with approximate samples and the generator RB. The curve makes clear that using approximate and exact samples leads to virtually indistinguishable estimates and

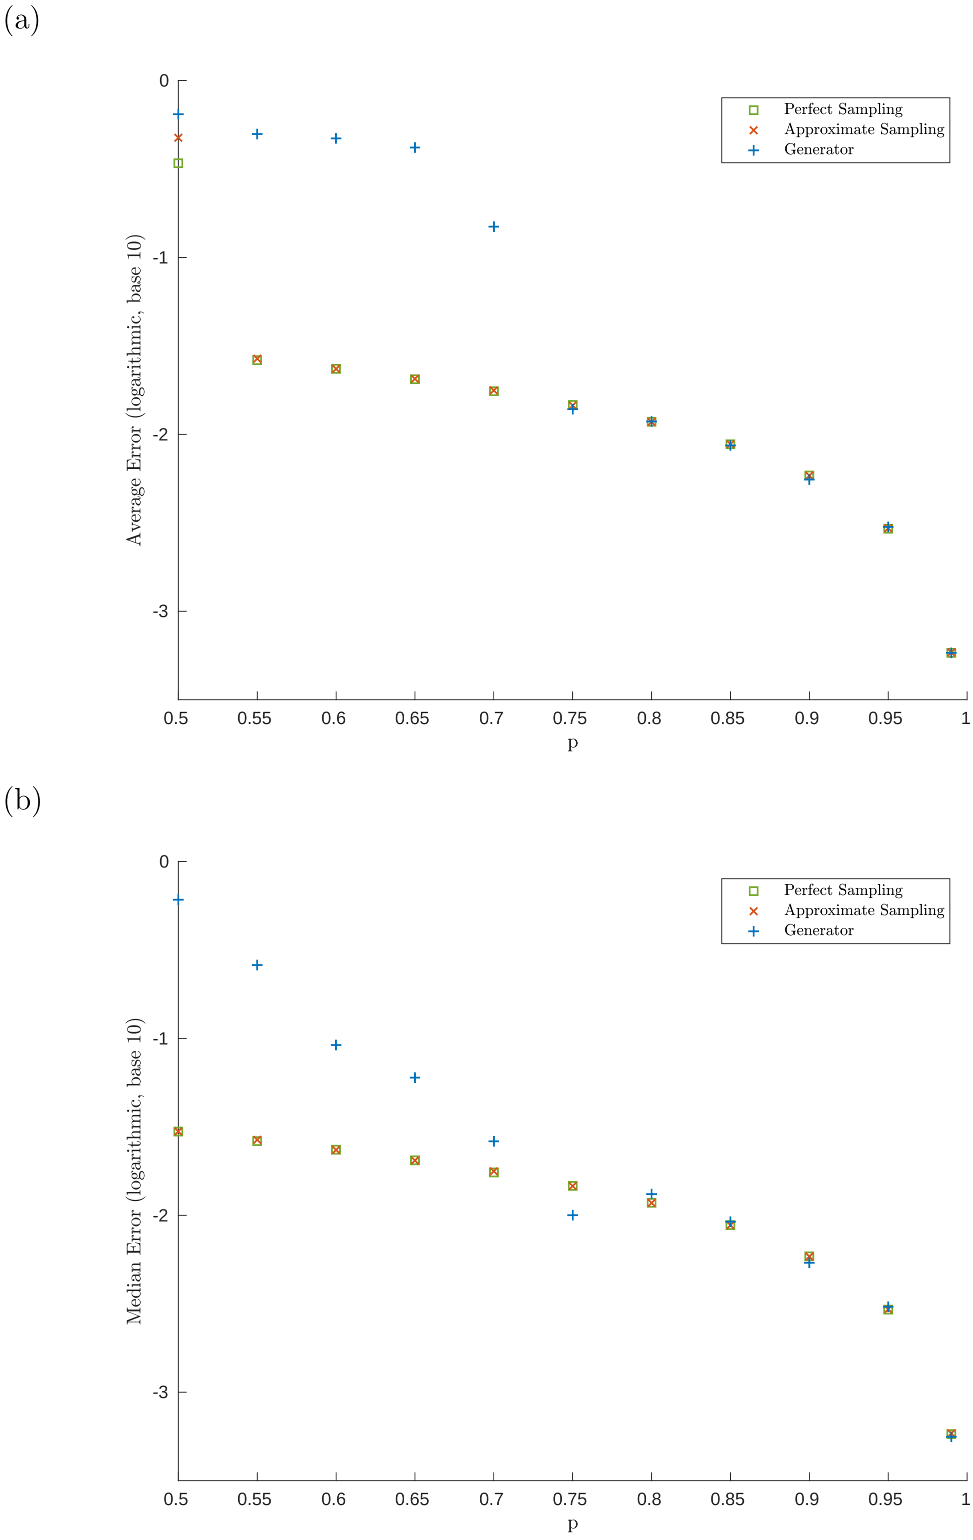

Figure 1. Plot of the average error (a) and mean error (b) as a function of p for different versions of the RB protocol. For each value of p we generated 20 with $M = 100$ and $m = 20$. For the generator RB we chose $b = 5$ and to obtain the approximate samples we ran the chain for 20 steps.

that all protocols have similar performance for p close to 1.

8. Conclusion and Open Problems

We have generalized the RB protocol to estimate the average gate fidelity of unitary representations of arbitrary finite groups. Our protocol is efficient when multiplying, inverting and sampling elements from the group can be done efficiently and we have shown some potential applications that go beyond the usual Clifford one. Moreover, we showed that using approximate samples instead of perfect ones from the Haar measure on the group does not lead to great errors. This can be seen as a stability result for RB protocols w.r.t. sampling which was not available in the literature and is also relevant in the Clifford case. We hope that this result can be useful in practice when one is not given a full description of the group but rather a set of generators. Moreover, we have shown how to perform RB by just implementing a set of generators and one arbitrary gate under some noise models. This protocol could potentially be more feasible for applications, as the set of gates we need to implement is on average simpler.

However, some questions remain open and require further work. It is straightforward to generalize the technique of interleaved RB to this more general scenario and this would also be a relevant development. It would be important to derive confidence intervals for the estimates as was done for the Clifford case in [7]. Moreover, it would be relevant to estimate not only the mean fidelity but also the variance of this quantity. The assumption that the noisy channel is the same for all gates is not realistic in many scenarios and should be seen as a 0-order approximation, as in [36]. It would be desirable to generalize our results to the case in which the channel depends weakly on the gate.

Acknowledgments

D.S.F. acknowledges support from the graduate program TopMath of the Elite Network of Bavaria, the TopMath Graduate Center of TUM Graduate School at Technische Universität München and by the Technische Universität München Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763.

A.K.H.'s work is supported by the Elite Network of Bavaria through the PhD programme of excellence *Exploring Quantum Matter*.

Appendix

A. Numerical Considerations

Here we gather some comments on the numerical issues associated with the RB procedure when estimating more than one parameter.

A.1. Fitting the Data to Several Parameters

In order to be able to estimate the average fidelity following the protocols discussed so far, it is necessary to fit noisy data points $\{x_i\}_{i=1}^m \subset \mathbb{R}$ to a curve $f : \mathbb{R} \to \mathbb{R}$ of the form

$$f(x) = a_0 + \sum_{k=1}^{n} a_k \mathrm{e}^{-b_k x}, \tag{A.1}$$

with $a_0, a_1, \ldots, a_n, b_1, \ldots, b_n \in \mathbb{C}$. Although this may look like an innocent problem at first sight, fitting noisy data to exponential curves is a difficult problem from the numerical point of view for large n. It suffers from many stability issues, as thoroughly discussed in [45]. Here we are going to briefly comment on some of the issues and challenges faced when trying to fit the data, although we admittedly only scratch the surface. For a more thorough analysis of some methods and issues, we refer to [45, 46].

We assume that we know the maximum number of different parameters, $2n + 1$, which we are fitting. This is given by the structure of the unitary representation at hand, as discussed in lemma 11. Luckily, significant progress has been made in the recent years to develop algorithms to overcome the issues faced in this setting and it is now possible to fit curves to data with a moderate number of parameters. It is also noteworthy that for $n = 2$ there exist stable algorithms based on geometric sums [46] which works for equispaced data, as is our case. For estimating more than two parameters one can use the algorithms proposed in [45], available at [47]. It should be said that the reliability and convergence of most algorithms found in the literature depends strongly on the choice of a good initial point. This tends not to be a problem, as we might have some assumptions where our fidelity approximately lies and choose the initial b_k accordingly. What could be another source of numerical instabilities is the fact that we have to input the model with a number of parameters, n. In case the eigenvalues of T are very close for different irreps, then this will lead to numerical instabilities. This is the case if the noise is described by a depolarizing channel, for example. Furthermore, it might be the case that the initial state in our protocol does not intersect with all eigenspaces of the channel. This may lead to some parameters a_k being 0 and we are not able to estimate some of the b_k from them.

Moreover, it is in principle not possible to tell which parameter corresponds to which irrep given the decomposition in lemma 11, which is again necessary to estimate the trace of the channel. So even in the case in which we have a small number of parameters, it is important that the different irreps associated to our parameters have a similar dimension or to assume that the spectrum of the twirled channel contains eigenvalues that are very close to each other. In this way, the most pessimistic estimate on the fidelity, as defined in (36), is not very far from the most optimistic, defined in (37). This is one of the reasons we focus on examples that only have a small number of parameters, say 1 or 2, and irreps of a similar dimension to avoid having numerical instabilities or estimates that range over an interval that is too large.

It is therefore important to develop better schemes to fit the data in the context of RB for more than one or two parameters. This is important from a statistical point of view, as it would be desirable to obtain confidence intervals for the parameters from the RB data. We will further develop this issue in appendix B. It would be worthwhile pursuing a Bayesian approach to this problem, as was done in [48] for the usual RB protocol.

A.2. Isolating the parameters

One way to possibly deal with issue is to isolate each parameter, that is, by preparing states that only have support on one of the irreps that are not the trivial one. In the case of non-degenerate unitary representations, discussed in theorem 5, we have the following:

Theorem 25 (Isolating parameters)**.** *Let $U : G \to \mathcal{M}_d$ be a simply irrep of a finite group G and $T : \mathcal{M}_d \to \mathcal{M}_d$ a channel which is covariant w.r.t. U. Then, for all eigenvalues λ_α there is a quantum state $\rho_\alpha = \frac{\mathbb{I}}{d} + X$, where $X = X^*$ and* $\mathrm{Tr}\,(X) = 0$*, such that*

$$T^m\left(\rho_\alpha\right) = \frac{\mathbb{I}}{d} + \lambda_\alpha^m X. \tag{A.2}$$

Proof. Consider the projections to the irreducible subspaces P_α defined in (11). For a self-adjoint operator $X \in \mathcal{M}_d$ we have that

$$P_\alpha(X)^* = \frac{\chi^\alpha(e)}{|G|} \sum_{g \in G} \chi^\alpha\left(g\right) U_g^* X U_g = P_\alpha(X),$$

as we are summing over the whole group and $\overline{\chi^\alpha\left(g^{-1}\right)} = \chi^\alpha\left(g\right)$. Therefore, we have that the P_α are hermiticity preserving. As the image of P_α is the eigenspace corresponding to the irreps, we thus only have to show that there exists a self-adjoint X such that $P_\alpha(X) \neq 0$. But the existence of such an X is clear, as we may choose a basis of $\mathcal{M}_d$ that consists of self-adjoint operators. Moreover, as for α not the trivial representation all eigenvectors are orthogonal to $\mathbb{I}$, it follows that $\mathrm{Tr}\,(X) = 0$ and that for $\epsilon > 0$ small enough $\frac{\mathbb{I}}{d} + \epsilon X$ is positive semidefinite. To finish the proof, note that simply irreducible channels always satisfy

$$T(\mathbb{I}) = \mathbb{I}. \qquad \square$$

Note that this also proves that the spectrum of irreducibly covariant channels is always real. That is, if we can prepare a state such as in (A.2), then we can perform the

RB with this as an initial state and estimate the eigenvalue corresponding to each irrep. This would bypass the problems discussed in appendix A.1. The proof of theorem 25 already hints a way of determining how to isolate the parameter: just apply the projector P_α to some states ρ_i. If the output is not 0, then we can in principle write down a state that "isolates" the parameter as in the proof of theorem 25. But we admit that it is not clear at this stage how to prepare such states in a simple and efficient way or how to extend such results to unitary representations which are not simply irreducible, although it is certainly a direction which could be further investigated. This approach would also lose one of the main advantages of the RB protocol, namely that it does not assume that we are able to prepare a specific initial state [7]. However, in the case of the monomial unitary matrices discussed in section 7.1, we can examine the projections and see how to isolate the parameters. To isolate the parameter α in (55), we can prepare initial states $\rho \in \mathcal{D}_d$ that have $1/d$ as their diagonal elements and at least one nonzero off-diagonal element, as then the projector corresponding to β vanishes on ρ and does not vanish on the one corresponding to α. To isolate the parameter β, one can prepare states ρ that are diagonal in the computational basis but are not the maximally mixed state, as can be seen by direct inspection.

B. Statistical Considerations

One of the main open questions left in our work is how to derive good confidence intervals for the average fidelity. For the case of the Clifford group, discussed in section 7.2, one can directly apply the results of [7, 8], but it is not clear how one should pick m and M for arbitrary finite groups. Especially in the case in which we are not working with Cliffords, it is not clear how many samples per point, M, we should gather and how big m should be, as it depends on the choice of the algorithm picked for fitting the curve. As noted in appendix A.1, this is not a trivial problem from a numerical point of view. However, it is possible to obtain estimates on how much the observed survival probability deviates from its expectation value by just using Hoeffding's inequality:

Theorem 26. *Let $\bar{F}(m,E,\rho)$ be the observed average fidelity with M samples and $F(m,E,\rho)$ the average fidelity for any of the protocols discussed before and $\epsilon > 0$. Then:*

$$\mathbb{P}(|F(m,E,\rho) - \bar{F}(m,E,\rho)| \geq \epsilon) \leq \mathrm{e}^{-2M\epsilon^2}. \tag{B.1}$$

Proof. This is just a straightforward application of Hoeffding's inequality [49], as $\bar{F}(m,E,\rho)$ is just the empirical average of a random variable whose value is contained in $[0,1]$ and whose expectation value is $F(m,E,\rho)$. □

This bound is extremely general, as we did not even have to use any property of the random variables or of the group at hand. One should not expect it to perform well for specific cases and the scaling it gives is still undesirable for applications. Indeed, to assure we are 10^{-4} close to the expectation value with probability of 0.95, we need around 6×10^8 samples, which is not feasible. Thus, it is necessary to derive more refined bounds for specific groups.

C. Proof of Theorem 17

Proof. From lemma 16 it suffices to show

$$\|\tilde{\nu} - \mu^{\otimes m}\|_1 \leq 2\sqrt{\frac{\log(|G|)}{1-|G|^{-1}}\sum_{k=1}^{m}\epsilon_k},$$

as the 1 → 1 norm contracts under quantum channels [39].

We will first show that

$$\|\tilde{\nu} - \mu^{\otimes m}\|_1 = \| \otimes_{k=1}^{m} \nu_k - \mu^{\otimes m}\|_1.$$

We may rewrite the distribution $\tilde{\nu}$ in terms of the ν_k as follows:

$$\begin{aligned}\mathbb{P}(\mathcal{D}_1 = g_1, \mathcal{D}_2 = g_2, \ldots, \mathcal{D}_m = g_m) =& \mathbb{P}(U_1 = g_1, U_2 = g_2 g_1^{-1}, \ldots, U_m = g_m g_{m-1}^{-1}) \\ =& \nu_1(g_1)\nu_2(g_2 g_1^{-1}) \ldots \nu_m(g_m g_{m-1}^{-1}),\end{aligned}$$

as the U_{g_i} are independent.

Note that the map $\sigma : G^m \to G^m$, $(g_1, \ldots, g_m) \mapsto \left(g_1, g_2 g_1^{-1}, \ldots, g_m g_{m-1}^{-1}\right)$ is bijective. Moreover, we have $\tilde{\nu} = \otimes_{k=1}^{m} \nu_k \circ \sigma$. As the total variation norm is invariant under compositions with bijections on the state space, we have

$$\|\tilde{\nu} - \mu^{\otimes m}\|_1 = \| \otimes_{k=1}^{m} \nu_k \circ \sigma - \mu^{\otimes m}\|_1 = \| \otimes_{k=1}^{m} \nu_k - \mu^{\otimes m} \circ \sigma^{-1}\|_1 = \| \otimes_{k=1}^{m} \nu_k - \mu^{\otimes m}\|_1,$$

where the last equality follows from the fact that the Haar measure is invariant under bijections. We will now bound $\| \otimes_{k=1}^{m} \nu_k - \mu^{\otimes m}\|_1$. By Pinsker's inequality [50], we have

$$\| \otimes_{k=1}^{m} \nu_k - \mu^{\otimes m}\|_1^2 \leq 4D\left(\otimes_{k=1}^{m}\nu_k || \mu^{\otimes m}\right) = 4\sum_{k=1}^{m} D\left(\nu_k||\mu\right). \tag{C.1}$$

Here D is the relative entropy. In [50, Theorem 1] they show that

$$D\left(\nu_k||\mu\right) \leq \frac{\log(|G|)}{1-|G|^{-1}}\|\mu - \nu_k\|_1$$

and from Equation (41) it follows that

$$D\left(\nu_k||\mu\right) \leq \frac{\log(|G|)}{1-|G|^{-1}}\epsilon_k. \tag{C.2}$$

Combining (C.2) with (C.1) and taking the square root yields the claim. □

D. Proof of Theorem 20

Proof. Let T_c and T_n be as in definition 19. Then we have

$$\mathcal{T}(T) = (1-\delta)T_c + \delta\mathcal{T}(T_n),$$

as T_c is already covariant, and

$$\begin{aligned}\mathcal{T}(T)^m =&(1-\delta)^m T_c^m + \delta(1-\delta)^{m-1}\sum_{j=0}^{m-1} T_c^j\mathcal{T}(T_n)T_c^{m-j-1}\\ &+\delta^2(1-\delta)^{m-2}\sum_{j_1+j_2+j_3=m-2} T_c^{j_1}\mathcal{T}(T_n)T_c^{j_2}\mathcal{T}(T_n)T_c^{j_3} + \mathrm{O}(\delta^3). \qquad \text{(D.1)}\end{aligned}$$

Moreover, as T_c is covariant with respect to this unitary representation, we have

$$\begin{aligned}\mathbb{E}(\mathcal{S}_{b,m+b+1}) = (1-\delta^m)T_c + \delta(1-\delta)^{m-1}\sum_{j=0}^{m-1}\mathbb{E}\left(T_c^j\mathcal{D}_{m-j}T_n\mathcal{D}^*_{m-j}T_c^{m-j-1}\right)\\ +\delta^2(1-\delta)^{m-2}\sum_{j_1+j_2+j_3=m-2}\mathbb{E}\left(T_c^{j_1}\mathcal{D}_{j_2+1}T_n\mathcal{D}^*_{j_2+1}T_c^{j_2}\mathcal{D}_{j_3+1}T_n\mathcal{D}^*_{j_3+1}T_c^{j_3}\right)+\mathrm{O}(\delta^3) \quad \text{(D.2)}\end{aligned}$$

It is clear that the terms of 0−order in δ in (D.1) and (D.2) coincide. Comparing each of the summands of first order we obtain:

$$\begin{aligned}&\mathbb{E}\left(T_c^j\mathcal{D}_{m-j}T_n\mathcal{D}^*_{m-j}T_c^{m-j-1}\right) - T_c^j\mathcal{T}(T_n)T_c^{m-j-1}\\ =&\sum_{g\in G}\left(\nu_{m-j}(g)-\frac{1}{|G|}\right)T_c^j\mathcal{U}_gT_n\mathcal{U}_g^*T_c^{m-j-1},\end{aligned}$$

where ν_{m-j} is the distribution of $\mathcal{D}_{m-j}$. Comparing the terms of second order we obtain:

$$\begin{aligned}&\mathbb{E}\left(T_c^{j_1}\mathcal{D}_{j_2+1}T_n\mathcal{D}^*_{j_2+1}T_c^{j_2}\mathcal{D}_{j_3+1}T_n\mathcal{D}^*_{j_3+1}T_c^{j_3}\right) - T_c^{j_1}\mathcal{T}(T_n)T_c^{j_2}\mathcal{T}(T_n)T_c^{j_3}\\ =&\sum_{g_1,g_2\in G}\left(\tau_{j_3+1,j_2+1}(g_1,g_2)-\frac{1}{|G|^2}\right)T_c^{j_1}\mathcal{U}_{g_1}T_n\mathcal{U}^*_{g_1}T_c^{j_2}\mathcal{U}_{g_2}T_n\mathcal{U}^*_{g_2}T_c^{j_3}.\end{aligned}$$

Here τ_{j_3+1,j_2+1} is the joint distribution of $\mathcal{D}_{j_3+1}$ and $\mathcal{D}_{j_2+1}$. Then, using arguments similar to those of theorem 17, we have that

$$\begin{aligned}&\|\mathcal{T}(T)^m - \mathbb{E}(S_{b,m+1})\|_{1\to1}\\ &\quad\leq\delta(1-\delta)^{m-1}\sum_{j=1}^{m}\|\nu_j-\mu\| + \delta^2(1-\delta)^{m-2}\sum_{j_1=1}^{m-1}\sum_{j_2=j_1+1}^{m}\|\tau_{j_1,j_2}-\mu^{\otimes2}\|_1+\mathrm{O}(\delta^3).\end{aligned}$$

Now, from our choice of b, we have $\|\nu_j-\mu\|_1\leq\frac{\epsilon}{m}$. Furthermore, we have that

$$\tau_{j_1,j_2}(g_1,g_2) = \mathbb{P}(\mathcal{D}_{j_1}=\mathcal{U}_{g_1},\mathcal{D}_{j_2}=\mathcal{U}_{g_2}) = \mathbb{P}(\mathcal{D}_{j_2}=\mathcal{U}_{g_2}|\mathcal{D}_{j_1}=\mathcal{U}_{g_1})\mathbb{P}(\mathcal{D}_{j_1}=\mathcal{U}_{g_1}).$$

By the construction of the $\mathcal{D}_j$, it holds that

$$\mathbb{P}(\mathcal{D}_{j_2} = \mathcal{U}_{g_2} | \mathcal{D}_{j_1} = \mathcal{U}_{g_1}) = \pi^{j_2 - j_1}(g_1, g_2),$$

where π is the stochastic matrix of the chain generated by A. From this we obtain

$$\begin{aligned}\sum_{g_1, g_2 \in G} \left| \tau_{j_1, j_2}(g_1, g_2) - \frac{1}{|G|^2} \right| &= \sum_{g_1, g_2 \in G} \left| \nu_{j_1}(g_1) \pi^{j_2 - j_1}(g_1, g_2) - \frac{1}{|G|^2} \right| \\ \leq \sum_{g_1, g_2 \in G} \left| \nu_{j_1}(g_1) - \frac{1}{|G|} \right| &\pi^{j_2 - j_1}(g_1, g_2) + \left| \frac{1}{|G|} \pi^{j_2 - j_1}(g_1, g_2) - \frac{1}{|G|^2} \right|. \qquad \text{(D.3)}\end{aligned}$$

As the matrix π is doubly stochastic, summing over g_2 first

$$\sum_{g_1, g_2 \in G} \left| \nu_{j_1}(g_1) - \frac{1}{|G|} \right| \pi^{j_2 - j_1}(g_1, g_2) = \sum_{g_1 \in G} \left| \nu_{j_1}(g_1) - \frac{1}{|G|} \right| \leq \epsilon m^{-1},$$

which again follows from our choice of b. We now estimate the other term in (D.3). From our assumption in (51) we have that

$$\sum_{j=1}^{m-1} \sum_{l=j+1}^{m} \sum_{g_1, g_2 \in G} \frac{1}{|G|} \left| \pi^{j_2 - j_1}(g_1, g_2) - \frac{1}{|G|} \right| \leq \sum_{j=1}^{m-1} \sum_{l=j+1}^{m} \lambda^{l-j}. \qquad \text{(D.4)}$$

Summing up both geometrical series which come up in the last expression of (D.4), we finally obtain

$$\sum_{j=1}^{m-1} \sum_{l=j+1}^{m} \sum_{g_1, g_2 \in G} \frac{1}{|G|} \left| \pi^{j_2 - j_1}(g_1, g_2) - \frac{1}{|G|} \right| \leq \frac{\lambda}{1 - \lambda}(m - 1).$$

Putting all inequalities together, we obtain the claim. □

E. Proof of Lemma 24

Proof. (2) ⇒ (1) can be seen by direct inspection. In order to prove the converse, we consider the Choi-Jamiolkowski state $\tau_T := \frac{1}{d} \sum_{i,j=1}^{d} T\big(|i\rangle\langle j| \big) \otimes |i\rangle\langle j|$. Then (1) is equivalent to the statement that τ_T commutes with all unitaries of the form $U \otimes \bar{U}$, $U \in MU(d, n)$. That is, we have

$$\sum_{i,j=1}^{d} U \otimes \bar{U} \left(T\big(|i\rangle\langle j| \big) \otimes |i\rangle\langle j| \right) (U \otimes \bar{U})^* = \sum_{i,j=1}^{d} T\big(|i\rangle\langle j| \big) \otimes |i\rangle\langle j| \,.$$

Restricting to the subgroup of diagonal unitaries in $MU(d, n)$, for which $U^* = \bar{U}$, we have

$$\sum_{i,j=1}^{d} \mathrm{e}^{\mathrm{i}(\phi_j - \phi_i)} U T\big(|i\rangle\langle j| \big) \bar{U} \otimes |i\rangle\langle j| = \sum_{i,j=1}^{d} T\big(|i\rangle\langle j| \big) \otimes |i\rangle\langle j| \,,$$

where $\mathrm{e}^{\mathrm{i}\phi_i}$ is the i-th diagonal entry of U. Comparing the tensor factors it follows that

$$\mathrm{e}^{\mathrm{i}(\phi_j-\phi_i)} U T\big(\,|i\rangle\langle j|\,\big)\bar{U} = T\big(\,|i\rangle\langle j|\,\big). \tag{E.1}$$

We will now show that we have

$$\tau_T = \sum_{i,j=1}^{d} A_{ij}\,|i\rangle\langle i| \otimes |j\rangle\langle j| + B_{ij}\,|i\rangle\langle j| \otimes |i\rangle\langle j|\,. \tag{E.2}$$

We have

$$T(|i\rangle\langle j|) = \sum_{k,l=1}^{d} a_{k,l}\,|k\rangle\langle l|$$

for some $a_{k,l} \in \mathbb{C}$. From (E.1) it follows that

$$\sum_{k,l=1}^{d} \mathrm{e}^{\mathrm{i}(\phi_k-\phi_l)} a_{k,l}\,|k\rangle\langle l| = \mathrm{e}^{\mathrm{i}(\phi_i-\phi_j)} \sum_{k,l=1}^{d} a_{k,l}\,|k\rangle\langle l| \tag{E.3}$$

for all diagonal unitaries. Again comparing both sides of (E.3) we have $a_{k,l}\mathrm{e}^{\mathrm{i}(\phi_i-\phi_j)} = \mathrm{e}^{\mathrm{i}(\phi_k-\phi_l)} a_{k,l}$. Suppose now $i \neq j$. For $a_{k,l} \neq 0$ we have

$$\mathrm{e}^{\mathrm{i}(\phi_i-\phi_j)} = \mathrm{e}^{\mathrm{i}(\phi_k-\phi_l)} \tag{E.4}$$

for all diagonal entries of diagonal unitaries. If k, l, i and j are all pairwise distinct, we have $i = k$ and $j \neq l$ or $i \neq k$ and $j = l$, then it is clear that we may always find a combination of ϕ_k, ϕ_l, ϕ_i and ϕ_j such that (E.4) is not satisfied, a contradiction. For $i = l$ and $k = j$, it is only possible to find such a combination for $n > 2$, as otherwise $\phi_i - \phi_j = -(\phi_i - \phi_j)$ always holds. This proves that we have

$$T\left(|i\rangle\langle j|\right) = B_{ij}\,|i\rangle\langle j| \tag{E.5}$$

for $i \neq j$. For $i = j$ we have analogously that

$$UT(|i\rangle\langle i|)\bar{U} = \sum_{k,l=1}^{d} \mathrm{e}^{\mathrm{i}(\phi_k-\phi_l)} a_{k,l}\,|k\rangle\langle l| = \sum_{k,l=1}^{d} a_{k,l}\,|k\rangle\langle l|\,.$$

In this case, we have $a_{k,l} = \mathrm{e}^{\mathrm{i}(\phi_k-\phi_l)} a_{k,l}$ for all possible phases of the form $e^{\mathrm{i}(\phi_k-\phi_l)}$. It is then clear that $a_{k,l} = 0$ unless $k = l$ by a similar argument as before. This gives

$$T\left(|i\rangle\langle i|\right) = \sum_{j=1}^{d} A_{ij}\,|j\rangle\langle j|\,. \tag{E.6}$$

Putting together (E.6) and (E.5) implies (E.2). Next, we will exploit that τ_T commutes in addition with permutations of the form $U_\pi \otimes U_\pi$ for all $\pi \in S_d$. For $i \neq j$ this implies that $A_{i,j} = A_{\pi(i),\pi(j)}$ and $B_{i,j} = B_{\pi(i),\pi(j)}$ so that there is only one independent off-diagonal element for each A and B. The case $i = j$ leads to a third parameter that is a coefficient in front of $\sum_\alpha |ii\rangle\langle ii|$. Translating this back to the level of projections then yields (55). The fact that the terms of (55) are projections can be seen by direct inspection. Note that the term corresponding to α is the difference of two projections, the identity and projection onto diagonal matrices. As the rank of the identity is d^2 and the space of diagonal matrices has dimension d, we obtain the claim. The same reasoning applies to the term corresponding to β, as it is the difference of the projection onto diagonal matrices and the projection onto the maximally mixed state. The latter is a projection of rank 1, which yields a rank of $d - 1$ for their difference. □

References

[1] Poyatos J F, Cirac J I and Zoller P 1997 *Phys. Rev. Lett.* **78** 390–393

[2] Chuang I L and Nielsen M A 1997 *J. Mod. Opt.* **44** 2455–2467

[3] Knill E, Leibfried D, Reichle R, Britton J, Blakestad R B, Jost J D, Langer C, Ozeri R, Seidelin S and Wineland D J 2008 *Phys. Rev. A* **77**(1) 012307

[4] Emerson J, Silva M, Moussa O, Ryan C, Laforest M, Baugh J, Cory D G and Laflamme R 2007 *Science* **317** 1893–1896

[5] Lévi B, López C C, Emerson J and Cory D G 2007 *Phys. Rev. A* **75**(2) 022314

[6] Emerson J, Alicki R and Zyczkowski K 2005 *J. Opt. B* **7** S347

[7] Wallman J J and Flammia S T 2014 *New J. Phys.* **16** 103032

[8] Helsen J, Wallman J J, Flammia S T and Wehner S 2017 *ArXiv e-prints* (*Preprint* `1701.04299`)

[9] Chow J M, Gambetta J M, Tornberg L, Koch J, Bishop L S, Houck A A, Johnson B R, Frunzio L, Girvin S M and Schoelkopf R J 2009 *Phys. Rev. Lett.* **102**(9) 090502

[10] Ryan C A, Laforest M and Laflamme R 2009 *New J. Phys.* **11** 013034

[11] Olmschenk S, Chicireanu R, Nelson K D and Porto J V 2010 *New J. Phys.* **12** 113007

[12] Brown K R, Wilson A C, Colombe Y, Ospelkaus C, Meier A M, Knill E, Leibfried D and Wineland D J 2011 *Phys. Rev. A* **84**(3) 030303

[13] Gaebler J P, Meier A M, Tan T R, Bowler R, Lin Y, Hanneke D, Jost J D, Home J P, Knill E, Leibfried D and Wineland D J 2012 *Phys. Rev. Lett.* **108**(26) 260503

[14] Barends R, Kelly J, Megrant A, Veitia A, Sank D, Jeffrey E, White T C, Mutus J, Fowler A G, Campbell B, Chen Y, Chen Z, Chiaro B, Dunsworth A, Neill C, O/'Malley P, Roushan P, Vainsencher A, Wenner J, Korotkov A N, Cleland A N and Martinis J M 2014 *Nature* **508** 500–503

[15] Xia T, Lichtman M, Maller K, Carr A W, Piotrowicz M J, Isenhower L and Saffman M 2015 *Phys. Rev. Lett.* **114**(10) 100503

[16] Muhonen J T, Laucht A, Simmons S, Dehollain J P, Kalra R, Hudson F E, Freer S, Itoh K M, Jamieson D N, McCallum J C, Dzurak A S and Morello A 2015 *J. Phys. Condens. Matter* **27** 154205

[17] Asaad S, Dickel C, Langford N K, Poletto S, Bruno A, Rol M A, Deurloo D and DiCarlo L 2016 *nph Quantum Inf.* **2** 16029

[18] Hashagen A K, Flammia S T, Gross D and Wallman J J 2018 *ArXiv e-prints* (*Preprint* `1801.06121`)

[19] Brown W G and Eastin B 2018 *ArXiv e-prints* (*Preprint* `1801.04042`)

[20] Cross A W, Magesan E, Bishop L S, Smolin J A and Gambetta J M 2016 *npj Quantum Inf.* **2** 16012

[21] Carignan-Dugas A, Wallman J J and Emerson J 2015 *Phys. Rev. A* **92**(6) 060302

[22] Van den Nest M 2011 *New J. Phys.* **13** 123004

[23] Nielsen M A and Chuang I L 2009 *Quantum Computation and Quantum Information* (Cambridge University Press)

[24] Heinosaari T and Ziman M 2012 *The Mathematical Language of Quantum Theory: From Uncertainty to Entanglement* (Cambridge University Press)

[25] Simon B 1996 *Representations of finite and compact groups* (*Graduate studies in mathematics* vol 10) (American Mathematical Society)

[26] Mendl C B and Wolf M M 2009 *Commun. Math. Phys.* **289** 1057–1086

[27] Mozrzymas M, Studziński M and Datta N 2017 *J. Math. Phys.* **58** 052204

[28] Levin D A and Peres Y 2017 *Markov chains and mixing times* vol 107 (American Mathematical Society)

[29] Saloff-Coste L 2004 Random walks on finite groups *Probability on discrete structures* (Springer) pp 263–346

[30] Nielsen M A 2002 *Phys. Lett. A* **303** 249–252

[31] Horn R A and Johnson C R 2009 *Matrix Analysis* (Cambridge University Press)

[32] Burgarth D, Chiribella G, Giovannetti V, Perinotti P and Yuasa K 2013 *New J. Phys.* **15** 073045

[33] Wallman J J 2018 *Quantum* **2** 47

[34] Proctor T, Rudinger K, Young K, Sarovar M and Blume-Kohout R 2017 *Phys. Rev. Lett.* **119** 130502

[35] Gambetta J M, Córcoles A D, Merkel S T, Johnson B R, Smolin J A, Chow J M, Ryan C A, Rigetti C, Poletto S, Ohki T A, Ketchen M B and Steffen M 2012 *Phys. Rev. Lett.* **109**(24) 240504

[36] Magesan E, Gambetta J M and Emerson J 2012 *Phys. Rev. A* **85**(4) 042311

[37] Magesan E, M G J and Emerson J 2011 *Phys. Rev. Lett.* **106** 180504

[38] Dankert C, Cleve R, Emerson J and Livine E 2009 *Phys. Rev. A* **80**(1) 012304

[39] Pérez-García D, Wolf M M, Petz D and Ruskai M B 2006 *J. Math. Phys.* **47** 083506–083506

[40] Harrow A W and Low R A 2009 *Commun. Math. Phy.* **291** 257–302

[41] Kliuchnikov V 2014 *New methods for quantum compiling* Ph.D. thesis University of Waterloo

[42] Leditzky F, Kaur E, Datta N and Wilde M M 2018 *Phys. Rev. A* **97**

[43] Gottesman D 1997 *ArXiv e-prints* (*Preprint* 9705052)

[44] Koenig R and Smolin J A 2014 *J. Math. Phys.* **55** 122202

[45] Hokanson J 2013 *Numerically stable and statistically efficient algorithms for large scale exponential fitting* Ph.D. thesis Rice University

[46] Holmström K and Petersson J 2002 *App. Math. Comput.* **126** 31–61

[47] Hokanson J 2014 Exponential fitting `https://github.com/jeffrey-hokanson/exponential_fitting_code`

[48] Hincks I, Wallman J J, Ferrie C, Granade C and Cory D G 2018 *ArXiv e-prints* (*Preprint* 1802.00401)

[49] Hoeffding W 1963 *J. Am. Stat. Assoc.* **58** 13–30

[50] Sason I 2015 *ArXiv e-prints* (*Preprint* 1503.07118)