Université Paris Cité
École Doctorale de Sciences Mathématiques de Paris Centre (ED386)
Institut de Recherche en Informatique Fondamentale (IRIF) – UMR 8243

Analog quantum algorithms for combinatorial optimization on neutral-atom devices: from graph embeddings to tunneling in quantum annealing

Par Elie Bermot

Thèse de doctorat en Informatique

Dirigée par Frédéric Magniez

Présentée et soutenue publiquement le 12 octobre 2026

Devant un jury composé de :

Thomas Ayral, Professeur	École polytechnique	*Rapporteur*
Andrea Simonetto, Enseignant-chercheur	ENSTA	*Rapporteur*
Sophie Laplante, Professeure	CNRS, Université Paris Cité	*Examinatrice*
Frédéric Magniez, Directeur de recherche	CNRS, Université Paris Cité	*Directeur de thèse*
Simon Apers, Chargé de recherche	CNRS, Université Paris Cité	*Membre invité*
Louis-Paul Henry, VP Quantum Applications	Pasqal	*Membre invité*

Acknowledgements

Summarizing three years of a PhD is quite a difficult task, but if there is one feeling that stands out above all others, it is gratitude. Sorry in advance: this is going to be a very long list of acknowledgments.

First, I would like to thank the members of my PhD jury, Thomas Ayral, Sophie Laplante, Andrea Simonetto, Frédéric Magniez, Simon Apers, and Louis-Paul Henry, for agreeing to evaluate my work. In particular, I would like to thank Thomas and Andrea for taking the time to review my manuscript, especially at such short notice. I am fully aware of the time and effort that such a task requires, and I am truly grateful to you both. I am deeply thankful for your insightful comments and constructive feedback, which helped me improve the manuscript and bring it to the highest standard I could hope for.

I would like to thank Simon Apers, my PhD supervisor, for absolutely everything he has brought me. I truly believe that I am a better scientist today thanks to him, to his rigour, and to the way he has constantly pushed me to improve and aim higher. Thank you for accepting me as your first PhD student, and for all our discussions, whether they were about science, record stores, musicals, coffee, or anything else. Above all, thank you for creating such a safe and supportive environment for me throughout these years, for always accepting my scientific ideas, or sometimes not-so-scientific ones, and for your constant trust. I felt respected and treated like a colleague, and even like your son sometimes. I will always cherish these last three years, and that is mostly because of you.

I would similarly like to thank Frédéric Magniez for agreeing to be my PhD supervisor. Thank you for all our scientific discussions, but also for our conversations about wine, food, life, and everything in between. I look up to you with great respect, and your constant drive to build an amazing quantum ecosystem is truly inspiring. Your support and guidance have been invaluable throughout this journey. Thank you also for giving me the opportunity to teach at the university. I think it was one of those opportunities that helped me glimpse what some of my future jobs could be.

I would also like to thank my PhD supervisor at Pasqal, Louis-Paul Henry, for being such a great supervisor. I do not think I would have become the scientist I am today without you. You taught me how to be a scientist inside a company, with all its aspects, both the good and the more difficult ones. Thank you for the countless conversations we have had, scientific or

otherwise: about how to make cheese, how to design state graphs, the best place to read your thesis, how to design SMART goals, and many other equally important topics. Thank you for listening, for your advice, and for your views on life. Thank you for everything.

I would also like to thank Loïc Henriet, first and foremost for helping to found such an amazing company and for creating such a stimulating place in which I could grow as a scientist. I do not think I would have been as lucky in my PhD if it had not been for you and your willingness to design such a challenging PhD project. Thank you for replying to my bold email. I cannot wait to see what the future holds at Pasqal.

I would particularly like to thank all my Pasqal co-authors: Wesley da Silva Coelho, Natalie Person, and Lucia Valor. I am really proud of the paper we wrote together, and grateful for your endless trust throughout this project. Wesley, thank you for your endless life and organizational advice; Natalie, for your scientific rigour and the best British jokes; and Lucia, for being an amazing scientist and friend ever since our time in Zurich.

I would also like to thank all my Pasqal colleagues, many of whom have now become close friends. First of all, thank you to Constantin Dalyac, my very first contact at Pasqal, for being such a great mentor. Filling your shoes is not an easy task, but having you around was a constant push and a constant motivation to become better. Louis Vignoli, thank you for being such a warm presence, and for always inspiring me to become a better scientist through rigour, precision, and intelligence. I am so glad to have met someone who gives so much love to science, with such dedication. Paul Rousset-Rouard, I have not known you for that long, but your presence is truly inspiring, and you are hilarious. I cannot wait to see what the future holds. François-Marie Le Régent, thank you for being such a good friend, someone I can talk to at 8:30 in Phénix about pantoufles, every other week, or simply joke around with. Thank you for inspiring me every single day. Thank you to Lucas Leclerc for being such a warm presence, while being insanely stylish, kind, and talented, and for agreeing to try random places around Joro or Patchwork, and, most importantly, for laughing at my random jokes.

Thank you as well to Yassine Naghmouchi, for being an insanely nice colleague, always trying to cheer up and help everyone around you; Mauro d'Arcangelo, for being such a great scientist and emotional support; Vittorio Vitale, for being an amazing Primitives leader, both for your scientific genius and for the endless restaurant chats, you are a joy to be around; Simone Tibaldi, for the great time we had when you were around and for the parties; Quentin Ma, for the great chats and jokes; Shaheen Acheche, for all the inspiring discussions about science and PhD life; and Annah Guandalino, for all our musical conversations and matcha breaks and Francesca Galasso for being a great friend. Thank you to Morgane Migliari for this amazing podcast adventure and for being such an inspiring person to be around; and to Isis Fradet for all our coffee breaks and your constant joy. I would also like to thank Diane Zink for the great discussions. Good luck with your PhD; I am sure you will do great at Pasqal. You are insanely talented, and I am confident you will thrive here! Thank you as well to Sergi Julia

Farre for the great discussions and the great time around, to Marie Wakim for the amazing discussions and the life thought, to Fadil Aidane for the great food chats, and life advices!

I would also like to thank Fabien Le Rolland for the Galika jokes and your constant respect for me; Angela Pulido, for being a real ray of sunshine in this company through your joy and kindness; Mehdi Djellabi, for the great times around beers and scientific chats ever since I joined Pasqal; Alexandre Dauphin, for being such an inspiring scientist and for the great jokes when we were in Brussels; Georges-Olivier Reymond, for co-founding such a great company; Mourad Beji, for leading the software department with such greatness; Matthias Hecker, for inspiring chats over your massive lunches; Lorenzo Moro, for being such a joy to be around and an amazing researcher; Igor Sokolov, for being such a great friend with whom I have followed the same path since Zurich; Monique Joaquim Witt Garzillo, for being such a warm presence at Pasqal; Victor Onofre, for being a great scientist ever since the Pasqal hackathon and for continuing to push toward a challenging quantum advantage; Francesco Musso, for being such a great scientist ever since I saw you during your internship, I have absolutely no doubt you will do amazing things during your PhD; Stefano Grava, for being a great Primitives colleague and for being so kind; Joseph Vovrosh, for such great times in Glasgow during AQC, your talent and your constant desire to push your knowledge to the maximum are truly inspiring, and they push your colleagues to want the best for themselves; and Emmanuelle Choyer, for being the first person to welcome me to Pasqal with your joyful personality and your constant desire to bring out the best in everyone. Thank you all for making Pasqal such a unique place to work. It is truly inspiring and challenging to help make the workplace even better while being part of such a thriving company.

I would like to thank Arthur Braida for the amazing time we spent working together. Thank you for always driving our work forward, for all the enjoyable conversations, and for inspiring me to become a better scientist. Ever since I met you, I knew you would be a great driving force, someone who constantly wants to learn more and push science forward, all in a very organized and calm way. I hope that one day I will get to try one of your meditation retreats!

I would also like to thank my colleagues at IRIF for making my time there so enjoyable. Étienne Objois, thank you for all the great times during conferences, coffee breaks, or just lunches; you are hilarious. Benjamin Mathieu-Bloise, thank you for being such a brother to me during this too short time with you. You are kind, talented, and warm, and I am proud to call you my friend. Roman Edenhofer, thank you for being an amazing office mate. You are insanely talented, and ever since I met you during your internship, you have constantly proven to me that you are a local genius, with a real desire to make your workplace even better. Thank you for accepting me, and sorry for all the noise I made on Tuesdays and Wednesdays. Daniel Szabo, thank you for being such a calm and warm presence ever since I joined, always joking around and creating a safe place at IRIF. I really enjoyed our numerous chats in Waterloo, Berlin, and so many other places. Arjan Cornelissen (*quel mec!*), thank you for being a great

friend and colleague, and for giving one of the best tutorials and talks I have ever seen, all while learning *Bande organisée*! Sebastian Zur, thank you for all our philosophical conversations, your amazing intelligence, humour, craziness, and simply your warm presence. Isabella Ziccardi, thank you for being an insanely nice colleague; Maël Luce, for being a great conference friend with whom I could talk; Émile Larroque, for the great chats; and Chandrima Kayal, for the great conversations in Riga.

Throughout my journey in quantum computing, I was also fortunate to meet wonderful people at conferences, including Joseph Cunningham and Galina Pass, as well as my friends from QuantX and Le Lab Quantique: Kenzo Bounegta, Elvira Shishenina, Alexandre Krajenbrink, and Jannes Stubbemann. Thank you for teaching me so much about organizing hackathons and for showing me how meaningful and rewarding science communication can be. I would also like to thank Christa Zoufal, my master's thesis supervisor at IBM, for truly inspiring me from the beginning of my master's degree and showing me what it means to be a great scientist, an amazing supervisor, and a great foodie.

I would now like to thank the friends I have met throughout the years, and especially those who helped me survive this PhD. I am beyond grateful to know such amazing people. I am a better person because I know you, and your love has always made me stronger.

To my childhood best friend Estelle Paris, thank you for your constant love and support, and for accepting me for who I am. I do not think I would be the man I am today without you and our countless Friday dinners. To my adoptive sister Assa Paris, thank you for Just Dance and for all the great times; being around you is a blessing. To my crazy best friend Amélie Liberge, thank you for all the love, laughter, trips, and the Jean-Michel à peu près. I do not think I would be who I am today without your love and our amazing times together. Who would have thought that we would still be laughing like crazy about Poroy today? I am speechless when trying to describe how much I love and care for you, my Melou. To my diva queen Mary Thirouin, thank you for the constant love, divaness, and for everything since MONIQUE and option musique. You are an icon, and I love you for pushing me into the world of music.

To my best friends Maxence Daddji Meunier, Valentine Aurélien Saue, Matteo Dedola, Adam Siegel, and Cédric Dussault: I am beyond grateful to know you. Maxence, I met you around the same time I started my PhD, but you very quickly became an ally and an acolyte, and I could not love you more than I do today. Valentine, my bestie who would move the world for me, thank you for being such a dedicated friend with the best humour I could ever hope for. Matteo, my mentor and mentee, you are a true inspiration to be around, with your endless talent, so thank you for being who you are. Adam, my bibou and C12 internship partner, you are a gold star, with your talent, wisdom, and foodieness. Cédric, my Odymus dance friend, I love you so much and I am beyond grateful to know you.

Adrien Braun, thank you for being an amazing friend in the short time I have known you; you truly deserve gallons of love for the generosity and love you spread in the world. Clémence

Couchot, my sister and the most comforting friend in the universe, your craziness should be studied, and you truly spread love and joy with your beauty and kindness. To all my best friends, thank you for your constant advice, your love, and your support. I truly do not think I would be the man I am today without you.

I am equally grateful for all the friends I have made through dance over the years. Once again, thank you to my best friend Estelle for introducing me to this world so long ago. Thank you to my queen Scarlleth Castro for bringing me to my first Groove class with Ralph. You dance like an icon, and you really push me to be my best and inspire me to become a better dancer and human. To my acolytes Hugo Deleaz and Benjamin Jouffret: you are insanely talented, and I am so glad to know you. I feel so free when I am with you, and I want to become a better dancer and weightlifter every day thanks to you. Thank you for being such rays of sunshine this past year. To my dance-viewing friend Jean-Guillaume Santi, thank you for your facial expressions when you dance and for trusting me to be a great pit crew with you. You truly value the people around you, so thank you for that, and for IOB. To my Beyoncé queen Malak Achahbar, thank you for being an amazing friend who radiates the passion of dance with your sunshine, passion, and energy. I am so inspired by you. I would also like to thank Elisa Biecher and Irene Saint-Louis for being such wonderful discoveries this past year in Delphine's class: your constant advice to help me become a better dancer, and the way you constantly lifted me up, whether through your amazing dancing abilities or your great life advice, made Mondays and Thursdays such joyful moments. I would also like to thank my amazing dance teacher, Delphine Lemaître, for creating such a safe environment, for giving me the will to keep pushing my abilities and my progress further, for being a breath of fresh air in my weeks, and for helping me cultivate what has become one of the greatest passions of my life.

Thank you as well to my safe-place friends: Yann Jestin; Anthony Da Silva, my Barbaloomis friend of the heart, whose constant love, joy, and dedication to helping me throughout my artistic journey make me love you more each day; Servane Demarquet, for being an amazing friend who pushed me to always stay true to myself and not try to become someone I am not, with your true kindness and our macareux; Antoine Palazzolo, for the constant memes, constant support, and constant desire to make everyone feel welcome in a true and insane community; Anne-Flavie Bastard, for the amazing dance creations – you are an amazing artist who truly deserves the world and more; Lucie Claris, for your constant kindness and desire to make the world a safer and more loving place; Melodie Turpin, for creating love around you; Vimala Maillard, for being a great friend and flatmate in Zurich and for all the constant motivation in sport, science, bobo life, and so much more; and Mariona Flores, for being an amazing friend – girl, you are crazy, and I love it.

I would particularly like to thank my Odymus family for giving me such a strong sense of freedom whenever I was not working this past year. To my beloved dancers: Deborah Aisenberg, my pissenlit; Léa Benquis, my Rheingold acolyte; Orhiane Hardouin, the yellow ray of sunshine;

Pauline Tricot (*pinky up les patounes!*); Alice Magat, the queen of expressions and one of the most loving people I know; Bastien Galant, my Rheingold partner; Cécile Ballorain, my coupé queen and future choreographer; Elise Lei, my Freak and Code Hays queen; Élodie Rochat, whose hugs are the best on Earth; Héloïse Rannou, my tango partner and queen; and Laure Villemonte, my make-up queen with a heart of gold.

I would also like to give a special thank-you to the friends I met there. First, Ileana Chocardelle, thank you for making me discover this association; you are a true love on this Earth, and I cannot wait to keep loving you every day. Anouk Thomas, one of my first friends there, you are insanely funny, loving, and warm, but also a *poucave quand même*. Sylvain Jouve, my twin, you truly want the best for everyone around you; you are insanely talented, and I cannot wait to see what the future holds for you. Gerson Cristovao, my Beyoncé twin, I love you so much, and thank you again for Freak and for the many more opportunities to come; you are such a talented diva. Dune Favero, your stage directions were insanely great. I love you and your constant love for art. I am inspired by you, and I cannot wait to work with you again. My twins, I love you so much that it should not even be allowed to love people this much.

I am also immensely grateful for all the wonderful people this association brought into my life, including Benjy Rouire (my favorite fraisita in the Gregorio multiverse, and for letting me into his world), Adrien Etheve (my sister), Jacques Boneff (the kindest person on this Earth, thank you for Star Ac!!), Noémie Perrot (the insanely talented Rita), Brune Dubost (you are a true inspiration), Philippine Coindre (please continue to shine), Constant Josse (I will try to susurrer to you more often), Ryan Privat (the kindest and most iconic), Alexis Bertrand (the best understudy), Galith Labade (my mazal), Guillemette Tesson (my Rheingold partner, I love you so much), Alice Liu (the diva herself!!), Romain Bollaert (the best croco, and you are so talented and deserve so much love), Théo Ferjani (I am obsessed with you and your talent; please continue to be such a light and kind heart), Nicolas Dovale (best Dali, and thank you for all the fun we had), Natalia Rest (the kindest soul with so much talent), Nada El Alami (the smile that could fill and inspire an entire room), Yann Teytaut (my research friend, I love you), and so many others.

Finally, I would like to thank my beautiful family for always giving me everything I have ever needed. Thank you, Mum and Dad, for giving me all the love and support I could ever have hoped for.

I would especially like to thank my amazing sister, Myriam, who is truly the best sister I could ever have wished for. You are the best part of me, and I know that I am the man I am today because of you. I would also like to thank her wonderful husband, Léo Hauchecorne, for always accepting me exactly as I am and for giving my sister all the love she has always deserved.

Thank you to my amazing cousins: Salomé Fuks, Manon Fuks, Ruben Guez, Lena May, Sarah May, Sacha Klezman, Oscar Blackstone, Max Blackstone, Raphaël Fettaya, Déborah and

Laura Levy, Clara Fettaya, Ethan Fettaya, Lior Hassine, and Noémie Beccache. Thank you as well to my wonderful aunts and uncles: Marielle and Richard Fettaya, Caroline May, David Fuks, Éric May, and Sonia.

And, naturally, I would like to thank everyone once again for accepting me for who I am, without censorship. I know I can be a lot, but you have given me everything I needed to become the best version of myself today, with all my foolishness, craziness, hyperactivity, infinite love, and slight insanity. I love you all from the bottom of my heart.

Abstract

Title: Analog quantum algorithms for combinatorial optimization on neutral-atom devices: from graph embeddings to tunneling in quantum annealing.

The goal of this thesis is to study quantum algorithms designed for neutral-atom devices. Conducted at IRIF in collaboration with Pasqal, it mainly focuses on quantum annealing, an analog approach to combinatorial optimization. It works by slowly evolving a system toward a ground state of a problem Hamiltonian. We study this approach from two complementary angles, which form the two parts of the thesis: which optimization problems can be encoded within hardware constraints, and under which conditions do the annealing dynamics prepare a solution efficiently.

The first part concerns the encoding of graph optimization problems on neutral-atom devices. Such devices natively implement the Maximum Independent Set (MIS) problem on unit-disk graphs through the Rydberg blockade mechanism. However, the class of unit-disk graphs that can currently be implemented on Pasqal's devices is restricted by hardware constraints, limiting the range of instances that can be addressed. To go further, we revisit the Rydberg blockade mechanism and study its limits, which allows us to extend the set of encodable instances to the broader class of disk graphs. This widens the family of graphs on which the MIS problem can be tackled on Pasqal's machines, beyond the native unit-disk setting.

The second part concerns the conditions under which quantum annealing prepares a solution efficiently. The runtime is governed by the minimum spectral gap encountered along the adiabatic evolution, but this gap is typically hard to compute exactly beyond a few special cases. We study a family of potentials that depend only on the Hamming weight, including sufficiently regular examples such as convex potentials. We identify conditions, phrased in terms of log-concavity of the ground-state distribution, that imply a unimodal ground state with well-controlled tails away from its mode. We then show that this preparation remains robust under the addition of a potential barrier. Finally, we apply this technique to quadratic potentials, which are well motivated by common physical implementations, including Ising-like models relevant to neutral-atom platforms.

Keywords: quantum annealing, adiabatic quantum computation, neutral atoms, Rydberg blockade, Maximum Independent Set, disk graphs, spectral gap, log-concavity, quantum tunneling, combinatorial optimization

Résumé

Titre : Algorithmes quantiques analogiques pour l'optimisation combinatoire sur plateformes à atomes neutres : plongements de graphes et l'effet tunnel en recuit quantique.

Le but de cette thèse est d'étudier des algorithmes quantiques conçus pour des dispositifs à atomes neutres. Menée à l'IRIF en collaboration avec Pasqal, elle porte principalement sur le recuit quantique, une approche analogique pour l'optimisation combinatoire. Cette approche consiste à faire évoluer lentement un système vers un état fondamental de l'Hamiltonien du problème. Nous étudions cette approche selon deux angles complémentaires, qui constituent les deux parties de la thèse : quels problèmes d'optimisation peuvent être encodés sous des contraintes expérimentales et technologiques, et dans quelles conditions la dynamique du recuit prépare efficacement une solution.

La première partie concerne l'encodage de problèmes d'optimisation sur graphes dans des dispositifs à atomes neutres. Ces dispositifs implémentent naturellement le problème du stable maximum, ou Maximum Independent Set (MIS), sur les graphes à disques unitaires grâce au mécanisme de blocage de Rydberg. Cependant, la classe des graphes à disques unitaires qui peuvent actuellement être réalisés sur les machines de Pasqal reste restreinte par des contraintes expérimentales et technologiques, ce qui limite l'éventail d'instances pouvant être traitées. Pour aller plus loin, nous réexaminons le mécanisme de blocage de Rydberg et en étudions les limites, ce qui nous permet d'étendre l'ensemble des instances encodables à la classe plus générale des graphes de disques. Cela élargit la famille de graphes sur lesquels le problème MIS peut être abordé sur les machines de Pasqal, au-delà du cadre natif des graphes à disques unitaires.

La seconde partie concerne les conditions dans lesquelles le recuit quantique prépare efficacement une solution. Le temps d'exécution est gouverné par le gap spectral minimal du Hamiltonien rencontré au cours de l'évolution adiabatique, mais ce gap est généralement difficile à calculer exactement en dehors de quelques cas particuliers. Nous étudions une famille de potentiels ne dépendant que du poids de Hamming, incluant des exemples suffisamment réguliers tels que des potentiels convexes. Nous identifions des conditions, formulées en termes de log-concavité de la distribution de l'état fondamental, qui impliquent un état fondamental unimodal avec une décroissance bien contrôlée loin du mode. Nous montrons ensuite que cette préparation reste robuste lorsque l'on ajoute une barrière de potentiel. Enfin, nous appliquons cette technique à des potentiels quadratiques, bien motivés par des implémentations physiques courantes, notamment des modèles de type Ising pertinents pour les plateformes à atomes neutres.

Mots-clés : recuit quantique, calcul quantique adiabatique, atomes neutres, blocage de Rydberg, ensemble indépendant maximum, graphes disques, gap spectral, log-concavité, effet tunnel quantique, optimisation combinatoire

Table of Contents

Acknowledgements **ii**

Abstract **x**

Résumé **xii**

Table of Contents **xiv**

List of Figures **xviii**

List of Tables **xviii**

Résumé substantiel en français **xxii**

1 Introduction **1**
1.1 Context 1
1.1.1 Quantum computation and quantum annealing 3
1.1.2 Neutral atoms and graph optimization 4
1.1.3 From embeddability to algorithmic guarantees 6
1.2 Overview of results 6
1.2.1 Disk graph embeddings on neutral-atom devices 6
1.2.2 Log-concavity of quantum ground states 7
1.3 Publications 7

2 Preliminaries **9**
2.1 Notations and basic definitions 9
2.2 Combinatorial optimization 10
2.3 Basic notions of quantum computing 12
2.3.1 Some linear algebra 12
2.3.2 The postulates of quantum computing 13
2.4 Quantum annealing and adiabatic evolution 15
2.5 Neutral-atom devices and Rydberg physics 18

2.5.1 Principles of neutral-atom devices 18
2.5.2 MIS preparation 21

3 Disk graphs **25**
3.1 Introduction 25
3.1.1 Quantum complexity of UD-MIS 25
3.1.2 From a sharp blockade radius to local blockade radii 27
3.1.3 Our contributions 27
3.1.4 Organization of the chapter 30
3.2 Preliminaries 30
3.2.1 Unit-disk graphs and disk graphs 30
3.2.2 Correlation matrix and independence violation 31
3.2.3 Sampling quality: exact and approximate MIS probabilities 32
3.3 Description of the Rydberg blockade phenomenon 34
3.3.1 Background 34
3.3.2 Methods to find the evolution operator 36
3.3.3 Sequential regime 37
3.3.4 Global regime 38
3.3.5 Local regime 41
3.3.6 Comparison between different regimes 44
3.4 Embedding disk graphs on neutral-atom devices 46
3.4.1 Embedding model and control knobs 46
3.4.2 Quench protocol: global vs local amplitudes 48
3.4.3 Evaluation metrics 48
3.4.4 Distance scaling and edge breaking 49
3.4.5 Illustrative instance and numerical results 50
3.5 Applications to maximum independent set problems 52
3.5.1 Endpoint Hamiltonians with finite global versus local drives 52
3.5.2 Evaluation: feasibility and sampling of exact/near-optimal solutions . . 53
3.5.3 Results on minimal non-unit-disk instances 53
3.5.4 Quantum state preparation 54
3.6 Conclusion and outlook 56

4 Log-concavity and tunneling: adiabatic quantum optimization for convex functions (with a spike) **57**
4.1 Introduction 57
4.1.1 From graph embeddings to adiabatic dynamics 57
4.1.2 Spectral gaps, tunneling, and approximate optimization 59
4.1.3 Main contributions 60

4.1.4 Organization of the chapter . . . 63
4.2 Preliminaries . . . 63
4.2.1 Symmetry reduction . . . 63
4.2.2 Binomial state and Hamming weight Hamiltonian with a spike . . . 64
4.2.3 Discrete log-concavity . . . 67
4.3 (Ultra-)Log-concavity for convex potentials (and beyond) . . . 70
4.3.1 Prior work . . . 70
4.3.2 Ultra-log-concavity from convexity . . . 71
4.3.3 γ-log-concavity from γ-relative monotonicity . . . 72
4.3.4 Proofs of γ-LC from γ-RM . . . 73
4.3.5 Examples of γ-LC from γ-RM . . . 74
4.4 Spectral gap bounds from (ultra-)log-concavity . . . 77
4.4.1 Hardy inequality . . . 77
4.4.2 Hardy constant from relative log-concavity . . . 78
4.4.3 Spectral gap bounds . . . 80
4.5 Log-concavity and tunneling . . . 82
4.6 Application: Quadratic Hamming weight with a spike . . . 83
4.6.1 Unperturbed quadratic: Overlap with the binomial ansatz . . . 83
4.6.2 Unperturbed quadratic: Improved spectral gap bound . . . 85
4.6.3 Perturbed quadratic and AQO . . . 86
4.6.4 Runtime for quadratic HWS . . . 87
4.7 Summary, related work and open questions . . . 87
4.7.1 Related work . . . 87
4.7.2 Open questions . . . 88

Conclusion and outlook **91**

A Log concavity supplementary results **95**
A.1 Proof of the binomial ansatz overlap . . . 95
A.1.1 Helper lemmas . . . 95
A.1.2 Binomial expectation lemmas . . . 96
A.1.3 Proof of claim . . . 97
A.2 Overlap bound for quadratic potentials . . . 98
A.3 Independent-set sampling for bipartite graphs . . . 100

Bibliography **107**

List of Figures

1.1 Example of the energy spectrum of the system during the annealing. 4
1.2 Encoding of a unit-disk graph from an atomic configuration 5

2.1 Illustration of the Rydberg blockade. 19
2.2 Example of a neutral-atom evolution $\Omega(t), \delta(t)$ with Pulser. 21

3.2 Blockade radius for locally driven two-atom systems with amplitudes Ω_0 and Ω_1, obtained by simulation 43
3.3 Gradient $\frac{\partial P_{11}^L}{\partial r}$ for locally driven systems 44
3.4 Maximal P_{11}^L for different Ω_0, Ω_1 45
3.5 Comparison of the double-excited probability for different drive configurations and Rydberg blockade 47
3.6 Graph embedding for unit-disk and disk graphs evaluated with the correlation matrix and maximal independence-violation probability 51
3.7 Relative sampling performance of local and global drives for non-unit-disk graphs from ground state analysis 54
3.8 Relative sampling performance of local and global drives from quantum simulation 55

4.1 Example of a non-monotone potential satisfying the RM condition 76
4.2 Plot of the function $m^*(a)/n$ obtained for $n = 40$ and $k_0 = -3$. 84
4.3 Overlap between the optimal binomial ansatz state $|\varphi_m\rangle$ and the ground state $|\psi_0(s)\rangle$ of $H(s) = -(1-s)A + sV$ for $n = 40$ and different potentials 89

A.1 Asymptotic log-density $\varphi_{AS}(\alpha, \beta)$ of independent sets in the random bipartite graph model. 101
A.2 Potential $V(x)$, V"(x), and two-dimensional plots of ϕ for $\Delta = 3$ and $\Delta = 4.80$. . . 104
A.3 Numerical solution for α as a function of Δ. 105
A.4 Numerical solution for the separation $1/2 - 2\alpha$ as a function of Δ. 105

List of Tables

3.1 Summary of the two-atom blockade regimes. 46
3.2 Violation probabilities for local and global drives 55

Résumé substantiel en français

L'optimisation est l'une des tâches centrales de l'informatique : l'objectif est de trouver une solution optimale parmi toutes les solutions admissibles. Cette formulation simple couvre un large éventail de problèmes, allant du routage et de l'ordonnancement à l'optimisation de portefeuille, à l'allocation de ressources, à la conception de réseaux sans fil, ou encore à l'analyse de grands graphes. Les problèmes d'optimisation peuvent faire intervenir des variables continues, telles que des paramètres à valeurs réelles, ou des variables discrètes, telles que des choix parmi des sous-ensembles, des affectations, des chemins ou des configurations de graphes. Ce dernier cadre est celui de l'optimisation combinatoire, où l'on recherche parmi un ensemble fini, mais souvent exponentiellement grand, d'objets admissibles. Dans de nombreux exemples, l'entrée peut être décrite par un graphe : les sommets, notés par les éléments d'un ensemble V, représentent des objets, des agents, des tâches ou des lieux physiques, tandis que les arêtes, notées par les éléments d'un ensemble E, encodent des interactions, des incompatibilités ou des contraintes.

Un exemple est le problème du stable maximum, ou Maximum Independent Set (MIS). Étant donné un graphe non orienté, $G = (V, E)$, un ensemble indépendant est un sous-ensemble de sommets ne contenant aucune paire de sommets adjacents, et l'objectif est d'en trouver un de taille maximale. MIS est un problème fondamental en optimisation de graphes et a été largement étudié [Erc13]. Il apparaît naturellement dans de nombreuses situations pratiques, tout en restant NP-difficile en général [Aro98]. Considérons, par exemple, le problème consistant à maximiser le nombre de talkies-walkies actifs simultanément. Si deux foyers proches utilisent la même fréquence, leurs signaux interfèrent, empêchant les deux communications d'être fiables. On représente chaque foyer par un sommet et l'on ajoute une arête entre deux foyers suffisamment proches pour interférer. Maximiser le nombre de talkies-walkies sans interférence revient alors à trouver un ensemble indépendant maximum dans ce graphe, correspondant à un plus grand ensemble d'émetteurs mutuellement non interférents. Des formulations similaires apparaissent dans la conception de réseaux sans fil [MC21], l'étiquetage cartographique [AvS98], l'ordonnancement [HS95] et l'allocation de ressources [ZWWZ17].

Depuis les propositions du début des années 1980 de Feynman et Benioff d'utiliser des systèmes quantiques comme dispositifs de calcul [Fey82, Ben80], le calcul quantique s'est imposé comme un modèle de calcul fondamentalement différent, exploitant les effets de la

mécanique quantique pour traiter l'information. Comme de nombreux problèmes d'optimisation sont difficiles pour les ordinateurs classiques, il est naturel de se demander si les ordinateurs quantiques peuvent en résoudre certains plus efficacement. Le calcul quantique se présente au moins sous deux formes principales : le modèle numérique, qui décrit un calcul comme une succession de portes élémentaires appliquées à des qubits, et le modèle analogique, dans lequel le système évolue continûment sous l'effet d'un Hamiltonien dépendant du temps. Dans le cadre numérique, plusieurs algorithmes quantiques sont connus pour surpasser les meilleurs algorithmes classiques connus, tels que l'algorithme de Shor pour la factorisation d'entiers [Sho94] et l'algorithme de Grover pour la recherche non structurée [Gro96]. Dans le cadre analogique, la simulation quantique est une application paradigmatique [Fey82, Dal23, GAN14] : les simulateurs quantiques analogiques sont naturellement adaptés aux systèmes de grande taille et fortement intriqués, avec des applications allant des propriétés de l'état fondamental [SSW+21, EWL+21, KBR+21, NHRDT24, KKO+20, LJFF+26, MCJ+17, SLK+21] aux dynamiques hors équilibre [MLE+25, AXB+24]. Plus généralement, la simulation quantique promet d'avoir un impact en chimie et en science des matériaux, pour des problèmes où la nature quantique du système rend la simulation classique exponentiellement coûteuse [DBK+22]. Un autre paradigme analogique important, central dans cette thèse, est le calcul quantique adiabatique, où la solution d'un problème est encodée dans l'état fondamental d'un Hamiltonien. Cependant, les ordinateurs quantiques sont actuellement très sensibles aux erreurs, tandis que les algorithmes présentant les avantages asymptotiques les plus nets nécessitent généralement des dispositifs de grande échelle, tolérants aux fautes, capables d'exécuter de longs calculs de manière fiable.

De telles machines pleinement tolérantes aux fautes ne sont pas encore disponibles. Cela motive l'étude des algorithmes quantiques à court terme, qui visent à exploiter les dispositifs actuellement disponibles malgré leur taille limitée, leur bruit et leurs contraintes matérielles. Comprendre ce qui peut être accompli dans ce régime, et où se situent ses limites, est une question clé du domaine.

Cette thèse étudie des algorithmes quantiques analogiques pour l'optimisation combinatoire. Nous nous concentrons sur le recuit quantique, une méthode analogique permettant de trouver les minima de fonctions de coût $V : \{0,1\}^n \to \mathbb{R}$ sur l'hypercube, qui exploite les fluctuations quantiques pour échapper aux minima locaux dans lesquels les algorithmes classiques peuvent rester piégés [FGG+01, KN98]. Le système est initialisé dans l'état fondamental d'un Hamiltonien simple, puis évolue lentement vers un Hamiltonien final dont l'état fondamental encode la solution recherchée. Ses performances sont gouvernées par l'écart spectral minimal le long de ce chemin : lorsque cet écart reste suffisamment grand, l'état fondamental du Hamiltonien final peut être préparé efficacement.

Cette thèse est organisée en deux parties, qui abordent deux aspects complémentaires de l'optimisation quantique analogique sur des dispositifs à atomes neutres à court terme. La première partie concerne l'encodage, c'est-à-dire l'ensemble des problèmes d'optimisation

de graphes qui peuvent être encodés sur une plateforme quantique à atomes neutres.Sur les dispositifs à atomes neutres, les atomes interagissent par le blocage de Rydberg, qui fournit une manière native d'encoder des problèmes de graphes tels que MIS. Cependant, les graphes qui émergent directement de la géométrie des atomes sont restreints par des contraintes expérimentales, telles que la distance minimale entre atomes. Comprendre et étendre cette classe native de graphes est donc essentiel pour évaluer la portée computationnelle des plateformes à atomes neutres.

La seconde partie concerne l'analyse théorique du processus de recuit : il s'agit de déterminer si la dynamique prépare effectivement l'état fondamental souhaité. Cela est gouverné par l'écart spectral, difficile à estimer pour des problèmes généraux, les analyses rigoureuses n'étant disponibles que dans des cas particuliers [AKR10, CPGW22, FGG+01]. Pour rendre l'analyse traitable, nous nous restreignons à des potentiels qui ne dépendent que du poids de Hamming, et nous demandons jusqu'où une préparation efficace de l'état fondamental peut encore être garantie.

Aperçu des résultats

Plongements de graphes de disques sur des dispositifs à atomes neutres Le chapitre 3 étudie des stratégies de plongement pour l'optimisation de graphes sur des dispositifs à atomes neutres. Le blocage de Rydberg natif encode directement des graphes de disques unitaires, mais cette classe est limitée par des contraintes géométriques et expérimentales, telles que la distance minimale entre atomes. Nous montrons que les graphes de disques deviennent accessibles en définissant le blocage à partir de la probabilité de double excitation. Plus précisément, nous définissons le rayon de blocage effectif comme la distance à laquelle la probabilité que deux atomes soient simultanément excités est égale à 1/2. Cette définition fournit une description précise du blocage de Rydberg.

À partir de cette perspective, nous caractérisons les plongements au moyen de la matrice de corrélation : des entrées proches de zéro indiquent des paires d'atomes bloquées, tandis que des entrées proches de un indiquent des paires pouvant être co-excitées. Comme le blocage de Rydberg limite la dynamique au sous-espace des ensembles indépendants [FSS04, LK12], nous évaluons en outre la qualité du plongement à travers la probabilité maximale de violer la condition d'ensemble indépendant.

Sur les instances étudiées, nos résultats numériques montrent que les schémas à adressage local préparent des ensembles indépendants optimaux ou proches de l'optimal plus efficacement que les protocoles globaux. Cela permet d'étendre l'optimisation par atomes neutres au-delà du cadre natif des graphes de disques unitaires, vers des familles de graphes plus larges et plus difficiles.

Log-concavité des états fondamentaux quantiques Le chapitre 4 étudie la structure de l'état fondamental instantané en optimisation quantique adiabatique. Pour une large famille de potentiels, incluant les potentiels convexes, nous démontrons que l'état fondamental de l'opérateur de Schrödinger discret correspondant est log-concave. Il s'agit d'un analogue discret du théorème continu de Brascamp–Lieb [BL76], qui fournit un nouvel outil structurel pour l'analyse du recuit quantique.

En particulier, cela implique des bornes inférieures polynomiales sur l'écart spectral pour une large gamme de potentiels, incluant des potentiels convexes, monotones et même non convexes. Nous montrons ensuite que l'analyse perturbative du pic due à Reichardt [Rei04] s'étend aux états fondamentaux log-concaves généraux. En exploitant des propriétés structurelles telles que la monotonie et la concentration, nous réduisons l'analyse de l'effet tunnel à travers un pic à l'estimation de la moyenne et de la variance de l'état fondamental sans pic, et nous donnons des conditions explicites sur les paramètres du pic sous lesquelles l'écart spectral reste inverse-polynomialement grand.

Comme application, nous analysons des potentiels quadratiques avec un pic. Bien que l'état fondamental n'ait pas de forme analytique, nous montrons qu'il conserve un grand recouvrement avec une distribution binomiale et utilisons cela pour caractériser précisément sa variance. Cela nous permet de démontrer que l'écart spectral reste constant sous des conditions sur le pic analogues à celles du problème original du poids de Hamming avec pic, généralisant directement ce résultat d'un potentiel linéaire à un potentiel quadratique. L'analyse suggère fortement une applicabilité plus large des accélérations fondées sur l'effet tunnel à des paysages présentant des pics.

Chapter 1

Introduction

1.1 Context

Optimization is one of the central tasks of computer science: the goal is to find an optimal solution among all feasible ones. This simple formulation covers a wide range of problems, from routing and scheduling to portfolio optimization, resource allocation, wireless network design, or the analysis of large graphs. Optimization problems can involve continuous variables, such as real-valued parameters, or discrete variables, such as choices among subsets, assignments, paths, or graph configurations. The latter setting is the domain of combinatorial optimization, where one searches over a finite, but often exponentially large, set of feasible objects. In many of these examples, the input can be described by a graph: vertices, denoted by elements of a set V, represent objects, agents, tasks, or physical locations, while edges, denoted by elements of a set E, encode interactions, incompatibilities, or constraints.

An example is the Maximum Independent Set problem (MIS). Given an undirected graph $G = (V, E)$, an independent set is a subset of vertices containing no pair of adjacent vertices, and the goal is to find one of maximum size. MIS is a fundamental problem in graph optimization and has been studied extensively [Erc13]. It arises naturally in many practical settings, while remaining NP-hard in general [Aro98]. Consider, for instance, the problem of maximizing the number of simultaneously active walkie-talkies. If two nearby households use the same frequency, their signals interfere, preventing both from communicating reliably. We represent each household by a vertex and add an edge between any two households that are close enough to interfere. Maximizing the number of non-interfering walkie-talkies then amounts to finding a maximum independent set in this graph, corresponding to a largest set of mutually non-interfering transmitters. Similar formulations arise in wireless network design [MC21], map

labelling [AvS98], scheduling [HS95], and resource allocation [ZWWZ17].

Since the early-1980s proposals by Feynman and Benioff to use quantum systems as computational devices [Fey82, Ben80], quantum computing has emerged as a fundamentally different model of computation, exploiting quantum mechanical effects to process information. Since many optimization problems are hard for classical computers, it is natural to ask whether quantum computers can solve some of them more efficiently. Quantum computation comes in at least two main forms: the digital model, in which a computation is decomposed into a sequence of elementary gates acting on qubits, and the analog model, in which the system evolves continuously under a time-dependent Hamiltonian. In the digital setting, several quantum algorithms are known to outperform the best known classical ones, such as Shor's algorithm for integer factorization [Sho94] and Grover's algorithm for unstructured search [Gro96]. In the analog setting, quantum simulation is a paradigmatic application [Fey82, Dal23, GAN14]: analog quantum simulators are naturally suited to large-scale, highly entangled systems, with applications ranging from ground-state properties [SSW+21, EWL+21, KBR+21, NHRDT24, KKO+20, LJFF+26, MCJ+17, SLK+21] to nonequilibrium dynamics [MLE+25, AXB+24]. More broadly, quantum simulation promises impact in chemistry and materials science, for problems where the quantum nature of the system makes classical simulation exponentially costly [DBK+22]. Another important analog paradigm, central to this thesis, is adiabatic quantum computation, where the solution of a problem is encoded in the ground state of a Hamiltonian. However, quantum computers are currently very sensitive to errors, while the algorithms with the clearest asymptotic advantages typically require large-scale, fault-tolerant devices capable of performing long computations reliably.

Such fully fault-tolerant machines are not yet available. This motivates the study of near-term quantum algorithms, which aim to exploit currently available devices despite their limited size, noise, and hardware constraints. Understanding what can be achieved in this regime, and where its limitations lie, is a key question in the field.

This thesis studies analog quantum algorithms for combinatorial optimization. We focus on quantum annealing, an analog method for finding the minima of cost functions $V : \{0,1\}^n \to \mathbb{R}$ on the hypercube, which leverages quantum fluctuations to escape the local minima that trap classical algorithms [FGG+01, KN98]. The system is initialized in the ground state of an easy Hamiltonian and then slowly evolved toward a final Hamiltonian whose ground state encodes the desired solution. Its performance is governed by the minimum spectral gap along this path: when the gap stays large enough, the target ground state can be prepared efficiently.

This thesis is organized in two parts, addressing two complementary aspects of analog quantum optimization on near-term neutral-atom devices. The first part concerns encoding, namely the range of graph optimization problems that can be represented on the hardware. On neutral-atom devices, atoms interact through the Rydberg blockade, which provides a native way to encode graph problems such as MIS. However, the graphs that arise directly from the

geometry of the atoms are restricted by hardware constraints, such as the minimum distance between atoms and the maximum spatial extent of the atomic layout. Understanding and extending this native graph class is therefore essential for assessing the computational reach of neutral-atom platforms.

The second part concerns the theoretical analysis of the annealing process: whether the dynamics actually prepares the desired ground state. This is governed by the spectral gap, which is difficult to estimate for general problems, with rigorous analyses available only in special cases [AKR10, CPGW22, FGG+01]. To make the analysis tractable, we restrict to potentials that depend only on the Hamming weight, and ask how far an efficient preparation of the ground state can still be guaranteed.

1.1.1 Quantum computation and quantum annealing

Having motivated quantum annealing as an analog approach to combinatorial optimization, we now review the relevant background in more detail, starting with a brief history of quantum computation.

Quantum computing was introduced in early-twentieth-century physics. Around 1900, Max Planck [Pla72] was one of the first to suggest that energy was quantized, thereby inaugurating the field of quantum mechanics. This was then followed by the works of Albert Einstein [EPR35], Niels Bohr [BJ25], Werner Heisenberg [Hei25], and Erwin Schrödinger [Sch26], among others. A quantum state of n two-level systems is described by a vector in a Hilbert space whose dimension grows exponentially with n. This exponential growth makes the simulation of quantum systems on classical computers computationally expensive, a difficulty famously emphasized by Feynman [Fey82], who argued that quantum systems might be more naturally simulated by other quantum systems. For near-term analog platforms, a particularly relevant framework is adiabatic quantum computation.

A central tool in this setting is the quantum adiabatic theorem, first proved by Born and Fock in 1928 [BF28, Kat50]: if a system is prepared in an instantaneous eigenstate and the Hamiltonian is varied slowly enough, the system stays in that eigenstate, provided it remains separated from the rest of the spectrum by a nonzero energy gap. Conversely, under a fast variation the system can be excited to a higher energy level [Lan32, Zen32].

Building on this theorem, Farhi introduced adiabatic quantum computation (AQC) in 2000 [FGGS00] as a framework for optimization. Instead of applying a sequence of gates, the algorithm prepares the system in the ground state of a simple, easy-to-prepare Hamiltonian. This Hamiltonian is then slowly evolved toward a target Hamiltonian whose ground state encodes the solution to the problem of interest. By the adiabatic theorem, if the evolution is slow relative to the spectral gap, the final state remains close to the target ground state, and hence to the solution it encodes.

Quantum annealing follows the same adiabatic principle, but is usually presented in the

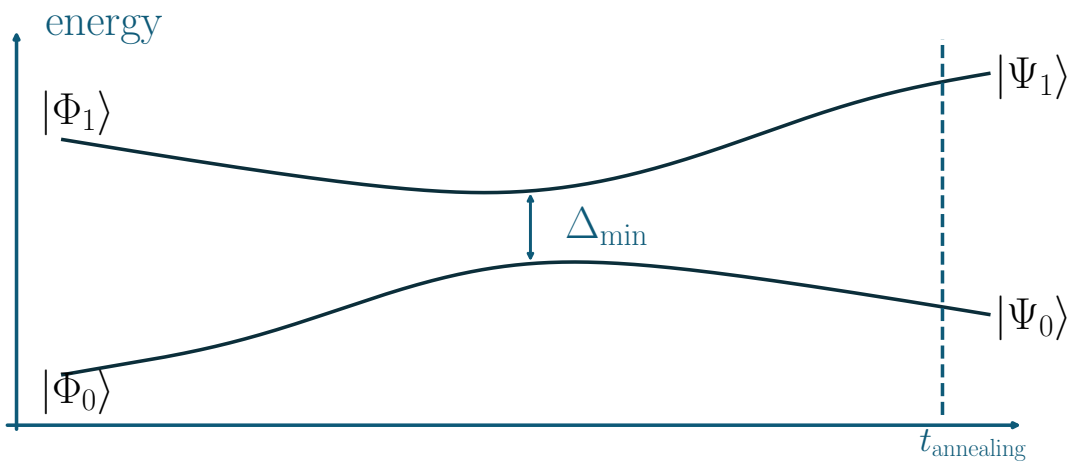


Figure 1.1: Example of the energy spectrum of the system during the annealing reproduced from [PKLV25] Initially, the two lowest energy levels are the ground state $|\Phi_0\rangle$ and the first excited state $|\Phi_1\rangle$. Over the annealing process, of total duration $t_{\text{annealing}}$, the system evolves and ends in the lowest eigenstates $|\Psi_0\rangle$ and $|\Psi_1\rangle$. The minimum spectral gap $\Delta_{\min}$ is the smallest energy separation between the two levels along the evolution. To follow the instantaneous ground state from $|\Phi_0\rangle$ to $|\Psi_0\rangle$, the evolution must be slow compared to $\Delta_{\min}$.

language of energy landscapes and quantum fluctuations. These fluctuations allow the system to explore classical configurations and potentially escape local minima that would trap classical algorithms. The hope, already articulated in the late 1980s [dFACB88, ACd89], is to exploit this mechanism to solve combinatorial optimization problems.

The performance of quantum annealing is controlled by the minimum spectral gap $\Delta_{\min}$ along the annealing path (see Figure 1.1). If this gap shrinks no faster than inverse-polynomially in the system size, the adiabatic theorem provides an efficient preparation guarantee. If it shrinks exponentially, the Landau-Zener analysis suggests that the required evolution time may become exponential. For hard optimization problems, this second regime is often the main obstruction to efficient adiabatic evolution.

1.1.2 Neutral atoms and graph optimization

Among near-term platforms, neutral-atom quantum processors provide a natural setting for analog quantum optimization: atoms can be viewed as vertices, and interacting pairs of atoms as edges of a graph. This thesis sits at the intersection of quantum algorithms and their implementation on such devices. Pasqal, a spin-off of the Institut d'Optique Graduate School in Palaiseau, builds quantum processors for computing and simulation by controlling individual rubidium atoms with lasers [SSW+21, BdLL+16]. In these devices, individual atoms are trapped with optical tweezers and excited to Rydberg states, which induce strong and controllable

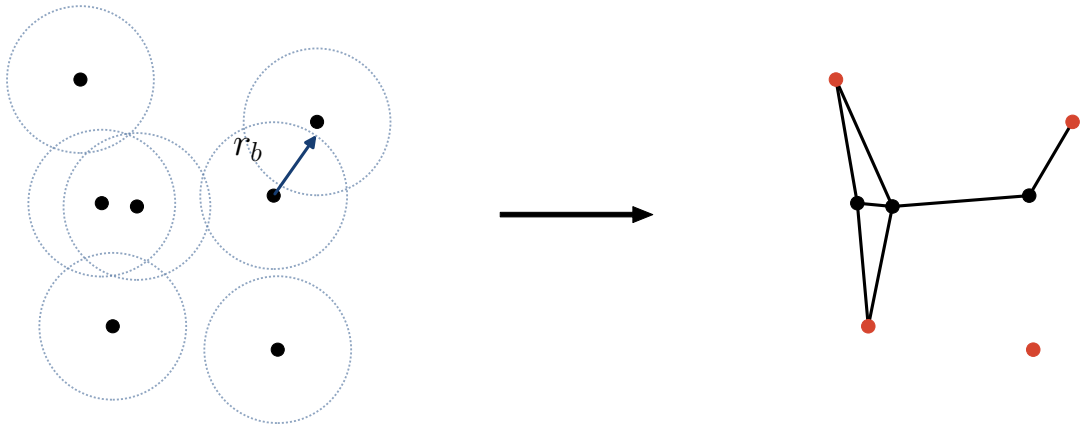


Figure 1.2: From a spatial configuration of Rydberg atoms to a unit-disk MIS graph problem. Two Rydberg atoms interact strongly when their separation is below a characteristic distance r_b (left), making it impossible for both to be excited at the same time. This is the Rydberg blockade, and it reproduces exactly the independent set constraint: in the graph defined by the atoms, an edge links any two atoms closer than r_b. On the right, we show the resulting graph, in which the red vertices form an independent set. Figure adapted from [HBS+20].

interactions between nearby atoms [HBS+20]. The key mechanism is the Rydberg blockade: two atoms that are sufficiently close cannot both be excited simultaneously, creating a direct link between the geometry of the atoms and a graph constraint.

This link is made explicit by associating a graph with an array of atoms: one places a vertex at each atom and joins two vertices by an edge whenever the corresponding atoms are separated by less than the blockade radius. Because of the Rydberg blockade, the set of atoms excited to the Rydberg state must form an independent set of this graph. Moreover, by making configurations with many excitations energetically favorable, the ground state of the Rydberg system can be made to encode the MIS problem [PWZ+18, DH22, EKC+22].

As one might expect, this type of constraint restricts the class of graphs that can be encoded. In particular, neutral-atom devices encode unit-disk graphs: a unit-disk graph is a graph that can be embedded in the two-dimensional Euclidean plane so that two vertices are adjacent if and only if they lie within a fixed unit distance of each other.

For this class of graphs, MIS admits a classical polynomial-time approximation scheme [CCJ90], so near-optimal solutions can be found efficiently and a quantum advantage seems limited. For general graphs, MIS is NP-hard in the worst case [Hås99, EJvL05] and even APX-complete for graphs of bounded maximum degree [PY91], so no efficient classical approximation is expected.

Since the hardware natively encodes only unit-disk graphs, exactly the regime where MIS is classically easy, a quantum advantage would require encoding harder graph classes.

A natural question is therefore how to extend neutral-atom devices beyond unit-disk graphs. The first part of this thesis addresses it by studying embeddings of the broader class of disk graphs [vL06, MM14], in which vertices correspond to disks of possibly different radii, with an edge between two vertices whenever the corresponding disks overlap.

1.1.3 From embeddability to algorithmic guarantees

Once a problem has been encoded into the physical hardware, we must still understand whether quantum annealing can succeed, which requires estimating the spectral gap.

One mechanism that makes quantum annealing appealing is tunneling. In classical simulated annealing [KGV83], the system explores the landscape through thermal transitions and must climb over energy barriers, which can lead to exponentially slow mixing; a quantum system may instead tunnel directly through a sufficiently narrow barrier. The canonical illustration is the Hamming-weight-with-a-spike problem [FGG02, Rei04], where one considers a potential of the form $V(x) = |x| + h(x), x \in \{0,1\}^n$, with h a localized perturbation, or "spike". In this setting, classical simulated annealing can fail to find the minimum in polynomial time, while quantum annealing succeeds. Because the unperturbed Hamming-weight Hamiltonian is exactly solvable, a substantial body of work has analyzed such perturbations [KC17, MAL16, BvD16, Ber20]. It is natural to ask how much of this tunneling advantage persists for more general optimization landscapes.

The second part of this thesis develops a structural approach to this question. In the continuous setting, log-concavity is a fundamental tool for proving functional inequalities and spectral-gap bounds, as illustrated for instance by recent works on log-concave sampling [Che22]. In the discrete setting relevant to adiabatic quantum optimization, however, this perspective has been much less explored. We consider permutation-symmetric potentials, i.e., those satisfying $V(x) = V(|x|)$, and prove that, under suitable conditions, the corresponding ground state is log-concave. This structural property makes the analysis of spikes tractable, placing the Hamming-weight-with-a-spike example within a broader theory of quantum tunneling.

1.2 Overview of results

1.2.1 Disk graph embeddings on neutral-atom devices

Chapter 3 studies embedding strategies for graph optimization on neutral-atom devices. The native Rydberg blockade directly encodes unit-disk graphs, but this class is limited by geometric and hardware constraints, such as the minimum distance between atoms. We show that disk graphs become accessible by defining the blockade through the two-atom excitation probability. More precisely, we define the effective blockade radius operationally as the distance at which

the probability that two atoms are simultaneously excited is equal to 1/2. This definition gives an accurate description of the physical blockade.

Using this perspective, we characterize embeddings through the correlation matrix: entries close to zero indicate pairs of atoms that are blockaded, while entries close to one indicate pairs that can be co-excited. Since the Rydberg blockade limits the dynamics to the independent-set subspace [FSS04, LK12], we further evaluate embedding quality through the maximum probability of violating the independent-set condition.

On the instances studied, our numerical results show that locally addressed driving schemes prepare larger independent sets on disk graphs than global protocols. By allowing individual atoms to be tuned independently, local addressing compensates for the non-uniformity of disk-graph embeddings and extends neutral-atom optimization beyond the native unit-disk setting, toward broader and harder graph families.

1.2.2 Log-concavity of quantum ground states

Chapter 4 studies the structure of the instantaneous ground state in adiabatic quantum optimization. For a broad family of potentials, including convex potentials, we prove that the ground state of the corresponding discrete Schrödinger operator is log-concave. This is a discrete counterpart of the continuous Brascamp-Lieb theorem [BL76] and provides a new structural tool for the analysis of quantum annealing.

In particular, this implies polynomial lower bounds on the spectral gap for a broad range of potentials, including convex, monotone, and even non-convex ones. We then show that the perturbative spike analysis of Reichardt [Rei04] extends from binomial states to general log-concave ground states. Exploiting structural properties such as monotonicity and concentration, we reduce the analysis of spike tunneling to estimating the mean and variance of the unperturbed ground state, and we give explicit conditions on the spike parameters under which the spectral gap remains inverse-polynomially large.

As an application, we analyze quadratic potentials with a spike. Although the ground state has no closed analytical form, we show that it maintains large overlap with a binomial distribution and use this to tightly characterize its variance. This allows us to prove that the spectral gap remains constant under spike conditions analogous to those in the original Hamming-weight-with-a-spike problem, directly generalizing that result from a linear to a quadratic potential. The analysis strongly suggests the broader applicability of tunneling-based speedups to landscapes with spikes.

1.3 Publications

This thesis is based on the following papers.

Chapter 3

- [BVC+25], Bermot, E., Valor, L., Coelho, W., Henry, L.-P., and Pearson, N. (2025). *Local Rydberg blockade regimes for disk graph embedding and quantum optimization*, submitted to the IOP Science proceedings, revisions submitted, arXiv:2506.13228.

Chapter 4

- [BBA26], Braida, A., Bermot, E., and Apers, S. (2026). *Log-concavity and tunneling: adiabatic quantum optimization for convex functions with a spike*, published in the European Symposium on Algorithms (ESA) 2026 proceedings and presented at Adiabatic Quantum Computing (AQC) 2026 conference, arXiv:2606.23614.

CHAPTER 2

PRELIMINARIES

In this chapter we introduce some notations, models and results that we will use throughout the thesis. In later chapters we will remind the reader of some of these before using them.

2.1 Notations and basic definitions

For $n \in \mathbb{N}$, we write $[n] = \{1, \dots, n\}$. We denote by $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$ the sets of nonnegative integers, integers, rational numbers, real numbers, and complex numbers, respectively. For a finite set S, its cardinality is denoted by $|S|$.

Let Σ be a finite alphabet. A string of length n over Σ is an element $x \in \Sigma^n$. In the special case $\Sigma = \{0, 1\}$, we call x a bitstring. The Hamming weight of a bitstring $x \in \{0, 1\}^n$ is

$$|x| := \Big|\{i \in [n] : x_i = 1\}\Big|.$$

Unless specified otherwise, log denotes the binary logarithm $\log_2$, and ln denotes the natural logarithm. For $f \in \mathbb{R}$, we write $\exp(f) = e^f$. The expectation of a random variable X is denoted by $\mathbb{E}[X]$, and its kth moment by $\mathbb{E}[X^k]$.

We use the asymptotic notations $O(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$, hiding constant factors and dependencies on parameters treated as constants. For functions $f, g : A \subseteq \mathbb{N} \to B \subseteq \mathbb{R}$, with $g(n) > 0$ for all sufficiently large n, we write

- $f(n) = O(g(n))$ if $\exists C > 0, \exists n_0$ s.t. $\forall n \geq n_0,\ f(n) \leq Cg(n)$,
- $f(n) = \Omega(g(n))$ if $\exists C > 0, \exists n_0$ s.t. $\forall n \geq n_0,\ f(n) \geq Cg(n)$,
- $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

Throughout the thesis, unless stated otherwise, a graph is denoted by $G = (V, E)$, with $V = [n]$. Thus $n = |V|$ denotes the number of vertices, which will also coincide with the

number of qubits or atoms in the corresponding encoding. We use $z_i \in \{0,1\}$ for classical binary variables.

Notation	Meaning
$G = (V, E)$	graph, usually with $V = [n]$
n	number of vertices/qubits/atoms
A_G	adjacency matrix of G
$z \in \{0,1\}^n$	classical bitstring
$\|z\|$	Hamming weight of z
$\mathcal{IS}(G)$	Set of independent sets of G
$\mathcal{MIS}(G)$	Set of maximum independent sets of G
$C_G(z)$	classical MIS cost function
H_C	diagonal cost Hamiltonian
A_{hc}	reduced hypercube adjacency operator
K	Hamming-weight operator, $K\|k\rangle = k\|k\rangle$
$\Delta(s)$	spectral gap of $H(s)$
$d_{\max}(G)$	maximum degree of G

2.2 Combinatorial optimization

An undirected graph $G = (V, E)$ consists of a vertex set V and an edge set $E \subseteq \{\{u, v\} : u, v \in V,\ u \neq v\}$. For $u, v \in V$, saying that u and v are connected by an edge means that $\{u, v\} \in E$. For a graph $G = (V, E)$ with $V = [n]$, we denote its adjacency matrix by $A_G \in \mathbb{R}^{n\times n}$, where

$$(A_G)_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$$

In this thesis, we will often work with graphs together with an explicit geometric representation. On a first approach, we study unit-disk graphs. A graph $G = (V, E)$ is a unit-disk graph if there exist points $\{\vec{R}_i\}_{i\in V} \subset \mathbb{R}^2$ and a radius $R > 0$ such that

$$\{i, j\} \in E \quad \Longleftrightarrow \quad \|\vec{R}_i - \vec{R}_j\| \leq R.$$

By rescaling all distances, one may equivalently take $R = 1$. More generally, unit-disk graphs belong to the broader class of intersection graphs, which appear naturally in applications such as wireless communication networks [HS95] and map labelling [AvS98, WDJ+06].

Combinatorial optimization problems consist in making choices from a finite set of discrete elements in order to optimize an objective function while satisfying constraints. Typically, the number of feasible solutions grows exponentially with the input size, which makes exhaustive search impractical. We illustrate this with a standard example [AB09].

Suppose you are organising a dance performance with a group of dancers, but you know that some pairs of dancers cannot perform simultaneously because they require the same stage area or prop. You can represent this situation by a graph $G = (V, E)$, where vertices represent dancers, and an edge $\{i, j\} \in E$ indicates that dancers i and j cannot perform at the same time. To feature as many dancers as possible in a single scene while avoiding conflicts, you wish to find a largest subset of vertices with no edges between them.

Maximum Independent Set (MIS). An independent set in G is a subset $S \subseteq V$ such that no two distinct vertices in S are adjacent, i.e.,

$$\forall \{u, v\} \subseteq S, \quad \{u, v\} \notin E. \tag{2.1}$$

A maximum independent set is an independent set of maximum cardinality, denoted $|MIS|$. We define $\mathcal{IS}(G)$ as the set of all independent sets of G, and $\mathcal{MIS}(G)$ as the set of all maximum independent sets of G. Checking all subsets of V requires $2^{|V|}$ steps in the worst case, which scales exponentially with $n = |V|$. It is believed that no polynomial-time algorithm solves MIS exactly on all instances; indeed MIS is NP-hard [GJ78].

A convenient formulation uses binary variables $z \in \{0, 1\}^n$, where $z_i = 1$ if vertex i is selected and $z_i = 0$ otherwise. One way to solve this problem is to keep the constraints explicitly and use the mixed-integer linear programming (MILP) formulation

$$\max \sum_{i \in V} z_i \quad \text{such that} \quad \forall \{u, v\} \in E,\ z_u + z_v \leq 1. \tag{2.2}$$

This formulation is standard in operations research and combinatorial optimization. Exact optimization methods, such as branch-and-bound [LW66, Moo91], can then be used to solve the constrained problem.

Another approach consists in removing the constraints by adding them as a penalty in the objective function. This leads to the cost function

$$C_G(z) = -\sum_{i=1}^{n} z_i + U \sum_{\{i,j\} \in E} z_i z_j, \tag{2.3}$$

where U is a penalty parameter such that $U \gg d_{\max}(G)$ and $d_{\max}(G)$ denotes the maximum degree of G. The first term favors selecting many vertices, while the second term penalizes selecting adjacent vertices. For any independent set, the second term is zero. Minimizers of C_G correspond to maximum independent sets.

This gives a quadratic unconstrained binary optimization problem (QUBO), of the general form

$$\sum_i q_{ii} z_i + \sum_{i<j} q_{ij} z_i z_j.$$

Such formulations can be tackled classically using heuristic methods such as simulated annealing [KGV83], genetic algorithms, or integer-programming methods.

The complexity of the MIS problem can also be studied through structural parameters of the input graph [CFH+25]. One example parameter is the treewidth, which captures how tree-like a graph is. A tree decomposition of width w represents the graph by a tree of bags, each containing at most $w+1$ vertices, while preserving the adjacency structure of the graph. Given a tree decomposition of width w of a graph with n vertices, MIS can be solved in time $O\left(2^w w^{O(1)} n\right)$. For unit-disk graphs, solving MIS is still NP-complete [CCJ90]. But a more geometric way to measure the complexity is through the thickness of the graph. Informally, it measures how many disk centers can be simultaneously contained in a horizontal strip of unit width, minimized over all possible orientations of the strip. When the thickness t is small, the graph has a nearly one-dimensional structure that can be exploited algorithmically: van Leeuwen [vL05] showed that MIS can be solved on such graphs in time $O\left(t^2 2^{2t} n\right)$.

The approximability of MIS follows a similar dichotomy between general and geometric instances. On general graphs, the problem is notoriously hard to approximate: for any $\varepsilon > 0$, no polynomial-time algorithm can achieve an approximation ratio of $n^{1-\varepsilon}$ unless $P = NP$ [Hås99]. On unit-disk graphs, by contrast, the geometric structure makes the problem much more tractable. For example, the greedy algorithm that iteratively selects the leftmost disk, removes its neighborhood, and continues recursively on the remaining graph, produces an independent set whose size is at least one third of the maximum independent set [HM85]. Since exact polynomial-time algorithms are not expected for NP-hard problems, a large body of work has focused on approximation algorithms. In particular, unit-disk graphs admit polynomial-time approximation schemes, based for instance on breadth-first-search sphere decompositions [NHK05]. More refined geometric methods yield a $(1+\varepsilon)$-approximation running in time $n^{O(1/\varepsilon)}$ for any $\varepsilon > 0$.

2.3 Basic notions of quantum computing

In classical computing, the basic unit of information is a bit, which takes one of the two values 0 or 1. Quantum mechanics allows physical systems to exist in superpositions of classical states. This motivates the notion of a quantum bit (qubit), which is a coherent superposition of the basis states $|0\rangle$ and $|1\rangle$. We recall the basic formalism below; the interested reader may consult [NC11].

2.3.1 Some linear algebra

A Hilbert space is a complete vector space equipped with an inner product. In this work, we only consider finite-dimensional Hilbert spaces over $\mathbb{C}$.

For $z \in \mathbb{C}$, we write $\overline{z}$ for the complex conjugate and $|z| = \sqrt{z\overline{z}}$ for the modulus.

We use Dirac (bra–ket) notation: a (column) vector is written $|v\rangle$, and its conjugate transpose

is written $\langle v|$. For $|u\rangle, |v\rangle \in \mathbb{C}^n$, their inner product is

$$\langle u|v\rangle = \sum_{i=1}^{n} \overline{u_i}\, v_i \in \mathbb{C}.$$

Their outer product is the rank-one matrix $|u\rangle\langle v| \in \mathbb{C}^{n\times n}$ with entries

$$(|u\rangle\langle v|)_{i,j} = u_i\, \overline{v_j}.$$

The l_2-norm of $|v\rangle \in \mathbb{C}^n$ is

$$\| |v\rangle \|_2 = \left(\sum_{i=1}^{n} |v_i|^2 \right)^{1/2}.$$

For a matrix $M \in \mathbb{C}^{n\times m}$, we denote by $\|M\|$ the induced operator norm:

$$\|M\| = \sup_{\||\psi\rangle\|_2=1} \|M|\psi\rangle\|_2.$$

The adjoint (conjugate transpose) of M is denoted $M^\dagger$. A matrix $M \in \mathbb{C}^{n\times n}$ is Hermitian if $M = M^\dagger$ and unitary if $MM^\dagger = M^\dagger M = I$.

The tensor product of vector spaces $\mathbb{C}^a$ and $\mathbb{C}^b$ is denoted $\mathbb{C}^a \otimes \mathbb{C}^b \simeq \mathbb{C}^{ab}$. The tensor product of vectors $|u\rangle$ and $|v\rangle$ is written $|u\rangle \otimes |v\rangle$, often abbreviated $|u\rangle|v\rangle$ or $|u, v\rangle$. For matrices $M \in \mathbb{C}^{a_1\times b_1}$ and $M' \in \mathbb{C}^{a_2\times b_2}$, their tensor product is $M \otimes M' \in \mathbb{C}^{a_1a_2\times b_1b_2}$.

2.3.2 The postulates of quantum computing

We recall the standard postulates in a form sufficient for our purposes.

Qubit. A qubit is a unit vector in $\mathbb{C}^2$. We denote the computational basis by

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Thus any qubit can be written as $\alpha|0\rangle + \beta|1\rangle$ with $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$.

Writing $\alpha = |\alpha|e^{i\varphi_\alpha}$ and $\beta = |\beta|e^{i\varphi_\beta}$, we have

$$\alpha|0\rangle + \beta|1\rangle = e^{i\varphi_\alpha}\left(|\alpha||0\rangle + |\beta|e^{i(\varphi_\beta-\varphi_\alpha)}|1\rangle\right).$$

The phase φ_α is a global phase and has no physical effect. The relative phase $\varphi = \varphi_\beta - \varphi_\alpha$, however, is physically meaningful because it affects interference, and can therefore change the probabilities of later measurement outcomes.

Register. A system of n qubits is represented by a unit vector in $(\mathbb{C}^2)^{\otimes n} \simeq \mathbb{C}^{2^n}$, with computational basis $\{|x\rangle : x \in \{0,1\}^n\}$. A pure n-qubit state $|s\rangle$ is fseparable if it can be written as a tensor product of single-qubit states,

$$|s\rangle = \bigotimes_{i=1}^{n} |s_i\rangle .$$

If no such decomposition exists, the state is entangled. For example,

$$\frac{1}{\sqrt{2}} (|00\rangle + |01\rangle) = |0\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle)$$

is separable, while

$$\frac{1}{\sqrt{2}} (|00\rangle + |11\rangle)$$

is entangled. More generally, the minimal number of product states needed to express a state provides one way to quantify its entanglement.

Time evolution. Quantum algorithms are implemented through unitary transformations applied to quantum states. If an n-qubit state $|\psi\rangle$ evolves under a unitary operator U, the resulting state is

$$|\phi\rangle = U |\psi\rangle .$$

There are different ways to realize such unitary evolutions. In the digital circuit model, the desired unitary is decomposed into a product of elementary gates chosen from a universal gate set. Common single-qubit gates include the Pauli X and Z gates and the Hadamard gate H:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

We also use the single-qubit number operator

$$N_i = \frac{I - Z_i}{2},$$

which acts as $\mathrm{diag}(0,1)$ on qubit i. On two qubits i and j,

$$N_i \otimes N_j = \mathrm{diag}(0,0,0,1), \qquad \mathrm{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

In the analog model, the evolution is specified directly by a Hamiltonian, which may depend on time. Physically, the Hamiltonian describes the energy of the system and determines how the state evolves. Mathematically, it is a Hermitian operator acting on the Hilbert space of the system. If the system is in the state $|\psi(t)\rangle$, its evolution is governed by Schrödinger's equation

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = H(t) |\psi(t)\rangle .$$

The corresponding unitary evolution is

$$|\psi(t)\rangle = \mathcal{T} \exp\left(-\frac{i}{\hbar}\int_0^t H(t')\,dt'\right)|\psi(0)\rangle\,,$$

where $\mathcal{T}$ is the time-ordering operator. Thus, while the digital model builds the evolution as a product of gates, the analog model generates it through continuous Hamiltonian dynamics. Under suitable assumptions, these viewpoints are computationally equivalent in the sense that both can realize universal quantum computation; see, for instance, the review [AL18a].

For a time-independent Hamiltonian H, the stationary states satisfy

$$H\,|\psi\rangle = E\,|\psi\rangle\,.$$

Since H is Hermitian, its eigenvalues are real and can be ordered $E_0 \le E_1 \le \cdots$. The lowest eigenvalue E_0 is the ground-state energy, and an associated eigenvector is called a ground state. If the initial state is

$$|\psi(0)\rangle = \sum_k \alpha_k\,|E_k\rangle\,, \qquad \sum_k |\alpha_k|^2 = 1,$$

where $H\,|E_k\rangle = E_k\,|E_k\rangle$, then the solution of Schrödinger's equation is

$$|\psi(t)\rangle = \sum_k \alpha_k e^{-iE_k t/\hbar}\,|E_k\rangle\,.$$

In the rest of the thesis, we often use units where $\hbar = 1$.

Measurement. If an n-qubit state is written in the computational basis as $\sum_{j=0}^{2^n-1}\alpha_j\,|j\rangle$, then measuring the register in this basis yields outcome $|j\rangle$ with probability $|\alpha_j|^2$ (Born rule).

2.4 Quantum annealing and adiabatic evolution

In this section, we present the quantum annealing framework and recall the adiabatic theorem in the form that is relevant for preparing the ground state of a target Hamiltonian $H(1)$.

Definition 2.1 (Adiabatic quantum computation[AL18a])**.** *Let H_0 and H_1 be two Hamiltonians, and let $s : [0,T] \to [0,1]$ be a schedule with $s(0) = 0$ and $s(T) = 1$. The associated adiabatic path is*

$$H(t) = (1 - s(t))H_0 + s(t)H_1.$$

The goal of adiabatic quantum computation is to start from the ground state of H_0, assumed easy to prepare, and to prepare a state close to the ground-state subspace of H_1.

A main use of this method is to encode the solution of a combinatorial optimization problem into the ground state of a Hamiltonian. For the MIS problem, given a graph $G = (V, E)$ with $V = [n]$, the cost function can be written as

$$C_G(z) = -\sum_{i=1}^{n} z_i + U\sum_{i<j}(A_G)_{ij} z_i z_j.$$

We associate to this classical cost function a diagonal Hamiltonian H_C defined by

$$H_C \left|z_1, \ldots, z_n\right\rangle = C_G(z_1, \ldots, z_n) \left|z_1, \ldots, z_n\right\rangle .$$

Finding a ground state of H_C is therefore equivalent to finding a minimizer of C_G, and hence a maximum independent set of G as long as $U \gg 1$.

Let $\{|E_k(s)\rangle\}_{k=0}^{2^n-1}$ denote the eigenstates of $H(s)$, with associated eigenvalues $E_k(s)$, ordered as $E_0(s) \leq E_1(s) \leq \cdots$. The algorithm starts from an easy-to-prepare ground state of $H(0)$. If $|\psi(T)\rangle$ denotes the state prepared at the end of the evolution, the goal is to obtain $|\langle\psi(T)|E_0(1)\rangle|^2 \geq 1 - \varepsilon^2$. The main theoretical guarantee is the adiabatic theorem, which states that the system remains close to the instantaneous ground state if the evolution is slow enough compared to the inverse spectral gap.

Assume that the ground state of $H(s)$ is non-degenerate for all $s \in [0,1]$, and let $\Delta(s)$ denote the spectral gap between the ground state and the first excited state of $H(s)$, with

$$\Delta_{\min} := \min_{s\in[0,1]} \Delta(s).$$

The quantum adiabatic algorithm starts in the ground state of an easily preparable initial Hamiltonian $H(0)$ and evolves according to Schrödinger's equation: Let $t \in [0,T]$ denote the physical time of the evolution. The unitary evolution $U(t)$ is governed by the time-dependent Schrödinger equation

$$i\frac{d}{dt}U(t) = H\left(\frac{t}{T}\right)U(t), \qquad t \in [0,T], \qquad U(0) = I. \tag{2.4}$$

Here, T is the total evolution time of the algorithm. Introducing the reduced time

$$s = \frac{t}{T} \in [0,1],$$

and writing $U(s) := U(t = Ts)$, we obtain

$$i\frac{d}{ds}U(s) = TH(s)U(s), \qquad s \in [0,1], \qquad U(0) = I. \tag{2.5}$$

The adiabatic theorem then guarantees that, provided that this evolution time T is sufficiently large, the evolved state

$$|\phi(s)\rangle = U(s)\,|\psi_0(0)\rangle$$

remains close to the instantaneous ground state $|\psi_0(s)\rangle$ throughout the evolution, and in particular to $|\psi_0(1)\rangle$ at the end.

Theorem 2.1 (Adiabatic theorem, special case of [Rei04])**.** *Let $H(s)$, $s \in [0,1]$, be a twice differentiable family of Hamiltonians. Assume that the corresponding ground state energy $E_0(s)$ is non-degenerate and can be defined continuously for all $s \in [0,1]$, and that the spectral gap satisfies*

$$\Delta(s) \geq \Delta_{\min} > 0.$$

If $\|H(s)\|, \|\dot{H}(s)\| \leq \gamma$ *and*

$$T \in \Omega\left(\frac{\gamma^2}{\varepsilon \Delta_{\min}^3}\right),$$

then $|\langle\psi_0(1)|\phi(1)\rangle|^2 \geq 1 - \varepsilon^2$.

In this work, we consider only linear interpolations of the form

$$H(s) = (1-s)H_0 + sH_1,$$

for which the ground state is assumed to remain non-degenerate along the path, and hence can be chosen continuously in s. The initial Hamiltonian H_0 is chosen so that its ground state is easy to prepare, while the final Hamiltonian H_1 is the target Hamiltonian whose ground state we wish to obtain. For a linear interpolation, the Hamiltonian is twice differentiable and

$$\dot{H}(s) = H_1 - H_0,$$

so that $\|\dot{H}(s)\| = O(\|H_0\| + \|H_1\|)$. For constant error ε, the runtime bound therefore reduces to

$$T = O\left(\frac{(\|H_0\| + \|H_1\|)^2}{\Delta_{\min}^3}\right).$$

Many works have tried to improve this bound by adding smoothness assumptions on the schedule $s(t)$ to obtain a bound of the form $T = O(\Delta_{\min}^{-2})$[EH12]. In all cases, the main challenge is thus to prove a good lower bound on $\Delta_{\min}$. Since in general practical settings, we have $\|H_0\|, \|H_1\| \in \text{poly}(n)$, the adiabatic algorithm is efficient provided that

$$\Delta_{\min} \in \Omega\left(\frac{1}{\text{poly}(n)}\right).$$

As a special case, we often look at

$$H_0 = -\sum_{i=1}^{n} X_i,$$

whose unique ground state is the uniform superposition

$$|\psi_0(0)\rangle = |+\rangle^{\otimes n} = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right)^{\otimes n}.$$

If, moreover, H_1 is diagonal in the computational basis, then the interpolation is stoquastic, and the Perron-Frobenius theorem implies that the ground state may be chosen to have positive real amplitudes. In this case, one may equivalently write

$$H(s) = -(1-s)\sum_{i=1}^{n} X_i + sV,$$

with V a diagonal cost Hamiltonian.

2.5 Neutral-atom devices and Rydberg physics

In this section, we briefly describe neutral-atom quantum processors, which constitute Pasqal's hardware platform. Neutral-atom arrays are a promising approach to quantum computing, supporting both digital and analog modes of operation and enabling programmable many-body Hamiltonians [BVC+13, HBS+20, BL20]. Further background and physics can be found in the references above.

2.5.1 Principles of neutral-atom devices

Qubit encoding. Neutral-atom devices consist of configurable arrays of individually trapped atoms, typically rubidium atoms. A cloud of atoms is first prepared, and individual atoms are then trapped using optical tweezers. This makes it possible to create registers in which each trap contains at most one atom and to rearrange the atoms into a desired geometry.

This trapping process imposes hardware constraints on the possible atomic layouts. In Pasqal's devices, for instance, the minimal distance between two atoms is of the order of a few micrometers. These constraints will be important when discussing which graphs can be embedded on the device.

Each atom encodes a qubit in two electronic states, denoted $|0\rangle$ and $|1\rangle$. In the Rydberg setting, $|0\rangle$ is chosen as a low-energy ground state, while $|1\rangle$ is a highly excited Rydberg state with principal quantum number n_{R}. The radiative lifetime of such states scales approximately as n_{R}^3, and can reach hundreds of microseconds for the states used in current neutral-atom platforms.

Rydberg Hamiltonian. Two Rydberg atoms separated by a distance R interact through a van der Waals interaction. For atoms located at positions $\vec{R}_i$ and $\vec{R}_j$, we denote the interaction strength by

$$J_{ij} = \frac{C_6}{\|\vec{R}_i - \vec{R}_j\|^6}.$$

The interaction Hamiltonian is then

$$H_{\mathrm{int}} = \sum_{i<j} J_{ij} N_i N_j.$$

The coefficient C_6 depends strongly on the chosen Rydberg level: for Rydberg states with principal quantum number n_{R}, it typically scales approximately as $C_6 \sim n_{\mathrm{R}}^{11}$. The interaction strength decays as $\|\vec{R}_i - \vec{R}_j\|^{-6}$. For instance, for $n_{\mathrm{R}} = 70$, $C_6 = 862\,\mathrm{GHz} \cdot \mu\mathrm{m}^6$, while for $n_{\mathrm{R}} = 82$, $C_6 = 5559\,\mathrm{GHz} \cdot \mu\mathrm{m}^6$ [ŠPAW17]. In this thesis, we use the Rydberg state with principal quantum number $n_{\mathrm{R}} = 70$.

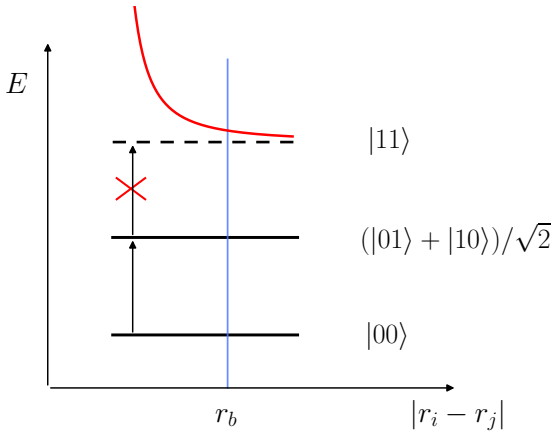


Figure 2.1: Illustration of the Rydberg blockade. Two atoms separated by a distance $R \geq r_b$ can both be excited to the Rydberg state $|1\rangle$ independently. When their distance r is lower than r_b, the van der Waals interaction shifts the doubly excited state $|11\rangle$ far out of resonance, preventing simultaneous excitation.

For each atom i, a coherent laser drive couples the two qubit states with amplitude $\Omega_i(t)$ and detuning $\delta_i(t)$, giving the effective Hamiltonian

$$H(t) = \frac{\hbar}{2} \sum_{i=1}^{n} \Omega_i(t) X_i - \hbar \sum_{i=1}^{n} \delta_i(t) N_i + \sum_{i<j} J_{ij} N_i N_j, \tag{2.6}$$

where X_i is the Pauli X operator on qubit i.

Rydberg blockade. For two atoms i and j separated by a distance R, the van der Waals interaction contributes a term $\frac{C_6}{R^6} N_i N_j$ to the Hamiltonian. Thus, when both atoms are in the Rydberg state, the energy of the doubly excited state $|11\rangle$ is shifted by an amount C_6/R^6. If the atoms are sufficiently close, this interaction-induced shift becomes large compared with the amplitude, and the laser drive is off-resonant with the transition to $|11\rangle$. This suppresses simultaneous excitation of the two atoms, a phenomenon known as the Rydberg blockade.

The corresponding blockade radius is commonly estimated as $r_b = \left(\frac{C_6}{\hbar\Omega}\right)^{1/6}$ where Ω is the amplitude. For two atoms separated by a distance smaller than r_b, simultaneous excitation is strongly suppressed. Conversely, because the van der Waals interaction decays as $1/R^6$, it can be neglected for atoms separated by distances $\|\vec{R}_i - \vec{R}_j\| \gg r_b$.

As a result, the low-energy dynamics naturally favors configurations in which no two atoms within the blockade radius are both in state $|1\rangle$. This is the physical mechanism behind the native encoding of independent sets on neutral-atom devices.

Embedding. The form of the Rydberg Hamiltonian naturally suggests an encoding of graph optimization problems. Consider a graph $G = (V, E)$ with $V = [n]$. To encode G on a neutral-atom device, each vertex $i \in V$ is assigned to an atom at position $\vec{R}_i$. The geometry of the

register then defines the interaction matrix

$$J_{ij} = \frac{C_6}{\|\vec{R}_i - \vec{R}_j\|^6}.$$

In the absence of laser driving, namely when $\Omega_i = 0$, and for a positive detuning $\delta_i = \delta > 0$, the Hamiltonian becomes

$$H_{\text{Ryd}} = -\hbar\delta \sum_{i=1}^{n} N_i + \sum_{i<j} J_{ij} N_i N_j.$$

This Hamiltonian has the same structure as the MIS cost function: the detuning term favors configurations with many excited atoms, while the interaction term penalizes simultaneous excitations of atoms that are close to each other.

The goal of the embedding is therefore to reproduce the adjacency matrix A_G: one wants J_{ij} to be large when $\{i, j\} \in E$, and sufficiently small when $\{i, j\} \notin E$. In the ideal blockade approximation, this gives the effective rule

$$\{i, j\} \in E \quad \Longleftrightarrow \quad \|\vec{R}_i - \vec{R}_j\| \leq r_b,$$

where r_b is the blockade radius. This is precisely the native unit-disk graph encoding induced by the Rydberg blockade.

The embedding problem is thus the task of choosing positions $\{\vec{R}_i\}_{i \in V}$ such that the induced interaction matrix reproduces, or approximates, the constraints of the target graph. In practice, this can be formulated as a continuous optimization problem over the atomic positions.

One way to evaluate the quality of an embedding is to study the correlation matrix

$$C_{ij} = \langle\psi| N_i N_j |\psi\rangle$$

for the system in the state $|\psi\rangle$. In this thesis, $|\psi\rangle$ will typically be either the final state of the evolution or an instantaneous state $|\psi(t)\rangle$. We also consider the maximal co-excitation probability along the dynamics,$\max_t \langle\psi(t)| N_i N_j |\psi(t)\rangle$. Small values of C_{ij} indicate that atoms i and j are unlikely to be simultaneously excited, which is the expected behavior for pairs that are blockaded. One may also consider the connected correlation matrix

$$C_{ij}^{\text{conn}} = \langle\psi| N_i N_j |\psi\rangle - \langle\psi| N_i |\psi\rangle \langle\psi| N_j |\psi\rangle .$$

This removes the contribution coming only from the marginal excitation probabilities of the two atoms. Comparing these correlation matrices with the adjacency matrix of the target graph gives a way to quantify whether the prepared state reflects the desired graph constraints.

The native blockade picture naturally encodes unit-disk graphs. One objective of this thesis is to investigate how this encoding can be extended beyond unit-disk graphs, in particular toward the broader class of disk graphs.

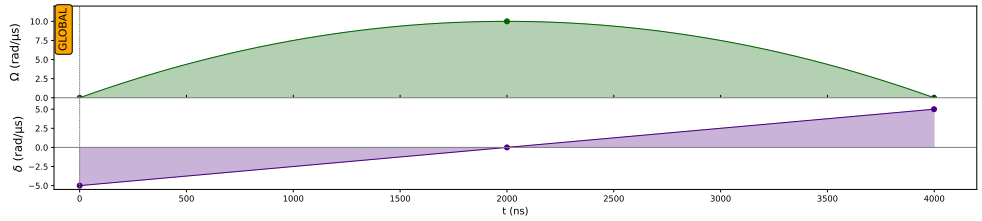


Figure 2.2: Example of a neutral-atom evolution $\Omega(t), \delta(t)$ generated with Pasqal's platform Pulser. The evolution is governed by $H(t) = \frac{\Omega(t)}{2} \sum_i X_i - \delta(t) \sum_i N_i + \sum_{i<j} \frac{C_6}{\|\vec{R}_i - \vec{R}_j\|^6} N_i N_j$. We usually impose $\Omega(0) = \Omega(T) = 0$, corresponding to the lasers being switched off at the beginning and end of the evolution, and choose $\delta(0) < 0$ so that the initial ground state is $|0 \dots 0\rangle$.

Evolution Once the target graph has been embedded into the atomic positions, the algorithm proceeds by choosing a time-dependent driving schedule, namely the functions $\Omega_i(t)$ and $\delta_i(t)$. This schedule determines the Hamiltonian $H(t)$, and the final state is

$$|\psi(T)\rangle = \mathcal{T} \exp\left(-i \int_0^T H(t)\, dt\right) |\psi_0\rangle\,.$$

where $\mathcal{T}$ is the time-ordering operator. If the system is initialized in the ground state and the parameters $\Omega_i(t)$, $\delta_i(t)$ are varied slowly, the evolution is adiabatic: the adiabatic theorem guarantees that the state remains close to the instantaneous ground state, provided the evolution is slow enough.

On the other hand, if the parameters are changed abruptly instead, the dynamics is a quench and is effectively governed by a time-independent Hamiltonian over each time interval. This is particularly relevant when studying the quantum approximate optimization algorithm (QAOA) [FGG14] where one studies states of the form

$$|\psi\rangle = e^{-i\alpha_p H_{mix}} e^{-i\beta_p H_C} \dots e^{-i\alpha_1 H_{mix}} e^{-i\beta_1 H_C} |+\rangle^{\otimes n}$$

where here $H_{mix} = -\sum_i X_i$ and H_C is the problem Hamiltonian as previously described.

Measurement. At the end of the evolution, the driving field is switched off, so that $\Omega(T) = 0$, and the atoms are measured in the computational basis. Each measurement outcome is a bitstring $z \in \{0,1\}^n$, where $z_i = 1$ means that atom i is found in the Rydberg state. The quality of the algorithm is then evaluated from the probability of sampling an independent set, or more specifically a maximum independent set.

2.5.2 MIS preparation

We now describe how to prepare a maximum independent set of a graph G on a neutral-atom device. The procedure consists of two steps:

- embedding the graph into the atomic register, so that the graph constraints are encoded in the atom positions;
- choosing a schedule $(\Omega_i(t), \delta_i(t))$ such that the final state has a large overlap with the MIS subspace.

By construction, the atom positions $\{\vec{R}_i\}_{i\in V}$ have been chosen so that the interaction term of the Rydberg Hamiltonian reproduces the adjacency matrix of the target graph. Consequently, setting $\Omega = 0$ and $\delta > 0$ yields a Hamiltonian whose ground-state subspace is the MIS subspace. It remains to design a schedule that drives the system into this ground state.

We follow the approach of [SSW+21], where adiabatic ramp sequences were used to prepare antiferromagnetic states on square and triangular lattices. This is a natural starting point because, in the Rydberg encoding, atoms excited to the Rydberg state represent occupied vertices, while the blockade forbids simultaneous excitation of neighboring atoms. On regular lattices, the resulting antiferromagnetic ordering corresponds to, or is closely related to, maximum independent-set configurations of the associated lattice graph. The schedule proceeds in three stages:

- linearly ramping up $\Omega(t)$ from 0 to $\Omega_{\max}$, while keeping $\delta(t) = \delta_0 < 0$, over a duration t_1;
- linearly ramping $\delta(t)$ from δ_0 to δ_f, while keeping $\Omega(t) = \Omega_{\max}$, over a duration t_2;
- linearly ramping down $\Omega(t)$ from $\Omega_{\max}$ to 0, while keeping $\delta(t) = \delta_f$, over a duration t_3.

The boundary conditions are chosen to ensure a valid adiabatic path. At $t = 0$, taking $\Omega(0) = 0$ and $\delta_0 < 0$ makes the all-zero state $|0\rangle^{\otimes n}$ the unique ground state, which is straightforward to prepare by optical pumping. At $t = t_f$, taking $\Omega(t_f) = 0$ and $\delta_f > 0$ makes the MIS subspace the ground-state subspace of the final Hamiltonian provided δ_f remains smaller than the relevant interaction penalties, so that configurations with adjacent excitations are not favored. Since the schedule varies slowly throughout, the adiabatic theorem guarantees that the system tracks the instantaneous ground state and ends close to the MIS at $t = t_f$. The full protocol is thus specified by six parameters: $\delta_0, \delta_f, \Omega_{\max}, t_1, t_2, t_3$.

There are several ways to choose the parameters of the pulse. A simple heuristic is to choose $\Omega_{\max}$ between the largest unwanted interaction and the smallest desired interaction. More precisely, define

$$U_{\min} = \min_{\{i,j\}\in E} \frac{C_6}{\|\vec{R}_i - \vec{R}_j\|^6}, \qquad U_{\max} = \max_{\{i,j\}\notin E} \frac{C_6}{\|\vec{R}_i - \vec{R}_j\|^6}.$$

One may then choose

$$\Omega_{\max} = U_{\max} + \alpha(U_{\min} - U_{\max}),$$

for some $\alpha \in [0, 1]$, so that $\Omega_{\max}$ lies between these two interaction scales. We also need to choose δ_f such that the MIS is the ground state of the final Hamiltonian. In practice, as inspired by [aD], we take

$$\delta_f = \frac{d_{\max}(G)}{d_{\max}(G) + 1} U_{\min},$$

where $d_{\max}(G)$ is the maximum degree of G. For the initial detuning, we take $\delta_0 = -\delta_f$. It is also possible to optimize the parameters using Bayesian optimization [CFH+25].

Chapter 3

Disk graphs

This chapter is partially based on results from [BVC+25], reproduced and adapted here with the authors' permission.

3.1 Introduction

In this chapter, we extend neutral-atom graph embeddings beyond the unit-disk setting. By revisiting the Rydberg blockade mechanism and the modelling assumptions behind the unit-disk approximation, we enlarge the class of graphs for which MIS can be implemented on a neutral-atom device.

3.1.1 Quantum complexity of UD-MIS

As discussed in the preliminary chapter, solving MIS on a unit-disk graph $G = (V, E)$ with a neutral-atom device can be viewed as a two-step pipeline. First, we choose atom positions $\{\vec{R}_i\}_{i \in V}$ such that the induced Rydberg interaction graph approximates the target adjacency structure. Second, we run a quantum evolution to prepare a state whose measurement yields a maximum independent set with high probability, most commonly via a quantum adiabatic algorithm that attempts to follow the instantaneous ground state of a time-dependent Hamiltonian. The final Hamiltonian is chosen so that its ground-state manifold corresponds to the set of maximum independent sets.

A number of works [EKC+22, CCL+23] have investigated the scaling of such annealing-based approaches for MIS on neutral-atom platforms, in particular on King's lattice instances. Throughout these studies, performance is typically characterized by the MIS size $|MIS|$, the success probability P_{MIS} of sampling an exact MIS after a fixed anneal, and the degeneracies

$D_{|MIS|}$ and $D_{|MIS|-1}$, i.e., the number of independent sets of size $|MIS|$ and $|MIS|-1$. In [EKC+22], the parameter $\mathcal{HP} = \frac{D_{|MIS|-1}}{|MIS| D_{|MIS|}}$, is used as a heuristic indicator of how difficult it is for the Rydberg dynamics to prepare an MIS. For King's lattice instances, $\mathcal{HP}$ is large precisely when the number of independent sets of size $|MIS|-1$ dominates that of size $|MIS|$, i.e., when $D_{|MIS|-1} \gg D_{|MIS|}$. For the hardest instances in [EKC+22], both simulated annealing and quantum annealing exhibit success probabilities of the form

$$P_{\mathrm{MIS}} = 1 - \exp\Big(-C\,\mathcal{HP}^{-a}\Big),$$

with $a = 1.03(4)$ for simulated annealing and $a = 0.95(15)$ for quantum annealing, indicating a broadly similar dependence on $\mathcal{HP}$ with a small advantage for the quantum approach in that benchmark. This qualitative picture is consistent with the follow-up study [ASM+23], which shows that state-of-the-art classical solvers can be competitive with the quantum protocol on the considered instances.

In addition, [CCL+23] provides a theoretical analysis of the spectral gap and identifies regimes in which the quantum adiabatic algorithm can be expected to outperform Markov-chain-based methods (including simulated annealing and parallel tempering). In particular, their results yield an explicit criterion in terms of the spectral gap Δ and the ratio $D_{|MIS|}/D_{|MIS|-1}$: when $\Delta \ll \frac{D_{|MIS|}}{D_{|MIS|-1}}$, quantum annealing is expected to suffer a slowdown relative to simulated annealing, whereas when $\Delta \gg \frac{D_{|MIS|}}{D_{|MIS|-1}}$, it can yield a speedup. Overall, these works show that quantum speedups over simulated annealing are possible, while a clear and robust advantage on existing UD-MIS benchmarks is not yet established.

This motivates a complementary perspective: even if one can in principle identify graph families for which quantum annealing achieves speedups for MIS, the class of graphs that can be realized on hardware remains a central bottleneck. Unit-disk graphs are already a geometrically constrained subclass of intersection graphs, and additional experimental and technological constraints (finite optical resolution, minimum atom spacing, finite system size, and connectivity constraints) further narrow the family of instances accessible in practice. A key challenge is therefore to enlarge the set of implementable graphs while retaining meaningful hardness guarantees for the resulting MIS instances.

Moreover, many UD-MIS benchmarks studied so far come from relatively narrow instance families, and in these regimes classical solution methods can outperform current quantum approaches [CFH+25], raising questions about the computational reach of existing protocols. In particular, [CFH+25] identifies graph parameters correlated with problem hardness within the unit-disk class and shows that both classical and quantum solvers display similar trends with respect to these hardness indicators. However, for currently accessible problem sizes, classical solvers can still outperform quantum experiments by orders of magnitude in time-to-solution, for example due to the overhead of measurement repetition on present-day hardware. Beyond hardware improvements, identifying experimentally realizable instances that are simultaneously

hard for classical solvers and easy for quantum ones remains challenging. A natural direction is therefore to expand the set of embeddable graphs toward families that support richer and harder instances, while still leveraging the analog structure of the Rydberg system.

Finally, it has become increasingly relevant to evaluate analog quantum protocols not only by their ability to reach exact ground states, but also by their capacity to produce high-quality approximate solutions [MBL25]. For MIS, this corresponds to reliably preparing independent sets close to optimal, which is the most relevant regime in the presence of noise and finite-time dynamics. This motivates the study of near-optimal independent-set sampling as a practical benchmark for analog quantum optimization protocols.

3.1.2 From a sharp blockade radius to local blockade radii

To extend the class of graphs that can be embedded in neutral-atom hardware, it is natural to revisit the modelling assumptions behind the standard embedding rule. In its simplest form, the Rydberg blockade picture draws an edge between two vertices whenever the corresponding atoms are placed at distance $r \leq r_b$. Equivalently, given positions $\{\vec{R}_i\}_{i\in V}$, the associated unit-disk graph connects i and j if $\|\vec{R}_i - \vec{R}_j\| \leq r_b$. This "hard-threshold" rule is effectively a sharp-threshold model: two atoms at distance r_b interact strongly enough to suppress double excitations, whereas atoms at distance $r_b + \varepsilon$ (for $\varepsilon > 0$) are treated as non-interacting.

Such a sharp cutoff would correspond to an idealized interaction Hamiltonian of the form

$$H_{\text{int}}^{\text{ideal}} = \sum_{i<j} h(\|\vec{R}_i - \vec{R}_j\|)\, N_i N_j, \qquad h(r) = \begin{cases} \infty & \text{if } r \leq r_b, \\ 0 & \text{otherwise.} \end{cases}$$

In real neutral-atom devices, however, interactions decay approximately as $1/r^6$, so pairs with separations close to r_b can still contribute non-negligible energy shifts. In particular, when embedding graphs (and especially when attempting dense or random embeddings), it is common to encounter pairs with $\|\vec{R}_i - \vec{R}_j\| \approx r_b$, where the hard-threshold approximation becomes inaccurate and the resulting dynamics may deviate from the ideal unit-disk model. This motivates a more refined theoretical treatment of the blockade mechanism, and in particular the move from a single sharp blockade radius to effective, potentially local blockade radii, as we develop in the remainder of this chapter.

3.1.3 Our contributions

Local Rydberg blockade. In this work, we introduce and characterize a two-atom theoretical framework for the Rydberg blockade that allows one to compare different amplitude settings. We consider three regimes. In the sequential regime, one atom is prepared in the excited state $|1\rangle$ and we then apply a quantum evolution to excite the second atom. In the global regime, both atoms undergo the same quantum evolution with a common amplitude Ω. In the local regime,

the two atoms undergo quantum evolutions with different amplitudes. We analyze the resulting two-atom dynamics in the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. For two atoms separated by a distance r, we denote by $P_{11}(r)$ the probability of occupying the doubly excited state $|11\rangle$. We then define the blockade radius as the distance r_b such that $P_{11}(r_b) = 1/2$. The following statement summarizes the two-atom blockade formulas that will be derived and validated in Sec. 3.3. The sequential, global, and local regimes are defined respectively in Section 3.3.3, Section 3.3.4 and Section 3.3.5.

Main Result 1 (Doubly excited-state probability)**.** *Let $P_{11}(r)$ be the doubly excited-state probability for two atoms separated by a distance r.*

- ***Sequential regime (Section 3.3.3).*** *We first consider the case where one atom is initially excited, and a quantum evolution is then applied to excite the second atom. In this setting, using exact diagonalization, we obtain*

$$P_{11}^{S}(r) = \frac{1}{1 + \left(\frac{\frac{C_6}{r^6} - \delta}{\Omega}\right)^2},$$

and the corresponding blockade radius is

$$r_b^S = \left(\frac{C_6}{\delta + \Omega}\right)^{1/6}.$$

- ***Global regime (Section 3.3.4).*** *We next consider the case where both atoms start in the ground state and are driven by a common amplitude Ω. By combining diagonalization with a reduction to the sequential regime, we obtain*

$$P_{11}^{G}(r) = \frac{1}{1 + \left[\frac{1}{\Omega} \cdot \frac{C_6}{(1.29\, r - 0.26\, r_b^S)^6}\right]^2},$$

which yields

$$r_b^G = 0.98\, r_b^S.$$

- ***Local regime (Section 3.3.5).*** *Finally, we consider the case where both atoms start in the ground state and are driven by local amplitudes Ω_0 and Ω_1. We define the effective amplitude*

$$\Omega_{\text{eff}} = \left(\frac{1}{2}\left[\Omega_0^{-1/6} + \Omega_1^{-1/6}\right]\right)^{-6}.$$

Numerically, we find

$$P_{11}^{L}(r) = \frac{1}{1 + \left[\frac{1}{\Omega_{\text{eff}}} \cdot \left(\frac{C_6}{(1.18\, r - 0.16\, r_b^S)^6}\right)\right]^2},$$

and the resulting effective blockade radius satisfies

$$r_b^L = \frac{1}{2}\left(r_b^G(\Omega_0) + r_b^G(\Omega_1)\right) \simeq 0.49\left(r_b^S(\Omega_0) + r_b^S(\Omega_1)\right).$$

These results provide a theoretical framework for characterizing the blockade radius, with the aim of improving the encoding of a broader class of graphs. As discussed above, extending the class of encodable graphs requires a more precise understanding of the Rydberg blockade mechanism. In particular, the global regime result yields an effective radius close to the sequential-regime value, supporting the consistency of the standard blockade interpretation in the homogeneous setting. In practice, the blockade effect is a soft, continuous phenomenon, sensitive not only to interatomic distance but also to the amplitude and detuning of the laser drive, especially in the presence of inhomogeneous driving conditions. As such, a more general theory is needed to describe how the Rydberg blockade operates under local driving, where each atom may experience a distinct laser amplitude.

Disk graph embedding. Building on the local blockade framework, we investigate the embedding of disk graphs, a natural generalization of unit disk graphs in which each vertex corresponds to a disk of arbitrary radius in the plane, and edges are present when two disks intersect. Disk graphs can represent more complex, heterogeneous structures and are known to encode classically challenging optimization problems [vLVL+06]. Crucially, embedding such graphs requires moving beyond the standard, idealized picture of the Rydberg blockade based on a global drive and a single fixed blockade radius. We introduce two complementary metrics to quantify the fidelity of a disk-graph encoding, given atom positions and local amplitudes. The first, a correlation matrix, captures joint excitation probabilities and reveals the effective adjacency structure induced by Rydberg interactions. The second, the maximum independence violation, provides a quantitative measure of how well the embedding respects the independent set properties of the desired graph. Together, these tools allow us to assess the fidelity of graph encodings in both global and local driving regimes.

Main Result 2. *Let $\epsilon_1, \epsilon_2 > 0$, $G = (V, E)$ be a disk graph and let $|\psi(t)\rangle$ be the quantum states associated with an encoding procedure. Let C denote the occupation correlation matrix, with entries defined as $C_{ij} = \max_t \langle\psi(t)|N_iN_j|\psi(t)\rangle$ and let $P_{\text{violation}}$ denote the probability of sampling a non-independent set. We say that the encoding of the graph G is accurate when:*

- *$P_{\text{violation}} < \epsilon_1$, and*
- *for all $(i, j) \in E$, the correlations $C_{ij} < \epsilon_2$.*

We further show that a quench corresponding to a π-pulse can efficiently prepare encodings beyond the unit-disk class: a global pulse with $\Omega = \pi$ prepares unit-disk instances, while local pulses (e.g., $\Omega_i = \pi$ on a subset of atoms and $\Omega_i = \pi/20$ on others) can realize disk-graph encodings.

Disk graph MIS. Finally, once disk graphs are embedded in a neutral-atom device, we show that local regimes with local amplitudes Ω_i can prepare MIS solutions more efficiently than a global regime. In particular, disk graphs that cannot be faithfully embedded under global driving become accessible under local addressing, and we observe substantially improved approximate optimization performance for instances outside the unit-disk class. These results extend the capabilities of neutral-atom processors to a broader class of geometrically and computationally complex optimization problems, and we also show how such instances can be treated within a quantum annealing procedure.

3.1.4 Organization of the chapter

We now give a brief overview of the remainder of the chapter. In Section 3.2, we collect preliminary facts on unit-disk and disk graphs, embedding metrics and on approximate optimization.

In Section 3.3, we introduce the local Rydberg blockade regime for two-atom systems and rederive the interaction structure for both sequential and global driving. We analyze how the resulting effective blockade radius depends on the driving configuration and detuning.

In Section 3.4, we extend this analysis to many-body systems and introduce the correlation matrix and maximum independence violation as tools to evaluate the faithfulness of graph embeddings under local blockade conditions.

In Section 3.5, we apply this framework to solving the Maximum Independent Set problem on disk graphs and show that local driving enhances approximate optimization performance relative to global approaches. We conclude by discussing the implications of our findings for scaling neutral-atom systems and tackling more computationally demanding problems in the analog regime.

3.2 Preliminaries

3.2.1 Unit-disk graphs and disk graphs

In this section, we introduce the two geometric graph classes that will be central throughout this work.

Definition 3.1 (Intersection graphs)**.** *Let S be a family of sets. A graph $G = (V, E)$ is an intersection graph of S if there exists a mapping $v \mapsto s_v \in S$ such that, for any two vertices $u, v \in V$,*

$$\{u, v\} \in E \quad \Longleftrightarrow \quad s_u \cap s_v \neq \emptyset.$$

Intersection graphs arise in concrete applications, including wireless networking [HS95] and map labeling [AvS98, WDJ+06]. Disk graphs are obtained by taking S to be a family of (closed) disks in the Euclidean plane.

Definition 3.2 (Disk graphs [DDK+15, CCJ90])**.** *A graph $G = (V, E)$ with $|V| = n$ is a disk graph if it is the intersection graph of a family of disks*

$$\mathcal{C} = \{C_1, \ldots, C_n\} \subset \mathbb{R}^2,$$

where disk C_i has radius $r_i > 0$. If all radii are equal (e.g. $r_i = 1$ for all i), then G is called a unit-disk graph.

Equivalently, if disk C_i has center $\vec{R}_i$ and radius $r_i > 0$, then

$$\{i, j\} \in E \quad \Longleftrightarrow \quad \|\vec{R}_i - \vec{R}_j\| \leq \frac{r_i + r_j}{2}.$$

Unit-disk graphs form a strict subclass of disk graphs. Moreover, every planar graph is a disk graph [TM79b]. Since MIS is NP-hard on unit-disk graphs [CCJ90], it is also NP-hard on disk graphs.

Although exact polynomial-time algorithms are therefore not expected for MIS on disk graphs (unless P = NP), approximation algorithms may still be possible. Disk graphs simultaneously generalize planar graphs and unit-disk graphs [TM79a], so one might expect a rich approximation theory for standard NP-hard problems on this class. Surprisingly, relatively few approximation schemes are known. To the best of our knowledge, PTAS-type results on disk graphs are known for Vertex Cover [EJS05, vL06], Dominating Set [GP10], and MIS [EJS05]. Maximum Clique admits approximation algorithms [Fei04], while determining its exact complexity on disk graphs remains an active topic.

3.2.2 Correlation matrix and independence violation

We now introduce two metrics used throughout this work to quantify the quality of a geometric embedding.

Definition 3.3 (Correlation matrix)**.** *Let $n \in \mathbb{N}$, let $|\psi\rangle$ be an n-qubit state, and let $\{N_i\}_{i\in[n]}$ denote the number operators on qubits $1, \ldots, n$. The disconnected correlation matrix is the matrix $C \in \mathbb{R}^{n\times n}$ with entries*

$$C_{ij} = \langle\psi| N_i N_j |\psi\rangle .$$

The connected correlation matrix is the matrix with entries

$$C_{ij} = \langle\psi| N_i N_j |\psi\rangle - \langle\psi| N_i |\psi\rangle \langle\psi| N_j |\psi\rangle .$$

For a time-dependent state $|\psi(t)\rangle$, we define the disconnected correlation

$$C_{ij}^{\max} = \max_{t\in[0,T]} \langle\psi(t)| N_i N_j |\psi(t)\rangle .$$

Since we only use this maximal disconnected correlation in the remainder of the chapter, we write $C_{ij} \equiv C_{ij}^{\max}$ from now on, unless explicitly stated otherwise. As discussed previously,

for a valid graph encoding we expect C_{ij} to be small on edges $(i,j) \in E$, since the Rydberg blockade should suppress simultaneous excitations of adjacent vertices. More precisely, a small value of $\max_t \langle\psi(t)| N_i N_j |\psi(t)\rangle$ indicates that atoms i and j are unlikely to be jointly excited during the dynamics. This is the characteristic signature of the blockade mechanism: when two atoms are separated by a distance smaller than the blockade radius, the doubly excited state is energetically suppressed.

As explained in Section 2.5.2, an excited atom corresponds to selecting the associated vertex in a candidate solution. Thus, simultaneous excitations on an edge correspond to selecting two adjacent vertices, which violates the independent-set constraint. A small value of C_{ij} on edges therefore indicates that the dynamics suppresses such invalid configurations. This quantity should, however, be interpreted together with the single-site excitation probabilities, since C_{ij} can also be small simply because one or both atoms remain mostly unexcited during the evolution. Since C_{ij} only probes pairwise correlations, we also introduce a global metric to quantify the total probability of violating the independent-set constraint. We associate to each subset $S \subseteq V$ the computational basis state $|z(S)\rangle = |z_1, \ldots, z_n\rangle$ where $z_i = 1$ iff $i \in S$. We write $|S\rangle$ for this state.

Definition 3.4 (Probability of violation)**.** *Let $G = (V, E)$ and let $\mathcal{IS}(G)$ denote the set of independent sets of G. Define the projector onto non-independent sets by*

$$\Pi_{\text{non-IS}} = I - \sum_{S \in \mathcal{IS}(G)} |S\rangle \langle S| .$$

For an n-qubit state $|\psi\rangle$, the probability of violation is

$$P_{\text{violation}} = \langle\psi| \Pi_{\text{non-IS}} |\psi\rangle .$$

Similarly, for time-dependent states we define

$$P^{\max}_{\text{violation}} = \max_{t \in [0,T]} \langle\psi(t)| \Pi_{\text{non-IS}} |\psi(t)\rangle .$$

Since this is the quantity used in the remainder of the chapter, we write $P_{\text{violation}} \equiv P^{\max}_{\text{violation}}$ when no ambiguity is possible.

Intuitively, embeddings that respect the intended blockade constraints should yield small values of $P_{\text{violation}}$, and simultaneously small correlations C_{ij} on edges $(i,j) \in E$.

3.2.3 Sampling quality: exact and approximate MIS probabilities

As discussed in the introduction, for unit-disk graphs the class of instances that can be embedded on current neutral-atom devices can be significantly narrower than the full unit-disk family, and exact MIS preparation may be challenging in practice.

A standard performance metric is the time-to-solution required to reach a target success probability (e.g. 99%). If each run takes time τ and succeeds with probability P_{MIS}, the number of repetitions needed to reach 99% success is

$$R_{99} = \frac{\log(1-0.99)}{\log(1-P_{\mathrm{MIS}})},$$

and the corresponding time-to-solution is $TTS_{99} = \tau R_{99}$.

Motivated by [MBL25], it is also natural to consider approximate optimization metrics. For $\varepsilon \in (0,1)$, let

$$p_\varepsilon = \Pr\Big[\, |IS| \geq (1-\varepsilon)|MIS| \,\Big]$$

denote the probability that the sampled configuration is an independent set within a $(1-\varepsilon)$ factor of optimal. One can then define an approximate time-to-solution by replacing P_{MIS} with p_ε in the above expression:

$$TT_\epsilon = \tau \frac{\log(1-0.99)}{\log\Big(1-\Pr\Big[\, |IS| \geq (1-\varepsilon)|MIS| \,\Big]\Big)}$$

This then computes the time to be ϵ-close to the MIS. In this work, we will also use $P_{|MIS|-k}$ to denote the cumulative probability

$$P_{|MIS|-k} = \Pr\big[\, |IS| \geq |MIS| - k \,\big], \qquad k \in \{0, \ldots, |MIS|\}.$$

We compute these quantities under two amplitude schemes: a global scheme, denoted by g, in which all atoms are driven with the same amplitude, and a local scheme, denoted by l, in which the amplitude can depend on the site. In both cases, the geometry $\{\vec{R}_i\}_{i\in V}$ and the detuning are kept fixed, so that the comparison only probes the effect of using uniform or local amplitudes. We evaluate the corresponding quantities in two settings: a ground-state calculation and a finite-time state-preparation protocol.

Ground-state search (more details in Section 3.5.2). For each driving scheme $\mathrm{dr} \in \{g, l\}$, we consider the Rydberg Hamiltonian

$$H^{\mathrm{dr}} = \sum_i \frac{\Omega_i^{\mathrm{dr}}}{2} X_i - \sum_i \delta N_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j.$$

Here,

$$\Omega_i^g = \Omega \qquad \text{and} \qquad \Omega_i^l = \Omega_i,$$

correspond respectively to the global and local amplitude schemes. We compute the ground state $\left|\psi_0^{\mathrm{dr}}\right\rangle$ of H^{dr} and evaluate the probability of obtaining an independent set whose size is within k vertices of the maximum independent-set size:

$$P^{\mathrm{dr}}_{|\mathrm{MIS}|-k} = \sum_{\substack{s\in\mathcal{IS}(G)\\ |s|\geq|\mathrm{MIS}|-k}} \left|\langle s|\psi_0^{\mathrm{dr}}\rangle\right|^2.$$

State preparation (more details in Section 3.5.4). Since the ground state is not necessarily prepared in a finite-time evolution, we also simulate the annealing dynamics. Starting from $|0\dots0\rangle$, the system evolves under the time-dependent Hamiltonian

$$H^{\mathrm{dr}}(t) = \sum_i \frac{\Omega_i^{\mathrm{dr}}(t)}{2} X_i - \sum_i \delta(t) N_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j.$$

The schedules $\Omega_i^{\mathrm{dr}}(t)$ and $\delta(t)$ are those described in Section 3.5.4. They are chosen such that, at the final time T,

$$\Omega_i^{\mathrm{dr}}(T) = \Omega_{i,f}^{\mathrm{dr}}, \qquad \delta(T) = \delta_f,$$

with final parameters matching those used in the corresponding ground-state search. From the final state $\left|\psi^{\mathrm{dr}}(T)\right\rangle$, we compute the same cumulative probability

$$P_{|\mathrm{MIS}|-k}^{\mathrm{dr}} = \sum_{\substack{s\in\mathcal{IS}(G)\\ |s|\geq|\mathrm{MIS}|-k}} \left|\langle s|\psi^{\mathrm{dr}}(T)\rangle\right|^2.$$

To compare the two schemes, we use the relative difference, defined as the difference normalized by the average:

$$\Delta_k = \frac{2\left(P_{|\mathrm{MIS}|-k}^{l} - P_{|\mathrm{MIS}|-k}^{g}\right)}{P_{|\mathrm{MIS}|-k}^{l} + P_{|\mathrm{MIS}|-k}^{g}}.$$

This metric provides a symmetric comparison between $P_{|\mathrm{MIS}|-k}^{l}$ and $P_{|\mathrm{MIS}|-k}^{g}$. It is preferable to the relative difference $(P_{|\mathrm{MIS}|-k}^{l} - P_{|\mathrm{MIS}|-k}^{g})/P_{|\mathrm{MIS}|-k}^{g}$, which can become numerically unstable when $P_{|\mathrm{MIS}|-k}^{g}$ is small. With this convention, $\Delta_k \simeq 0$ indicates comparable performance between the two schemes, while positive values indicate an advantage of the local scheme. In particular, values close to 2 correspond to the regime where $P_{|\mathrm{MIS}|-k}^{l} \gg P_{|\mathrm{MIS}|-k}^{g}$.

3.3 Description of the Rydberg blockade phenomenon

The goal of this section is to characterize the Rydberg blockade quantitatively in a more general setting. We first focus on a local blockade regime, in which atoms experience distinct local amplitudes. We then study how blockade behavior changes across different driving configurations and introduce consistent, amplitude-dependent definitions of blockade strength and blockade radius. These results provide the foundation for the many-body analysis and graph-embedding applications developed in the following sections.

3.3.1 Background

The Rydberg blockade mechanism [UJH+09, GMW+09] can be understood from the energy-level structure in the computational (Z) basis. In the van der Waals interaction regime and in

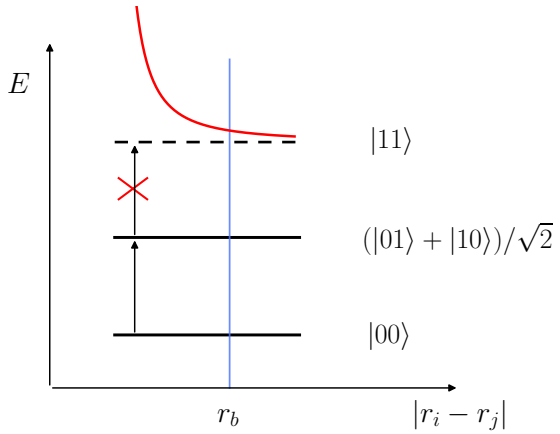


Figure 3.1: Illustration of the Rydberg blockade. Two atoms separated by a distance $R \geq r_b$ can both be excited to the Rydberg state $|1\rangle$ independently. When their distance r is lower than r_b, the van der Waals interaction shifts the doubly excited state $|11\rangle$ far out of resonance, preventing simultaneous excitation.

units where $\hbar = 1$, the Rydberg Hamiltonian (see Eq. (2.6)) can be written as

$$H(t) = \frac{1}{2}\sum_{i=1}^{n} \Omega_i(t) X_i - \sum_{i=1}^{n} \delta_i(t) N_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j,$$

where $N_i = (1 - Z_i)/2 = |1\rangle\langle 1|$ is the occupation of the Rydberg state of atom i, and X_i is the corresponding Pauli-X operator. The parameters $\Omega_i(t)$ and $\delta_i(t)$ correspond to the (possibly local) amplitude and detuning of the laser driving the $|0\rangle \leftrightarrow |1\rangle$ transition . The interaction term, characterized by the van der Waals coefficient C_6, describes the Rydberg–Rydberg interaction which decays as $1/r^6$.

For a pair of atoms, the interaction shifts the energy of the doubly excited state $|11\rangle$ by $C_6/|\vec{R}_i - \vec{R}_j|^6$. In the resonant case $\delta = 0$, when the atoms are close enough that $C_6/r^6 \gg \Omega$, the state $|11\rangle$ becomes strongly off-resonant and the state $|11\rangle$ cannot be populated. In this regime, the dynamics predominantly remains in the single-excitation manifold, and under symmetric driving can populate the entangled state $(|01\rangle + |10\rangle)/\sqrt{2}$.

This standard picture motivates a notion of blockade in terms of these two regimes.

Definition 3.5 (Blockaded pair of atoms (heuristic criterion) see Fig. 3.1)**.** *Consider two neutral atoms separated by a distance R, with global amplitude Ω and detuning δ. We say that the pair is in a blockaded regime when double excitation $|11\rangle$ is strongly suppressed, which heuristically corresponds to*

$$\Omega + \delta \ll \frac{C_6}{R^6}.$$

We say that the pair is in a non-blockaded regime when

$$\Omega + \delta \gg \frac{C_6}{R^6}.$$

This criterion is commonly summarized by introducing a blockade radius [JCZ+00, LFC+01],

$$r_b = \left(\frac{C_6}{\Omega + \delta}\right)^{1/6}. \tag{3.1}$$

This heuristic criterion captures the regime in which the interaction shift dominates the relevant driving and detuning scales. However, this sharp-radius picture does not quantify the crossover between partially blockaded and non-blockaded regimes, namely the intermediate regime where $\Omega + \delta \sim \frac{C_6}{R^6}$. To obtain a quantitative measure of blockade strength, we use the maximal population of the doubly excited state.

Definition 3.6 (Doubly excited-state probability and effective blockade radius)**.** *Consider two neutral atoms separated by a distance r, with Hamiltonian $H(t)$ and evolved state $|\psi(t)\rangle$. Define*

$$P_{11}(r) = \max_t \left|\langle 11|\psi(t)\rangle\right|^2.$$

We then define an effective blockade radius as a critical distance r_b such that

$$P_{11}(r_b) = \frac{1}{2}.$$

This definition is more useful in practice than a heuristic condition such as $r \ll (C_6/(\Omega + \delta))^{1/6}$, because the latter only gives an approximate length scale and does not quantify how small the distance must be for double excitations to be effectively suppressed. By contrast, $P_{11}(r)$ is directly accessible in numerical simulations and can be estimated on a QPU by measuring the frequency of joint excitations. It therefore provides a way to infer the blockade radius from the actual dynamics, rather than from a heuristic comparison of interaction, drive, and detuning scales.

For sufficiently long evolutions, we expect P_{11} to approach a value of 1 if the atoms are very far apart where the interactions are negligible. In the blockaded regime, $P_{11}(r)$ remains bounded away from 1. This motivates using $P_{11}(r)$ as a continuous blockade indicator, and defining r_b through the threshold $P_{11}(r_b) = 1/2$. Analytically, the maximum is taken over all times. In numerical simulations, we approximate it by a maximum over the simulated interval $t \in [0, T]$ for T to be the duration of the simulation.

We now study how $P_{11}(r)$ and the associated effective blockade radius depend on three configurations: sequential, global, and local.

3.3.2 Methods to find the evolution operator

Under the Hamiltonian $H(t)$, the evolved quantum state is $|\psi(t)\rangle = U(t)\,|\psi_0\rangle$, where $|\psi_0\rangle$ is the initial state. To quantify how blockaded a pair of atoms is, we study the maximal population of the doubly excited state $|11\rangle$ during the evolution, i.e. the maximum of $|\,\langle 11|\psi(t)\rangle\,|^2$ over time. This quantity can be obtained from the transition amplitude $\langle 11|U(t)|\psi_0\rangle$. When the

Hamiltonian is time-independent, there are two standard and equivalent ways to compute $U(t) = e^{-iHt}$. First, using the spectral decomposition $H|\psi_i\rangle = E_i|\psi_i\rangle$, one has

$$U(t) = \sum_i e^{-iE_i t}|\psi_i\rangle\langle\psi_i|. \tag{3.2}$$

Alternatively, one can diagonalize H as $H = VDV^{-1}$, where $D = \mathrm{diag}(E_i)$ is the diagonal matrix of eigenvalues. This gives

$$U(t) = Ve^{-iDt}V^{-1}, \tag{3.3}$$

where e^{-iDt} is obtained by exponentiating the diagonal entries of D.

In what follows, we will use the second representation, which is convenient for analytic manipulations.

3.3.3 Sequential regime

We first consider a sequential two-atom configuration, labeled S, with associated blockade radius r_b^S and doubly excited-state probability $P_{11}^S(r)$. In this regime, one atom is initially prepared in the Rydberg state, and we then apply a quantum evolution with constant amplitude Ω and detuning δ to drive the second atom. This setting provides the simplest way to build intuition for the blockade mechanism, since the dynamics reduces to an effective two-level system in which the interaction with the already excited atom appears as an additional fixed detuning term.

Theorem 3.1 (Sequential regime)**.** *Consider two neutral atoms separated by a distance r. Assume that one atom is excited and we then attempt to excite the second atom with amplitude Ω and detuning δ. Then*

$$P_{11}^S(r) = \frac{1}{1 + \left(\frac{\frac{C_6}{r^6} - \delta}{\Omega}\right)^2}. \tag{3.4}$$

Proof. In the sequential setting, $P_{11}^S(r)$ follows from an explicit diagonalization of an effective 2×2 Hamiltonian. Introduce the notation

$$A = -\delta, \qquad B = \Omega/2, \qquad C = \frac{C_6}{r^6} - 2\delta,$$

so that the effective Hamiltonian can be written as

$$H = \begin{pmatrix} A & B \\ B & C \end{pmatrix}. \tag{3.5}$$

Let $D = \sqrt{(A-C)^2 + 4B^2}$ and $R = A - C$. Diagonalizing $H = VDV^{-1}$ yields

$$U(t) = Ve^{-iDt}V^{-1} = \begin{pmatrix} \frac{R-D}{2B} & \frac{R+D}{2B} \\ 1 & 1 \end{pmatrix} \begin{pmatrix} e^{-i(R-D)t/2} & 0 \\ 0 & e^{-i(R+D)t/2} \end{pmatrix} \begin{pmatrix} \frac{-B}{D} & \frac{R+D}{2D} \\ \frac{B}{D} & \frac{-(R+D)}{2D} \end{pmatrix}. \tag{3.6}$$

For an initial state $|\psi_0\rangle = |\psi_{10}\rangle$, the transition amplitude to $|\psi_{11}\rangle$ is

$$U_{10\to 11} = \langle\psi_{11}|U(t)|\psi_0\rangle = \frac{B}{D}\left[e^{-i(R+D)t/2} - e^{-i(R-D)t/2}\right]. \tag{3.7}$$

Therefore the population of $|11\rangle$ is

$$P_{11}(t) = |U_{10\to 11}|^2 = \left(\frac{B}{D}\right)^2 \left[2 - 2\cos(Dt)\right], \tag{3.8}$$

whose maximum is attained when $\cos(Dt) = -1$, giving

$$P_{11}^S(r) = \max_t P_{11}(t) = \frac{4B^2}{D^2} = \frac{1}{1 + \left(\frac{\frac{C_6}{r^6} - \delta}{\Omega}\right)^2}, \tag{3.9}$$

as claimed. □

3.3.4 Global regime

We next consider the global two-atom configuration (labeled G), with associated blockade radius r_b^G and doubly excited-state probability $P_{11}^G(r)$. In this regime, both atoms start in the ground state and undergo the same quantum evolution with a common amplitude Ω and detuning δ. Compared to the sequential case, the relevant dynamics now lives in the full two-qubit Hilbert space, and computing $P_{11}^G(r)$ amounts to diagonalizing a 4×4 Hamiltonian. We perform this diagonalization using Wolfram Mathematica [Inc24]. For fixed (r, Ω), we find that $P_{11}^G(r)$ is larger than in the sequential case, which leads to a smaller effective blockade radius, as illustrated in Fig. 3.5(b).

Theorem 3.2 (Global regime, approximate expression)**.** *Consider two atoms separated by a distance r, evolving under the two-atom Rydberg Hamiltonian with common amplitude Ω and detuning δ. Assume the initial state is $|0\ldots0\rangle$. Then the doubly excited-state probability is*

$$P_{11}^G(r) = \frac{1}{1 + \left[\frac{1}{\Omega}\cdot\left(\frac{C_6}{(1.29\,r - 0.26\,r_b^S)^6}\right)\right]^2}, \tag{3.10}$$

where $r_b^S = \left(\frac{C_6}{\Omega+\delta}\right)^{1/6}$. In particular, solving $P_{11}^G(r_b^G) = \frac{1}{2}$ yields

$$r_b^G = 0.98\, r_b^S. \tag{3.11}$$

Proof. We diagonalize the Hamiltonian of the two-atom system with global drive and use the resulting eigendecomposition to compute the population of the doubly excited state. Let

$$A = -\delta, \qquad B = \Omega/2, \qquad C = \frac{C_6}{r^6} - 2\delta.$$

In the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, the Hamiltonian can be written as

$$\mathcal{H} = \begin{pmatrix} 0 & B & B & 0 \\ B & A & 0 & B \\ B & 0 & A & B \\ 0 & B & B & C \end{pmatrix}. \tag{3.12}$$

To quantify blockade, we study the maximum population of $|11\rangle$ over time, starting from an initial state $|\psi_0\rangle$. Let $\{(|\psi_i\rangle, E_i)\}_i$ denote an eigenbasis of H. We expand

$$|\psi_0\rangle = \sum_i \lambda_i |\psi_i\rangle, \tag{3.13}$$

$$|11\rangle = \sum_i \beta_i |\psi_i\rangle, \tag{3.14}$$

$$U(t) = \sum_i e^{-iE_i t} |\psi_i\rangle \langle\psi_i|. \tag{3.15}$$

Then the transition amplitude to $|11\rangle$ is

$$U_{11}(t) = \langle 11|U(t)|\psi_0\rangle = \sum_i e^{-iE_i t}\lambda_i\beta_i. \tag{3.16}$$

Hence the population of $|11\rangle$ is

$$\begin{aligned} |U_{11}(t)|^2 &= \left(\sum_i e^{-iE_i t}\lambda_i\beta_i\right)\left(\sum_j e^{-iE_j t}\lambda_j\beta_j\right)^* \\ &= \sum_i (\lambda_i\beta_i)^2 + 2\sum_{i>j}(\lambda_i\beta_i)(\lambda_j\beta_j)\cos\Big((E_i - E_j)t\Big). \end{aligned} \tag{3.17}$$

Since all entries of H are real, we can choose real eigenvectors and therefore take $\lambda_i, \beta_i \in \mathbb{R}$.

To upper-bound the maximal population, we use that $\cos(\cdot) \le 1$, which gives

$$\max_t |U_{11}(t)|^2 \;\le\; \sum_i (\lambda_i\beta_i)^2 + 2\sum_{i>j} \left|(\lambda_i\beta_i)(\lambda_j\beta_j)\right|. \tag{3.18}$$

In the cases considered here (and in particular for the parameter regimes used in Fig. 3.5), this upper bound is observed numerically to be tight, and we use it as the expression for the maximal population.

The eigenvectors of H can be written as

$$v_1 = (0, -1, 1, 0), \tag{3.19}$$

$$v_i = \left(\frac{E_i - C}{E_i}, \frac{E_i - C}{2B}, \frac{E_i - C}{2B}, 1\right), \qquad \text{for } i = 2, 3, 4. \tag{3.20}$$

For instance, if we start in $|\psi_0\rangle = |00\rangle$, only the components along v_2, v_3, v_4 contribute, and we obtain

$$\left[\max_t |U_{11}(t)|^2\right]_{|\psi_0\rangle=|00\rangle} = \sum_{i=2}^{4}\left(\frac{E_i - C}{E_i}\cdot\frac{1}{L_i^2}\right)^2 + 2\sum_{i>j\ge 2}\left|\left(\frac{E_i - C}{E_i}\cdot\frac{1}{L_i^2}\right)\left(\frac{E_j - C}{E_j}\cdot\frac{1}{L_j^2}\right)\right|, \tag{3.21}$$

where L_i denotes the Euclidean norm of v_i, i.e.

$$L_i = \sqrt{\left(\frac{E_i - C}{E_i}\right)^2 + 2\left(\frac{E_i - C}{2B}\right)^2 + 1}. \tag{3.22}$$

The eigenvalues can be written as

$$E_2 = \frac{1}{3}\left(A + C - 2\sqrt{P}\cos\left(\frac{\tan^{-1}(Q/M)}{3}\right)\right), \tag{3.23}$$

$$E_3 = \frac{1}{3}\left(A + C + \sqrt{P}\cos\left(\frac{\tan^{-1}(Q/M)}{3}\right) - \sqrt{3P}\sin\left(\frac{\tan^{-1}(Q/M)}{3}\right)\right), \tag{3.24}$$

$$E_4 = \frac{1}{3}\left(A + C + \sqrt{P}\cos\left(\frac{\tan^{-1}(Q/M)}{3}\right) + \sqrt{3P}\sin\left(\frac{\tan^{-1}(Q/M)}{3}\right)\right), \tag{3.25}$$

where

$$P = -3AC + 12B^2 + (A + C)^2, \tag{3.26}$$

$$M = 27B^2C + \frac{9}{2}(A + C)(AC - 4B^2) - (A + C)^3, \tag{3.27}$$

$$Q = \sqrt{P^3 - M^2}. \tag{3.28}$$

We observe that for $\delta = 0$ the transition region of $P_{11}^G(r)$ around the threshold value $P_{11}^G(r) = 1/2$ has a shape similar to that of the sequential curve. Motivated by this, we approximate the global curve by shifting and rescaling the sequential curve around their respective threshold points. Writing Δ^π (resp. Δ^G) for the slope at the threshold of the sequential (resp. global) curve, one obtains

$$\Delta^\pi = \left.\frac{\partial P_{11}}{\partial r}\right|_{r=r_b} = \frac{3}{r_b}, \qquad \Delta^G = \left.\frac{\partial P_{11}^G}{\partial r}\right|_{r=r_b^G} = \frac{3.864}{r_b}, \tag{3.29}$$

and the mapping

$$r \longmapsto \frac{\Delta^G}{\Delta^\pi}(r - r_b^G) + r_b \tag{3.30}$$

yields the simplified expression

$$P_{11}^G(r) = \frac{1}{1 + \left[\frac{1}{\Omega}\cdot\left(\frac{C_6}{(1.29\,r - 0.26\,r_b)^6}\right)\right]^2}. \tag{3.31}$$

Solving $P_{11}^G(r) = 1/2$ gives $r_b^G = 0.98\,r_b^S$. □

We observe that the approximation $r_b^G \approx 0.98\,r_b^S$ is in good agreement with simulation results (Fig. 3.5(c)). The small difference between r_b^G and r_b^S indicates that the global-pulse dynamics leads to an effective blockade radius very close to the sequential estimate. This supports the use of the usual heuristic blockade radius as a practical length scale in this regime, while the definition based on $P_{11}^G(r)$ provides a more operational way to extract it from the actual dynamics.

3.3.5 Local regime

Finally, we consider the local regime (labeled L, with associated blockade radius r_b^L and doubly excited-state probability P_{11}^L), in which the two atoms evolve simultaneously but with different local amplitudes. This regime captures amplitude asymmetry in a minimal two-atom setting and serves as the starting point for the many-body dynamics considered later in this work. In the numerical studies below, we choose the same average amplitude as in the previous regimes,

$$\Omega = \frac{\Omega_0 + \Omega_1}{2} \in \{1, 3\} \text{ rad}/\mu\text{s},$$

and fix the ratio $\Omega_0/\Omega_1 = 3$.

For simplicity, we restrict to the zero-detuning case. The corresponding Hamiltonian reads

$$H = \frac{1}{2}\sum_i \Omega_i X_i + \sum_{i<j} \frac{C_6}{r_{ij}^6} N_i N_j. \tag{3.32}$$

The difficulty compared with the uniform case is not the absence of a closed-form diagonalization per se, but rather the loss of permutation symmetry: site-dependent amplitudes prevent the reduction to symmetric and antisymmetric sectors, even for two atoms. As a result, unlike the sequential and global regimes, we do not derive a closed-form expression directly from exact diagonalization. Instead, we obtain an accurate approximation for $P_{11}^L(r)$ by combining numerical simulations with the same rescaling strategy used in the global regime.

Theorem 3.3 (Local regime, empirical expression)**.** *Consider two neutral atoms separated by a distance r, evolving at zero detuning ($\delta = 0$) with local amplitudes Ω_0 and Ω_1. Define the effective amplitude*

$$\Omega_{\text{eff}} = \left(\frac{1}{2}\left[\Omega_0^{-1/6} + \Omega_1^{-1/6}\right]\right)^{-6}.$$

Then, assuming the initial state is $|0\ldots0\rangle$, we have

$$P_{11}^L(r) \sim \frac{1}{1 + \left[\frac{1}{\Omega_{\text{eff}}} \cdot \left(\frac{C_6}{(1.18\, r - 0.16\, r_b^S)^6}\right)\right]^2}, \tag{3.33}$$

and the resulting effective blockade radius satisfies

$$r_b^L = \frac{1}{2}\left(r_b^G(\Omega_0) + r_b^G(\Omega_1)\right) \simeq 0.49\left(r_b^S(\Omega_0) + r_b^S(\Omega_1)\right).$$

Proof. We use the same observation of near-symmetry around the threshold $P_{11} = 1/2$ as in the global regime, and approximate $P_{11}^L(r)$ by a shifted and rescaled version of the sequential curve. Concretely, we model the local curve by the rescaling

$$r \longmapsto \left(\Delta^L/\Delta^\pi\right)\left[r - r_b^L\right] + r_b^S, \tag{3.34}$$

where r_b^L is the local blockade radius and Δ^L is the slope of $P_{11}^L(r)$ at $r = r_b^L$.

We choose r_b^L as the simplest symmetric function of (Ω_0, Ω_1) that reduces to the global blockade radius when $\Omega_0 = \Omega_1$. Specifically, we define

$$r_b^L = \frac{1}{2}\left(r_b^G(\Omega_0) + r_b^G(\Omega_1)\right), \tag{3.35}$$

where $r_b^G(\Omega)$ denotes the global blockade radius obtained in the homogeneous two-atom system driven with amplitude Ω. Using $r_b^G(\Omega) \simeq 0.98\, r_b^S(\Omega)$, Eq. (3.35) can equivalently be written as

$$r_b^L \simeq 0.49\left(r_b^S(\Omega_0) + r_b^S(\Omega_1)\right). \tag{3.36}$$

We compare Eq. (3.36) with the values of r_b^L obtained from simulations for different pairs (Ω_0, Ω_1) in Fig. 3.2. For each instance, we sample the two-atom distance r in a narrow window around the expected value of r_b^L, in order to obtain high precision in the extracted radius. Each data point, shown as a circle, is obtained by evolving the two-atom system under a time-independent Hamiltonian with constant amplitudes Ω_0 and Ω_1 and zero detuning for a duration of $15\,\mu$s. From this evolution, we compute P_{11}^L and infer r_b^L as the distance satisfying $P_{11}^L(r) = 1/2$. We then compare these numerical values with the expected behavior from Eq. (3.36). Explicitly, this gives $r_b^L = 0.49\left[\left(\frac{C_6}{\Omega_0}\right)^{1/6} + \left(\frac{C_6}{\Omega_1}\right)^{1/6}\right]$. As shown in Fig. 3.2, the simulated points follow the predicted trends for the different values of Ω_0 and ratios Ω_1/Ω_0, with deviations mainly appearing for small values of Ω_1/Ω_0. Since $P_{11}^L(r)$ is defined as a maximum over time, a finite simulation time can underestimate the true maximum, which in turn leads to a slight overestimate of the inferred blockade radius. We therefore expect Eq. (3.36) to provide a lower bound for the values obtained in finite-time simulations, consistent with Fig. 3.2. Finally, because small ratios Ω_1/Ω_0 rarely satisfy the criteria required for the rescaling approximation to be accurate, we restrict the comparison to $\Omega_1/\Omega_0 \in [0.4, 1]$.

Equation (3.36) also motivates introducing an effective amplitude for the local system as the generalized mean with exponent $-1/6$,

$$\Omega_{\text{eff}} = \left(\frac{1}{2}\left[\Omega_0^{-1/6} + \Omega_1^{-1/6}\right]\right)^{-6}. \tag{3.37}$$

We estimate Δ^L from simulation data by applying the rescaling (3.34) to each dataset and choosing Δ^L to minimize the least-squares distance between the simulated $P_{11}(r)$ curve and the rescaled model curve. This procedure is repeated over all combinations of (Ω_0, Ω_1). The resulting values of Δ^L are shown as a function of $1/r_b(\Omega = \Omega_{\text{eff}})$ in Fig. 3.3, and are well fitted by

$$\Delta^L = \left.\frac{\partial P_{11}^L}{\partial r}\right|_{r=r_b^L} = \frac{3.475}{r_b^\pi(\Omega = \Omega_{\text{eff}})}. \tag{3.38}$$

For a small number of simulations, the estimated curve $P_{11}^L(r)$ exhibits local non-monotonic deviations, despite the expected monotonic dependence of the double-excitation probability on the interatomic distance. We interpret these deviations as numerical fluctuations and account

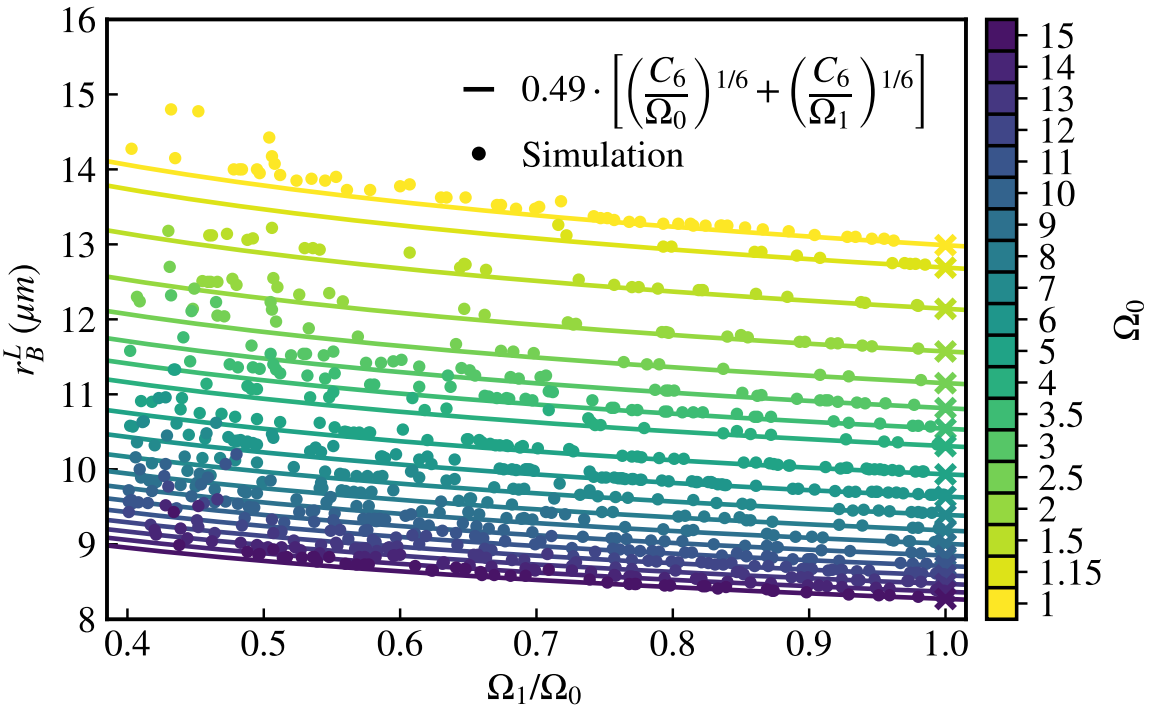


Figure 3.2: Blockade radius for locally driven two-atom systems with amplitudes Ω_0 and Ω_1. Circles show simulation results obtained by evolving the two-atom system under a time-independent Hamiltonian with constant amplitudes Ω_0 and Ω_1 and zero detuning for $15\,\mu$s. The blockade radius r_b^L is inferred from the condition $P_{11}^L(r) = 1/2$, where $P_{11}^L(r)$ is computed as the maximum double-excitation probability during the evolution. Crosses mark the case $\Omega_0 = \Omega_1$, and solid lines show the prediction of Eq. (3.36).

for them by assigning a larger uncertainty to the corresponding simulations. These uncertainties are then included in the fitting procedure of Fig. 3.3, so that simulations with stronger non-monotonic deviations contribute less to the fit. More precisely, for each simulation we define a non-monotonicity score F_i by measuring the variations of $P_{11}^L(r)$ within the regions where the data violate monotonicity. This score is used to weight the linear fit: data points with larger F_i are assigned smaller weights. In practice, the fit is obtained by minimizing a weighted residual of the form $\sum_i w_i d_i^2$, where d_i is the distance between data point i and the fitted line, and w_i is a decreasing function of the non-monotonicity score F_i.

Combining the rescaling model with the effective amplitude Ω_{eff} yields

$$P_{11}^L(r) = \frac{1}{1 + \left[\frac{1}{\Omega_{\text{eff}}} \cdot \left(\frac{C_6}{(1.18\,r - 0.16\,r_b^S)^6}\right)\right]^2}, \tag{3.39}$$

which is plotted in Fig. 3.4(a) for five representative simulations and shows good agreement. For all simulations, the residuals between Eq. (3.39) and the simulated curves are shown in

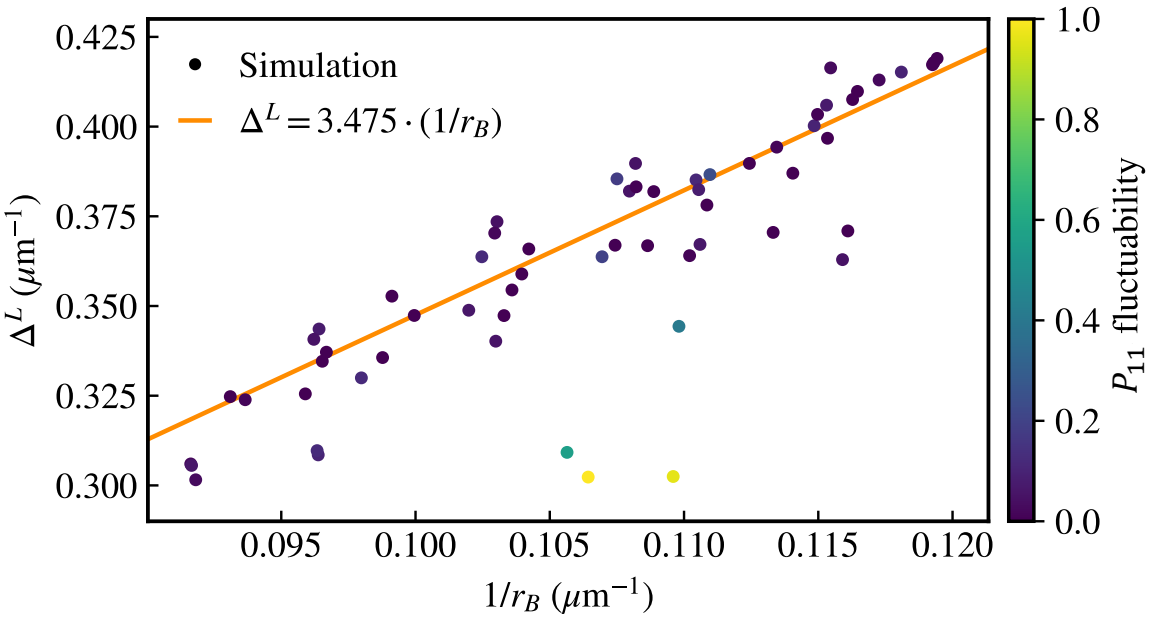


Figure 3.3: Gradient for locally driven systems obtained by minimizing the distance from the sequential model to simulated data shifted and rescaled according to $r \to (\Delta^L/\Delta^\pi) \cdot [r - r_b^L] + r_b^S$. Points with larger non-monotonicity scores are shown in lighter colors.

Fig. 3.4(b), where each residual is computed as the average absolute difference between the simulated doubly excited-state probability and the predicted one over the sampled distances. None exceeds 2%.

The simulated and predicted values of r_b for all regimes are shown in Fig. 3.5(c). As in the global scheme, the local curve $P_{11}^L(r)$ exhibits a dependence on r similar to that of the sequential curve, and the simulations follow the predicted trend. The local-drive simulations, shown as pink triangles, deviate slightly from the pink solid lines, consistently with the interpretation that the predicted curves provide lower bounds on the blockade radius. Overall, this suggests that our estimate of the blockade radius in the local-amplitude setting captures the effective blockade scale reasonably well. □

3.3.6 Comparison between different regimes

We now summarize and compare the three two-atom configurations introduced above through their corresponding doubly excited-state probabilities and effective blockade radii.

Theorem 3.4 (Doubly excited-state probability)**.** *Let $P_{11}(r)$ denote the maximal probability of occupying the doubly excited state $|11\rangle$ for two atoms separated by a distance r.*

- ***Sequential regime.*** *Assume that one atom is excited first, and we then apply a quantum*

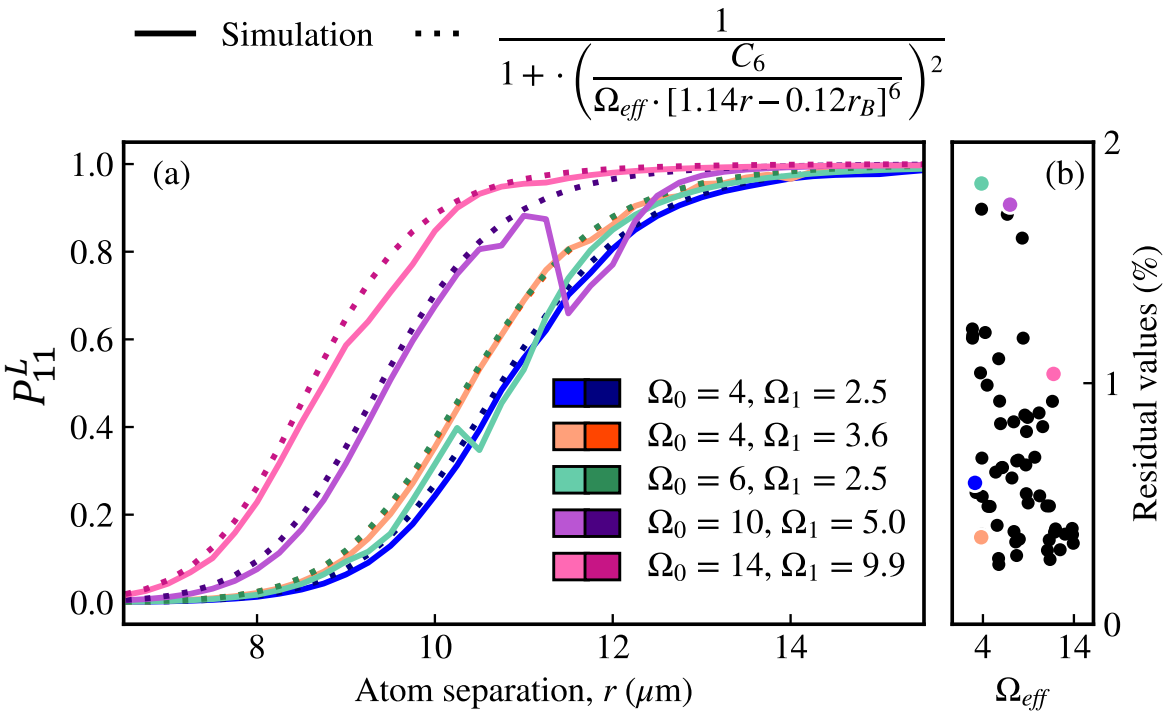


Figure 3.4: (a) Maximal occupation of the doubly excited state for five simulations of locally driven systems, compared with Eq. (3.39). (b) Residuals for all simulations. Each residual is computed as the average absolute difference between the simulated doubly excited-state probability and the predicted one over the sampled distances; the five instances shown in (a) are highlighted in color.

evolution that aims to excite the second atom as well. In this setting,

$$P_{11}^S(r) = \frac{1}{1+\left(\frac{\frac{C_6}{r^6}-\delta}{\Omega}\right)^2},$$

and the corresponding blockade radius is $r_b^S = \left(\frac{C_6}{\delta+\Omega}\right)^{1/6}$.

- ***Global regime.*** *Assume that both atoms start in the ground state. We then apply a global quantum evolution that aims to excite both atoms using a common Rabi frequency* Ω*. In this setting,*

$$P_{11}^G(r) = \frac{1}{1+\left(\frac{1}{\Omega}\cdot\frac{C_6}{(1.29\,r-0.26\,r_b^S)^6}\right)^2},$$

which yields $r_b^G = 0.98\,r_b^S$.

- ***Local regime.*** *Assume that both atoms start in the ground state. We then apply a quantum evolution that aims to excite both atoms, but with local amplitudes* Ω_0 *and* Ω_1*. Define the*

Regime	Initial state	Drive	Effective radius
Sequential	$\|10\rangle$	$\Omega_0 = 0, \Omega_1 = \Omega$	$r_b^S = (C_6/(\Omega + \delta))^{1/6}$
Global	$\|00\rangle$	$\Omega_0 = \Omega_1 = \Omega$	$r_b^G \simeq 0.98 r_b^S$
Local	$\|00\rangle$	$\Omega_0 \neq \Omega_1$	$r_b^L = \frac{1}{2}(r_b^G(\Omega_0) + r_b^G(\Omega_1)) \simeq 0.49(r_b^S(\Omega_0) + r_b^S(\Omega_1))$

Table 3.1: Summary of the two-atom blockade regimes.

effective amplitude

$$\Omega_{\text{eff}} = \left(\frac{1}{2}\left[\Omega_0^{-1/6} + \Omega_1^{-1/6}\right]\right)^{-6}.$$

Then

$$P_{11}^L(r) = \frac{1}{1 + \left[\frac{1}{\Omega_{\text{eff}}} \cdot \left(\frac{C_6}{(1.18\, r - 0.16\, r_b^S)^6}\right)\right]^2},$$

and the resulting effective blockade radius satisfies

$$r_b^L = \frac{1}{2}\left(r_b^G(\Omega_0) + r_b^G(\Omega_1)\right) \simeq 0.49\left(r_b^S(\Omega_0) + r_b^S(\Omega_1)\right).$$

Figure 3.5(a) illustrates the three configurations, and Fig. 3.5(b) compares the corresponding doubly excited-state probabilities as functions of the separation $r \equiv r_{12}$ for several amplitudes. Taken together, these results provide two outcomes. First, they give a quantitative way to compare blockade strength across configurations, which will be useful for graph-embedding applications. Second, they yield a consistent definition of the pairwise blockade radius, including in the presence of local amplitudes. In the next sections, we investigate how these two-atom results extend to graph embedding in many-body systems.

3.4 Embedding disk graphs on neutral-atom devices

Having developed a quantitative description of the Rydberg blockade in the presence of local amplitudes (Sec. 3.3), we now apply this framework to graph embedding on a neutral-atom quantum processing unit (QPU). Our goal is to move beyond unit-disk graph embeddings and encode disk graphs (see Sec. 3.2) by combining atom positions with local amplitudes. As explained in Section 3.3.5, when two atoms are driven with local amplitudes Ω_0 and Ω_1, the resulting effective blockade radius depends on these amplitudes and can therefore differ from the standard global-amplitude blockade radius. Consequently, local amplitudes can be used to tune the effective blockade radius associated with a pair of atoms, even at fixed interatomic distance.

3.4.1 Embedding model and control knobs

We consider a register of atoms placed at positions $\{\vec{R}_i\}_{i \in V}$. The pairwise blockade condition (Theorem 3.3) depends on (Ω_i, Ω_j) and becomes non-uniform across the register. A disk graph

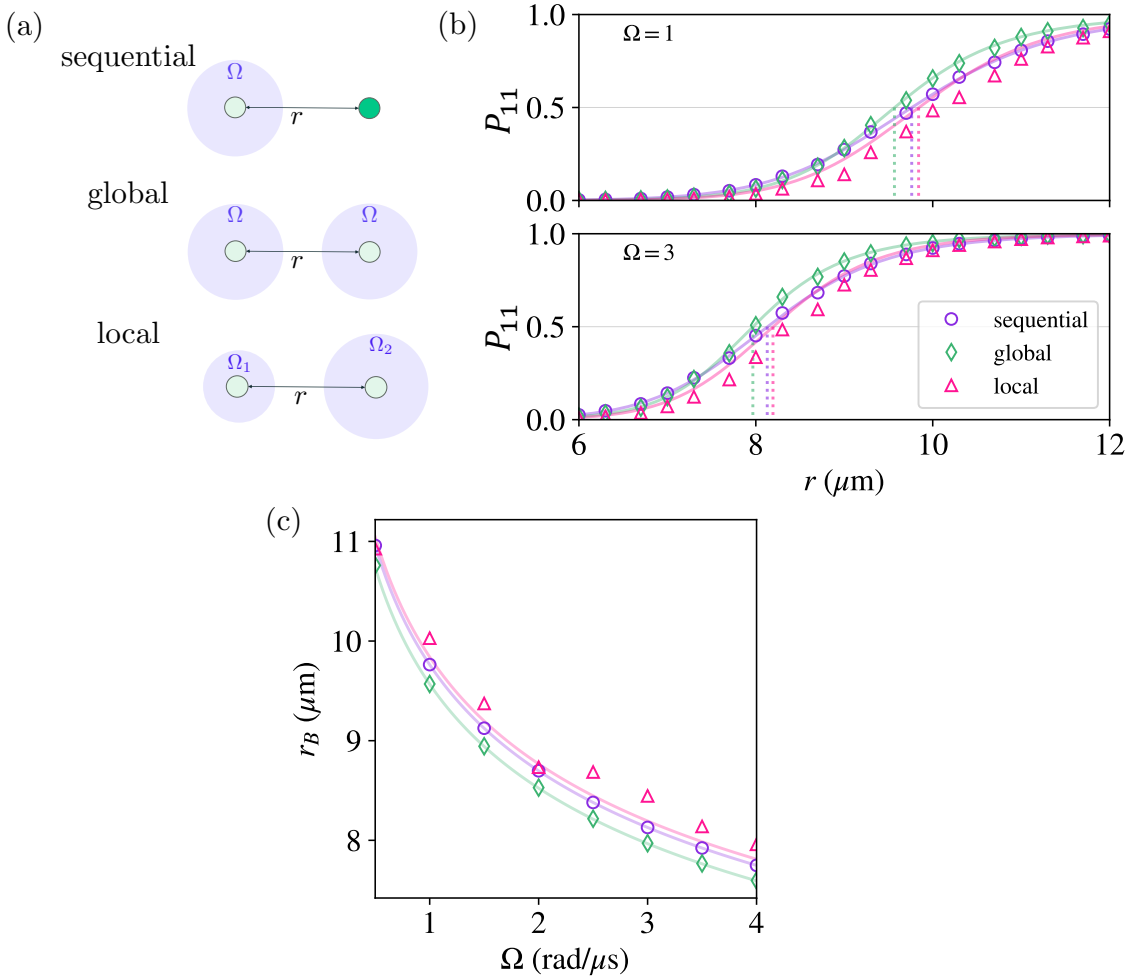


Figure 3.5: (a) The three driving scenarios are illustrated, where the light (dark) green circles correspond to atoms initialized in the ground (excited) state and driven with lasers with amplitudes Ω. (b) The maximum probability of measuring both atoms in the Rydberg state during a quench with average amplitudes 1 rad/μs (upper) and 3 rad/μs (lower) for each driving scenario. Markers indicate simulated results over 50μs, solid lines are calculated using the equations mentioned in the main text. Dotted lines show where the calculated values cross $P_{RR} = 0.5$ and hence define the blockade radius r_B. (c) The simulated (markers) and calculated (solid lines) values of r_B for a range of amplitudes. For the locally driven case the average amplitude is Ω and the ratio is $\Omega_2/\Omega_1 = 3$

with vertex positions $\vec{R}_i$ and radius r_i satisfies

$$\{i,j\} \in E \quad \Longleftrightarrow \quad \|\vec{R}_i - \vec{R}_j\| \leq \frac{r_i + r_j}{2}.$$

The local blockade model gives an effective pairwise blockade radius

$$r_{b,ij}^{L} = \frac{1}{2}\left(r_b^G(\Omega_i) + r_b^G(\Omega_j)\right) \simeq 0.49\left(r_b^S(\Omega_i) + r_b^S(\Omega_j)\right).$$

Thus, by identifying

$$r_i = r_b^G(\Omega_i)$$

or equivalently $r_i \simeq 0.98\, r_b^S(\Omega_i)$, the local blockade condition reproduces the disk-graph intersection condition. In the global setting, this recovers the usual unit-disk picture: an edge is expected when $\|\vec{R}_i - \vec{R}_j\|$ lies within the blockade regime associated with Ω. In the local setting, the constraint becomes pair-dependent through local blockade radii, enabling non-uniform connectivity that cannot be captured by a single global blockade radius and hence enabling disk-graph embeddings.

3.4.2 Quench protocol: global vs local amplitudes

To test the disk-graph embedding strategy, we perform a quench of duration $T = 100\,\mu$s using the Rydberg Hamiltonian. We compare two configurations:

- a global pulse with uniform amplitude Ω, corresponding to a global blockade radius r_b (hence giving a unit-disk embedding). This corresponds to the Hamiltonian

$$H = \frac{\Omega}{2}\sum_{i\in V} X_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j$$

- a local-amplitude scheme with amplitudes $\{\Omega_i\}_i$, corresponding to local blockade radii $r_{b,i}$(hence giving a disk-graph embedding). This corresponds to the Hamiltonian

$$H = \sum_{i\in V} \frac{\Omega_i}{2} X_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j$$

The resulting states $|\psi(T)\rangle$ encode graph structure through the blockade behavior described in Section 3.3. We evaluate this encoding using the metrics introduced below and illustrated in Fig. 3.6.

3.4.3 Evaluation metrics

We evaluate embedding quality using the two metrics introduced in Sec. 3.2: the maximal correlation matrix

$$C_{ij} = \max_t \langle\psi(t)| N_i N_j |\psi(t)\rangle,$$

and the maximal independent-set violation probability

$$P_{\text{violation}} = \max_t \langle\psi(t)| \Pi_{\text{non-IS}} |\psi(t)\rangle.$$

Definition 3.7 (Embedding-fidelity criterion)**.** *Let $\epsilon_1, \epsilon_2 > 0$, let $G = (V, E)$ be a target disk graph, and let $|\psi(t)\rangle$ be the state produced by an embedding procedure. We say that the embedding represents G faithfully if:*

- $P_{\text{violation}} < \epsilon_1$*, and*
- *for all edges $(i, j) \in E$, the correlations satisfy $C_{ij} = \max_t \langle\psi(t)|N_i N_j|\psi(t)\rangle < \epsilon_2$.*

This criterion is motivated by the blockade interpretation of graph constraints. Edges $(i, j) \in E$ correspond to pairs of atoms for which simultaneous excitation should be suppressed by the Rydberg blockade, which is captured by small values of $C_{ij} = \max_t\langle\psi(t)|N_i N_j|\psi(t)\rangle$ on those pairs. In parallel, $P_{\text{violation}}$ measures the maximal weight, over the evolution, of configurations that violate the independent-set constraint, namely configurations containing at least one adjacent excited pair. Requiring both quantities to be small therefore tests the embedding both locally, through the edge correlations C_{ij}, and globally, through the total probability of violating the independent-set constraint.

Small correlations on edges are necessary but not sufficient for a faithful embedding: non-edge pairs should not be systematically suppressed by unintended blockade effects. We therefore also inspect C_{ij} on non-edges as a metric for spurious constraints. The role of the thresholds ϵ_1 and ϵ_2 is to make this criterion quantitative. In practice, the thresholds are chosen according to the graph and numerical setting under consideration. The threshold ϵ_2 is selected so as to best separate edge pairs from non-edge pairs in the correlation matrix: edges should have $C_{ij} < \epsilon_2$, while non-edges should typically have larger correlations. The threshold ϵ_1 controls the tolerated probability of violating the independent-set constraint; we choose it to be small enough to emphasize embeddings that give little weight to non-independent configurations.

3.4.4 Distance scaling and edge breaking

An encoding of a graph in a neutral-atom device is controlled by both the register geometry and the pulse parameters. More precisely, the interaction strength C_6 and the driving amplitude Ω determine a characteristic blockade length. In the global case, this blockade length is fixed for all atom pairs, while in the local case it depends on the local amplitudes. To probe the robustness of our embedding strategy, we fix the pulse, in particular the amplitudes, and rescale the atomic coordinates by introducing a scaling parameter λ, $\vec{R}_i \mapsto \lambda\vec{R}_i$. Increasing λ increases all interatomic distances while keeping the target graph fixed. Eventually, at least one intended edge no longer satisfies the corresponding blockade condition; we say that this edge is broken. We refer to this as edge breaking, and define the critical scaling λ_c as the value at which the first intended edge breaks.

The behavior of λ_c differs between unit-disk and disk-graph embeddings. In the unit-disk case, all edges are governed by a single threshold radius. As λ increases, the first broken edge at λ_c is the longest intended edge; shorter edges remain present until larger values of λ. In the

disk-graph case, edges may break in a different order because blockade radii vary across atoms, so the first broken edge is not necessarily the longest one.

3.4.5 Illustrative instance and numerical results

As shown in Fig. 3.6, we consider a star-like arrangement of atoms and compare two target graphs defined on the same positions.

We fix the global amplitude to $\Omega = \pi$ rad/μs and infer the corresponding blockade radius r_b^G from Eq. (3.11). We then construct the disk graph by assigning a larger effective blockade radius to atoms 0, 4, and 5. In practice, this is achieved by setting their local amplitudes to $\Omega_i = \pi/20$ rad/μs, which increases the corresponding effective blockade radius, denoted by $r_{b,c}$, again estimated using Eq. (3.11). The atoms are arranged with uniform separation $0.9\, r_b^G$.

We analyze three protocols:

1. a global pulse applied to the unit-disk embedding (i.e. with $\Omega = \pi$);
2. a local pulse applied to the disk-graph embedding, with $\Omega_i = \pi/20$ for atoms $0, 4, 5$ and $\Omega_i = \pi$ for all others;
3. a local drive with randomly shuffled amplitude radii, also applied to the disk graph, to act as a control case. In the randomly shuffled protocol, each atom i is independently assigned a local amplitude $\Omega_i = \pi/20$ with probability $3/7$, and $\Omega_i = \pi$ otherwise.

Figure 3.6 shows both the maximum independence violation as a function of the scaling parameter λ/λ_c for the global, local, and randomly shuffled driving protocols, and the corresponding correlation matrices for $\lambda/\lambda_c = 0.8$ and $\lambda/\lambda_c = 2$ in the former two cases. After trying several thresholds, we fix $\epsilon_1 = 0.5$ and $\epsilon_2 = 0.1$. In the correlation matrices, red squares with diagonal crosses denote edges present in the encoded graph. For $\lambda < \lambda_c$, independence is expected to be preserved. This is confirmed for both the global pulse (green diamonds) and local pulse (pink unfilled triangles), where the maximum independence violation remains below $0.2 < \epsilon_1$, as expected for a correct embedding before an edge is broken. In contrast, the randomly shuffled protocol (pink filled triangles) exhibits significantly larger maximum independence violations, around $0.8 > \epsilon_1$, indicating a breakdown of independence and a failure of the embedding. In this regime, the correlation matrix elements C_{ij} corresponding to edges in both the unit-disk and disk graphs are below ϵ_2, consistent with the Rydberg blockade phenomenon and with the embedding. Notably, the disk graph includes three additional edges, $(0, 4)$, $(0, 5)$, and $(4, 5)$, relative to the unit-disk graph, yet these also show low correlations below ϵ_2 in the presence of the local drive designed to encode them, confirming that independence is maintained even for these added connections. For larger λ/λ_c, atoms are spaced farther apart, making double excitations more likely as the conditions for edge encoding are violated. Accordingly, the maximum independence violation is above the threshold ϵ_1 across all protocols, and the structure of the correlation matrices changes, with many entries increasing beyond ϵ_2.

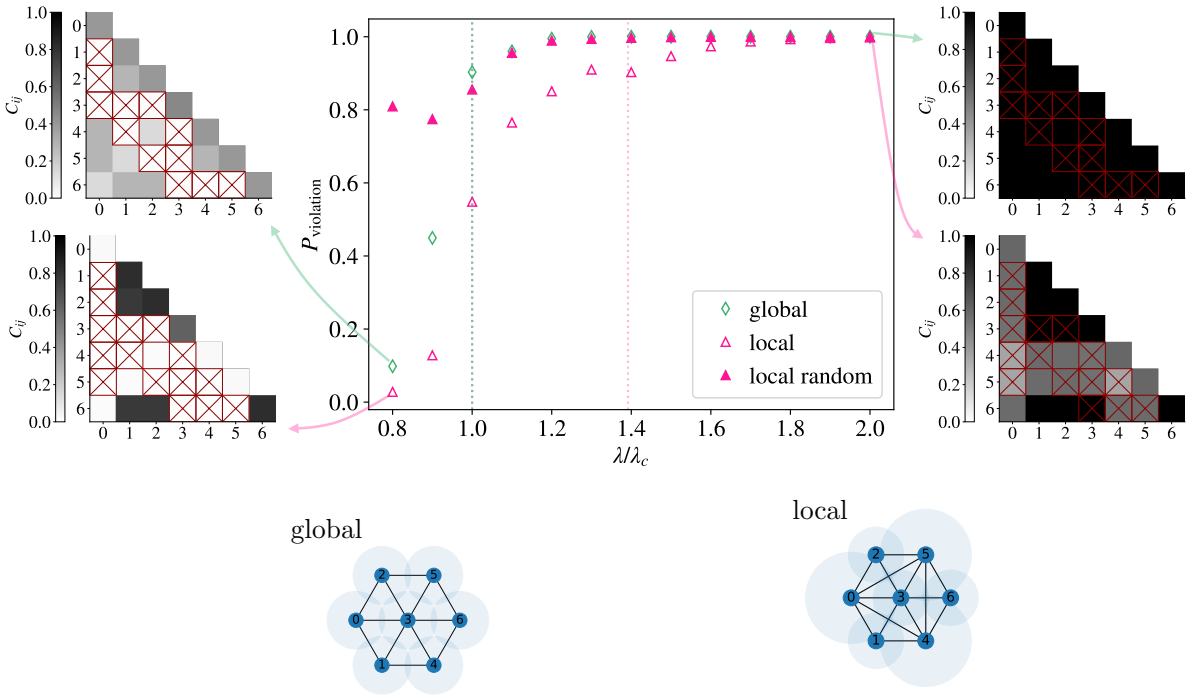


Figure 3.6: Graph embeddings for a unit-disk target (top) and a disk-graph target (bottom), evaluated using the maximal correlation matrix and the maximal independence-violation probability. Both cases use the same atomic positions; local radii induce different connectivity. A 100 μs quench is performed with the atomic register scaled by λ, and $\lambda = \lambda_c$ (black dotted line) denotes the point at which the first intended edge is expected to break. The maximal correlation matrix is shown for a connected regime ($\lambda/\lambda_c = 0.8$) and a more separated regime ($\lambda/\lambda_c = 2$); red crosses indicate edges in the target graph.

These observations demonstrate that both the maximum independence violation metric and the correlation matrix can be used to assess how faithfully a given graph has been encoded. An important limitation is visible in the locally driven case: some non-edge pairs also display low correlations. This indicates that the dynamics may impose additional effective constraints beyond the target graph. Thus, while the local drive successfully suppresses simultaneous excitations on intended edges, it does not necessarily guarantee uniform exploration of the full independent-set manifold. A more faithful embedding would suppress correlations on edges while allowing comparable excitation correlations across non-edge pairs. Future work could use the thresholds ϵ_1 and ϵ_2 to quantify this separation more systematically. For example, one could construct a correlation-based score comparing edge and non-edge pairs, such as

$$\sum_{(i,j)\notin E} C_{ij} - \sum_{(i,j)\in E} C_{ij}.$$

A large value of this score would indicate a clear separation between non-edge and edge

correlations. This could then be used to identify, for a given prepared state $|\psi(t)\rangle$, which target graphs are best represented by the dynamics. In addition, the embedding may be improved by optimizing the quantum evolution, for instance by using an annealing schedule or by varying the evolution time.

3.5 Applications to maximum independent set problems

As we previously explained, Rydberg atom platforms are a natural candidate for encoding and solving the MIS problem via analog quantum dynamics and quantum annealing [FGG+01, DC08, AL18b]. This is done by encoding the target graph for which we want to solve the MIS problem in the final Hamiltonian of the quantum evolution (see Section 3.4). In this section, we focus on the particular case of solving the MIS on disk graphs and we investigate whether a quantum annealing strategy can prepare high-quality independent sets on disk graphs.

3.5.1 Endpoint Hamiltonians with finite global versus local drives

As shown in the previous section, local amplitudes can be used to encode disk-graph constraints in neutral-atom devices. There, we studied the embedding through quench dynamics using the maximum violation probability and the correlation matrix. We now ask whether the corresponding Hamiltonian can be used as the endpoint of an annealing strategy for preparing MIS solutions.

Let $G = (V, E)$ be a disk graph with associated local radii $\{r_i\}_{i=1}^{|V|}$. We assume that we assign to each node i a position vector $\vec{R}_i$. For parameters (κ, δ_f), we define the parametrized Hamiltonian

$$H^{dr}(\kappa, \delta_f) = \kappa \sum_{i \in V} \frac{C_6}{2r_i^6} X_i - \delta_f \sum_{i \in V} N_i + \sum_{i<j} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} N_i N_j \tag{3.40}$$

where $\mathrm{dr} \in \{g, l\}$ denotes the global or local driving scheme. In the case of a local drive, we consider the local amplitudes $\Omega_i = \kappa \frac{C_6}{r_i^6}$ which we use to encode disk graphs. On the contrary, in the global-drive comparison, we replace the local amplitudes by a uniform amplitude chosen from the average radius, $r_{\mathrm{avg}} = \frac{1}{|V|} \sum_{i=1}^{|V|} r_i$, which gives a global amplitude $\Omega = \kappa \frac{C_6}{r_{\mathrm{avg}}^6}$.

Unlike the conventional MIS problem Hamiltonian, this endpoint Hamiltonian retains nonzero amplitudes, with $\Omega_i \neq 0$. Its eigenstates are therefore not computational-basis states. Consequently, MIS performance must be evaluated through measurement probabilities in the Z-basis, rather than by directly identifying the ground-state bitstring.

3.5.2 Evaluation: feasibility and sampling of exact/near-optimal solutions

For $dr \in \{l, g\}$, let $\left|\psi_0^{dr}(\kappa, \delta_f)\right\rangle$ denote the ground state of $H^{dr}(\kappa, \delta_f)$. We define the exact MIS overlap as

$$P_{\mathrm{MIS}}^{dr}(\kappa, \delta_f) = \sum_{z \in \mathcal{MIS}(G)} |\langle z|\psi_0^{dr}(\kappa, \delta_f)\rangle|^2, \tag{3.41}$$

where $\mathcal{MIS}(G)$ denotes the set of maximum independent sets of G each encoded as a state $|S\rangle$ just like previously described. We then optimize over the parameters,

$$P_{\mathrm{MIS}}^{dr} = \max_{(\kappa, \delta_f)} P_{\mathrm{MIS}}^{dr}(\kappa, \delta_f). \tag{3.42}$$

Before comparing optimization performance, we first verify that the corresponding ground states have most of their weight in the independent-set subspace. To this end, we evaluate the violation probability $P_{\mathrm{violation}}$ introduced in Sec. 3.4. This ensures that the sampled configurations have small weight outside the independent-set subspace.

To assess approximate optimization, we consider for each $k \geq 0$ the probability of sampling an independent set of size at least $|MIS| - k$, which we denote by $P_{\mathrm{MIS}-k}^{dr}$. To quantify the relative improvement provided by local addressing, we use the relative sampling enhancement

$$\Delta_k = \frac{2\left(P_{\mathrm{MIS}-k}^{l} - P_{\mathrm{MIS}-k}^{g}\right)}{P_{\mathrm{MIS}-k}^{l} + P_{\mathrm{MIS}-k}^{g}}, \tag{3.43}$$

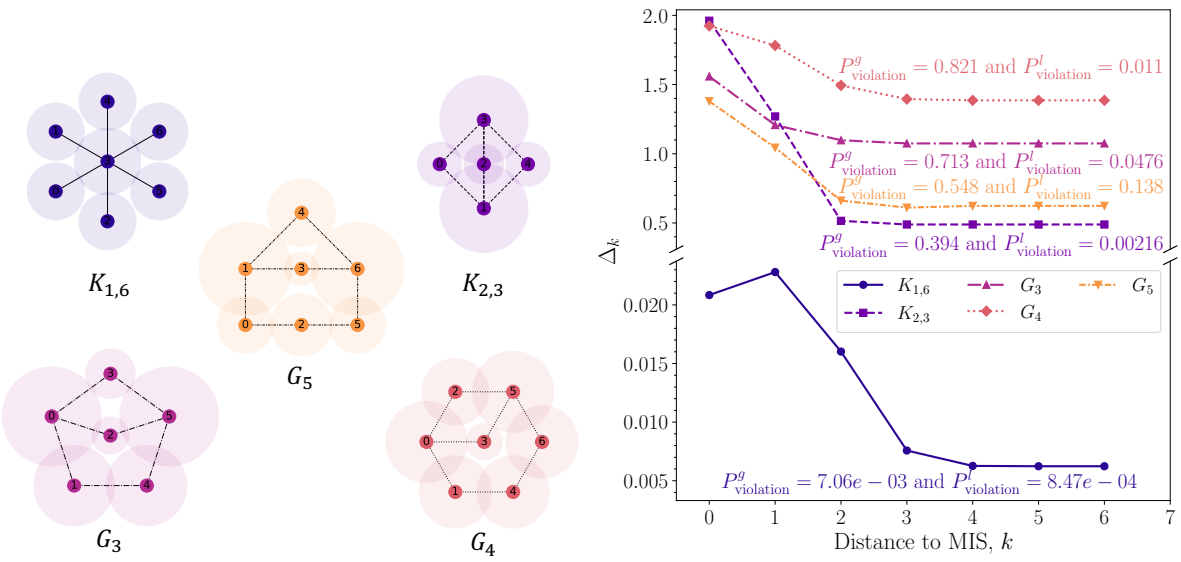
which satisfies $|\Delta_k| \leq 2$. As explained in Section 3.2.3, we want Δ_k to be as large as possible, indicating that the local scheme outperforms the global one for preparing near-optimal solutions.

3.5.3 Results on minimal non-unit-disk instances

We apply this analysis to several minimal non-unit-disk graph instances [AZ18], shown in Fig. 3.7(a). For each instance and each drive configuration, we optimize (κ, δ_f) to maximize $P_{\mathrm{MIS}}^{\mathrm{dr}}$. Figure 3.7(b) reports the resulting relative enhancement Δ_k, together with the corresponding violation probabilities $P_{\mathrm{violation}}^{g}$ and $P_{\mathrm{violation}}^{l}$. A good configuration is characterized by a large positive value of Δ_k, combined with a large ratio $P_{\mathrm{violation}}^{g}/P_{\mathrm{violation}}^{l}$.

Across these non-unit-disk instances, local driving systematically reduces $P_{\mathrm{violation}}$. For example, for the non-unit-disk graph G_4 shown in red in Fig. 3.7b(a), we find $P_{\mathrm{violation}}^{g}/P_{\mathrm{violation}}^{l} = 75$, indicating that the global scheme produces a much larger weight on non-independent configurations. As expected, this is accompanied by a relatively high positive enhancement, close to 1.5. More generally, across all non-unit-disk instances, the relative enhancements are positive, showing that the local scheme is better suited than the global one for preparing near-optimal solutions. This indicates that local addressability is beneficial both for maintaining feasibility and for approximate MIS preparation beyond the unit-disk class.

Overall, these findings suggest that retaining local drive amplitudes as part of the embedding can improve the approximate optimization landscape, allowing Rydberg systems to explore high-quality solutions on graph classes that are not naturally accessible under global driving.



(a) (b)

Figure 3.7: Relative sampling performance of local and global drives obtained from the ground state analysis. Panel (a) displays the minimal non-unit-disk graph instances considered. For each instance, the parameters (κ, δ_f) are optimized to maximize $P^{\text{dr}}_{\text{MIS}}$. Panel (b) shows the relative sampling enhancement $\Delta_k = \frac{2\left(P^l_{\text{MIS}-k} - P^g_{\text{MIS}-k}\right)}{P^l_{\text{MIS}-k} + P^g_{\text{MIS}-k}}$, which compares the probabilities of preparing an independent set of size at least $|\text{MIS}|-k$ under local and global driving. Annotated values indicate the corresponding violation probabilities $P^g_{\text{violation}}$ and $P^l_{\text{violation}}$, which quantify the probability of leaving the independent-set subspace. Note the difference in scale between the curves for $K_{1,6}$ and those for the other graphs; this choice highlights the qualitative similarity of the trends while acknowledging their differing ranges. Across these non-unit-disk graphs, local driving systematically reduces $P_{\text{violation}}$ and improves the sampling of near-optimal solutions up to $k \approx 2$. This illustrates that local addressability is important both for maintaining dynamics within the feasible independent-set subspace and for enhancing approximate MIS preparation.

3.5.4 Quantum state preparation

The previous subsection analyzed the ground state of the optimized endpoint Hamiltonian $H(\kappa_{\text{opt}}, \delta_{f,\text{opt}})$. We now ask whether this state can be approximately prepared by an explicit annealing schedule. To do so, we consider the following schedule

$$\Omega_i(t) = \kappa_{\text{opt}} \frac{C_6}{r_i^6} \frac{t}{T}, \qquad \delta(t) = \delta_0 + \frac{t}{T}(\delta_{f,\text{opt}} - \delta_0) \tag{3.44}$$

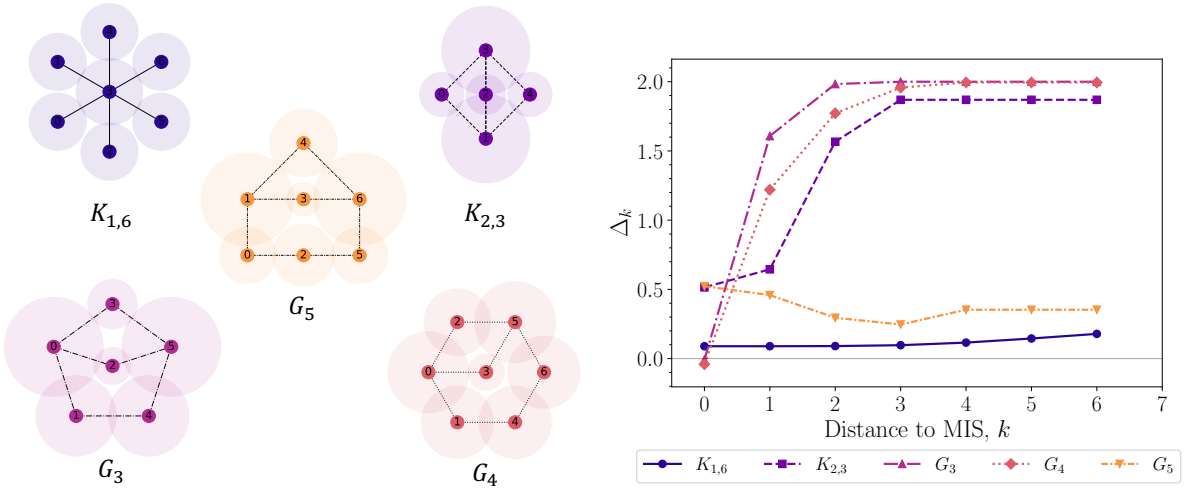
for the local case and otherwise for the global case

$$\Omega(t) = \kappa_{\text{opt}} \frac{C_6}{r_{\text{avg}}^6} \frac{t}{T}, \qquad \delta(t) = \delta_0 + \frac{t}{T}(\delta_{f,\text{opt}} - \delta_0). \tag{3.45}$$

In both cases, we choose δ_0 sufficiently negative so that $|0\dots0\rangle$ is the initial ground state. We choose T to be the maximum duration given by the hardware that we consider. In practice, as in Section 2.5.2 we can take

$$\delta_0 = -\frac{d_{\max}(G)}{d_{\max}(G)+1} \min_{\{i,j\}\in E} \frac{C_6}{|\vec{R}_i - \vec{R}_j|^6} \tag{3.46}$$

where $d_{\max}(G)$ is the maximum degree of G, and $\vec{R}_i$ is the position of vertex i in the Euclidean plane.



(a) (b)

Figure 3.8: Relative sampling performance of local and global drives obtained by quantum simulation. Panel (a) displays the minimal non-unit-disk graph instances considered. Panel (b) shows the relative sampling enhancement Δ_k. The corresponding violation probabilities are reported in Table 3.2.

Graph	$P^g_{\text{violation}}$	$P^l_{\text{violation}}$
$K_{1,6}$	0.107	0.0239
$K_{2,3}$	0.41	1.16e-03
G_3	1.22e-03	8.99e-03
G_4	0.0328	0.0715
G_5	0.432	0.0296

Table 3.2: Violation probabilities $P^g_{\text{violation}}$ and $P^l_{\text{violation}}$ for the graph instances considered.

Figure 3.8 shows the relative enhancement for the different non-unit-disk graphs, computed from the quantum simulation protocol of Eqs. (3.44)–(3.45). The finite-time annealing results differ from the endpoint ground-state predictions. This is expected: the linear schedule does not

perfectly follow the instantaneous ground state, and nonadiabatic transitions near small spectral gaps can populate excited states. This is accompanied by a larger probability of sampling non-independent configurations in some cases, as shown in Table 3.2. For instance, for G_3 and G_4, $P^l_{\text{violation}}$ is larger than $P^g_{\text{violation}}$, although both remain below 7%. Nevertheless, the local protocol still exhibits a positive relative enhancement in several cases, suggesting that the benefit of local addressing can persist beyond the endpoint ground-state analysis.

3.6 Conclusion and outlook

In this work, we revisited the Rydberg blockade phenomenon beyond the conventional model based on a global, fixed blockade radius. Motivated by the need to embed and solve combinatorial problems on more general graph families, we introduced and analyzed the local Rydberg blockade regime, in which atoms experience different drive amplitudes. We characterized this regime using a quantitative framework based on the double-excitation probability, and proposed two metrics: the correlation matrix and the maximum independence violation to assess how well a physical configuration respects graph-theoretic independence constraints. We proposed and tested a method to solve the MIS problem for this class of graphs using local amplitudes and projective measurements, and compared it favorably to a set of control cases of unit-disk graphs with the same average disk radii.

Leveraging this framework, we showed that disk graphs, generalizations of unit-disk graphs with local radii, can be deterministically embedded into Rydberg atom arrays through the use of local driving amplitudes. We applied this embedding to the MIS problem in the context of quantum annealing and demonstrated that finite local drives significantly enhance the probability of preparing high-quality approximate solutions on disk-graph instances while staying in the independent set subspace. These results indicate that locally tailored driving schemes can meaningfully extend the computational capabilities of Rydberg platforms. We note that this approach is also directly applicable to unit-disk graphs, recovering previously established results, and can in principle be used to prepare approximate MIS configurations for such graphs within the same embedding framework.

The many-body implementation of this embedding remains qualitative, and would benefit from further study in order to quantitatively assess its effectiveness in useful applications. Experimental challenges in annealing to a state with a finite drive may limit the feasibility of the specific MIS algorithm suggested, even for approximate solutions. The extent to which the projective measurement requirement reduces the efficiency also warrants further investigation.

Chapter 4

Log-concavity and tunneling: adiabatic quantum optimization for convex functions (with a spike)

This chapter is based on results from [BBA26], reproduced and adapted here with the authors' permission.

4.1 Introduction

In this chapter, we derive conditions under which quantum annealing efficiently prepares ground states of permutation-symmetric cost functions, including convex potentials. The central structural notion is discrete log-concavity: we prove that broad classes of one-dimensional Hamiltonians have log-concave ground states, and we use this property to obtain spectral-gap bounds. We then show that log-concavity also provides a robust way to analyze tunneling through perturbations. As a main application, we extend the Hamming-weight-with-a-spike analysis from linear to quadratic potentials.

4.1.1 From graph embeddings to adiabatic dynamics

In the previous chapter, we studied the encoding side of neutral-atom quantum optimization. Starting from the Rydberg blockade mechanism, we showed how going beyond the unit-disk approximation enlarges the class of graph instances that can be represented on quantum hardware, from unit-disk graphs to more general disk graphs. The quality of such an embedding was assessed through quantities such as the probability of violating the independent-set constraints.

Encoding, however, is only the first step. Once a graph or cost function has been implemented, one must still choose a quantum strategy that prepares a useful low-energy state of the final Hamiltonian. In the adiabatic setting, this question is governed by the spectral gap along the annealing path. As reviewed in Section 2.4, if the gap remains inverse-polynomially large, the adiabatic theorem gives an efficient preparation guarantee, whereas an exponentially small gap can lead to exponentially long annealing times.

It becomes particularly important to design conditions to get an inverse-polynomially large gap. A comprehensive review of these mechanisms is given in [AL18b], which also discusses techniques to widen an exponentially small gap, such as changing the initial Hamiltonian, adding a catalyst Hamiltonian, or modifying the final Hamiltonian. Here, we are instead interested in conditions on the final Hamiltonian itself that guarantee a large gap. In particular, for the MIS preparation with neutral atom devices, existing studies of unit-disk MIS, especially on King's-lattice instances, suggest that quantum annealing can exhibit speedups in some regimes, but that a clear and robust advantage over classical methods is not yet established for currently accessible benchmarks [EKC+22, CCL+23, ASM+23, CFH+25]. Moreover, in realistic finite-time and noisy experiments, the relevant output of these procedures is often not an exact maximum independent set but a high-quality, near-optimal independent set. This motivates a broader approximate-optimization viewpoint [MBL25], in which one studies not only exact ground-state preparation, but also the structure of the low-energy landscape. For MIS, this means moving from the exact task of finding a maximum independent set to the approximate task of preparing or sampling large independent sets. A natural way to formalize this sampling viewpoint is through the hard-core model.

Given a graph $G = (V, E)$, the hard-core distribution with parameter $\lambda > 0$ is

$$\pi_\lambda(I) = \frac{\lambda^{|I|}}{Z_G(\lambda)}, \qquad Z_G(\lambda) = \sum_{J \in \mathcal{IS}(G)} \lambda^{|J|}, \qquad I \in \mathcal{IS}(G),$$

where $\mathcal{IS}(G)$ denotes the set of independent sets of G. At $\lambda = 1$, the partition function counts independent sets, while large values of λ favor large independent sets. In the large-λ limit, the distribution concentrates on maximum independent sets. Classical algorithms for this model exhibit a rich complexity landscape. For a graph of maximum degree $d_{\max}$, there exist polynomial-time algorithms when $\lambda < \lambda_c := \frac{(d_{\max}-1)^{d_{\max}-1}}{(d_{\max}-2)^{d_{\max}}}$, while no polynomial-time approximation is possible when $\lambda > \lambda_c$ unless $NP = RP$ [LV97, Wei06, Sly10]. These questions are also closely related to independent-set sampling in random bipartite graphs [DFJ02]: for graphs of size $2n$ and maximum degree $d_{\max} \geq 6$, it is shown that any Markov chain algorithm updating at most n nodes at a time has mixing time $\Omega(e^{\gamma n})$ for some $\gamma > 0$. Since neutral-atom dynamics naturally produces distributions over independent sets, this suggests a possible connection between hard-core sampling and neutral-atom state preparation. Exploring whether such sampling problems can be implemented or benchmarked on neutral-atom devices is an interesting direction for future work.

The present chapter does not address the full independent-set sampling problem on embeddable graphs. Instead, we study a simplified but analytically tractable setting: cost functions on the hypercube that depend only on the Hamming weight, for which the adiabatic dynamics reduces to a one-dimensional Hamiltonian. Such potentials are not meant to be native neutral-atom MIS Hamiltonians; an arbitrary function of the Hamming weight would generally require highly non-local interactions. Their role is rather to isolate the mechanisms that determine whether quantum annealing can efficiently prepare low-energy states. This setting is nonetheless connected to the independent-set sampling motivation above: as shown in Section A.3, the independent-set density in the random bipartite graph model of [DFJ02] can be reduced, in a suitable symmetric regime, to a convex potential depending only on the Hamming weight.

In particular, this Hamming-weight setting gives a toy model in which we can analytically obtain rigorous sufficient conditions for the success of quantum annealing. The canonical example is the Hamming-weight-with-a-spike problem, which we review in detail below, and which will serve as the main testbed for our techniques.

4.1.2 Spectral gaps, tunneling, and approximate optimization

Adiabatic quantum optimization (AQO) was introduced as a method for finding the minima of complicated cost functions $V : \{0,1\}^n \to \mathbb{R}$ on the hypercube, leveraging quantum fluctuations to escape local minima that trap classical algorithms [FGS+94, KN98, FGG+01]. Its theoretical analysis reduces to understanding the spectral gap of the interpolating Hamiltonian

$$H(s) = -(1-s)\sum_{i=1}^{n} X_i + sV \tag{4.1}$$

for $s \in [0,1]$, where X_i is the Pauli-X operator acting on the i-th qubit, and V is the diagonal operator satisfying $V|x\rangle = V(x)|x\rangle$ for $x \in \{0,1\}^n$. If the gap remains at least inverse-polynomially large in n throughout the annealing path, then the adiabatic theorem guarantees that the target ground state can be prepared efficiently. However, establishing lower bounds on the gap is in general a hard problem, and the cases where it can be done rigorously are relatively few – see e.g. [FGG+01, RC02, AKR10, Rei04, HF13, GHV21, BCC+25, Lee25] for some examples.

A large part of the results concerns the restrictive yet still interesting family of cost functions that is called "symmetric", satisfying $V(x) = V(|x|)$. A canonical such function is the Hamming weight function $V(x) = |x|$. Due to the exact solvability of the adiabatic Hamiltonian for this potential, a significant body of work was devoted to the performance of AQO for perturbations of this function [FGG02, KC17, MAL16, Rei04, BvD16, Ber20]. The resulting "Hamming weight with a spike" problem (HWS) concerns potentials of the form $V(x) = |x| + h(x)$, for some perturbation or "spike" h. While such a spike can create local minima and obstruct classical methods such as simulated annealing, the aforementioned works derive precise conditions on the shape of the spike under which the gap remains open, and

hence AQO still succeeds. E.g., Reichardt [Rei04] uses a perturbative argument based on the overlap between the unperturbed ground state and the spike to establish a spectral gap bound. Subsequent works [KC17, MAL16, BvD16] extended this analysis, some of them using approximation methods such as the spin-coherent path-integral instanton approach [MAL16] and the discrete WKB method [KC17].

These analyses critically hinge on the fact that the unperturbed linear potential $V(x) = |x|$ yields an exactly solvable Hamiltonian, with explicit expressions for the eigenvalues and eigenvectors, and indeed there seems to be no work considering more general convex potentials with a spike.[1] In this work, we go beyond linear potentials and consider a larger family of potentials, including but not limited to the broadly relevant family of convex symmetric potentials. While exact expressions for the corresponding ground states seem elusive, we establish strong structural properties based on "log-concavity" that give us a useful handle on them. As our main proof-of-concept, we extend the HWS analysis from linear to quadratic potentials (for which already we have no closed-form ground state), showing that the tunneling mechanism persists even in this more challenging setting. Quadratic potentials are the textbook workhorse for convex optimization, and we see our work as opening up the study of quantum tunneling in the more general setting of discrete convex optimization with AQO.

4.1.3 Main contributions

In the following we summarize our contributions in more detail.

(Ultra-)Log-concavity for a broad class of potentials. A discrete sequence $\{p_k\}_{k\in[n]}$ is *log-concave* if for all k it satisfies

$$p_k^2 \geq p_{k+1}p_{k-1}.$$

A distribution $\{p_k\}$ is *ultra-log-concave* (of order n) if the sequence $\{p_k/\binom{n}{k}\}_{k\in[n]}$ is log-concave.

As part of our first contribution, we introduce log-concavity as a critical tool for analyzing (discrete) AQO. In the continuous setting, log-concavity is well-known to be instrumental to proving spectral gap bounds, as is witnessed e.g. by all the recent works on log-concave sampling [Che22]. In the discrete setting however, log-concavity seems to have been less explored in this context. As an example, a key result by Brascamp and Lieb [BL76] in the continuous setting established that the ground state of a Schrödinger operator with convex potential is (ultra-)log-concave. Turning to the discrete setting, we were unable to find an analog of this key result. As our first main result, we establish discrete log-concavity of the ground state of a rather large family of Hamiltonians, including Schrödinger operators on the discrete line with single-well potential, and the AQO Hamiltonian (Eq. (4.1)) with symmetric potentials that are convex (and even certain multi-well potentials).

[1] There are some exceptions in the *continuous* setting, see e.g. [LSL23] and references therein.

We prove all of our results for 1-dimensional Hamiltonians of the form $H = -A + V$, where A is band-diagonal with $A_{k,k-1} = A_{k-1,k} = a_k$ for log-concave $\{a_k\}$. This includes path Hamiltonians, but also AQO Hamiltonians (cf. Eq. (4.1)) with a symmetric potential, for which there is a well-known symmetry reduction from the hypercube to the path. We establish a condition called "relative monotonicity" (RM) of the potential V with respect to the coefficients $\{a_k\}$ that ensures log-concavity of the ground state. The RM-condition constrains ratios of consecutive values of the potential compared to the growth of a_{k+1}/a_k.

Main Result 3 (Theorem 4.1 and Theorem 4.2)**.**
• *If a 1-dimensional Hamiltonian $H = -A + V$ satisfies the RM-condition then its ground state is* log-concave*.*
• *If H is the symmetry-reduced AQO Hamiltonian with symmetric and convex V, then the ground-state is* ultra-log-concave*.*

The second bullet yields a discrete version of the continuous result by Brascamp and Lieb [BL76]. The first bullet allows us to go beyond convexity. E.g., if the a_k's are uniform (i.e., H is the Schrödinger operator on a path), then the RM-condition merely requires that the potential be single-well (i.e., $-V$ must be unimodal) rather than convex. On the other hand, in the hypercube case, the RM condition even captures certain multi-well potentials, having local minima. This demonstrates how log-concavity can already capture certain quantum tunneling effects.

Spectral gap for log-concave ground state. As a first demonstration of the utility of log-concavity, we show that log-concavity (resp. ultra-log-concavity) of the ground state implies an inverse-polynomial (resp. constant) lower bound on the spectral gap of H. More specifically, we use a discrete Hardy inequality [Mic99] to prove the following theorem.

Main Result 4 (Theorem 4.4)**.** *Let $H = -A + V$ be a 1-dimensional Schrödinger operator with log-concave ground state $\{\psi_k\}$ and spectral gap Δ. Then*

$$\Delta \in \Omega\left(\frac{1}{nS_n}\right)$$

with $S_n = \sum_k \frac{1}{a_k}$. In the case of the AQO Hamiltonian (Eq. (4.1)) with V symmetric and convex we get

$$\Delta \geq \frac{1}{4}.$$

We note that $S_n = n$ on the line, while $S_n \leq \pi$ for the symmetry-reduced hypercube.

Combined with our characterization of potentials yielding log-concave ground states, this implies polynomial lower bounds for a broad range of (convex, single-well, and even multi-well) potentials. We compare our results with existing bounds, e.g. those by Jarret and Jordan on a discrete version of the fundamental gap lemma for convex potentials [JJ14]. Our bounds either recover their bounds (up to constants), or extend them beyond convex potentials.

Log-concavity enables tunneling through spikes. As a second demonstration, we show that the perturbative spike analysis of Reichardt [Rei04] can be easily extended from binomial states to general log-concave ground states. We exploit their structural properties, such as unimodality and concentration, to establish the following result.

Main Result 5 (Theorem 4.5)**.** *Consider a 1-dimensional Hamiltonian $H = -A + V$ with spectral gap Δ and a log-concave ground state with variance σ and mean μ. Consider a non-negative perturbation h supported on an interval $[\ell, u)$ and satisfying an upper bound $h(k) \leq \bar{h}$. If either*

$$\bar{h}\frac{u-\ell}{\sigma} \in O(\Delta) \quad \text{or} \quad \frac{\max\{\ell-\mu, \mu-u\}}{\sigma} \in \Omega(\Delta),$$

then the perturbed Hamiltonian with potential $V + h$ has spectral gap $\Omega(\Delta)$.

This result relies on log-concavity to reduce the analysis of spike tunneling to estimating the mean and variance of the unperturbed ground state, and it provides explicit conditions on the spike parameters under which tunneling occurs.

A tractable case: quadratic potentials. As our final demonstration (and our initial motivation), we deploy these results to study the "quadratic Hamming weight with a spike" (quadratic HWS) problem as a key proof of concept. Quadratic HWS provides a direct extension to Reichardt's result on the (linear) HWS problem [Rei04], and it shows that tunneling happens for the quadratic case under the exact same conditions.

Specifically, we specialize to quadratic potentials $V_q(k) = a(k - k_0)^2$ for $k \in [n]$ and $k_0 \leq 0$, so that the minimizer remains at $k = 0$. While for linear potentials, the ground state is known to follow a binomial distribution $|\varphi_m\rangle$ with mean m, there seems to be no closed-form expression available for the ground state under a quadratic potential. Nonetheless, we show the somewhat surprising result that the ground state maintains a constant overlap of at least $1/2$ with a binomial state $|\varphi_m\rangle$. Combining this with the log-concavity of the ground state that we established, such overlap bound yields tight bounds on the mean and variance of the quadratic ground state, throughout the whole annealing path. Plugging these estimates into the spike bound above, we obtain our final result.

Main Result 6 (Theorem 4.7)**.** *Consider the AQO Hamiltonian (Eq. (4.1)) with symmetric potential $V(k) = a(k - k_0)^2$ for $k_0 \leq 0$, and spike perturbation h, supported on $[\ell, u) \subseteq [0, n]$ and satisfying $h(k) \leq \bar{h}$. If*

$$\bar{h}\,\frac{u-\ell}{\sqrt{\ell}} \in O(1), \tag{4.2}$$

then AQO prepares an approximation of $|\psi_{qs}(1)\rangle = |0\rangle$ from $|\psi_{qs}(0)\rangle$, the binomial state with mean at $n/2$, in time $O(n^4)$.

The spike condition Eq. (4.2) is identical to the one derived by Reichardt for the linear HWS problem [Rei04], confirming that the tunneling mechanism persists for quadratic potentials and strongly suggesting the broader applicability of AQO to convex potentials with spikes.

4.1.4 Organization of the chapter

We now give a brief overview on the remainder of the chapter.

In Section 4.2, we collect preliminary facts on adiabatic quantum optimization and discrete log-concavity. In Section 4.3 we prove our main results on log-concave ground states. In Section 4.4 we use these to establish new spectral gap bounds based on Hardy's inequality. In Section 4.5, we use log-concavity to analyze quantum tunneling through spike perturbations. Finally, in Section 4.6, we analyze the quadratic HWS problem.

4.2 Preliminaries

4.2.1 Symmetry reduction

In this section, we describe the reduction from the bitstring space $\{0,1\}^n$ of dimension 2^n to the Hamming weight subspace of dimension $n+1$ spanned by uniform superpositions over bitstrings of a given Hamming weight.

Definition 4.1. *The Hamming weight subspace is spanned by the vectors*

$$|k\rangle = \frac{1}{\sqrt{\binom{n}{k}}} \sum_{x\in\{0,1\}^n : |x|=k} |x\rangle$$

The key point is that, for a symmetric potential V satisfying $V|x\rangle = V(|x|)\,|x\rangle$, the adiabatic algorithm remains restricted to the Hamming weight basis.

Claim 1. *In the Hamming weight subspace, the mixing Hamiltonian $\sum_i X_i$ becomes the tridiagonal operator:*

$$(A_{\mathrm{hc}})_{k-1,k} = (A_{\mathrm{hc}})_{k,k-1} := a_k := \sqrt{k(n+1-k)} \tag{4.3}$$

for $0 < k \le n$.

Proof. Let $k, l \in \{0,\dots,n\}$. We have

$$\langle k| \sum_{i=1}^{n} X_i |l\rangle = \frac{1}{\sqrt{\binom{n}{k}\binom{n}{l}}} \sum_{\substack{x\in\{0,1\}^n \\ |x|=k}} \sum_{\substack{y\in\{0,1\}^n \\ |y|=l}} \sum_{i=1}^{n} \langle x| X_i |y\rangle .$$

Now $\langle x| X_i |y\rangle = 1$ if and only if x is obtained from y by flipping its i-th bit, and it is 0 otherwise. Hence the matrix element counts the number of pairs (x,y) with $|x|=k$, $|y|=l$, and Hamming distance 1, normalized by $\sqrt{\binom{n}{k}\binom{n}{l}}$.

Flipping one bit can only change the Hamming weight by ± 1, so the matrix element is zero unless $k = l \pm 1$.

Assume first that $k = l + 1$. For each bitstring y of Hamming weight l, there are $n - l$ zero bits that can be flipped to obtain a bitstring x of Hamming weight $l + 1 = k$. Therefore

$$\begin{aligned}\langle k| \sum_{i=1}^{n} X_i |l\rangle &= \frac{\binom{n}{l}(n-l)}{\sqrt{\binom{n}{k}\binom{n}{l}}} \\ &= \sqrt{\frac{\binom{n}{l}}{\binom{n}{l+1}}}(n-l) \\ &= \sqrt{(l+1)(n-l)}.\end{aligned}$$

Since $k = l + 1$, this is

$$\sqrt{k(n+1-k)}.$$

By symmetry, the same value is obtained when $l = k + 1$. Therefore, in the basis $\{|k\rangle\}_{k=0}^{n}$, the operator $\sum_i X_i$ is tridiagonal with

$$(A_{\mathrm{hc}})_{k-1,k} = (A_{\mathrm{hc}})_{k,k-1} = \sqrt{k(n+1-k)}$$

for $0 < k \leq n$. □

Since we consider symmetric potentials, we will only need symmetry-reduced Hamiltonians of the form

$$H(s) = -(1-s)A_{\mathrm{hc}} + sV.$$

Definition 4.2. *On the Hamming-weight subspace, we denote by K the diagonal operator*

$$K = \sum_{k=0}^{n} k\,|k\rangle\langle k|\,.$$

Equivalently, $K|k\rangle = k|k\rangle$.

4.2.2 Binomial state and Hamming weight Hamiltonian with a spike

In this section, we introduce the Hamming weight with a spike problem. Our goal is to adapt the analysis of Reichardt [Rei04] to a symmetry-reduced Hamiltonian with a spike perturbation.

Definition 4.3 (Hamming weight with a spike)**.** *Let $n \in \mathbb{N}$, let $0 \leq l \leq u \leq n$, and let $h > 0$. On the Hamming-weight subspace, we define the spike projector*

$$\Pi_{[l,u]} = \sum_{k=l}^{u} |k\rangle\langle k|\,.$$

The Hamming weight with a spike potential is

$$V_h = K + h\Pi_{[l,u]},$$

where $K = \sum_{k=0}^{n} k \left|k\right\rangle \left\langle k\right|$. *Equivalently, in the full bitstring space, this potential is given by*

$$V_h(z) = |z| + h\,\mathbf{1}_{\{l,\dots,u\}}(|z|), \qquad z \in \{0,1\}^n.$$

The unique minimizer of V_h *is the bitstring* 0^n.

We are interested in the symmetry-reduced Hamiltonian

$$H_h(\alpha) = -A_{\mathrm{hc}} + \alpha V_h = -A_{\mathrm{hc}} + \alpha\left(K + h\Pi_{[l,u]}\right), \qquad \alpha \geq 0.$$

with eigenvectors $\left|\psi_k(\alpha,h)\right\rangle$ for $k \in \{0,\dots,n\}$ and associated eigenvalues $E_k(\alpha,h)$. This is equivalent, up to a positive rescaling, to the usual annealing interpolation

$$\widetilde{H}_h(s) = -(1-s)A_{\mathrm{hc}} + sV_h,$$

with $\alpha = s/(1-s)$. We will therefore work directly with the family $H_h(\alpha)$.

The case $h = 0$ gives the linear Hamming-weight Hamiltonian

$$H_0(\alpha) = -A_{\mathrm{hc}} + \alpha K.$$

This Hamiltonian is the unperturbed reference point for the spike problem. Its ground state is a binomial state, which can be written explicitly.

Definition 4.4 (Hamming weight ground state)**.** *For* $\alpha \geq 0$*, define*

$$\rho_\alpha = \frac{1}{2}\left(1 - \frac{\alpha}{\sqrt{\alpha^2+4}}\right).$$

For $h = 0$*, the binomial state of parameter* ρ_α *is*

$$\left|\psi_0(\alpha,0)\right\rangle = \sum_{k=0}^{n} \sqrt{\binom{n}{k}} \rho_\alpha^{k/2} (1-\rho_\alpha)^{(n-k)/2} \left|k\right\rangle.$$

Incidentally, as discussed for instance in [Rei04], the binomial state is also the ground state of the linear Hamming weight Hamiltonian $H_0(\alpha)$. In this framework, α is the free parameter, but it is also possible to pick α such that the binomial distribution has a fixed mean $m \in [0, n/2]$. We define the binomial state as follows.

Definition 4.5 (Binomial state $\left|\varphi_m\right\rangle$)**.** *For* $m \in [0, n/2]$*, define the binomial state* $\left|\varphi_m\right\rangle$ *as*

$$\left|\varphi_m\right\rangle = \sum_{k=0}^{n} \sqrt{\binom{n}{k}} \left(\frac{m}{n}\right)^{k/2} \left(1 - \frac{m}{n}\right)^{\frac{n-k}{2}} \left|k\right\rangle.$$

This corresponds to the state $\left|\psi_0(\alpha,0)\right\rangle$ *with*

$$\alpha = \frac{n-2m}{\sqrt{m(n-m)}}. \tag{4.4}$$

The distribution $|\langle k|\varphi_m\rangle|^2$ describes a binomial law $\mathcal{B}(n, m/n)$, and so the mean and variance are

$$\langle\varphi_m|K|\varphi_m\rangle = m \quad \text{and} \quad \langle\varphi_m|(K-m)^2|\varphi_m\rangle = \frac{m(n-m)}{n},$$

where K denotes the linear potential, $K|k\rangle = k|k\rangle$.

Lemma 4.1 (Linear Hamming-weight Hamiltonian)**.** *For every $\alpha \geq 0$, the Hamiltonian*

$$H_0(\alpha) = -A_{\text{hc}} + \alpha K$$

has ground state $|\psi_0(\alpha, 0)\rangle$, ground-state energy

$$E_0(\alpha) = \frac{n}{2}\left(\alpha - \sqrt{\alpha^2+4}\right),$$

and spectral gap

$$\Delta_0(\alpha) = \sqrt{\alpha^2+4}.$$

Claim 2 (Adapted from [Rei04])**.** *Let $H_h(\alpha) = -A_{\text{hc}} + \alpha\left(K + h\Pi_{[l,u]}\right)$, $\alpha \geq 0$, with $1 \leq l \leq u \leq n$. If $h\frac{u-l}{\sqrt{l}} = o(1)$, then the spectral gap of $H_h(\alpha)$ is $\Omega(1)$. Consequently, using the adiabatic theorem stated above, the corresponding linear adiabatic interpolation has runtime $O(n^2)$.*

Proof. Let $H_0(\alpha) = -A_{\text{hc}} + \alpha K$ be the linear Hamming-weight Hamiltonian, corresponding to the case $h = 0$. We denote by $E_j(\alpha, h)$ the j-th eigenvalue of $H_h(\alpha)$, associated with eigenstate $|\psi_j(\alpha, h)\rangle$ and by $E_j(\alpha, 0)$ the j-th eigenvalue of $H_0(\alpha)$, both ordered increasingly.

By Lemma 4.1, the spectral gap of $H_0(\alpha)$ is

$$\Delta_0(\alpha) = E_1(\alpha, 0) - E_0(\alpha, 0) = \sqrt{\alpha^2+4}.$$

Moreover, its ground state is the binomial state $|\psi_0(\alpha, 0)\rangle$, so that

$$|\langle k|\psi_0(\alpha, 0)\rangle|^2 = \binom{n}{k}\rho_\alpha^k(1-\rho_\alpha)^{n-k}$$

with ρ_α as defined above.

The spike perturbation is positive semidefinite $\alpha h\Pi_{[l,u]} \geq 0$. Therefore, by Weyl's monotonicity lemma, adding a positive semidefinite perturbation increases all eigenvalues,

$$E_1(\alpha, h) \geq E_1(\alpha, 0).$$

On the other hand, using the variational principle with the unperturbed ground state $|\psi_0(\alpha, 0)\rangle$, we obtain

$$\begin{aligned} E_0(\alpha, h) &\leq \langle\psi_0(\alpha, 0)| H_h(\alpha) |\psi_0(\alpha, 0)\rangle \\ &= \langle\psi_0(\alpha, 0)| H_0(\alpha) |\psi_0(\alpha, 0)\rangle + \alpha h \langle\psi_0(\alpha, 0)| \Pi_{[l,u]} |\psi_0(\alpha, 0)\rangle \\ &= E_0(\alpha, 0) + \alpha h \sum_{k=l}^{u} |\langle k|\psi_0(\alpha, 0)\rangle|^2 . \end{aligned}$$

Using the binomial estimate from [Rei04],

$$\sum_{k=l}^{u} |\langle k|\psi_0(\alpha,0)\rangle|^2 = O\left(\frac{u-l}{\sqrt{l}}\right).$$

Therefore

$$E_0(\alpha,h) - E_0(\alpha,0) \leq \alpha h\, O\left(\frac{u-l}{\sqrt{l}}\right).$$

Combining this with Weyl monotonicity gives

$$\Delta_h(\alpha) \geq \sqrt{\alpha^2+4} - \alpha\, O\left(h\frac{u-l}{\sqrt{l}}\right).$$

Since $\alpha \leq \sqrt{\alpha^2+4}$, if $h(u-l)/\sqrt{l} = o(1)$, then

$$\Delta_h(\alpha) \geq (1-o(1))\sqrt{\alpha^2+4} = \Omega(1).$$

For the corresponding adiabatic interpolation

$$\widetilde{H}_h(s) = -(1-s)A_{\mathrm{hc}} + sV_h,$$

we have

$$\|\widetilde{H}_h(s)\| = O(n), \qquad \|\dot{\widetilde{H}}_h(s)\| = O(n),$$

under the same assumptions. Since the gap is $\Omega(1)$, the adiabatic theorem stated earlier Theorem 2.1 gives a runtime $O(n^2)$. □

4.2.3 Discrete log-concavity

A key concept in our work is that of (discrete) log-concavity. In this section we recall some properties of log-concave sequences.

Definition 4.6 (Discrete log-concavity [Joh07])**.** *A non-negative sequence $\{p_k\}_{k\in\mathbb{Z}}$ is log-concave on interval $I \subseteq \mathbb{Z}$ if*

$$p_k^2 \geq p_{k+1}p_{k-1} \tag{4.5}$$

for all $k \in I$. If $I = \mathbb{Z}$ then we simply say that $\{p_k\}_{k\in\mathbb{Z}}$ is log-concave.

Equivalently, defining the score sequence $\{r_k\}$ by $r_k = p_{k+1}/p_k$, log-concavity at k is equivalent to $\{r_k\}$ being non-increasing from $k-1$ to k. It turns out that many classical discrete distributions belong to this class: the discrete uniform, Bernoulli, binomial, Poisson, geometric, and negative binomial distributions, as well as convolutions of Bernoulli distributions with arbitrary parameters (see [JG06] and references therein).

In this list, the binomial distribution will play a privileged role throughout this work. In fact, one can base a strengthened notion of log-concavity on the binomial distribution.

Definition 4.7 (Discrete ultra-log-concavity [Lig97])**.** *Fix integer $n \geq 0$ and let $q_k = \binom{n}{k}$ for $k \in [n]$. A non-negative sequence $\{p_k\}_{k\in[n]}$ is ultra-log-concave (of order n) if*

$$\left(\frac{p_k}{q_k}\right)^2 \geq \frac{p_{k+1}}{q_{k+1}}\frac{p_{k-1}}{q_{k-1}} \qquad (0 < k < n). \tag{4.6}$$

4.2.3.1 Useful facts and bounds

We will use the following two elementary facts.

Lemma 4.2. *Fix an interval $I \subseteq \mathbb{Z}$ and a non-negative sequence $\{p_k\}_{k\in\mathbb{Z}}$.*
- *If $\{p_k\}$ is concave on I, then it is log-concave on I.*
- *If $\{p_k\}$ is log-concave on I, then so is $\{c^k p_k\}$ for any $c \geq 0$.*

Proof. For the first bullet we combine concavity ($2p_k \geq p_{k+1} + p_{k-1}$) with the AM-GM inequality ($p_{k-1} + p_{k+1} \geq 2\sqrt{p_{k-1}p_{k+1}}$) to yield

$$p_k^2 \geq \left(\frac{p_{k+1} + p_{k-1}}{2}\right)^2 \geq p_{k-1}p_{k+1}.$$

The second bullet trivially follows from the fact that the rescaling by c^k drops out in the log-concavity condition Eq. (4.5). □

Remark 1. *A consequence of the first bullet is that the hypercube coefficients $a_k = \sqrt{k(n+1-k)}$ from Eq. (4.3) are log-concave, since they are concave:*

$$a_{k+1} + a_{k-1} - 2a_k = \sqrt{(k+1)(n-k)} + \sqrt{(k-1)(n+2-k)} - 2\sqrt{k(n+1-k)} < 0$$

We continue with some crucial yet less elementary claims. A first key property of a log-concave distribution is that it is *unimodal*, meaning that there exists a mode k^* such that $p_{k-1} \leq p_k$ for $k \leq k^*$ and $p_k \geq p_{k+1}$ for $k \geq k^*$, and so in particular $\max_k p_k = p_{k^*}$.

Claim 3 ([KG71, Theorem 3])**.** *A log-concave distribution $\{p_k\}$ is unimodal.*

Moreover, the mode of a log-concave distribution can be related to its variance.

Claim 4 ([BMM22, Theorem 1.1])**.** *Let σ^2 be the variance of a log-concave distribution $\{p_k\}$. Then*

$$\frac{1}{\sqrt{1+12\sigma^2}} \leq p_{k^*} \leq \frac{2}{\sqrt{1+4\sigma^2}}.$$

Another important property of log-concave distributions is that they exhibit geometric decay away from their mode. This is captured by the following tail bound.

Claim 5 ([MP25, Lemma 2.3])**.** *If μ is the mean of a log-concave distribution $\{p_k\}$, then for all $\alpha \geq 0$ it holds that*

$$\sum_{|k-\mu|\geq\alpha} p_k \;\leq\; 2\exp\left(-\frac{\alpha}{2\Big(2\sum_k |k-\mu|\, p_k + 1\Big)}\right).$$

Remark 2. *From this bound we can derive the more convenient bound*

$$\sum_{|k-\mu|\geq\alpha} p_k \leq 2\exp\left(-\frac{\alpha}{2(2\sigma+1)}\right).$$

This follows from Jensen's inequality applied to the concave function $\sqrt{\cdot}$*:*

$$\sum_k |k-\mu|p_k = \sum_k \sqrt{(k-\mu)^2}p_k \leq \sqrt{\sum_k (k-\mu)^2 p_k} = \sigma.$$

4.2.3.2 Overlapping log-concave distributions

We will prove a final bound that relates the mean and variance of two log-concave distributions that have a large overlap.

Lemma 4.3. *Let* $\{p_k\}, \{q_k\}$ *be two log-concave distributions on* $\mathbb{Z}$ *with means* μ_p *and* μ_q *and variances* σ_p^2 *and* σ_q^2 *such that* $\left(\sum_k \sqrt{p_k q_k}\right)^2 \geq \frac{1}{2}$*. Then*

$$\mu_q = \mu_p + O(\sigma_p) + O(1), \qquad \sigma_q = \Theta(\sigma_p) + O(1).$$

Proof. We first show that for any $I \subset \mathbb{Z}$,

$$q(I) \geq \left(\frac{1}{\sqrt{2}} - \sqrt{p(I^c)}\right)^2. \tag{4.7}$$

Indeed, by Cauchy–Schwarz,

$$\sum_{k\in I}\sqrt{p_k q_k} \leq \sqrt{\sum_{k\in I} p_k}\sqrt{\sum_{k\in I} q_k} = \sqrt{p(I)\,q(I)} \leq \sqrt{q(I)},$$

and similarly $\sum_{k\notin I}\sqrt{p_k q_k} \leq \sqrt{p(I^c)\,q(I^c)} \leq \sqrt{p(I^c)}$. Since $|\langle p|q\rangle| \geq 1/\sqrt{2}$, this yields

$$\frac{1}{\sqrt{2}} \leq \sum_k \sqrt{p_k q_k} = \sum_{k\in I}\sqrt{p_k q_k} + \sum_{k\notin I}\sqrt{p_k q_k} \leq \sqrt{q(I)} + \sqrt{p(I^c)},$$

which implies Eq. (4.7).

Now let $I_p = [\mu_p - 10\sigma_p,\ \mu_p + 10\sigma_p] \cap \mathbb{Z}$. By Chebyshev's inequality, $p(I_p^c) \leq \frac{1}{100}$. By Eq. (4.7) this implies that

$$q(I_p) \geq \left(\frac{1}{\sqrt{2}} - \frac{1}{10}\right)^2 \geq \frac{9}{25}.$$

Letting $q_{\max} = \max_{k\in\mathbb{Z}} q(k)$ we then get that

$$q_{\max} \geq \frac{1}{|I_p|}\sum_{k\in I_p} q(k) = \frac{q(I_p)}{|I_p|} \geq \frac{9}{25}\frac{1}{1+20\sigma_p}.$$

By log-concavity and Claim 4 we also have

$$q_{\max} \leq \frac{2}{\sqrt{1+4\sigma_q^2}}.$$

Combining both inequalities gives $\sqrt{1+4\sigma_q^2} \leq \frac{25}{18}(1+20\sigma_p)$, and hence $\sigma_q \leq \frac{25}{36}(1+20\sigma_p)$. By symmetry (interchanging p and q) we also obtain $\sigma_p \leq \frac{25}{36}(1+20\sigma_q)$. Together these imply that $\sigma_q = \Theta(\sigma_p) + O(1)$.

Finally, define $I_q = [\mu_q - 10\sigma_q,\ \mu_q + 10\sigma_q] \cap \mathbb{Z}$. Chebyshev's inequality gives $q(I_q) \geq 99/100$. Since $q(I_p) \geq 9/25$, the intervals I_p and I_q must overlap. Hence

$$\mu_q - 10\sigma_q < \mu_p + 10\sigma_p \qquad \text{and} \qquad \mu_q + 10\sigma_q > \mu_p - 10\sigma_p,$$

which implies $\mu_p - 10(\sigma_p + \sigma_q) < \mu_q < \mu_p + 10(\sigma_p + \sigma_q)$. Using the previous bounds on the variances then yields $\mu_q = \mu_p + O(\sigma_p) + O(1)$. □

4.3 (Ultra-)Log-concavity for convex potentials (and beyond)

As our first contribution, we establish a large family of potentials V for which the corresponding ground state is discretely log-concave or even ultra-log-concave. We do so for 1-dimensional stoquastic Hamiltonians of the form

$$H = -A + V$$

where V encodes the potential on $\{0, 1, ..., n\}$, and A is defined by $\langle k-1|A|k\rangle = \langle k|A|k-1\rangle = a_k$ for a concave, positive sequence $\{a_k \in \mathbb{R}\}_{k=1}^n$. When $a_k = \sqrt{k(n+1-k)}$ this corresponds to the symmetry-reduced adjacency matrix A_{hc} of the hypercube from Section 4.2.1.

4.3.1 Prior work

In the continuous setting it was established by Brascamp and Lieb [BL76, Theorem 6.1] that the ground state of a continuous Schrödinger operator $-\partial^2 + V$ on the real line with a convex potential is log-concave. Stronger even, they showed that if $V(x) = \alpha^2 x^2 + W(x)$ for a convex W, then the ground state ψ has the form $\psi(x) = e^{-\alpha x^2}\varphi(x)$ for some log-concave function φ. Clearly, this is a continuous analogue to discrete ultra-log-concavity (Definition 4.7). Apart from being a useful structural property, this was also used later to establish strong bounds on the spectral gap of these operators [AC11, YZ86].

Transferring this picture to the discrete setting turns out to be non-trivial. Indeed, a clear bottleneck when transferring the proof in [BL76] to the discrete setting is that discrete log-concavity loses some of the nice properties of its continuous variant. E.g., in the continuous setting, multivariate log-concavity is well-understood and enjoys strong stability properties: in particular, [BL76] uses that marginalization of a log-concave distribution remains log-concave. In the discrete setting, however, defining a meaningful notion of multivariate log-concavity is more delicate, and key properties such as preservation under marginalization fail [BN14, Chapter 6].

One exception that we found is the work by Johnson [Joh07]. As explained in e.g. [GLM+23], Johnson proves a discrete version of the Brascamp-Lieb result, but only for one particular A matrix on the half-line $\mathbb{N}$ that corresponds to a particular queuing process on $\mathbb{N}$.

4.3.2 Ultra-log-concavity from convexity

As our first result in this direction, we establish a discrete version of the aforementioned result by Brascamp and Lieb [BL76]. Let $|\psi\rangle = \sum_{k=0}^{n} \psi_k |k\rangle$ denote the ground state of $H = -A + V$ with energy E_0, and define the function

$$F(k) := V(k) - E_0.$$

By Perron–Frobenius and positivity of the a_k's, we can assume $\psi_k > 0$ for all $0 \leq k \leq n$. Projecting the eigenvalue equation $H|\psi\rangle = E_0|\psi\rangle$ onto $\langle k|$ gives a key recurrence relation:

$$F(k)\psi_k = a_{k+1}\psi_{k+1} + a_k\psi_{k-1}, \qquad 0 \leq k \leq n, \tag{4.8}$$

with the boundary convention $\psi_{-1} = \psi_{n+1} = 0$. Given that the a_k's and ψ_k's are positive, this implies that also $F(k) > 0$ for all $0 \leq k \leq n$. We use this recurrence relation to prove our theorem.

Theorem 4.1 (Ultra-log-concavity for convex potentials)**.** *If $\{a_k = \sqrt{k(n+1-k)}\}$ and V is convex, then $\{\psi_k^2\}$ is ultra-log-concave.*

Proof. We will prove that $\phi_k = \psi_k/\sqrt{\binom{n}{k}}$ is log-concave, which is equivalent to proving that the ratios $r_k = \frac{\phi_k}{\phi_{k-1}}$ for $1 \leq k \leq n$ are non-increasing. By contradiction, we will suppose $r_k < r_{k+1}$ for some k.

We start by rewriting Eq. (4.8) as

$$F(k) = (n-k)r_{k+1} + k/r_k.$$

Now let $a \geq 0$ be the first index such that $r_a < r_{a+1}$, and let $n \geq b \geq a+1$ be the first index after a such that $r_b > r_{b+1}$, so we have

$$r_{a-1} \geq r_a < r_{a+1} \leq \ldots \leq r_{b-1} \leq r_b > r_{b+1},$$

where we formally set $r_0 = \infty$ and $r_{n+1} = 0$. Using this, we can lower bound $F(a) - F(a-1)$ and upper bound $F(b) - F(b-1)$ by

$$F(a) - F(a-1) > \frac{1}{r_a} - r_a, \qquad F(b) - F(b-1) < \frac{1}{r_b} - r_b.$$

Since the function $r \mapsto \frac{1}{r} - r$ is decreasing for positive r, this implies that

$$F(a) - F(a-1) > F(b) - F(b-1).$$

However, our convexity assumption on V implies convexity of F, and so this contradicts our assumption. □

4.3.3 γ-log-concavity from γ-relative monotonicity

As our second contribution in this direction, we describe a more technical framework for establishing log-concavity of the ground state. Critically, it allows us to go beyond convex potentials, showing that log-concavity can persist even in the occurrence of local minima. While somewhat technically involved, we show that the framework is tight in (and motivated by) capturing limiting cases such as the uniform and binomial ground state.

In the most general case, we will focus on proving a weighted log-concavity property of the ground-state, i.e.

$$\psi_k^2 \geq \gamma_k \, \psi_{k-1}\psi_{k+1} \tag{4.9}$$

where the weights $\{\gamma_k\}$ are defined such that:

$$\gamma_1, \ldots, \gamma_{n-1} > 0, \qquad \gamma_0 = \gamma_n = 1.$$

More precisely, we will say that $\{\psi_k\}$ is γ-log-concave (γ-LC) on an interval $I \subseteq \{0, \ldots, n\}$ if Eq. (4.9) is satisfied for all $k \in I$. As an example, ultra-log-concavity of $\{\psi_k^2\}$ (Definition 4.7) corresponds to the case where $\gamma_k = \sqrt{\frac{k+1}{k}\frac{n-k+1}{n-k}}$.

To prove γ-LC, we assume that $\{a_k\}_{k=1}^n$ is positive (while setting $a_0 = a_{n+1} = 0$) and satisfies a related log-concavity condition:

$$a_k^2 \gamma_k \gamma_{k-1} \geq a_{k-1}a_{k+1}, \qquad \forall 1 \leq k \leq n. \tag{4.10}$$

Recalling the function $F(k) = V(k) - E_0 > 0$, we also need to define the cumulative weights

$$\Gamma_k^- := \frac{a_1}{F(0)} \prod_{\ell=0}^{k-1} \gamma_\ell, \qquad \Gamma_k^+ := \frac{a_n}{F(n)} \prod_{\ell=k+1}^{n} \gamma_\ell,$$

with the convention that empty products are equal to 1. These weights are motivated by the following bound.

Lemma 4.4 (Cumulative ratio)**.**
- *If* $\{\psi_\ell\}$ *is* γ*-LC on* $\{0, \ldots, k-1\}$ *then* $\frac{\psi_{k-1}}{\psi_k} \geq \Gamma_k^-$.
- *If* $\{\psi_\ell\}$ *is* γ*-LC on* $\{k+1, \ldots, n\}$ *then* $\frac{\psi_{k+1}}{\psi_k} \geq \Gamma_k^+$.

Proof. The first bullet follows from propagating γ-LC from $\ell = k-1$ to $\ell = 1$:

$$\frac{\psi_{k-1}}{\psi_k} \geq \gamma_{k-1}\frac{\psi_{k-2}}{\psi_{k-1}} \geq \cdots \geq \frac{\psi_0}{\psi_1} \prod_{l=1}^{k-1} \gamma_l \overset{(4.8)}{=} \frac{a_1}{F(0)} \prod_{l=1}^{k-1} \gamma_l = \Gamma_k^-.$$

The second bullet follows from propagating γ-LC from $\ell = k+1$ to $\ell = n-1$:

$$\frac{\psi_{k+1}}{\psi_k} \geq \gamma_{k+1}\frac{\psi_{k+2}}{\psi_{k+1}} \geq \cdots \geq \frac{\psi_n}{\psi_{n-1}} \prod_{l=k+1}^{n-1} \gamma_l \overset{(4.8)}{=} \frac{a_n}{F(n)} \prod_{l=k+1}^{n-1} \gamma_l = \Gamma_k^+. \qquad \square$$

We can now define a condition linking the potential V to $\{a_k\}$ and $\{\gamma_k\}$ under which we will prove γ-LC of $\{\psi_k\}$.

Definition 4.8 (γ-relative monotonicity (γ-RM))**.** *The Hamiltonian $H = -A + V$ satisfies the left γ-RM condition on $\{0, \dots, k_0\}$ if, for all $1 \le k \le k_0$,*

$$a_k \gamma_k F(k) - a_{k+1} F(k-1) \le \Gamma_{k-1}^{-} \left(a_k^2 \gamma_k \gamma_{k-1} - a_{k+1} a_{k-1} \right). \tag{4.11}$$

It satisfies the right γ-RM condition on $\{k_1, \dots, n-1\}$ if, for all $k_1 \le k \le n-1$,

$$a_{k+1} \gamma_k F(k) - a_k F(k+1) \le \Gamma_{k+1}^{+} \left(a_{k+1}^2 \gamma_{k+1} \gamma_k - a_{k+2} a_k \right). \tag{4.12}$$

Under this condition, we can prove weighted log-concavity of $\{\psi_k\}$ on the interval I.

Theorem 4.2 (γ-LC from γ-RM)**.** *Assume that $\{a_k\}$ satisfies Eq. (4.10).*
• *If $H = -A + V$ satisfies the left γ-RM condition on $\{1, \dots, k_0\}$ then $\{\psi_k\}$ is γ-LC on $\{0, \dots, k_0\}$.*
• *If $H = -A + V$ satisfies the right γ-RM condition on $\{k_1, \dots, n-1\}$ then $\{\psi_k\}$ is γ-LC on $\{k_1, \dots, n\}$.*

In particular, if these conditions hold for some $k_1 = k_0 + 1$ then $\{\psi_k\}$ is γ-LC on $[n]$.

4.3.4 Proofs of γ-LC from γ-RM

The proofs of Theorem 4.2 are by strong induction.

Strong induction from the left. The boundary initialization at $k = 0$ is trivial because we set $\psi_{-1} = 0$. Now let $k \le k_0$ and assume γ-LC for all $\ell \in \{0, \dots, k-1\}$. By Lemma 4.4 we have

$$\frac{\psi_{k-1}}{\psi_k} \ge \Gamma_k^{-}. \tag{4.13}$$

To prove γ-LC at k, we start from the recurrence relation, Eq. (4.8). After multiplying by $a_{k+1} \gamma_{k-1} \psi_k$ this is

$$a_{k+1} \gamma_{k-1} F(k-1) \psi_k \psi_{k-1} = a_k a_{k+1} \gamma_{k-1} \psi_k^2 + a_{k-1} a_{k+1} \gamma_{k-1} \psi_k \psi_{k-2}.$$

By left γ-RM at k and γ-LC at $k-1$ this implies that

$$a_k \gamma_k \gamma_{k-1} F(k) \psi_k \psi_{k-1} - \gamma_{k-1} \Gamma_{k-1}^{-} \left(a_k^2 \gamma_k \gamma_{k-1} - a_{k+1} a_{k-1} \right) \psi_k \psi_{k-1} \le a_k a_{k+1} \gamma_{k-1} \psi_k^2 + a_{k-1} a_{k+1} \psi_{k-1}^2.$$

Again using Eq. (4.8) and absorbing γ_{k-1} this is equivalent to

$$a_k a_{k+1} \gamma_k \gamma_{k-1} \psi_{k-1} \psi_{k+1} + a_k^2 \gamma_k \gamma_{k-1} \psi_{k-1}^2 - \Gamma_k^{-} \left(a_k^2 \gamma_k \gamma_{k-1} - a_{k+1} a_{k-1} \right) \psi_k \psi_{k-1} \le a_k a_{k+1} \gamma_{k-1} \psi_k^2 + a_{k-1} a_{k+1} \psi_{k-1}^2.$$

Regrouping yields

$$\left(a_k^2\gamma_k\gamma_{k-1} - a_{k+1}a_{k-1}\right)\psi_k\psi_{k-1}\left(\frac{\psi_{k-1}}{\psi_k} - \Gamma_k^-\right) \le a_k a_{k+1}\gamma_{k-1}(\psi_k^2 - \gamma_k\psi_{k-1}\psi_{k+1}).$$

By our condition on $\{a_k\}$ and Eq. (4.13) we have positivity of the left hand side, and so

$$\gamma_k\psi_{k-1}\psi_{k+1} \le \psi_k^2. \qquad \square$$

Strong induction from the right. We proceed in an almost identical manner. The boundary initialization at $k = n$ is trivial because $\psi_{n+1} = 0$. Now let $k \ge k_1$ and assume γ-LC for all $\ell \in \{k+1, \dots, n\}$. By Lemma 4.4 we have

$$\frac{\psi_{k+1}}{\psi_k} \ge \Gamma_k^+. \tag{4.14}$$

To prove γ-LC at k, we start from the recurrence relation, Eq. (4.8). After multiplying by $a_k\gamma_{k+1}\psi_k$ this is

$$a_k\gamma_{k+1}F(k+1)\psi_k\psi_{k+1} = a_k a_{k+1}\gamma_{k+1}\psi_k^2 + a_k a_{k+2}\gamma_{k+1}\psi_k\psi_{k+2}$$

By right γ-RM at k and γ-LC at $k+1$ this implies that

$$a_{k+1}\gamma_k\gamma_{k+1}F(k)\psi_k\psi_{k+1} - \gamma_{k+1}\Gamma_{k+1}^+\left(a_{k+1}^2\gamma_k\gamma_{k+1} - a_k a_{k+2}\right)\psi_k\psi_{k+1} \le a_k a_{k+1}\gamma_{k+1}\psi_k^2 + a_k a_{k+2}\psi_{k+1}^2.$$

Again using Eq. (4.8) and absorbing γ_{k+1} this is equivalent to

$$a_k a_{k+1}\gamma_k\gamma_{k+1}\psi_{k-1}\psi_{k+1} + a_{k+1}^2\gamma_k\gamma_{k+1}\psi_{k+1}^2 - \Gamma_k^+\left(a_{k+1}^2\gamma_k\gamma_{k+1} - a_k a_{k+2}\right)\psi_k\psi_{k+1} \le a_k a_{k+1}\gamma_{k+1}\psi_k^2 + a_k a_{k+2}\psi_{k+1}^2.$$

Regrouping yields

$$\left(a_{k+1}^2\gamma_k\gamma_{k+1} - a_k a_{k+2}\right)\psi_k\psi_{k+1}\left(\frac{\psi_{k+1}}{\psi_k} - \Gamma_k^+\right) \le a_k a_{k+1}\gamma_{k+1}(\psi_k^2 - \gamma_k\psi_{k-1}\psi_{k+1}).$$

By our condition on $\{a_k\}$ and Eq. (4.14) we have positivity of the left hand side, and so

$$\gamma_k\psi_{k-1}\psi_{k+1} \le \psi_k^2. \qquad \square$$

4.3.5 Examples of γ-LC from γ-RM

To better grasp the full reach of this framework, we make explicit some particularly interesting cases.

4.3.5.1 Uniform γ, uniform a_k

In this case, γ-LC (cf. Eq. (4.9)) reduces to normal log-concavity, and the γ-RM conditions (cf. Definition 4.8) elegantly reduce to

$$V(k) \le V(k-1) \qquad (1 \le k \le k_0), \qquad\qquad V(k) \le V(k+1) \qquad (k_1 \le k \le n-1).$$

In particular, if V is a single-well potential (equivalently, $-V$ is unimodal) then the ground state of H is log-concave on the full interval $[n]$.

4.3.5.2 Uniform γ, log-concave a_k

For uniform $\{\gamma_k\}$, the condition Eq. (4.10) reduces to log-concavity of the $\{a_k\}$. By Remark 1 this setting includes as a particular case of interest the symmetry-reduced hypercube Hamiltonian, in which case $\{a_k = \sqrt{k(n+1-k)}\}$. We will see that such non-trivial $\{a_k\}$ can yield log-concave ground states even for potentials with multiple wells.

A simple yet important limit case is when $V(k) = a_k + a_{k+1}$ (where we set $a_0 = a_{n+1} = 0$). This makes the Hamiltonian H correspond to the Laplacian of the weighted path, the ground state of which we know to be the uniform state with energy $E_0 = 0$. In some sense this is the limit case of log-concavity, but it is also the limit case of our theorem because it saturates the γ-RM conditions with equality.

Going beyond this canonical example, we observe that the conditions are satisfied if

$$\frac{F(k)}{F(k-1)} \leq \frac{a_{k+1}}{a_k} \qquad (1 \leq k \leq k_0), \qquad\qquad \frac{F(k+1)}{F(k)} \geq \frac{a_{k+1}}{a_k} \qquad (k_1 \leq k \leq n-1), \quad (4.15)$$

and so F has to satisfy a sort of relative monotonicity condition as compared to $\{a_k\}$ (hence the term γ-RM). It is this additional slack that allows us to prove log-concavity of the ground state for numerous non-monotone potentials. The following lemma captures one such family, of which we display an instance in Fig. 4.1.

Lemma 4.5. *If $\{a_k = \sqrt{k(n+1-k)}\}$ and $\langle\varphi_{n/2}|V|\varphi_{n/2}\rangle \leq n$, the conditions*

$$\frac{V(k)}{V(k-1)} \leq \frac{a_{k+1}}{a_k} \qquad (1 \leq k \leq n/2), \qquad\qquad \frac{V(k+1)}{V(k)} \geq \frac{a_{k+1}}{a_k} \qquad (n/2 < k \leq n),$$

are sufficient for the corresponding H to satisfy Eq. (4.15) and hence to have a log-concave ground state.

Remark 3. *The state $|\varphi_{n/2}\rangle$ is the binomial state with mean at $n/2$ (see Section 4.2.2). The condition on the expected value of V at $n/2$ is implied e.g. when $\max_k V(k) \leq n$. In the extreme case, V increases from $k = 0$ to $k = n/2$. This only imposes that $V(0) \leq 2\sqrt{n}$ because $\max V(k) = V(n/2) \leq \frac{a_{n/2+1}}{a_1} V(0) \simeq \frac{\sqrt{n}}{2} V(0)$. We get the same argument on the right side, i.e. $V(n) \leq 2\sqrt{n}$.*

Proof. For $k \leq n/2$, we have $1 - \frac{a_{k+1}}{a_k} \leq 0$. Since $E_0 \leq \langle\varphi_{n/2}|H|\varphi_{n/2}\rangle = -n + \langle\varphi_{n/2}|V|\varphi_{n/2}\rangle \leq 0$, it follows that $E_0\left(1 - \frac{a_{k+1}}{a_k}\right) \geq 0$. Combined with our assumption on V this yields

$$V(k) - V(k-1)\frac{a_{k+1}}{a_k} \leq 0 \leq E_0\left(1 - \frac{a_{k+1}}{a_k}\right),$$

and this implies the first inequality in Eq. (4.15). Similarly, for $k \geq n/2$, we have $(1 - \frac{a_{k+1}}{a_k}) \geq 0$, and therefore $E_0\left(1 - \frac{a_{k+1}}{a_k}\right) \leq 0$. Combined with our assumption on V this yields

$$V(k+1) - V(k)\frac{a_{k+1}}{a_k} \geq 0 \geq E_0\left(1 - \frac{a_{k+1}}{a_k}\right),$$

and this implies the second inequality in Eq. (4.15). □

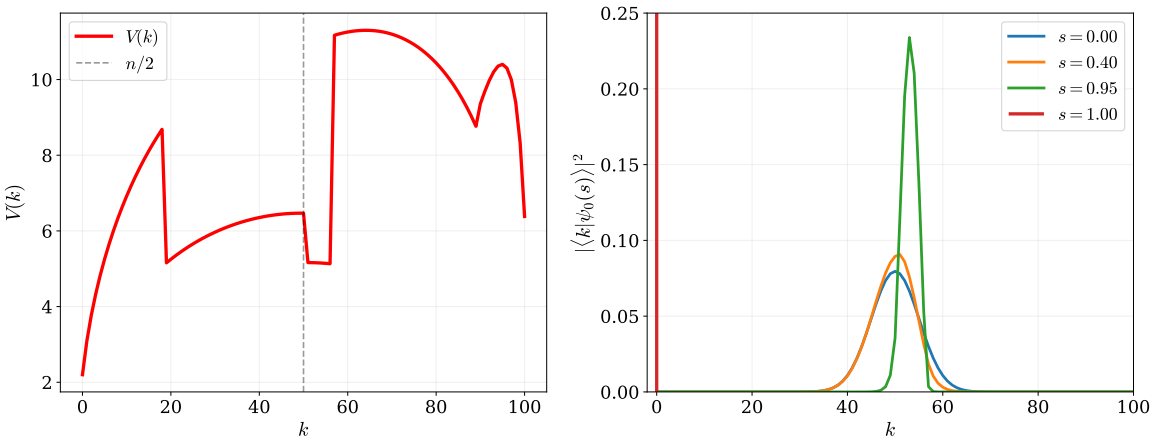


Figure 4.1: Example of a non-monotone potential satisfying the RM condition for $n = 100$ with $a_k = \sqrt{k(n+1-k)}$ and $\max V(k) \le n$. Panel (a) shows the potential, which satisfies the one-sided ratio bounds $\frac{V(k)}{V(k-1)} \le \frac{a_{k+1}}{a_k}$ for $k \le n/2$ and $\frac{V(k+1)}{V(k)} \ge \frac{a_{k+1}}{a_k}$ for $k \ge n/2+1$, where E_0 is the ground-state energy of $H = -A+V$. Panel (b) shows the shape of the log-concave ground state of $H(s) = -(1-s)A+sV$ for different values of s.

4.3.5.3 Binomial γ, binomial a_k

As a final example of our framework, we consider the case where $\left\{\gamma_k = \sqrt{\frac{(k+1)(n-k+1)}{k(n-k)}}\right\}$. In this case, γ-LC corresponds to ultra-log-concavity (ULC) of $\{\psi_k^2\}$ (Definition 4.7).

We will argue that our framework again tightly captures the limiting case, which is the binomial distribution. Indeed, in Section 4.2.2 we already saw that the squared ground state of the Hamming weight Hamiltonian (with linear potential $V(k) = \alpha k$) is a binomial distribution. In the following we show that this example satisfies our γ-RM conditions with equality.

We first observe that $\gamma_k \frac{a_k}{a_{k+1}} = 1 + \frac{1}{n-k}$ and $\gamma_k \frac{a_{k+1}}{a_k} = 1 + \frac{1}{k}$, so the $\{a_k\}$ satisfy Eq. (4.10). We then use a telescoping product to rewrite the product

$$\prod_{\ell=1}^{k-1} \gamma_\ell = \sqrt{\frac{kn}{n-k+1}}, \qquad \prod_{\ell=k+1}^{n-1} \gamma_\ell = \sqrt{\frac{n(n-k)}{k+1}}.$$

This simplifies the γ-RM conditions to

$$(n-k)\,[V(k) - V(k-1)] + V(k) - E_0 \le \frac{n^2}{V(0) - E_0} \qquad (1 \le k \le k_0),$$

$$k[V(k) - V(k+1)] + V(k) - E_0 \le \frac{n^2}{V(n) - E_0} \qquad (k_1 \le k \le n-1).$$

Now, for the linear potential $V(k) = \alpha k$ we have the exact expression $E_0 = \frac{n}{2}(\alpha - \sqrt{\alpha^2 + 4})$ (cf. Lemma 4.1), and this is easily checked to satisfy both inequalities with equality.

4.4 Spectral gap bounds from (ultra-)log-concavity

To demonstrate the power of log-concavity, we show that it implies useful lower bounds on the spectral gap Δ of H. The argument builds on a discrete Hardy inequality described by Miclo [Mic99].

4.4.1 Hardy inequality

We again consider a 1-dimensional Hamiltonian $H = -A + V$ on $\{0, \dots, n\}$ with diagonal potential V and A band-diagonal with $\langle k-1|A|k\rangle = \langle k|A|k-1\rangle = a_k$. We will assume that $\{a_k\}$ is positive, in which case there is a unique ground state $|\psi_0\rangle = \sum_{\ell=0}^{n} \psi_\ell |\ell\rangle$ that can be chosen positive. We denote by $\Delta > 0$ the spectral gap of H.

We can connect such a stoquastic Hamiltonian to a Markov chain through the so-called ground state transformation. Indeed, the operator

$$L = -\mathrm{Diag}(\psi_\ell)^{-1}(H - E_0)\mathrm{Diag}(\psi_\ell)$$

is a valid Markov chain generator, satisfying $\langle k|L|\ell\rangle \geq 0$ for $k \neq \ell$ and $L\vec{1} = 0$. More precisely, it corresponds to a birth-and-death process on $\{0, \dots, n\}$ with stationary probabilities

$$p_k = \psi_k^2$$

and ergodic flows

$$b_\ell = p_{\ell-1}L_{\ell-1,\ell} = p_\ell L_{\ell,\ell-1} = a_\ell\sqrt{p_\ell p_{\ell-1}}.$$

Following [Mic99], we associate to L a family of discrete Hardy constants: for any $0 < i < n$,

$$B_+(p,i) := \max_{i<k\leq n} p([k,n]) \sum_{\ell=i+1}^{k} \frac{1}{b_\ell}, \qquad B_-(p,i) := \max_{0\leq k<i} p([0,k]) \sum_{\ell=k+1}^{i} \frac{1}{b_\ell}.$$

The following lemma shows that these Hardy constants tightly characterize the spectral gap of H.

Lemma 4.6 (Discrete Hardy inequality)**.** *For any index $0 < i < n$,*

$$\frac{1}{4}\min\left(\frac{1}{B_+(p,i)}, \frac{1}{B_-(p,i)}\right) \leq \Delta \leq \min\left(\frac{1}{p([0,i])B_+(p,i)}, \frac{1}{p([i,n])B_-(p,i)}\right).$$

Proof. The lower bound follows from the lower bound in [Mic99, Proposition 3]. The upper bound is not stated explicitly as such, but it follows from the proof of [Mic99, Proposition 2]. Indeed, replacing the median m by i in that proof recovers the bound that we claim. □

Remark 4. *A particularly relevant case is when $i = m$ a median of p, so that $p([0, m]) \geq 1/2$ and $p([m, n]) \geq 1/2$. Then*

$$\frac{1}{4B} \leq \Delta \leq \frac{2}{B}$$

with $B = \max\{B_+(p, m), B_-(p, m)\}$. This corresponds to [Mic99, Proposition 2].

4.4.1.1 Hardy constant for uniform and binomial distribution

We first treat some cases of special interest. Consider the uniform distribution over $\{0, \ldots, n\}$.

Lemma 4.7 (Uniform Hardy constant)**.** *Let $S_n = \sum_{l=1}^{n} \frac{1}{a_l}$ and $\{p_k = \frac{1}{n+1}\}$. Then, for any i,*

$$\max\{B_+(p, i), B_-(p, i)\} \leq (n + 1)S_n.$$

Proof. Explicitly writing out the constants yields

$$B_+(p, i) = \max_{i<k\leq n} \frac{n - k + 1}{n + 1} \sum_{l=i+1}^{k} \frac{n + 1}{a_l} \leq (n - i + 1)S_n$$

and

$$B_-(p, i) = \max_{0\leq k<i} \frac{k + 1}{n + 1} \sum_{l=k+1}^{i} \frac{n + 1}{a_l} \leq (i + 1)S_n. \qquad \square$$

We will also need a bound on the Hardy constants for the binomial distribution $p_k \propto c^k \binom{n}{k}$ with $\{a_k = \sqrt{k(n + 1 - k)}\}$ derived from the hypercube symmetry reduction. Rather than deriving it explicitly, we can use that the corresponding Hamiltonian is precisely the symmetry reduction of the Hamming weight Hamiltonian $-A + \alpha K$ as described in Section 4.2.2, for the particular choice of $\alpha = \frac{n-2\mu}{\sqrt{\mu(n-\mu)}}$ (with μ the mean of q). From Lemma 4.1 we know that this Hamiltonian has a spectral gap $\Delta = \sqrt{\alpha^2 + 4} \geq 2$. Combined with the upper bound on the gap from Lemma 4.6, we get the lemma below.

Lemma 4.8 (Binomial Hardy constant)**.** *Let $a_k = \sqrt{k(n - k + 1)}$, $p_k \propto c^k \binom{n}{k}$, and m_p a median of p satisfying $p([0, m_p]) \geq 1/2$ and $p([m_p, n]) \geq 1/2$.*

- *If $i \geq m_p$, then $B_+(p, i) \leq 1$.*
- *If $i \leq m_p$, then $B_-(p, i) \leq 1$.*

4.4.2 Hardy constant from relative log-concavity

In this section we show one way of bounding the Hardy constant: if a distribution p is "log-concave relative to" some other distribution q, then we can upper bound the Hardy constant of p using that of q. For distributions p and q, we say that p is log-concave relative to q if the distribution $\{p_k/q_k\}$ is log-concave. It is easily seen that any log-concave distribution p is relative log-concave with respect to $q_k \propto c^k$ (Definition 4.6). More interestingly, if $q_k \propto \binom{n}{k}$ then

we recover the notion of ultra-log-concavity (Definition 4.7). We will be using the following technical lemma.

Lemma 4.9 (Single crossing lemma)**.** *Suppose that p is log-concave relative to q, and let k^* denote the mode of $\{r_k = p_k/q_k\}$. Let $b_\ell = a_\ell\sqrt{p_{\ell-1}p_\ell}$ and $c_\ell = a_\ell\sqrt{q_{\ell-1}q_\ell}$.*

1. *For all k, ℓ such that $k^* < \ell \leq k$, we have*

$$p([k,n]) \leq r_k q([k,n]) \qquad \text{and} \qquad b_\ell \geq r_k c_\ell.$$

2. *Assume there exists $0 < i < k^*$ such that either $q([0,i]) = p([0,i])$ or $q([i,n]) = p([i,n])$. Then for all k, ℓ such that $k > i$ and $i < \ell \leq k^*$, we have*

$$p([k,n]) \leq q([k,n]) \qquad \text{and} \qquad b_\ell \geq c_\ell.$$

Proof. To prove the first point, we note that log-concavity (and hence unimodality) of $\{r_k\}$ implies that the sequence is non-increasing for $k \geq k^*$. This implies that

$$p([k,n]) = \sum_{j=k}^{n} r_j q_j \leq r_k \sum_{j=k}^{n} q_j = r_k q([k,n]),$$

but also that $p_\ell \geq r_k q_\ell$ for $k^* \leq \ell \leq k$, and hence $b_\ell \geq r_k c_\ell$ for $k^* < \ell \leq k$.

For the second point, we assume that $p([0,s]) = q([0,s])$ for either $s = i$ or $s = i-1$ with $i < k^*$. Equivalently, this is $\sum_{j=0}^{s}(r_j - 1)q_j = 0$. Since the sequence $\{r_j\}$ is non-decreasing on $\{0, \ldots, k^*\}$, this implies that $r_s \geq 1$, and moreover $1 \leq r_s \leq r_{s+1} \leq \cdots \leq r_{k^*}$. This already implies that $b_\ell \geq c_\ell$ for $i < \ell \leq k^*$.

By unimodality and the fact that $r_i \geq 1$, the sequence $r_j - 1$ on $\{i, \ldots, n\}$ must be nonnegative up to some point, and nonpositive past that point (hence "single crossing"). By our assumption, the weighted sum over the interval $\sum_{j=i+1}^{n}(r_j - 1)q_j$ is nonpositive.[2] This implies that any suffix $\sum_{j=k}^{n}(r_j - 1)q_j$ of the sum must be nonpositive for any $k > i$, and hence $p([k,n]) \leq q([k,n])$ for $k > i$. □

The following theorem then shows how to bound the Hardy constant of a distribution p by using the Hardy constant of a distribution q relative to which it is log-concave and with which it shares the same tail mass. Note that we can always enforce this condition by replacing $\{q_k\}$ with $\{c^k q_k\}$ for an appropriate choice of c (indeed, log-concavity relative to $\{q_k\}$ implies log-concavity relative to $\{c^k q_k\}$).

Theorem 4.3 (Relative Hardy bound)**.** *Suppose that p is log-concave relative to q, and let k^* be the mode of $\{p_k/q_k\}$. Let $0 < i < n$ be such that either*

$$q([0,i]) = p([0,i]) \qquad \text{or} \qquad q([i,n]) = p([i,n]) \qquad \text{or} \qquad i = k^*.$$

Then

$$B_+(p,i) \leq B_+(q,i) \qquad \text{and} \qquad B_-(p,i) \leq B_-(q,i).$$

[2] It is zero if $q([0,i]) = p([0,i])$. It is nonpositive if $q([i,n]) = p([i,n])$, which follows from the fact that $r_i \geq 1$.

Proof. We only prove the bound on B_+, since the bound on B_- follows by symmetry.

Let $b_\ell = a_\ell\sqrt{p_{\ell-1}p_\ell}$ and $c_\ell = a_\ell\sqrt{q_{\ell-1}q_\ell}$. The bound on B_+ follows if we prove

$$\frac{p([k,n])}{b_\ell} \leq \frac{q([k,n])}{c_\ell} \quad (i < \ell \leq k).$$

This follows from point 1. in Lemma 4.9 if either $k^* \leq i$, or $k^* > i$ and $k^* < \ell \leq k$. The remaining case where $k^* > i$ and $i < \ell \leq k^*$ follows from point 2. in Lemma 4.9, where we invoke the assumption that either $q([0,i]) = p([0,i])$ or $q([i,n]) = p([i,n])$. □

4.4.3 Spectral gap bounds

Combined with our criteria for log-concavity of the ground state, we get our main result of the section.

Theorem 4.4. *Consider 1-dimensional $H = -A + V$ with ground state distribution p and spectral gap Δ.*

1. ***Uniform, log-concave:*** *if $\{a_\ell = 1\}$ and p log-concave then*

$$\Delta \in \Omega(1/n^2).$$

   This holds e.g. when $\{a_\ell = 1\}$ and V is single-well (Section 4.3.5).

2. ***Hypercube, log-concave:*** *if $\{a_\ell = \sqrt{\ell(n+1-\ell)}\}$ and p log-concave then*

$$\Delta \in \Omega(1/n).$$

   This holds e.g. when $\{a_\ell = \sqrt{\ell(n+1-\ell)}\}$ and H satisfies the 1-RM condition (Section 4.3.5).

3. ***Hypercube, ultra-log-concave:*** *if $\{a_\ell = \sqrt{\ell(n+1-\ell)}\}$ and p ultra-log-concave then*

$$\Delta \in \Omega(1).$$

   This holds e.g. when $\{a_\ell = \sqrt{\ell(n+1-\ell)}\}$ and V is convex (Theorem 4.1).

Proof. All proofs go via the discrete Hardy inequality (Lemma 4.6), stating that $\Delta \geq \frac{1}{4B(p,i)}$ with $B(p,i) = \max\{B_+(p,i), B_-(p,i)\}$. To bound $B(p,i)$, we combine our relative Hardy bound (Theorem 4.3) with our bounds on the Hardy constants of the geometric (Lemma 4.7) and binomial distribution (Lemma 4.8).

Towards proving points 1. and 2., we note that if p is log-concave with mode k^* then it is log-concave relative to the uniform distribution $\{q_k \propto 1\}$ and the mode of p/q is at k^*. We apply Theorem 4.3 with $i = k^*$. Hence

$$B(p,i) \overset{\text{Theorem 4.3}}{\leq} B(q,i) \overset{\text{Lemma 4.7}}{\leq} S_n(n+1). \tag{4.16}$$

Point 1. then follows by noting that $S_n = n$ in that case. Point 2. follows by showing that in this case $S_n \leq \pi$ for every $n \geq 1$. Define

$$f(x) = \frac{1}{\sqrt{x(n+1-x)}} \qquad (0 < x < n+1)$$

so that $f(\ell) = 1/a_\ell$. Since $f(x) = f(n+1-x)$, f is symmetric about the midpoint. Also, f is positive and decreasing on $\left(0, \frac{n+1}{2}\right]$, so for $k \leq \frac{n+1}{2}$, $f(k) \leq \int_{k-1}^{k} f(x)dx$. Hence we may bound the sum by an integral:

$$S_n \leq 2 \sum_{k=1}^{\lfloor (n+1)/2 \rfloor} f(k) \leq 2 \int_0^{(n+1)/2} \frac{dx}{\sqrt{x(n+1-x)}}.$$

Now make the substitution $x = (n+1)\sin^2\theta$ so that $dx = 2(n+1)\sin\theta\cos\theta\, d\theta$ and

$$\sqrt{x(n+1-x)} = \sqrt{(n+1)\sin^2\theta \cdot (n+1)\cos^2\theta} = (n+1)\sin\theta\cos\theta.$$

Therefore

$$\int_0^{(n+1)/2} \frac{dx}{\sqrt{x(n+1-x)}} = \int_0^{\pi/4} 2\, d\theta = \frac{\pi}{2}$$

and so $S_n \leq 2 \cdot \frac{\pi}{2} = \pi$.

Finally we prove point 3., which corresponds to the case where p is log-concave relative to the binomial distribution q. We apply Theorem 4.3 with $i = m_p$ a median of p, satisfying $p([0, m_p]) \geq 1/2$ and $p([m_p, n]) \geq 1/2$.
If $B(p, m_p) = B_+(p, m_p)$, then we choose c such that $q([0, m_p]) = p([0, m_p])$. Then if m_q is the leftmost median of q, we have that $m_q \leq m_p$. Using our bound on the Hardy constant of any binomial (Lemma 4.8), we get

$$B(p, m_p) = B_+(p, m_p) \overset{\text{Theorem 4.3}}{\leq} B_+(q, m_p) \overset{\text{Lemma 4.8}}{\leq} 1.$$

If $B(p, m_p) = B_-(p, m_p)$, then we choose c such that $q([m_p, n]) = p([m_p, n])$. Letting m'_q denote the rightmost median of q, we get that $m'_q \geq m_p$. Using our bound on the Hardy constant of any binomial (Lemma 4.8), we get

$$B(p, m_p) = B_-(p, m_p) \overset{\text{Theorem 4.3}}{\leq} B_-(q, m_p) \overset{\text{Lemma 4.8}}{\leq} 1. \qquad \square$$

It is interesting to compare our bounds with the results by Jarret and Jordan in [JJ14]:

- For the 1-dimensional Hamiltonian $H = -A + V$ with uniform $\{a_\ell\}$ and V convex, they show a bound $\Omega(1/n^2)$ on the spectral gap. Point 1. in Theorem 4.4 generalizes this bound to the much weaker case where $-V$ is merely unimodal.
- For the hypercube Hamiltonian $H = -A_{hc} + V$ with a symmetric potential V, we can use the symmetry reduction in Section 4.2.1. Point 3. in Theorem 4.4 then shows a constant gap when V is convex, in correspondence to the findings in [JJ14]. Point 2. in Theorem 4.4

goes beyond this, and shows an $\Omega(1/n)$ lower bound when H satisfies the 1-RM condition (Definition 4.8). This includes potentials that are non-convex and can even have local minima.

We do mention that the bounds in Jarret and Jordan [JJ14] are tight down to the constant, in analogy to the continuous fundamental gap proof [AC11] which their work is based on. This implies that points 1. and 3. in Theorem 4.4 are tight up to a constant.

Regarding point 2. in Theorem 4.4, we note that this bound is also tight up to a constant. Indeed, for the aforementioned potential $V(k) = a_k + a_{k+1}$, the ground state is uniform and hence log-concave. If the a_k's are derived from the hypercube then point 2. in Theorem 4.4 gives a bound $\Omega(1/n)$. To see that this is tight, note that the ground state energy in this case is 0, and the first excited energy can be shown to be $O(1/n)$ by using an ansatz state proportional to $\sum_k (k - n/2)|k\rangle$, orthogonal to the uniform ground state.

4.5 Log-concavity and tunneling

We now demonstrate the utility of log-concavity in analyzing quantum tunneling through localized spike perturbations. Following a variation of Reichardt's argument for the HWS problem [Rei04], we show that for log-concave ground states we can easily control the spectral gap when adding a spike perturbation. In the following we define the mean μ and variance σ^2 of a ground state $|\psi_0\rangle$ as the mean and variance of the distribution $p_k = \psi_k^2$. The key point will be that for a log-concave ground state, the mass placed on a short interval is uniformly controlled by the variance and the distance of the interval from the mean.

Theorem 4.5. *Consider a Hamiltonian $H = -A + V$ with spectral gap Δ and a log-concave ground state with variance σ and mean μ. Consider a perturbation $h : \{0, 1, \ldots, n\} \to \mathbb{R}_{\geq 0}$ supported on an interval $[\ell, u - 1] \subseteq [0, n]$ and satisfying an upper bound $h(k) \leq \bar{h}$. If $\alpha = \max\{\mu - u, \ell - \mu\}$ and*

$$\bar{h} \cdot \min\left\{ \frac{2(u-\ell)}{\sqrt{1+4\sigma^2}}, 2\exp\left(-\frac{\alpha}{2(2\sigma+1)}\right) \right\} \leq \Delta/2$$

then the perturbed Hamiltonian $H' = H + V'$ with $V'|k\rangle = h(k)|k\rangle$ has spectral gap $\Delta' \geq \Delta/2$.

Proof. Let λ_0, λ_1 (resp. λ_0', λ_1') denote the two lowest eigenvalues of H (resp. H'). By Weyl's monotonicity theorem, we have that $\lambda_1' \geq \lambda_1$. Conversely, with $p_k = \psi_k^2$, the variational principle with $|\psi_0\rangle$ as a test function implies that

$$\lambda_0' - \lambda_0 \leq \langle \psi_0 | V' | \psi_0 \rangle \leq \bar{h} \sum_{k=\ell}^{u-1} p_k.$$

Since the p_k's are log-concave, we know from Claim 4 that

$$p_{k^*} \leq \frac{2}{\sqrt{1+4\sigma^2}}.$$

This yields a bound $\lambda_0' - \lambda_0 \leq \bar{h} \sum_{k=\ell}^{u-1} p_k \leq \bar{h}(u-\ell)p_{k^*} \leq \bar{h}\frac{2(u-\ell)}{\sqrt{1+4\sigma^2}}$, which is good when σ is sufficiently large.

We get an alternative bound when σ is sufficiently small as compared to the distance between the mean μ and the spike. Here we use the second bound based on log-concavity, Claim 5, which states that

$$\sum_{|k-\mu|\geq\alpha} p_k \leq 2\exp\left(\frac{-\alpha}{2(2\sigma+1)}\right)$$

for any $\alpha \geq 0$. If either $\mu \leq \ell - \alpha$ or $\mu \geq u + \alpha$, then this implies that

$$\lambda_0' - \lambda_0 \leq \bar{h} \sum_{k=\ell}^{u-1} p_k \leq 2\bar{h}\exp\left(\frac{-\alpha}{2(2\sigma+1)}\right). \qquad \square$$

4.6 Application: Quadratic Hamming weight with a spike

In this section we develop the "quadratic Hamming weight with a spike" (quadratic HWS) problem. For this we will specialize our analysis to quadratic potentials perturbed by a spike. Specifically we consider a family of Hamiltonians

$$H_{qs}(s) = -(1-s)A_{hc} + sV_{qs}$$

for $s \in [0,1]$, with $V_{qs} = V_q + h$ where $V_q(k) = a(k-k_0)^2$ for $k_0 \leq 0$ (ensuring that the minimizer is at $k = 0$) and h is the spike perturbation. We aim to use Theorem 4.5 to bound the spectral gap $\Delta_{qs}(s)$ of H_{qs}, which requires a bound on the spectral gap $\Delta_q(s)$ of the *unperturbed* quadratic Hamiltonian $H_q(s) = -(1-s)A_{hc} + sV_q$, as well as estimates of the mean and variance of its ground state $|\psi_q(s)\rangle$. While for a linear potential, as in the HWS problem, the ground state is known explicitly as a binomial distribution, this is not so for quadratic potentials. We address this key challenge by showing that $|\psi_q(s)\rangle$ does maintain a constant overlap with the binomial ground state of a linear potential. Combined with log-concavity of the ground state, this yields explicit control of the mean and variance of $|\psi_q(s)\rangle$. This allows us to invoke the perturbation bound from Section 4.5 and generalize the HWS problem to quadratic potentials.

4.6.1 Unperturbed quadratic: Overlap with the binomial ansatz

We now consider the Hamiltonian $H = -A_{hc} + V$ for a convex potential V. The following claim quantifies the overlap of the ground state of H with an ansatz state $|\varphi_m\rangle$. The overlap depends on the second difference of the potential, which is defined as $\Delta^2 V(k) = V(k+1) + V(k-1) - 2V(k)$. Smaller $\Delta^2 V(k)$ leads to higher overlaps with the binomial state. We prove the claim in Appendix A.1.

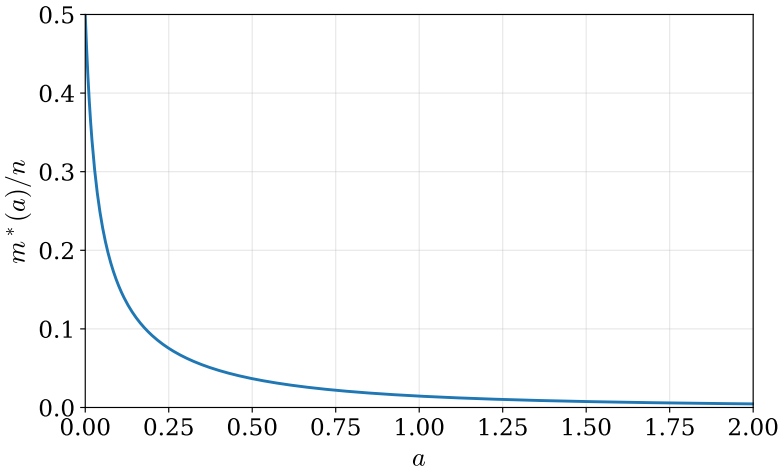


Figure 4.2: Plot of the function $m^*(a)/n$ obtained for $n = 40$ and $k_0 = -3$. $m^*(a)$ is obtained by solving the equation $a = \frac{n-2m^*}{\sqrt{m^*(n-m^*)}} \cdot \frac{1}{1-2k_0+2(n-1)m^*/n}$. As expected, m^* is located at $n/2$ for $a = 0$ and it decreases to 0 as a increases.

Claim 6 (Overlap bound)**.** *Let $|\psi_0\rangle$ denote the ground state of $H = -A_{hc} + V$ with V a convex function with bounded second difference $\Delta^2 V(k) \leq C$. Let m be chosen so as to minimize $\langle\varphi_m|H|\varphi_m\rangle$. Then it holds that*

$$|\langle\psi_0|\varphi_m\rangle|^2 \geq 1 - \left(1 - \frac{2m}{n}\right)\left((\mu - m)\left(1 - \frac{\Delta V(\lfloor m \rfloor)}{\mathbb{E}[\Delta V(Y)]}\right) + \frac{C\, m(1-m/n)}{2\mathbb{E}[\Delta V(Y)]}\right),$$

where $Y \sim \mathcal{B}(n-1, m/n)$ and $\mu = \langle\psi_0|K|\psi_0\rangle$.

In the quadratic case $V_q(k) = a(k-k_0)^2$, we have $\Delta V_q(k) = a(2(k-k_0)+1)$, and so the lower bound becomes simpler. We prove the below theorem in Appendix A.2, and we plot the optimal mean $m^*(a)$ as a function of a in Fig. 4.2.

Theorem 4.6 (Overlap with the binomial ansatz for quadratic potentials)**.** *Let $|\psi_{q,a}\rangle$ denote the ground state of $H = -A_{hc} + V_q$ with $V_q(k) = a(k-k_0)^2$ for $k_0 \leq 0$ and $a \in \mathbb{R}_{\geq 0}$. Then there exists $m^*(a) \in [0, n/2]$ such that*

$$\left|\langle\psi_{q,a}|\varphi_{m^*(a)}\rangle\right|^2 \geq \frac{1}{2}.$$

Now denote by μ_a and σ_a^2 the mean and variance associated to $|\psi_{q,a}\rangle$, which are given by

$$\mu_a = \langle\psi_{q,a}|K|\psi_{q,a}\rangle\,, \qquad \sigma_a^2 = \langle\psi_{q,a}|(K-\mu_a)^2|\psi_{q,a}\rangle\,.$$

Using the fact that both $|\psi_{q,a}\rangle$ and $|\varphi_m\rangle$ are log-concave, we can relate the associated mean and variance using Lemma 4.3. Applying this lemma to $\{p_k = |\langle k|\varphi_{m^*(a)}\rangle|^2\}$ and $\{q_k = |\langle k|\psi_{q,a}\rangle|^2\}$,

we obtain

$$\begin{aligned}\mu_a &= m^*(a) + O\Big(\sqrt{m^*(a)}\Big) + O(1),\\ \sigma_a &= \Theta\Big(\sqrt{m^*(a)}\Big) + O(1).\end{aligned} \tag{4.17}$$

Thus, although the quadratic ground state does not admit a closed form, the binomial ansatz provides a tractable approximation throughout the annealing path.

4.6.2 Unperturbed quadratic: Improved spectral gap bound

We now turn back to the interpolating, adiabatic Hamiltonian with quadratic potential,

$$H_q(s) = -(1-s)A_{hc} + sV_q.$$

From Theorem 4.4 (or [JJ14]), we get a bound $\Delta(H_q(s)) \in \Omega(1-s)$ on the spectral gap of the unperturbed, quadratic Hamiltonian $H_q(s)$. Unfortunately, that bound becomes useless around the end of the annealing path, when $s \to 1$. Luckily, a finer reading shows that the bound of [JJ14] can actually be extended until the end of the annealing path.

Lemma 4.10. *Consider $H(s) = -(1-s)A_{hc}+sV$ with V convex, and define $\beta = \min_{0\le i<j\le n}|\frac{V(j)-V(i)}{j-i}|$. Then*

$$\Delta(H(s)) \ge \sqrt{s^2\beta^2 + 4(1-s)^2}.$$

In particular, if there exists $c > 0$ such that $\beta \ge c$, then

$$\Delta(H(s)) \ge \frac{2c}{\sqrt{c^2+4}} \qquad \textit{uniformly for all } s \in [0,1].$$

Proof. Rewrite

$$H(s) = (1-s)\widetilde{H}_\lambda, \qquad \widetilde{H}_\lambda = -A + \lambda V, \qquad \lambda = \frac{s}{1-s},$$

so that $\Delta(H(s)) = (1-s)\Delta(\widetilde{H}_\lambda)$. By Lemma 10 of [JJ14], for convex V it holds that

$$\Delta(\widetilde{H}_\lambda) \ge \Delta(-A + \lambda\beta K) = \sqrt{(\lambda\beta)^2 + 4},$$

where $\beta = \min_{i<j}|\frac{V(j)-V(i)}{j-i}|$ and the equality is from Lemma 4.1. This proves that $\Delta(H(s)) \ge (1-s)\sqrt{\left(\frac{s}{1-s}\beta\right)^2 + 4} = \sqrt{s^2\beta^2 + 4(1-s)^2}$.

We have $\Delta(H(s)) \ge \sqrt{c^2s^2 + 4(1-s)^2}$. This is minimized for $s = \frac{4}{4+c^2}$ which then gives $\Delta(H(s)) \ge \frac{2c}{\sqrt{c^2+4}}$ uniformly for $s \in [0,1]$. □

In particular, for V with unique minimizer at 0, we have $\beta = V(1) - V(0) > 0$ which satisfies by definition this lemma. For the quadratic potential $V_q(k) = a(k-k_0)^2$ that we consider, we have that $\Delta V_q(k) \ge a$. Assuming $a \in \Omega(1)$, this yields a uniform bound $\Delta_q(s) \in \Omega(1)$ on the spectral gap of the unperturbed quadratic Hamiltonian $H_q(s)$.

4.6.3 Perturbed quadratic and AQO

We will now combine the bound on the spectral gap of the quadratic Hamiltonian $H_q(s)$, as well as the estimates on the mean and variance of its ground state, with the perturbative tunneling argument in Section 4.5. This will allow us to establish our analysis of the quadratic HWS problem.

4.6.3.1 Spectral gap of quadratic HWS

We now combine the general spike bound with the overlap analysis. The key point is that the variance of the true ground state is comparable to that of the binomial ansatz.

Theorem 4.7 (Quadratic potential with a spike)**.** *Consider the quadratic Hamming weight with a spike Hamiltonian*

$$H_{qs}(s) = -(1-s)A_{hc} + sV_{qs} = -(1-s)A_{hc} + s(V_q + h)$$

for $s \in [0,1]$, where the spike h is supported on $[\ell, u-1] \subseteq [0,n]$ and satisfies an upper bound $h(k) \leq \bar{h}$. If

$$\bar{h}\,\frac{u-\ell}{\sqrt{\ell}} \in O(1)$$

then the spectral gap of $H_{qs}(s)$ satisfies $\Delta_{qs}(s) \in \Omega(1)$.

Proof. Let $|\psi_q(s)\rangle$ denote the ground state and $\Delta_q(s)$ the spectral gap of the unperturbed Hamiltonian $H_q(s) = -(1-s)A_{hc} + sV_q$. By Lemma 4.10 we know that $\Delta_q(s) \in \Omega(1)$. By Theorem 4.5 we get that the perturbed gap $\Delta_{qs}(s) \in \Omega(1)$ provided that

$$\min\left\{\frac{2(u-\ell)}{\sqrt{1+4\sigma^2(s)}}, 2\exp\left(-\frac{\alpha(s)}{2(2\sigma(s)+1)}\right)\right\} \leq 1/\bar{h}, \tag{4.18}$$

where $\sigma^2(s)$ is the variance of the unperturbed ground state $|\psi_q(s)\rangle$ and $\alpha(s) = \max\{\ell - \mu(s), \mu(s) - u\}$ is the distance of the mean $\mu(s)$ of $|\psi_q(s)\rangle$ from the support $[\ell, u-1]$ of the spike h. By applying Theorem 4.6 and Eq. (4.17) to the Hamiltonian $-A_{hc} + \frac{s}{1-s}V_q$, there exists $m^*(s)$ such that $\mu(s) = m^*(s) + O(\sqrt{m^*(s)}) + O(1)$ and $\sigma(s) = \Theta(\sqrt{m^*(s)}) + O(1)$. Our proof distinguishes two cases.

First, assume that $m^*(s) \geq \ell/2$. Then $\sigma(s) \in \Omega(\sqrt{\ell})$ and

$$\frac{2(u-\ell)}{\sqrt{1+4\sigma^2(s)}} = \frac{u-\ell}{\sqrt{\ell}} \cdot \frac{2\sqrt{\ell}}{\sqrt{1+4\sigma^2(s)}} \in O(1/\bar{h})$$

by our assumption that $\bar{h}\,\frac{u-\ell}{\sqrt{\ell}} \in O(1)$.

Then, assume that $m^*(s) < \ell/2$. Then $\alpha(s) \in \Omega(\ell)$ and $\sigma(s) \in O(\sqrt{\ell}) + O(1)$. In this case, we get

$$\exp\left(-\frac{\alpha(s)}{2(2\sigma(s)+1)}\right) \leq \exp\left(-\Omega\left(\sqrt{\ell}\right)\right).$$

It remains to note that $\exp\left(-\Omega\left(\sqrt{\ell}\right)\right) \in O(1/\sqrt{\ell})$, and to use that $O(1/\sqrt{\ell}) \in O(1/\bar{h})$ by our assumption. □

The spike condition $\bar{h}\,\frac{u-\ell}{\sqrt{\ell}} \in O(1)$ is identical to the condition derived by Reichardt for the HWS problem [Rei04]. When for instance $u, \ell \in \Omega(n)$, the condition admits a spike of height and width $\Omega(n^{1/4})$.

4.6.4 Runtime for quadratic HWS

Finalizing our analysis, we combine our spectral gap bound with the AQO algorithm. By Section 2.4, AQO approximately follows the ground state $|\psi_{qs}(s)\rangle$ of $H_{qs}(s)$ and prepares the target state $|\psi_{qs}(1)\rangle = |0\rangle$ in time

$$O\left(\frac{(n + \max_k |V(k)|)^2}{(\min_s \Delta(s))^3}\right).$$

This yields the following corollary.

Corollary 4.1 (Informal version of Theorem 4.7)**.** *Let $V_q(k) = a(k-k_0)^2$ with $k_0 \leq 0$ and consider the perturbed Hamiltonian $H_{qs}(s) = -(1-s)A_{hc} + s(V_q + h)$, where the spike h is supported on $[\ell, u-1] \subseteq [0,n]$ and satisfies $h(k) \leq \bar{h}$. If $\bar{h}\,\frac{u-\ell}{\sqrt{\ell}} = O(1)$, then AQO prepares an approximation of $|\psi_{qs}(1)\rangle = |0\rangle$ from $|\psi_{qs}(0)\rangle$, the binomial state with mean at $n/2$, in time $O(n^4)$.*

We note that the setup is symmetric with respect to shifting the minimizer $|0\rangle$ to some other, hidden string $|z^*\rangle$, $z^* \in \{0,1\}^n$. In that case, the potential V would satisfy the symmetry condition $V(x) = V(|x \oplus z^*|)$ and the same analysis goes through.

4.7 Summary, related work and open questions

Our work establishes discrete log-concavity as a critical tool for analyzing discrete optimization with the quantum adiabatic algorithm. We proved log-concavity for a large family of potentials (including and beyond convex potentials), and used it to establish new spectral gap bounds. As a concrete application we proved that AQO can efficiently solve the quadratic HWS problem, yielding a direct extension to the results by Reichardt [Rei04] and Farhi, Goldstone and Gutmann [FGG02] on the (linear) HWS problem.

4.7.1 Related work

In this section we compare our results to some related works.

4.7.1.1 Spectral gap analysis of spike Hamiltonians

There exist more refined analyses of the linear HWS Hamiltonian. For spike position $u, \ell \approx n/4$, spike width $u - \ell = n^\alpha$ and spike height $h = n^\beta$, it was shown in [Rei04] that $\alpha + \beta < 1/2$ implies a constant gap for linear potentials, and our result extends this bound to quadratic potentials. In later works, Brady and van Dam [BvD16] (improving on Kong and Crosson [KC17]) showed that if $\alpha + \beta > 1/2$ but $2\alpha + \beta < 1$, then the gap scales as $n^{1/2-\alpha-\beta}$, while for $2\alpha + \beta > 1$, the gap becomes exponentially small. It would be interesting to extend these results to other potentials, as is suggested in [BvD16], and our result seems to be the first one in this direction.

4.7.1.2 Classical benchmark algorithms

Our polynomial runtime bound for the quadratic HWS problem should be held against comparable classical algorithms (which, similarly to AQO, do not exploit any particular structure of the problem). To this end, we note that a typical approach based on classical Markov chains and simulated annealing would take exponential time to solve this problem. This can be demonstrated in an identical fashion to the arguments in [FGG02, CH16] for the linear HWS problem. Let h be the height and $w = u - \ell$ the width of the spike, and assume $h, w \geq n^c$ for $c \in \Omega(1)$. If the Markov chain flips less than w bits in its proposal, then it can only cross the spike by accepting a move that jumps into the spike, but this will be rejected with probability $1 - O(\exp(-n^c))$. Conversely, if the Markov chain flips more than w bits, then the probability of a proposal ever jumping to the minimizer is $O(\exp(-n^c))$.

Notably, it was shown by Crosson and Harrow [CH16] and Bergamaschi [Ber20] that a more advanced classical Markov chain algorithm called simulated quantum annealing can solve the (linear) HWS problem under exactly the same assumptions on the spike perturbations. This algorithm bypasses the former argument by using non-local "wordline updates". It seems reasonable to expect that this algorithm would also solve the quadratic HWS problem, and so we only claim an advantage against the simplest simulated annealing algorithms.

4.7.2 Open questions

Finally, the results of this work suggest many directions for future research.

Sufficient and necessary condition for log-concavity. The $\gamma-$RM condition is a sufficient condition. We leave the question open of whether it is also a necessary condition. We do prove that at least two important limiting cases of (ultra-)log-concavity (the uniform and the binomial distribution) satisfy the $\gamma-$RM condition with equality.

Can we extend the quadratic HWS problem to general convex potentials? We specialized to quadratic potentials with minimizer at 0 because we can prove a strong overlap

bound between the corresponding ground states and our binomial ansatz states. A natural question is whether this approach extends to other convex potentials. First, note that if the minimizer is not at 0 then the ansatz is not well suited. We show this in Fig. 4.3 where we plot the overlap with the best binomial ansatz for a shifted quadratic potential $V(k) = (k - k_*)^2$ with minimizer at $k_* = n/3$. More unexpectedly, the overlap can decrease even if the minimizer is at 0. We show this in the figure for a piecewise-linear potential.

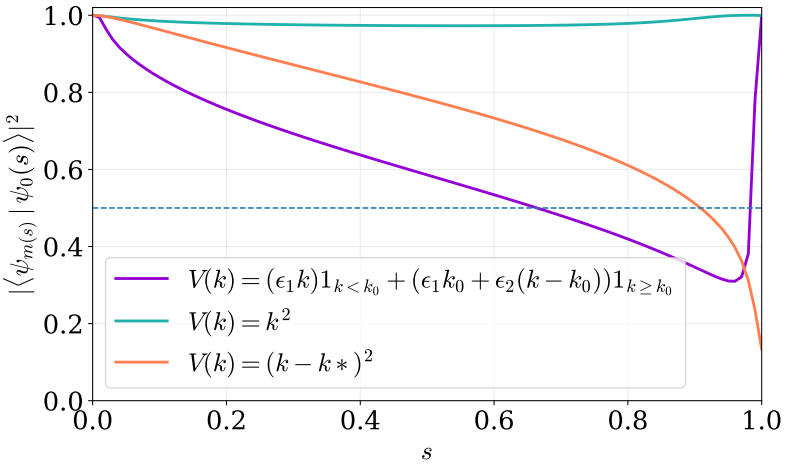


Figure 4.3: Overlap between the optimal binomial ansatz state $|\varphi_m\rangle$ and the ground state $|\psi_0(s)\rangle$ of $H(s) = -(1-s)A + sV$ for $n = 40$ and different potentials: the quadratic potential $V(k) = k^2$ (used in Section 4.6), a shifted quadratic potential $V(k) = (k - k_*)^2$ with minimizer at $k_* = n/3$, and a piecewise-linear potential $V(k) = \epsilon_1 k\,\mathbf{1}_{k<k_0} + (\epsilon_1 k_0 + \epsilon_2(k - k_0))\,\mathbf{1}_{k\geq k_0}$, with $n = 40$, $\epsilon_1 = 1/n$, $\epsilon_2 = 2n$, and $k_0 = n/4$.

Can this framework be applied to study tunneling for local Hamiltonians? A key point is that a localized spike perturbation corresponds to a highly non-local Hamiltonian (coupling many qubits together). It would be interesting to extend our discussion to local target Hamiltonians. E.g., the Curie-Weiss model considers target Hamiltonian $H_1 = -\sum_{i,j\in[n]} Z_i Z_j$. As discussed for instance in [BS12], the spectral gap of $H(s) = -(1-s)\sum_i X_i + sH_1$ as a function of s undergoes a phase transition from constant to exponentially small in n. Such a phase transition makes the Curie-Weiss model a natural benchmark for our techniques. We leave it as an open question whether the log-concavity framework developed here can be adapted to handle such local models.

Can this framework be extended to non-permutation-symmetric potentials? A key ingredient of our analysis is the reduction from the hypercube to the $(n+1)$-dimensional Hamming weight subspace, which allows us to work on a weighted path graph. Extending our

analysis to more general potentials on the hypercube is a clear open question. It raises interesting points such as what the correct notions of convexity and log-concavity are on the hypercube. We expect this to be non-trivial. E.g., it was shown in [JJ15] that single-basin potentials on general graphs, a condition slightly weaker than convexity, can already yield exponentially small spectral gaps. This contrasts with the 1-dimensional case where monotonicity of the potential on a uniform path was enough (Section 4.3.3). Additionally, multivariate log-concavity is less trivial to define in the discrete case than in the continuous case [BN14, KL19], and identifying the right structural substitute in higher dimensions should be interesting.

Finally, a slightly different approach was taken by Montanaro and Zhou [MZ24], who argued that certain quantum algorithms (QAOA in particular) for symmetric functions can be shown to be robust against bounded, asymmetric perturbations.

Conclusion and outlook

This thesis has investigated analog quantum algorithms for combinatorial optimization on near-term neutral-atom devices. The central motivation has been twofold. On the one hand, neutral-atom platforms provide a natural physical mechanism, the Rydberg blockade, for encoding graph optimization problems such as Maximum Independent Set. On the other hand, even when a problem can be encoded into the hardware, the success of quantum annealing depends on the spectral properties of the associated Hamiltonian path, and in particular on the minimum spectral gap. The results presented in this thesis address these two complementary aspects: first, by enlarging the class of graphs for which Maximum Independent Set can be represented on neutral-atom hardware, and second, by deriving rigorous guarantees for the annealing dynamics within families of symmetric cost functions.

The first contribution of the thesis concerns graph embeddings on neutral-atom devices. The Rydberg blockade provides a native way to encode Maximum Independent Set on unit-disk graphs. However, although Maximum Independent Set remains hard on unit-disk graphs in general, the particular instances that can be implemented directly on current neutral-atom hardware are geometrically restricted and often classically tractable in practice. We therefore investigated how neutral-atom devices can be extended beyond this native setting.

In Chapter 3, we introduced an operational definition of the Rydberg blockade in terms of the two-atom excitation probability. This made it possible to embed disk graphs, for which a local definition of the blockade radius is needed. We then characterized embeddings using correlation matrices and the probability of violating the independent-set constraint. Finally, we showed that near-optimal independent sets on disk graphs can be prepared using locally addressed driving schemes. Our numerical results show that local driving yields better performance than global protocols. Local control therefore extends the computational reach of neutral-atom processors beyond unit-disk graphs.

Several questions remain open. A first direction is to develop a more systematic theory for enlarging the classes of graphs that can be embedded in neutral-atom devices, in particular by further exploring embedding metrics. A second direction is to incorporate more realistic experimental constraints, such as imperfect atom placement, laser noise, limited addressability, and decoherence, in order to assess the robustness of these embeddings. Finally, it would be valuable to identify graph families that are both physically accessible and classically challenging.

Such instances could provide more targeted benchmarks for testing the potential advantage of neutral-atom quantum optimization.

The second contribution of the thesis concerns the theoretical analysis of quantum annealing. Even when an optimization problem is successfully encoded in a Hamiltonian, the success of the algorithm is governed by the minimum spectral gap along the annealing path. Since estimating this gap is difficult for general optimization landscapes, we focused on permutation-symmetric potentials, i.e., potentials that depend only on the Hamming weight. This setting retains nontrivial tunneling phenomena while reducing the analysis to a structured one-dimensional problem.

In Chapter 4, we proved that, for a broad family of potentials including convex potentials, the ground state of the corresponding discrete Schrödinger operator is log-concave. This provides a structural tool for the analysis of adiabatic quantum optimization, since log-concavity implies regularity properties of the ground-state distribution that can be used to lower bound the spectral gap. We then applied this framework to spike perturbations, inspired by the Hamming-weight-with-a-spike problem, where a localized perturbation is added to the Hamming-weight function. Our results show that the perturbative analysis of this problem extends to more general log-concave ground states. By reducing the problem to estimates on the mean, variance, and concentration of the unperturbed ground state, we obtained explicit conditions under which the spectral gap remains inverse-polynomially large. As an application, we showed that for quadratic potentials with a spike, the spectral gap remains constant under spike conditions analogous to those of the original linear Hamming-weight problem. This suggests that tunneling-based speedups may persist beyond exactly solvable models.

Several directions remain open. A first question is whether the log-concavity framework can be extended beyond symmetric potentials, for example to fully connected ferromagnetic systems or graph-based cost functions with local structure. A second direction is to sharpen the spectral-gap bounds obtained from log-concavity, and to identify the precise boundary between regimes where tunneling is efficient and regimes where the gap becomes exponentially small. Finally, it would be interesting to develop a more general theory of quantum tunneling in discrete optimization landscapes, including landscapes with multiple barriers, asymmetric wells, plateaus, or competing local minima.

A broader long-term direction is to develop a more systematic methodology for designing quantum algorithms on neutral-atom platforms. This is particularly relevant in an industrial context, where the goal is not only to study isolated examples, but to identify problem classes, embeddings, and algorithmic routines that can realistically be implemented on a given device. A natural next step would be to combine these perspectives into a framework that starts from the physical constraints of the platform, derives families of implementable Hamiltonians, and studies whether the resulting algorithms admit rigorous performance guarantees.

Such a framework should not be restricted to quantum annealing. Although annealing is a

natural paradigm for analog neutral-atom devices, other routines may provide complementary ways of exploiting the same hardware. Standard primitives from digital quantum algorithms, such as quantum phase estimation, could be used to analyze spectra, certify gaps, or design algorithms for Hamiltonians that can be embedded on neutral-atom platforms. More generally, hybrid schemes combining analog evolution, local control, measurements, and digital subroutines could lead to algorithms that are less heuristic and more amenable to theoretical analysis.

This would help bridge the two sides of the thesis: the hardware-oriented question of what can be encoded, and the algorithmic question of what can be solved with theoretical guarantees. Developing such a platform-level theory could provide more principled benchmarks for neutral-atom processors and clarify where practical quantum advantages may realistically arise.

APPENDIX A

Log concavity supplementary results

A.1 Proof of the binomial ansatz overlap

In this appendix, we prove the overlap bound in Claim 6 for a convex potential V. Throughout, we let $\bar{V}$ denote the extension of V to the continuous interval $[0, n]$ by linear interpolation, so that

$$\bar{V}(x) = V(\lfloor x \rfloor) + \{x\}\Big(V(\lceil x \rceil) - V(\lfloor x \rfloor)\Big), \tag{A.1}$$

where we used the notation $\{x\} = (x - \lfloor x \rfloor)$.

A.1.1 Helper lemmas

We first prove some helper lemmas.

Lemma A.1 (Variational overlap bound)**.** *Let H be a Hamiltonian with non-degenerate ground state $|\psi_0\rangle$, ground state energy λ_0 and spectral gap Δ. Then for any quantum state $|\phi\rangle$ it holds that*

$$|\langle\phi|\psi_0\rangle|^2 \geq 1 - \frac{\langle\phi|H|\phi\rangle - \lambda_0}{\Delta}.$$

Proof. Let $(\lambda_k, |\psi_k\rangle)_k$ be the eigenvalues and eigenvectors of H, ordered so that $\lambda_0 < \lambda_1 \leq ....$ Then

$$\langle\phi|(H - \lambda_0)|\phi\rangle = \sum_{k>0}(\lambda_k - \lambda_0)|\langle\phi|\psi_k\rangle|^2 \geq \Delta \sum_{k>0} |\langle\phi|\psi_k\rangle|^2 = \Delta\Big(1 - |\langle\phi|\psi_0\rangle|^2\Big),$$

which rearranges to give the result. □

Lemma A.2 (Convex upper bound)**.** *Let $V : \mathbb{Z} \to \mathbb{R}$ be a discrete convex function. Then for all $a, b \in \mathbb{R}$,*

$$\bar{V}(a) - \bar{V}(b) \leq (a-b)\Delta V(\lfloor a \rfloor).$$

Proof. For the continuous, piecewise differentiable extension $\bar{V}$ we have the standard fact that

$$\bar{V}(a) - \bar{V}(b) \leq (a-b)\bar{V}'(a)$$

when $\bar{V}$ is differentiable at a. This happens whenever $a \notin \mathbb{Z}$, in which case $\bar{V}'(a) = \Delta V(\lfloor a \rfloor)$. The statement for $a \in \mathbb{Z}$ follows by continuity and taking the limit $\lim_{\varepsilon \to 0_+} a + \varepsilon$. □

A.1.2 Binomial expectation lemmas

Next we prove some technical lemmas on binomial expectations.

Lemma A.3 (Binomial energy)**.** *Let $X \sim \mathcal{B}(n, m/n)$ for $m \in [0, n]$. If $\Delta^2 V(k) \leq C$ for all $1 \leq k \leq n-1$ then*

$$\mathbb{E}[V(X)] \leq \bar{V}(m) + \frac{C}{2} m(1 - m/n).$$

Proof. Let $W(k) = V(k) - \frac{C}{2}k^2$. Using that $\Delta^2\left(\frac{C}{2}k^2\right) = C$ we get that $\Delta^2 W(k) = \Delta^2 V(k) - C \leq 0$, and so W is discrete concave. Now let $\bar{W}$ denote the concave extension of W to the continuous interval $[0, n]$ by linear interpolation. Jensen's inequality then gives

$$\mathbb{E}[\bar{W}(X)] \leq \bar{W}(\mathbb{E}[X]) = \bar{V}(m) - \frac{C}{2}\mathbb{E}[X]^2.$$

On the other hand,

$$\mathbb{E}[\bar{W}(X)] = \mathbb{E}[\bar{V}(X)] - \frac{C}{2}\mathbb{E}[X^2].$$

The claim follows by using that $\mathbb{E}[X^2] - \mathbb{E}[X]^2 = \mathrm{Var}[X] = m(1 - m/n)$. □

Lemma A.4 (Binomial derivative)**.** *Let $m \in [0, n/2]$, $X \sim \mathcal{B}(n, m/n)$ and $Y \sim \mathcal{B}(n-1, m/n)$. Then*

$$\frac{\mathrm{d}}{\mathrm{d}m}\mathbb{E}[V(X)] = \mathbb{E}[\Delta V(Y)].$$

Proof. Let $\Pr[X = k] = q_{m,n}(k) = \binom{n}{k}\left(\frac{m}{n}\right)^k\left(1 - \frac{m}{n}\right)^{n-k}$. Then (e.g. by the derivative rule of Bernstein polynomials)

$$\frac{\mathrm{d}}{\mathrm{d}m} q_{m,n}(k) = q_{m,n-1}(k-1) - q_{m,n-1}(k)$$

where we set $q_{m,n-1}(-1) = 0$. Plugging in and reshuffling shows that

$$\frac{\mathrm{d}}{\mathrm{d}m}\sum_{k=0}^{n} q_{m,n}(k)V(k) = \sum_{k=0}^{n} q_{m,n-1}(k)\Big(V(k+1) - V(k)\Big).$$ □

Lemma A.5 (Optimal m)**.** *Let $m \in [0, n/2]$, and choose m such that*

$$\frac{d}{dm} \langle \varphi_m | - A + V | \varphi_m \rangle = 0.$$

Then $\frac{d}{dm} \langle \varphi_m | V | \varphi_m \rangle = \alpha(m)$ with $\alpha(m) = \frac{n-2m}{\sqrt{m(n-m)}}$ (cf. (4.4)). In particular, m satisfies

$$\alpha(m) = \mathbb{E}[\Delta V(Y)] \quad \text{for} \quad Y \sim \mathcal{B}(n-1, m/n).$$

Proof. The variational energy is $-2\sqrt{m(n-m)} + \langle \varphi_m | V | \varphi_m \rangle$. Differentiating and setting to zero,

$$-\frac{n-2m}{\sqrt{m(n-m)}} + \frac{d}{dm} \langle \varphi_m | V | \varphi_m \rangle = 0.$$

Using $\alpha(m) = (n-2m)/\sqrt{m(n-m)}$ and $\frac{d}{dm} \langle \varphi_m | V | \varphi_m \rangle = \mathbb{E}[\Delta V(Y)]$ (Lemma A.4) gives the result. □

A.1.3 Proof of claim

Finally we prove the claim, which we restate below.

Claim 6 (Overlap bound)**.** *Let $|\psi_0\rangle$ denote the ground state of $H = -A_{hc} + V$ with V a convex function with bounded second difference $\Delta^2 V(k) \le C$. Let m be chosen so as to minimize $\langle \varphi_m | H | \varphi_m \rangle$. Then it holds that*

$$|\langle \psi_0 | \varphi_m \rangle|^2 \ge 1 - \left(1 - \frac{2m}{n}\right) \left((\mu - m) \left(1 - \frac{\Delta V(\lfloor m \rfloor)}{\mathbb{E}[\Delta V(Y)]} \right) + \frac{C\, m(1 - m/n)}{2\mathbb{E}[\Delta V(Y)]} \right),$$

where $Y \sim \mathcal{B}(n-1, m/n)$ and $\mu = \langle \psi_0 | K | \psi_0 \rangle$.

Proof of Claim 6. Let $\alpha = \alpha(m)$ as in Lemma 4.1. Then $H_\alpha = -A_{hc} + \alpha K$ has ground state $|\varphi_m\rangle$, ground state energy $\lambda_\alpha = \frac{n}{2}\left(\alpha - \sqrt{4 + \alpha^2}\right)$ and spectral gap $\Delta_\alpha = \sqrt{4 + \alpha^2}$. Applying Lemma A.1 with $H = H_\alpha$ and $|\phi\rangle = |\psi_0\rangle$ we get

$$|\langle \psi_0 | \varphi_m \rangle|^2 \ge 1 - \frac{\langle \psi_0 | H_\alpha | \psi_0 \rangle - \lambda_\alpha}{\Delta_\alpha}.$$

We expand the numerator:

$$\begin{aligned} \langle \psi_0 | H_\alpha | \psi_0 \rangle - \lambda_\alpha &= \langle \psi_0 | H_\alpha | \psi_0 \rangle - \langle \varphi_m | H_\alpha | \varphi_m \rangle \\ &= \alpha(\mu - m) - \Big(\langle \psi_0 | A_{hc} | \psi_0 \rangle - \langle \varphi_m | A_{hc} | \varphi_m \rangle \Big). \end{aligned}$$

Since $|\psi_0\rangle$ is the ground state of $H = -A_{hc} + V$, we have $\langle \psi_0 | H | \psi_0 \rangle \le \langle \varphi_m | H | \varphi_m \rangle$ and hence

$$\langle \psi_0 | A_{hc} | \psi_0 \rangle - \langle \varphi_m | A_{hc} | \varphi_m \rangle \ge \langle \psi_0 | V | \psi_0 \rangle - \langle \varphi_m | V | \varphi_m \rangle .$$

Therefore

$$|\langle\psi_0|\varphi_m\rangle|^2 \geq 1 - \frac{\alpha}{\sqrt{4+\alpha^2}}(\mu - m) + \frac{1}{\sqrt{4+\alpha^2}}\Big(\langle\psi_0|V|\psi_0\rangle - \langle\varphi_m|V|\varphi_m\rangle\Big).$$

By Jensen's inequality, $\langle\psi_0|V|\psi_0\rangle \geq \bar{V}(\mu)$ for $\mu = \langle\psi_0|K|\psi_0\rangle$. Combining this with Lemma A.3, $\langle\varphi_m|V|\varphi_m\rangle \leq \bar{V}(m) + \frac{C}{2}m(1-m/n)$, yields

$$\langle\psi_0|V|\psi_0\rangle - \langle\varphi_m|V|\varphi_m\rangle \geq \bar{V}(\mu) - \bar{V}(m) - \frac{C}{2}m\Big(1 - \frac{m}{n}\Big) \tag{A.2}$$

$$\geq (\mu - m)\Delta V(\lfloor m\rfloor) - \frac{C}{2}m\Big(1 - \frac{m}{n}\Big), \tag{A.3}$$

where the last inequality used Lemma A.2, $\bar{V}(m) - \bar{V}(\mu) \leq (m-\mu)\Delta V(\lfloor m\rfloor)$.

Substituting, and using that $\alpha = \mathbb{E}[\Delta V(Y)]$ from Lemma A.5 and $\frac{\alpha}{\sqrt{4+\alpha^2}} = 1 - \frac{2m}{n}$, gives the result. □

A.2 Overlap bound for quadratic potentials

In this appendix, we specialize the overlap bound to the quadratic case, proving Theorem 4.6.

Claim 7. *Consider $H = -A_{hc} + V_q$ with $V_q(k) = a(k-k_0)^2$ and ground state $|\psi_0\rangle$. If m^* minimizes $\langle\varphi_m|H|\varphi_m\rangle$ then it satisfies*

$$a = \frac{n - 2m^*}{\sqrt{m^*(n-m^*)}} \cdot \frac{1}{1 - 2m^*/n + 2(m^* - k_0)},$$

and

$$|\langle\psi_0|\varphi_{m^*}\rangle|^2 \geq 1/2.$$

Proof. By Lemma A.5 we have that the optimal m^* satisfies $\frac{n-2m^*}{\sqrt{m^*(n-m^*)}} = \mathbb{E}[\Delta V(Y)]$ with $Y \sim \mathcal{B}(n-1, m^*/n)$. Since $\Delta V(k) = a(2(k-k_0)+1)$ we can write out the right hand side as

$$\mathbb{E}[\Delta V(Y)] = a\Big(1 - 2m^*/n + 2(m^* - k_0)\Big),$$

and this yields the first identity.

Towards proving the second bound, we combine this expression with Eq. (A.2) in the proof of Claim 6 to get

$$|\langle\psi_0|\varphi_m\rangle|^2 \geq 1 - \left(1 - \frac{2m}{n}\right) f(\mu, m),$$

where $\mu = \langle\psi_0|K|\psi_0\rangle$ and

$$f(\mu, m) = \mu - m - \frac{\bar{V}_q(\mu) - \bar{V}_q(m) - am(1-m/n)}{a\,(1 - 2m/n + 2(m-k_0))}.$$

By Lemma A.6 we can rewrite $\bar{V}_q(x) = a(x-k_0)^2 + a\{x\}(1-\{x\})$. Using the bound $0 \le \{x\}(1-\{x\}) \le 1/4$, this implies that

$$\bar{V}_q(\mu) - \bar{V}_q(m) \ge a\left[(\mu-k_0)^2 - (m-k_0)^2 - 1/4\right].$$

Substituting and combining terms shows that

$$\begin{aligned} f(\mu,m) &\le \mu - m - \frac{(\mu-k_0)^2 - (m-k_0)^2 - 1/4 - m(1-m/n)}{1-2m/n+2(m-k_0)} \\ &= \frac{(1-2m/n)(\mu-m) - (\mu-m)^2 + 1/4 + m(1-m/n)}{1-2m/n+2(m-k_0)} \\ &\le \frac{(1-2m/n)(\mu-m) - (\mu-m)^2 + 1/4 + m(1-m/n)}{1-2m/n+2m}. \end{aligned}$$

We further simplify this bound by noting that

$$(1-2m/n)(\mu-m) - (\mu-m)^2 \le 1/4,$$

which follows from the fact that for $y \in \mathbb{R}$ and $q \in [0,1]$ it holds that $qy - y^2 = -(y-q/2)^2 + q^2/4 \le q^2/4 \le 1/4$. This yields the bound

$$f(\mu,m) \le \frac{1/2 + m(1-m/n)}{1-2m/n+2m}.$$

It remains to prove that

$$(1-2m/n)f(\mu,m) \le \frac{(1-2m/n)(1/2+m(1-m/n))}{1-2m/n+2m} \le 1/2,$$

but this is equivalent to the condition $(1-2m/n)(1-m/n) \le 1$, which is satisfied by our assumption that $m \le n/2$. □

Lemma A.6. *For $V_q(x) = a(x-k_0)^2$ and $x \in [0,n]$ we have*

$$\bar{V}_q(x) = a(x-k_0)^2 + a\{x\}(1-\{x\}).$$

Proof. This follows by using the expressions for $\bar{V}_q(x)$ (Eq. (A.1)) and $\Delta V_q(\lfloor x \rfloor)$ to write out

$$\bar{V}_q(x) = V_q(\lfloor x \rfloor) + a\{x\}(2(\lfloor x \rfloor - k_0) + 1)$$

and compare it with

$$a(x-k_0)^2 = a(\lfloor x \rfloor + \{x\} - k_0)^2 = V_q(\lfloor x \rfloor) + a\{x\}(2(\lfloor x \rfloor - k_0) + \{x\}).$$ □

A.3 Independent-set sampling for bipartite graphs

We consider a random bipartite graph G with bipartition $L \sqcup R$, where $|L| = n$ and $|R| = \gamma n$. The graph is obtained by drawing Δ independent uniformly random matchings between L and R. For $0 \le \alpha, \beta \le 1$, let $I(\alpha, \beta)$ denote the set of independent sets containing αn vertices in L and $\beta\gamma n$ vertices in R. We write

$$\mathcal{E}_{AS}(\alpha, \beta) = \mathbb{E}_G\Big[|I(\alpha, \beta)|\Big]$$

for the expected number of such independent sets.

Then

$$\mathcal{E}_{AS}(\alpha, \beta) = \binom{n}{\alpha n}\binom{\gamma n}{\beta\gamma n}\left(\frac{\binom{(1-\beta)\gamma n}{\alpha n}}{\binom{\gamma n}{\alpha n}}\right)^{\Delta}. \tag{A.4}$$

Indeed, after choosing the αn vertices on the left and the $\beta\gamma n$ vertices on the right, each matching must avoid connecting the selected vertices in L to the selected vertices in R. For a single matching, the images of the selected left vertices form a uniformly random αn-subset of R. The probability that this subset avoids the $\beta\gamma n$ selected right vertices is therefore

$$\frac{\binom{(1-\beta)\gamma n}{\alpha n}}{\binom{\gamma n}{\alpha n}}.$$

Since the Δ matchings are independent, this probability is raised to the power Δ. The expression is understood on the domain where the binomial coefficients are well-defined, in particular where $\alpha \le (1-\beta)\gamma$; outside this domain, the corresponding probability is zero.

Using Stirling's approximation, we obtain

$$\mathcal{E}_{AS}(\alpha, \beta) = \exp\Big(n\varphi_{AS}(\alpha, \beta) + o(n)\Big),$$

where

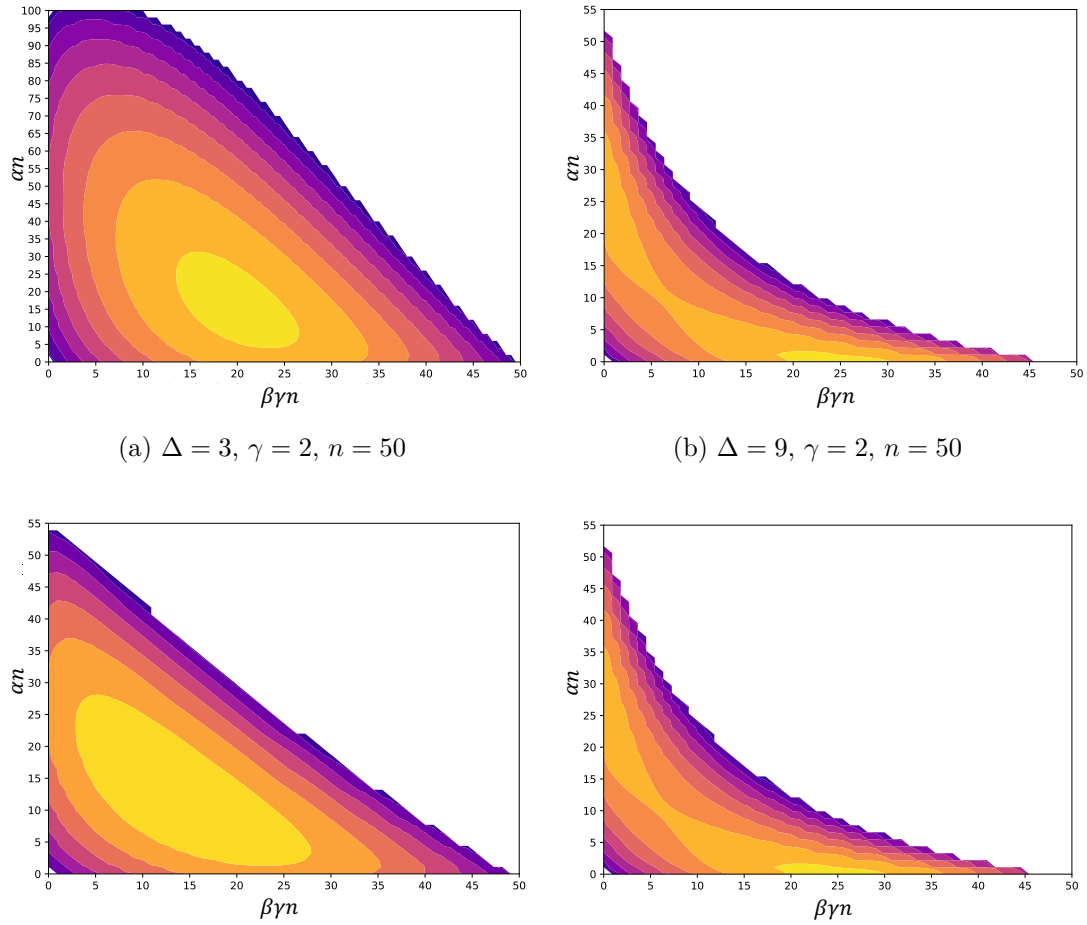
$$\begin{aligned}\varphi_{AS}(\alpha, \beta) = {} & -\alpha \ln \alpha - (1-\alpha)\ln(1-\alpha) + \gamma\Big[-\beta \ln \beta - (1-\beta)\ln(1-\beta)\Big] \\ & + \Delta\Big[(\gamma-\alpha)\ln(\gamma-\alpha) + \gamma(1-\beta)\ln(1-\beta) \\ & \quad - \Big((1-\beta)\gamma - \alpha\Big)\ln\Big((1-\beta)\gamma - \alpha\Big) - \gamma\beta\ln\gamma\Big].\end{aligned}$$

Equivalently, writing

$$H(x) = -x\ln x - (1-x)\ln(1-x)$$

for the binary entropy, this can be expressed as

$$\begin{aligned}\varphi_{AS}(\alpha, \beta) = {} & H(\alpha) + \gamma H(\beta) \\ & + \Delta\gamma\left[\left(1 - \frac{\alpha}{\gamma}\right)\ln\left(1 - \frac{\alpha}{\gamma}\right) + (1-\beta)\ln(1-\beta)\right. \\ & \quad \left. - \left(1 - \frac{\alpha}{\gamma} - \beta\right)\ln\left(1 - \frac{\alpha}{\gamma} - \beta\right)\right].\end{aligned}$$



(a) $\Delta = 3$, $\gamma = 2$, $n = 50$

(b) $\Delta = 9$, $\gamma = 2$, $n = 50$

(c) $\Delta = 3$, $\gamma = 1.1$, $n = 50$

(d) $\Delta = 9$, $\gamma = 1.1$, $n = 50$

Figure A.1: Asymptotic log-density $\varphi_{AS}(\alpha, \beta)$ of independent sets in the random bipartite graph model. The graph is obtained by drawing Δ independent random matchings between vertex sets L and R with $|L| = n$ and $|R| = \gamma n$.

We now restrict this potential to a fixed total density. For $\rho \in [0, 1]$, define

$$\mathcal{V}_\rho = \{(x, y) \in [0, 1]^2 : x + y = \rho\}.$$

We focus on the case $\rho = 1/2$. In the symmetric case $\gamma = 1$, we define the one-dimensional restriction

$$f(x) = \varphi_{AS}\left(x, \frac{1}{2} - x\right), \qquad x \in \left[0, \frac{1}{2}\right].$$

This gives

$$\begin{aligned} f(x) = &-x \ln x - \left(\frac{1}{2} - x\right) \ln\left(\frac{1}{2} - x\right) + \frac{\Delta}{2} \ln 2 \\ &+ (\Delta - 1)\left[(1 - x)\ln(1 - x) + \left(\frac{1}{2} + x\right) \ln\left(\frac{1}{2} + x\right)\right]. \end{aligned} \tag{A.5}$$

To determine the shape of f, we compute

$$f'(x) = -\ln x + \ln\left(\frac{1}{2} - x\right) + (\Delta - 1)\left[\ln\left(\frac{1}{2} + x\right) - \ln(1 - x)\right], \tag{A.6}$$

$$f''(x) = -\frac{1}{x} - \frac{1}{\frac{1}{2} - x} + (\Delta - 1)\left[\frac{1}{\frac{1}{2} + x} + \frac{1}{1 - x}\right] \tag{A.7}$$

$$= \frac{-2(3\Delta - 4)x^2 + (3\Delta - 4)x - 1}{x(2x + 1)(1 - x)(1 - 2x)}. \tag{A.8}$$

For $x \in (0, 1/2)$, the denominator is positive. Hence the sign of f'' is determined by

$$Q(x) = -2(3\Delta - 4)x^2 + (3\Delta - 4)x - 1.$$

The discriminant of Q is

$$D = (3\Delta - 4)(3\Delta - 12).$$

Thus, for $\Delta \leq 4$, $D \leq 0$ and $Q(x) \leq 0$ on $(0, 1/2)$, so f is concave. Since f is symmetric around $x = 1/4$, it has a unique maximum at $x = 1/4$.

For $\Delta > 4$, the two inflection points are

$$x_\pm = \frac{1 \pm \sqrt{1 - \frac{8}{3\Delta - 4}}}{4}.$$

In this case, f is concave on $(0, x_-) \cup (x_+, 1/2)$ and convex on (x_-, x_+).

It remains to determine the qualitative shape of f when $\Delta > 4$. Writing

$$x_- = \frac{1}{4}(1 - u), \qquad u = \sqrt{\frac{3\Delta - 12}{3\Delta - 4}} \in (0, 1),$$

one has

$$\Delta = \frac{12 - u^2}{3(1 - u^2)}.$$

A direct substitution gives

$$f'(x_-) = \ln\frac{1 + u}{1 - u} + \frac{9 - u^2}{3(1 - u^2)} \ln\frac{3 - u}{3 + u}.$$

To show that this quantity is negative, define

$$g(u) = 3(1 - u^2) \ln\frac{1 + u}{1 - u} + (9 - u^2) \ln\frac{3 - u}{3 + u}.$$

Then $f'(x_-) < 0$ is equivalent to $g(u) < 0$. Moreover,

$$g'(u) = -2u\left(3 \ln\frac{1 + u}{1 - u} + \ln\frac{3 - u}{3 + u}\right).$$

The term in parentheses is nonnegative on $[0, 1)$, since its derivative is

$$\frac{48}{(3 - u)(1 - u)(1 + u)(3 + u)} > 0$$

and it vanishes at $u = 0$. Hence g is decreasing, and therefore $g(u) < g(0) = 0$ for $u > 0$. Thus $f'(x_-) < 0$. By symmetry, $f'(x_+) = -f'(x_-) > 0$.

Since $f'(0) = +\infty$ and $f'(1/2) = -\infty$, this implies that for $\Delta > 4$ the function f has two symmetric local maxima, located at α and $1/2 - \alpha$, and one local minimum at $x = 1/4$. The parameter $\alpha \in (0, 1/4)$ is determined by the equation

$$\frac{\frac{1}{2} - \alpha}{\alpha} = \left(\frac{1-\alpha}{\frac{1}{2}+\alpha}\right)^{\Delta-1}.$$

The qualitative behavior can therefore be summarized as follows:

- For $\Delta \leq 4$, f is concave on $(0, 1/2)$ and has a unique maximum at $x = 1/4$.
- For $\Delta > 4$, f has two symmetric local maxima at α and $1/2 - \alpha$, separated by a local minimum at $x = 1/4$.

Near the bifurcation point $\Delta = 4$, we can estimate the location of the two maxima perturbatively. Write

$$\alpha = \frac{1}{4}(1 - \beta), \qquad \beta \ll 1.$$

Then the defining equation becomes

$$\frac{1+\beta}{1-\beta} = \left(\frac{3+\beta}{3-\beta}\right)^{\Delta-1}.$$

Taking logarithms and expanding up to order β^3 gives

$$\beta + \frac{\beta^3}{3} = (\Delta - 1)\left(\frac{\beta}{3} + \frac{\beta^3}{81}\right).$$

For the nonzero branch, this yields

$$\beta^2 = \frac{27(\Delta - 4)}{28 - \Delta},$$

and hence, close to $\Delta = 4$,

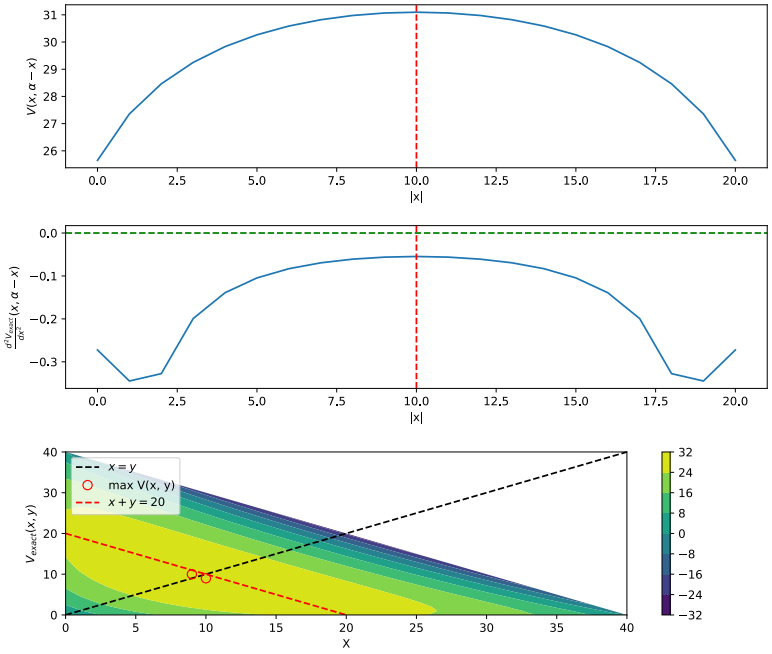
$$\alpha = \frac{1}{4}\left(1 - \sqrt{\frac{27(\Delta-4)}{28-\Delta}}\right).$$

This approximation describes the splitting of the symmetric maximum $x = 1/4$ into two local maxima as Δ crosses 4.

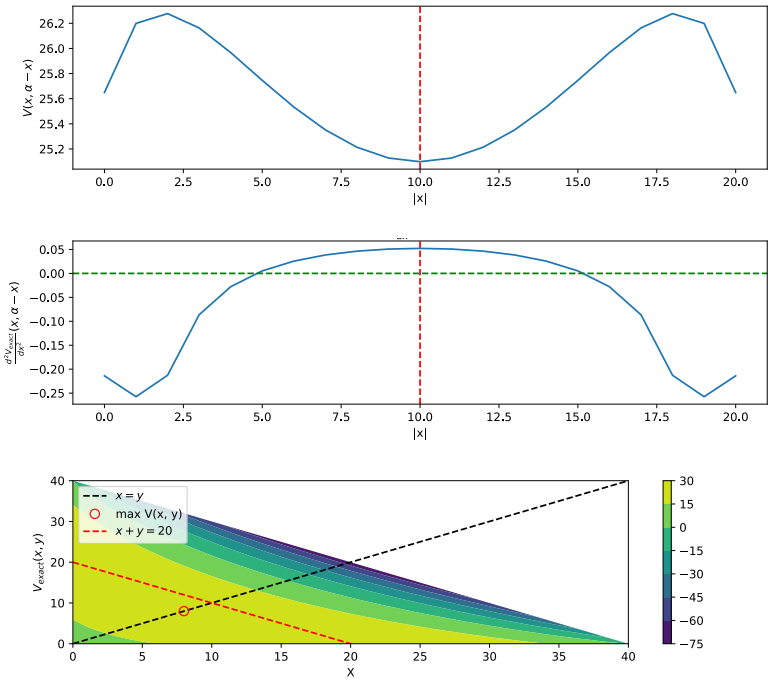
Away from the perturbative regime, the equation

$$\frac{\frac{1}{2} - \alpha}{\alpha} = \left(\frac{1-\alpha}{\frac{1}{2}+\alpha}\right)^{\Delta-1}$$

can be solved numerically. The resulting dependence of α on Δ is shown in Fig. A.3. The separation between the two local maxima is $1/2 - 2\alpha$, whose numerical behavior is shown in Fig. A.4.



(a) $\Delta = 3$, $n = 40$



(b) $\Delta = 4.80$, $n = 40$

Figure A.2: Potential $V(x) = \phi(x, 1/2 - x)$, second derivative of V, and two-dimensional plots of ϕ for $\Delta = 3$ and $\Delta = 4.80$.

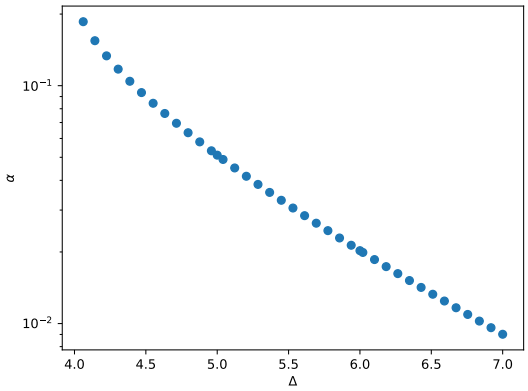


Figure A.3: Numerical solution for α as a function of Δ.

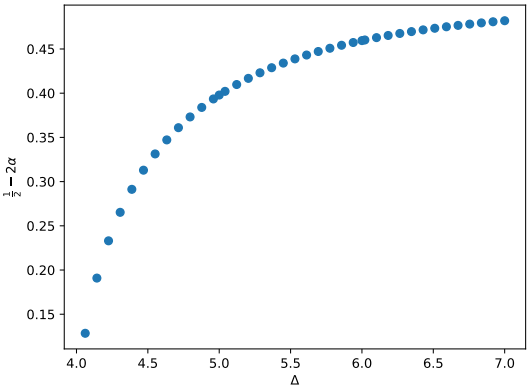


Figure A.4: Numerical solution for the separation $1/2 - 2\alpha$ as a function of Δ.

A.3.0.1 Effective quantum model

We now describe the effective one-dimensional quantum model suggested by the preceding classical potential. Consider the Hamiltonian

$$H(s) = -(1-s)\sum_{i=1}^{2n} X_i - sn\sum_{k,l}\phi(k/n, l/n)\,|k,l\rangle\,\langle k,l|\,,$$

where $|k,l\rangle$ denotes the symmetric state with Hamming weight k on the left and l on the right.

We are interested in the restriction to a fixed total Hamming weight $k+l=r$. The diagonal part then reduces to

$$-sn\sum_{l=0}^{r}\phi(l/n,(r-l)/n)\,|l\rangle\,\langle l|\,.$$

The transverse-field term does not preserve the total Hamming weight, so this restriction should be understood as an effective model rather than as an invariant subspace of the original transverse-field Hamiltonian. Within this effective description, transitions correspond to moving

one excitation from one side to the other. This gives the Ehrenfest-chain kinetic term on $\{0, \ldots, r\}$.

For $r = n/2$, define

$$g(l) = -nf(l/n).$$

The effective Hamiltonian takes the form

$$H_{\text{eff}}^{n/2}(s) = -(1-s)K_{\text{Ehr}} + s\sum_{l=0}^{n/2} g(l)\,|l\rangle\,\langle l|\,,$$

where K_{Ehr} is the hopping operator associated with the Ehrenfest chain. In particular, maxima of f become minima of the potential g.

The shape analysis above suggests two different regimes. For $\Delta \leq 4$, the function f has a single maximum, so $g = -nf$ has a single well. In this regime, one expects the effective Hamiltonian to have a polynomially large, and possibly constant, spectral gap. For $\Delta > 4$, the function f has two symmetric local maxima separated by a local minimum, so g becomes a double-well potential. In this regime, one expects the minimum gap to be controlled by tunneling between the two wells, and it may become exponentially small depending on the barrier height and width.

Bibliography

[AB09] S. Arora and B. Barak.
Computational complexity: a modern approach.
Cambridge University Press, 2009.
(page 10)

[AC11] B. Andrews and J. Clutterbuck.
Proof of the fundamental gap conjecture.
Journal of the American Mathematical Society, 24(3):899–916, July 2011.
doi:10.1090/S0894-0347-2011-00699-1.
(pages 70 and 82)

[ACd89] B. Apolloni, C. Carvalho, and D. de Falco.
Quantum stochastic optimization.
Stochastic Processes and their Applications, 33(2):233–244, 1989.
doi:https://doi.org/10.1016/0304-4149(89)90040-9.
(page 4)

[aD] C. and Dalyac.
Quantum many-body dynamics for combinatorial optimisation and machine learning.
PhD thesis, Agence Bibliographique de l'Enseignement Supérieur.
doi:10.70675/d8d31cecza123z48adz86b7ze8a30aaa7c58.
(page 22)

[AKR10] B. Altshuler, H. Krovi, and J. Roland.
Anderson localization makes adiabatic quantum optimization fail.
Proceedings of the National Academy of Sciences, 107(28):12446–12450, 2010, https://www.pnas.org/doi/pdf/10.1073/pnas.1002116107.
doi:10.1073/pnas.1002116107.
(pages xxiv, 3, and 59)

[AL18a] T. Albash and D. A. Lidar.
Adiabatic quantum computation.
Reviews of Modern Physics, 90(1):015002, 2018.
(page 15)

[AL18b] T. Albash and D. A. Lidar.
Adiabatic quantum computation.
Reviews of Modern Physics, 90(1), Jan. 2018.
doi:10.1103/revmodphys.90.015002.
(pages 52 and 58)

[Aro98] S. Arora.
The approximability of NP-hard problems.
In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, page 337–348, New York, NY, USA, 1998. Association for Computing Machinery.
doi:10.1145/276698.276784.
(pages xxii and 1)

[ASM+23] R. S. Andrist, M. J. A. Schuetz, P. Minssen, R. Yalovetzky, S. Chakrabarti, D. Herman, N. Kumar, G. Salton, R. Shaydulin, Y. Sun, M. Pistoia, and H. G. Katzgraber.
Hardness of the maximum-independent-set problem on unit-disk graphs and prospects for quantum speedups.
Physical Review Research, 5(4), Dec. 2023.
doi:10.1103/physrevresearch.5.043277.
(pages 26 and 58)

[AvS98] P. K. Agarwal, M. van Kreveld, and S. Suri.
Label placement by maximum independent set in rectangles.
Computational Geometry, 11(3):209–218, 1998.
doi:https://doi.org/10,1016/S0925-7721(98)00028-5.
(pages xxii, 2, 10, and 30)

[AXB+24] A. Ali, H. Xu, W. Bernoudy, A. Nocera, A. D. King, and A. Banerjee.
Quantum quench dynamics of geometrically frustrated ising models.
Nature Communications, 15(1), Dec. 2024.
doi:10.1038/s41467-024-54701-4.
(pages xxiii and 2)

[AZ18] A. Atminas and V. Zamaraev.
On forbidden induced subgraphs for unit disk graphs.
Discrete Comput. Geom., 60(1):57–97, July 2018.
doi:10.1007/s00454-018-9968-1.
(page 53)

[BBA26] A. Braida, E. Bermot, and S. Apers.
Log-concavity and tunneling: adiabatic quantum optimization for convex functions (with a spike), 2026, 2606.23614.
URL https://arxiv.org/abs/2606.23614.
(pages 8 and 57)

[BCC+25] A. Braida, S. Chakraborty, A. Chaudhuri, J. Cunningham, R. Menavlikar, L. Novo, and J. Roland.
Unstructured Adiabatic Quantum Optimization: Optimality with Limitations.
Quantum, 9:1790, July 2025.
doi:10.22331/q-2025-07-11-1790.
(page 59)

[BdLL+16] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys.
An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays.
Science, 354(6315):1021–1023, nov 2016.
doi:10.1126/science.aah3778.
(page 4)

[Ben80] P. Benioff.
The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines.
J. Stat. Phys., 22(5):563–591, May 1980.
(pages xxii and 2)

[Ber20] T. Bergamaschi.
Simulated quantum annealing is efficient on the spike hamiltonian.
arXiv preprint arXiv:2011.15094, 2020.
(pages 6, 59, and 88)

[BF28] M. Born and V. Fock.
Beweis des Adiabatensatzes.
Zeitschrift fur Physik, 51(3-4):165–180, Mar. 1928.
doi:10.1007/BF01343193.
(page 3)

[BJ25] M. Born and P. Jordan.

Zur quantenmechanik.
Eur. Phys. J. A, 34(1):858–888, Dec. 1925.
(page 3)

[BL76] H. J. Brascamp and E. H. Lieb.
On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation.
Journal of functional analysis, 22(4):366–389, 1976.
(pages xxv, 7, 60, 61, 70, and 71)

[BL20] A. Browaeys and T. Lahaye.
Many-body physics with individually controlled Rydberg atoms.
Nature Physics, 16(2):132–142, Feb 2020.
doi:10.1038/s41567-019-0733-z.
(page 18)

[BMM22] S. G. Bobkov, A. Marsiglietti, and J. Melbourne.
Concentration functions and entropy bounds for discrete log-concave distributions.
Combinatorics, Probability and Computing, 31(1):54–72, 2022.
(page 68)

[BN14] O. Barndorff-Nielsen.
Information and exponential families: in statistical theory.
John Wiley & Sons, 2014.
(pages 70 and 90)

[BS12] V. Bapst and G. Semerjian.
On quantum mean-field models and their quantum annealing.
Journal of Statistical Mechanics: Theory and Experiment, 2012(06):P06007, June 2012.
doi:10.1088/1742-5468/2012/06/p06007.
(page 89)

[BVC+13] L. Beguin, A. Vernier, R. Chicireanu, T. Lahaye, and A. Browaeys.
Direct measurement of the van der Waals interaction between two Rydberg atoms.
Physical review letters, 110(26):263201, 2013.
(page 18)

[BVC+25] E. Bermot, L. Valor, W. Coelho, L.-P. Henry, and N. Pearson.
Local Rydberg blockade regimes for disk graph embedding and quantum optimization, 2025, 2506.13228.
URL https://arxiv.org/abs/2506.13228.
(pages 8 and 25)

[BvD16] L. T. Brady and W. van Dam.
Spectral-gap analysis for efficient tunneling in quantum adiabatic optimization.
Physical Review A, 94(3):032309, 2016.
(pages 6, 59, 60, and 88)

[CCJ90] B. N. Clark, C. J. Colbourn, and D. S. Johnson.
Unit disk graphs.
Discrete Mathematics, 86(1):165–177, 1990.
doi:https://doi.org/10.1016/0012-365X(90)90358-O.
(pages 5, 12, and 31)

[CCL+23] M. Cain, S. Chattopadhyay, J.-G. Liu, R. Samajdar, H. Pichler, and M. D. Lukin.
Quantum speedup for combinatorial optimization with flat energy landscapes, 2023, 2306.13123.
(pages 25, 26, and 58)

[CFH+25] P. Cazals, A. François, L. Henriet, L. Leclerc, M. Marin, Y. Naghmouchi, W. da Silva Coelho, F. Sikora, V. Vitale, R. Watrigant, M. W. Garzillo, and C. Dalyac.
Identifying hard native instances for the maximum independent set problem on neutral atoms quantum processors, 2025, 2502.04291.
URL https://arxiv.org/abs/2502.04291.

(pages 12, 23, 26, and 58)

[CH16] E. Crosson and A. W. Harrow.
Simulated quantum annealing can be exponentially faster than classical simulated annealing.
In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 714–723. IEEE, 2016.
(page 88)

[Che22] S. Chewi.
Log-concave sampling.
Book draft available at https://chewisinho.github.io/, 2022.
(pages 6 and 60)

[CPGW22] T. Cubitt, D. Perez-Garcia, and M. M. Wolf.
Undecidability of the spectral gap.
Forum of Mathematics, Pi, 10, 2022.
doi:10.1017/fmp.2021.15.
(pages xxiv and 3)

[Dal23] A. J. Daley.
Twenty-five years of analogue quantum simulation.
Nature Reviews Physics, 5(12):702–703, Nov. 2023.
doi:10.1038/s42254-023-00666-0.
(pages xxiii and 2)

[DBK+22] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller.
Practical quantum advantage in quantum simulation.
Nature, 607(7920):667–676, July 2022.
(pages xxiii and 2)

[DC08] A. Das and B. K. Chakrabarti.
Colloquium: Quantum annealing and analog quantum computation.
Reviews of Modern Physics, 80(3):1061–1081, Sept. 2008.
doi:10.1103/revmodphys.80.1061.
(page 52)

[DDK+15] G. K. Das, M. De, S. Kolay, S. C. Nandy, and S. Sur-Kolay.
Approximation algorithms for maximum independent set of a unit disk graph.
Information Processing Letters, 115(3):439–446, 2015.
doi:https://doi.org/10.1016/j.ipl.2014.11.002.
(page 31)

[dFACB88] D. de Falco, B. Apolloni, and N. Cesa-Bianchi.
A numerical implementation of quantum annealing.
07 1988.
(page 4)

[DFJ02] M. Dyer, A. Frieze, and M. Jerrum.
On counting independent sets in sparse graphs.
SIAM Journal on Computing, 31(5):1527–1541, 2002, https://doi.org/10.1137/S0097539701383844.
doi:10.1137/S0097539701383844.
(pages 58 and 59)

[DH22] C. Dalyac and L. Henriet.
Embedding the MIS problem for non-local graphs with bounded degree using 3D arrays of atoms, 2022.
doi:10.48550/ARXIV.2209.05164.
(page 5)

[EH12] A. Elgart and G. A. Hagedorn.
A note on the switching adiabatic theorem.
Journal of Mathematical Physics, 53(10), 2012.
doi:10.1063/1.4748968.

(page 17)

[EJS05] T. Erlebach, K. Jansen, and E. Seidel.
Polynomial-time approximation schemes for geometric intersection graphs.
SIAM Journal on Computing, 34(6):1302–1323, 2005, https://doi.org/10.1137/S0097539702402676.
doi:10.1137/S0097539702402676.
(page 31)

[EJvL05] D. K. E. J. van Leeuwen.
Graph-Theoretic Concepts in Com- puter Science.
Springer Berlin Heidel- berg, Berlin, Heidelberg), 2005.
(page 5)

[EKC+22] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler, S.-T. Wang, M. Greiner, V. Vuletic, and M. D. Lukin.
Quantum optimization of maximum independent set using Rydberg atom arrays.
Science, 0(0):eabo6587, 2022, https://www.science.org/doi/pdf/10.1126/science.abo6587.
doi:10.1126/science.abo6587.
(pages 5, 25, 26, and 58)

[EPR35] A. Einstein, B. Podolsky, and N. Rosen.
Can quantum-mechanical description of physical reality be considered complete?
Physical Review, 47(10):777–780, May 1935.
doi:10.1103/PhysRev.47.777.
(page 3)

[Erc13] K. Erciyes.
Maximal Independent Sets, pages 135–155.
Springer London, London, 2013.
doi:10.1007/978-1-4471-5173-9_10.
(pages xxii and 1)

[EWL+21] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin.
Quantum phases of matter on a 256-atom programmable quantum simulator.
Nature, 595(7866):227–232, 2021.
doi:10.1038/s41586-021-03582-4.
(pages xxiii and 2)

[Fei04] U. Feige.
Approximating Maximum Clique by Removing Subgraphs.
SIAM J. Discrete Math., 18:219–225, 10 2004.
doi:10.1137/S089548010240415X.
(page 31)

[Fey82] R. P. Feynman.
Simulating physics with computers.
International Journal of Theoretical Physics, 21(6-7):467–488, jun 1982.
doi:10.1007/bf02650179.
(pages xxii, xxiii, 2, and 3)

[FGG+01] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda.
A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem.
Science, 292(5516):472–475, Apr. 2001.
doi:10.1126/science.1057726.
(pages xxiii, xxiv, 2, 3, 52, and 59)

[FGG02] E. Farhi, J. Goldstone, and S. Gutmann.
Quantum adiabatic evolution algorithms versus simulated annealing.
arXiv preprint quant-ph/0201031, 2002.
(pages 6, 59, 87, and 88)

[FGG14] E. Farhi, J. Goldstone, and S. Gutmann.
A quantum approximate optimization algorithm, 2014, 1411.4028.
URL https://arxiv.org/abs/1411.4028.
(page 21)

[FGGS00] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser.
Quantum computation by adiabatic evolution.
arXiv:quant-ph/0001106., 2000.
(page 3)

[FGS+94] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll.
Quantum annealing: A new method for minimizing multidimensional functions.
Chemical physics letters, 219(5-6):343–348, 1994.
(page 59)

[FSS04] P. Fendley, K. Sengupta, and S. Sachdev.
Competing density-wave orders in a one-dimensional hard-boson model.
Phys. Rev. B, 69:075106, Feb 2004.
doi:10.1103/PhysRevB.69.075106.
(pages xxiv and 7)

[GAN14] I. M. Georgescu, S. Ashhab, and F. Nori.
Quantum simulation.
Rev. Mod. Phys., 86:153–185, Mar 2014.
doi:10.1103/RevModPhys.86.153.
(pages xxiii and 2)

[GHV21] A. Gilyén, M. B. Hastings, and U. Vazirani.
(Sub)exponential advantage of adiabatic quantum computation with no sign problem.
In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1357–1369, 2021.
(page 59)

[GJ78] M. R. Garey and D. S. Johnson.
Strong NP-completeness results: Motivation, examples, and implications.
Journal of the ACM (JACM), 25(3):499–508, 1978.
(page 11)

[GLM+23] N. Gozlan, X.-M. Li, M. Madiman, C. Roberto, and P.-M. Samson.
Log-Hessian and deviation bounds for Markov semi-groups, and regularization effect in L^1.
Potential Analysis, 58(1):123–158, 2023.
(page 71)

[GMW+09] A. Gaëtan, Y. Miroshnychenko, T. Wilk, A. Chotia, M. Viteau, D. Comparat, P. Pillet, A. Browaeys, and P. Grangier.
Observation of collective excitation of two individual atoms in the Rydberg blockade regime.
Nature Physics, 5(2):115–118, 2009.
doi:https://doi.org/10.1038/nphys1183.
(page 34)

[GP10] M. Gibson and I. A. Pirwani.
Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier.
In M. de Berg and U. Meyer, editors, *Algorithms – ESA 2010*, pages 243–254, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
(page 31)

[Gro96] L. K. Grover.
A fast quantum mechanical algorithm for database search.
In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
doi:10.1145/237814.237866.
(pages xxiii and 2)

[Hås99] J. Håstad.
Clique is hard to approximate within $n^{1-\epsilon}$.
Acta Mathematica, 182(1):105 – 142, 1999.
doi:10.1007/BF02392825.
(pages 5 and 12)

[HBS+20] L. Henriet, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak.
Quantum computing with neutral atoms.
Quantum, 4:327, Sept. 2020.
doi:10.22331/q-2020-09-21-327.
(pages 5 and 18)

[Hei25] W. Heisenberg.
Uber quantentheoretische umdeutung kinematischer und mechanischer beziehungen.
Eur. Phys. J. A, 33(1):879–893, Dec. 1925.
(page 3)

[HF13] M. B. Hastings and M. H. Freedman.
Obstructions to classically simulating the quantum adiabatic algorithm.
Quantum Information and Computation, 13(11&12):1038–1076, 2013.
(page 59)

[HM85] D. S. Hochbaum and W. Maass.
Approximation schemes for covering and packing problems in image processing and vlsi.
Journal of the ACM, 32(1):130–136, Jan. 1985.
doi:10.1145/2455.214106.
(page 12)

[HS95] M. Huson and A. Sen.
Broadcast scheduling algorithms for radio networks.
In *Proceedings - IEEE Military Communications Conference MILCOM*, volume 2, pages 647–651. IEEE, 1995.
Proceedings of the 1995 Military Communications Conference (MILCOM). Part 1 (of 3) ; Conference date: 05-11-1995 Through 08-11-1995.
(pages xxii, 2, 10, and 30)

[Inc24] W. R. Inc.
Mathematica, Version 14.2, 2024.
URL https://www.wolfram.com/mathematica.
Champaign, IL, 2024.
(page 38)

[JCZ+00] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin.
Fast quantum gates for neutral atoms.
Phys. Rev. Lett., 85:2208–2211, Sep 2000.
doi:10.1103/PhysRevLett.85.2208.
(page 36)

[JG06] O. Johnson and C. Goldschmidt.
Preservation of log-concavity on summation.
ESAIM: Probability and Statistics, 10:206–215, Apr. 2006.
doi:10.1051/ps:2006008.
(page 67)

[JJ14] M. Jarret and S. P. Jordan.
The fundamental gap for a class of Schrödinger operators on path and hypercube graphs.
Journal of Mathematical Physics, 55(5), May 2014.
doi:10.1063/1.4878120.
(pages 61, 81, 82, and 85)

[JJ15] M. Jarret and S. P. Jordan.

Adiabatic optimization without local minima.
Quantum Information and Computation, 15(3–4):181–199, Mar. 2015.
(page 90)

[Joh07] O. Johnson.
Log-concavity and the maximum entropy property of the Poisson distribution.
Stochastic Processes and their Applications, 117(6):791–802, 2007.
doi:https://doi.org/10.1016/j.spa.2006.10.006.
(pages 67 and 71)

[Kat50] T. Kato.
On the adiabatic theorem of quantum mechanics.
Journal of the Physical Society of Japan, 5(6):435–439, 1950, https://doi.org/10.1143/JPSJ.5.435.
doi:10.1143/JPSJ.5.435.
(page 3)

[KBR+21] A. D. King, C. D. Batista, J. Raymond, T. Lanting, I. Ozfidan, G. Poulin-Lamarre, H. Zhang, and M. H. Amin.
Quantum annealing simulation of out-of-equilibrium magnetization in a spin-chain compound.
PRX Quantum, 2:030317, Jul 2021.
doi:10.1103/PRXQuantum.2.030317.
(pages xxiii and 2)

[KC17] L. Kong and E. Crosson.
The performance of the quantum adiabatic algorithm on spike hamiltonians.
International Journal of Quantum Information, 15(02):1750011, 2017.
(pages 6, 59, 60, and 88)

[KG71] J. Keilson and H. Gerber.
Some results for discrete unimodality.
Journal of the American Statistical Association, 66(334):386–389, 1971, https://www.tandfonline.com/doi/pdf/10.1080/01621459.1971.10482273.
doi:10.1080/01621459.1971.10482273.
(page 68)

[KGV83] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi.
Optimization by simulated annealing.
Science, 220(4598):671–680, 1983, https://www.science.org/doi/pdf/10.1126/science.220.4598.671.
doi:10.1126/science.220.4598.671.
(pages 6 and 11)

[KKO+20] P. Kairys, A. D. King, I. Ozfidan, K. Boothby, J. Raymond, A. Banerjee, and T. S. Humble.
Simulating the shastry-sutherland ising model using quantum annealing.
PRX Quantum, 1:020320, Dec 2020.
doi:10.1103/PRXQuantum.1.020320.
(pages xxiii and 2)

[KL19] B. Klartag and J. Lehec.
Poisson processes and a log-concave Bernstein theorem.
Studia Mathematica, 247(1), 2019.
(page 90)

[KN98] T. Kadowaki and H. Nishimori.
Quantum annealing in the transverse ising model.
Physical Review E, 58(5):5355–5363, Nov. 1998.
doi:10.1103/physreve.58.5355.
(pages xxiii, 2, and 59)

[Lan32] L. D. Landau.
A theory of energy transfer. 2.
Phys. Z. Sowjetunion, 2, 1932.

doi:10.1016/B978-0-08-010586-4.50014-6.
(page 3)

[Lee25] E. Lee.
A quantum analogue of convex optimization.
arXiv preprint arXiv:2510.02151, 2025.
(page 59)

[LFC+01] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller.
Dipole blockade and quantum information processing in mesoscopic atomic ensembles.
Phys. Rev. Lett., 87:037901, Jun 2001.
doi:10.1103/PhysRevLett.87.037901.
(page 36)

[Lig97] T. M. Liggett.
Ultra logconcave sequences and negative dependence.
Journal of Combinatorial Theory, Series A, 79(2):315–325, 1997.
(page 68)

[LJFF+26] L. Leclerc, S. Julià-Farré, G. S. Freitas, G. Villaret, B. Albrecht, L. Béguin, L. Bourachot, C. Briosne-Frejaville, D. Claveau, A. Cornillot, J. de Hond, D. Diallo, C. Dupays, R. Dupont, T. Eritzpokhoff, E. Gottlob, L. Henriet, M. Kaicher, L. Lassablière, A. Lindberg, Y. Machu, H. Mamann, T. Pansiot, J. Ripoll, E. S. Choi, A. Signoles, J. Vovrosh, B. Ximenez, V. Zapf, S. Zhang, H. Zhou, M. Lee, T. Mendes-Santos, C. Dalyac, A. Browaeys, and A. Dauphin.
One-to-one quantum simulation of a frustrated magnet with 256 qubits, 2026, 2603.20372.
URL https://arxiv.org/abs/2603.20372.
(pages xxiii and 2)

[LK12] I. Lesanovsky and H. Katsura.
Interacting Fibonacci anyons in a Rydberg gas.
Phys. Rev. A, 86:041601(R), Oct 2012.
doi:10.1103/PhysRevA.86.041601.
(pages xxiv and 7)

[LSL23] Y. Liu, W. J. Su, and T. Li.
On quantum speedups for nonconvex optimization via quantum tunneling walks.
Quantum, 7:1030, 2023.
(page 60)

[LV97] M. Luby and E. Vigoda.
Approximately counting up to four (extended abstract).
In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '97, page 682–687, New York, NY, USA, 1997. Association for Computing Machinery.
doi:10.1145/258533.258663.
(page 58)

[LW66] E. L. Lawler and D. E. Wood.
Branch-and-bound methods: A survey.
Operations Research, 14(4):699–719, Aug. 1966.
doi:10.1287/opre.14.4.699.
(page 11)

[MAL16] S. Muthukrishnan, T. Albash, and D. A. Lidar.
Tunneling and speedup in quantum optimization for permutation-symmetric problems.
Physical Review X, 6(3), July 2016.
doi:10.1103/physrevx.6.031010.
(pages 6, 59, and 60)

[MBL25] H. Munoz-Bauza and D. Lidar.
Scaling advantage in approximate optimization with quantum annealing.
Physical Review Letters, 134(16), Apr. 2025.

doi:10.1103/physrevlett.134.160601.
(pages 27, 33, and 58)

[MC21] Z. Mlika and S. Cherkaoui.
Massive iot access with noma in 5g networks and beyond using online competitiveness and learning.
IEEE Internet of Things Journal, 8(17):13624–13639, 2021.
doi:10.1109/jiot.2021.3068061.
(pages xxii and 1)

[MCJ+17] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner.
A cold-atom fermi–hubbard antiferromagnet.
Nature, 545(7655):462–466, May 2017.
doi:10.1038/nature22362.
(pages xxiii and 2)

[Mic99] L. Miclo.
An example of application of discrete Hardy's inequalities.
Markov Processes and Related Fields, 5(3):319–330, 1999.
(pages 61, 77, and 78)

[MLE+25] T. Manovitz, S. H. Li, S. Ebadi, R. Samajdar, A. A. Geim, S. J. Evered, D. Bluvstein, H. Zhou, N. U. Koyluoglu, J. Feldmeier, P. E. Dolgirev, N. Maskara, M. Kalinowski, S. Sachdev, D. A. Huse, M. Greiner, V. Vuletić, and M. D. Lukin.
Quantum coarsening and collective dynamics on a programmable simulator.
Nature, 638(8049):86–92, Feb. 2025.
doi:10.1038/s41586-024-08353-5.
(pages xxiii and 2)

[MM14] C. McDiarmid and T. Müller.
The number of disk graphs.
European Journal of Combinatorics, 35:413–431, 2014.
(page 6)

[Moo91] R. E. Moore.
Global optimization to prescribed accuracy.
Computers & Mathematics with Applications, 21(6):25–39, 1991.
doi:https://doi.org/10.1016/0898-1221(91)90158-Z.
(page 11)

[MP25] A. Marsiglietti and P. Pandey.
A note on statistical distances for discrete log-concave measures.
Statistics & Probability Letters, 216:110257, 2025.
(page 68)

[MZ24] A. Montanaro and L. Zhou.
Quantum speedups in solving near-symmetric optimization problems by low-depth QAOA.
arXiv preprint arXiv:2411.04979, 2024.
(page 90)

[NC11] M. A. Nielsen and I. L. Chuang.
Quantum Computation and Quantum Information: 10th Anniversary Edition.
Cambridge University Press, New York, NY, USA, 10th edition, 2011.
doi:10.1017/CBO9780511976667.
(page 12)

[NHK05] T. Nieberg, J. Hurink, and W. Kern.
A robust ptas for maximum weight independent sets in unit disk graphs.
In *Graph-Theoretic Concepts in Computer Science: 30th International Workshop, WG 2004, Bad Honnef, Germany, June 21-23, 2004. Revised Papers 30*, pages 214–221. Springer, 2005.
(page 12)

[NHRDT24] P. Narasimhan, S. Humeniuk, A. Roy, and V. Drouin-Touchette.
Simulating the transverse-field ising model on the kagome lattice using a programmable quantum annealer.
Phys. Rev. B, 110:054432, Aug 2024.
doi:10.1103/PhysRevB.110.054432.
(pages xxiii and 2)

[PKLV25] E. Pavarini, E. Koch, A. Lichtenstein, and D. Vollhardt.
Understanding Correlated Materials with DMFT, volume correl25.
Forschungzentrum Jülich, 2025.
doi:10.34734/FZJ-2025-03722.
(page 4)

[Pla72] M. Planck.
Planck's Original papers in quantum physics.
Classic papers in physics ; v. 1. Taylor and Francis ; Halsted Press Division Wiley, London : New York, german and english ed. edition, 1972.
(page 3)

[PWZ+18] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin.
Quantum Optimization for Maximum Independent Set Using Rydberg Atom Arrays, 2018, 1808.10816.
URL https://arxiv.org/abs/1808.10816.
(page 5)

[PY91] C. H. Papadimitriou and M. Yannakakis.
Optimization, approximation, and complexity classes.
Journal of Computer and System Sciences, 43(3):425–440, 1991.
doi:https://doi.org/10.1016/0022-0000(91)90023-X.
(page 5)

[RC02] J. Roland and N. J. Cerf.
Quantum search by local adiabatic evolution.
Physical Review A, 65(4):042308, 2002.
(page 59)

[Rei04] B. W. Reichardt.
The quantum adiabatic optimization algorithm and local minima.
In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 502–510, 2004.
(pages xxv, 6, 7, 16, 59, 60, 62, 63, 64, 65, 66, 67, 82, 87, and 88)

[Sch26] E. Schrödinger.
An undulatory theory of the mechanics of atoms and molecules.
Phys. Rev., 28:1049–1070, Dec 1926.
doi:10.1103/PhysRev.28.1049.
(page 3)

[Sho94] P. W. Shor.
Algorithms for quantum computation: discrete logarithms and factoring.
In *Proceedings 35th annual symposium on foundations of computer science*, pages 124–134. Ieee, 1994.
doi:10.1109/sfcs.1994.365700.
(pages xxiii and 2)

[SLK+21] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin.
Probing topological spin liquids on a programmable quantum simulator.
Science, 374(6572):1242–1247, Dec. 2021.
doi:10.1126/science.abi8794.
(pages xxiii and 2)

[Sly10] A. Sly.
Computational transition at the uniqueness threshold.

In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, page 287–296, USA, 2010. IEEE Computer Society.
doi:10.1109/FOCS.2010.34.
(page 58)

[ŠPAW17] N. Šibalić, J. Pritchard, C. Adams, and K. Weatherill.
ARC: An open-source library for calculating properties of alkali Rydberg atoms.
Computer Physics Communications, 220:319–331, Nov. 2017.
doi:10.1016/j.cpc.2017.06.015.
(page 18)

[SSW+21] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys.
Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms.
Nature, 595(7866):233–238, jul 2021.
doi:10.1038/s41586-021-03585-1.
(pages xxiii, 2, 4, and 22)

[TM79a] W. P. Thurston and J. W. Milnor.
The Geometry and Topology of Three-Manifolds.
Princeton University Department of Mathematics, 1979.
URL https://slmath.org.
Princeton Lecture Notes. Chapter 7 by John W. Milnor.
(page 31)

[TM79b] W. P. Thurston and J. W. Milnor.
The geometry and topology of three-manifolds, 1979.
(page 31)

[UJH+09] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman.
Observation of Rydberg blockade between two atoms.
Nature Physics, 5(2):110–114, Jan. 2009.
doi:10.1038/nphys1178.
(page 34)

[vL05] E. J. van Leeuwen.
Approximation Algorithms for Unit Disk Graphs, page 351–361.
Springer Berlin Heidelberg, 2005.
doi:10.1007/11604686_31.
(page 12)

[vL06] E. J. van Leeuwen.
Better approximation schemes for disk graphs.
In *Algorithm Theory–SWAT 2006: 10th Scandinavian Workshop on Algorithm Theory, Riga, Latvia, July 6-8, 2006. Proceedings 10*, pages 316–327. Springer, 2006.
(pages 6 and 31)

[vLVL+06] E. J. van Leeuwen, J. Van Leeuwen, et al.
On the representation of disk graphs.
UU WINFI Informatica en Informatiekunde, 2006.
(page 29)

[WDJ+06] W. Wu, H. Du, X. Jia, Y. Li, and S. C.-H. Huang.
Minimum connected dominating sets and maximal independent sets in unit disk graphs.
Theoretical Computer Science, 352(1):1–7, 2006.
doi:https://doi.org/10.1016/j.tcs.2005.08.037.
(pages 10 and 30)

[Wei06] D. Weitz.
Counting independent sets up to the tree threshold.

In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 140–149, New York, NY, USA, 2006. Association for Computing Machinery.
doi:10.1145/1132516.1132538.
(page 58)

[YZ86] Q. Yu and J. Q. Zhong.
Lower bounds of the gap between the first and second eigenvalues of the Schrödinger operator.
Transactions of the American Mathematical Society, 294:341–349, 1986.
URL https://api.semanticscholar.org/CorpusID:119924254.
(page 70)

[Zen32] C. Zener.
Non-Adiabatic Crossing of Energy Levels.
Proceedings of the Royal Society of London Series A, 137(833):696–702, Sept. 1932.
doi:10.1098/rspa.1932.0165.
(page 3)

[ZWWZ17] J. Zhou, L. Wang, W. Wang, and Q. Zhou.
Efficient graph-based resource allocation scheme using maximal independent set for randomly- deployed small star networks.
Sensors, 17(11), 2017.
doi:10.3390/s17112553.
(pages xxii and 2)