

The bulk spectral gap is semi-decidable: a convergent family of certified upper bounds

Xiangling Xu^{1†}, Matthias Schötz², Jie Wang^{3*}, Victor Magron⁴, Igor Klep^{5,6}, Omar Fawzi⁷, Marc-Olivier Renou¹

¹Inria, CPHT, LIX, CNRS, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

²Universität Würzburg, Institute of Mathematics, Emil-Fischer-Straße 31, 97074 Würzburg, Germany

³State Key Laboratory of Mathematical Sciences, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China

⁴Université de Toulouse; LAAS-CNRS, 7 avenue du colonel Roche, F-31400 Toulouse, France

⁵University of Ljubljana, Faculty of Mathematics and Physics, Jadranska 21, 1000 Ljubljana, Slovenia

⁶University of Primorska, Faculty of Mathematics, Natural Sciences and Information Technologies, Glagoljaška 8, 6000 Koper, Slovenia

⁷Inria, ENS Lyon, UCBL, LIP, F-69342 Lyon Cedex 07, France

† xu.xiangling@inria.fr

* wangjie212@amss.ac.cn

Abstract. Determining spectral gaps in the thermodynamic limit is a central challenge in quantum many-body physics. Existing rigorous methods are largely limited to special settings, while variational numerical approaches typically provide estimates rather than certified bounds. Here we introduce a complete family of certified upper bounds on the bulk spectral gap of quantum many-body systems. These upper bounds are obtained by solving a series of semidefinite programs and they become arbitrarily tight at the cost of more computational resources. This shows that the bulk spectral gap is semi-decidable, in contrast to undecidability results for alternative notions of spectral gap based on sequences of finite systems with prescribed boundary conditions. As a proof of principle, we apply our algorithm to the spin- $\frac{1}{2}$ kagome lattice Heisenberg antiferromagnet and obtain, to our knowledge, the first nontrivial certified upper bounds on its bulk spectral gap.

The spectral gap, the energy difference between the ground state and the first excited state, is fundamental in quantum many-body physics. A nonzero gap often implies exponential decay of correlations [Has07] and phase stability [BHM10; Den+02; Sac99] while a vanishing gap is associated with criticality. In quantum computation, the spectral gap governs the runtime of adiabatic quantum algorithms [AL18].

Determining the spectral gap is notoriously difficult, and many open many-body-physics problems can be formulated as spectral gap problems. For example, the Haldane conjecture states that the antiferromagnetic Heisenberg model in 1D with integer spins is gapped [Hal83]. Rigorous bounds on the spectral gap are confined to special settings, most notably frustration-free systems, tracing back to ideas developed for models such as the AKLT chain [Aff+88; Aff+87; FNW92; Kna88; Lem19; Nac96; You23]. Outside such structured settings, however, the ground states and low-energy excitations are not known in general. In fact, from a computational point of view, the spectral gap in the thermodynamic limit is undecidable [Bau+20; CPW15].

Hence, for general settings, one usually relies on variational numerical approaches, including tensor-network methods [Orú14; Sch11], variational Monte Carlo [BS17], and neural states Ansätze [CT17]. These have achieved remarkable success, but provide estimates rather than rigorous certificates. A paradigmatic example is the spin- $\frac{1}{2}$ kagome lattice Heisenberg antiferromagnet (cf. Fig. 1), a

canonical frustrated model whose thermodynamic-limit gap remains debated despite extensive numerical work, with competing results supporting both gapped and gapless scenarios [ZGS25].

A missing central ingredient is a general method for obtaining *rigorous upper bounds* on the spectral gap. Lower bounds certify gappedness, whereas upper bounds rule out proposed finite gaps and provide a route to proving gaplessness when they can be driven to zero. Such bounds are known in structured settings, for example from symmetry-based constraints such as Lieb–Schultz–Mattis-type theorems [Has04] and from Goldstone-type arguments [Nac06]. Alternatively, in special models such as the AKLT chain and related spin-1 valence-bond-solid chains, rigorous upper bounds can also be obtained variationally from explicit low-energy trial states [AAH88; Kna88]. Outside these special cases, however, a generally applicable certified framework is still lacking.

In this letter, we introduce the *first general algorithm* for rigorously upper bounding *bulk spectral gaps* in the thermodynamic limit. It produces a hierarchy of increasingly precise, but computationally more costly, upper bounds. These upper bounds converge to the bulk spectral gap, the gap of the infinite system itself as probed by local excitations, independent of any prescribed boundary condition. The same framework can also incorporate prescribed symmetries as additional constraints, yielding certified bounds within the corresponding symmetry sector. As a proof of principle, we apply the method

to the spin- $\frac{1}{2}$ kagome lattice Heisenberg antiferromagnet. We obtain a fully general certified upper bound $2.18J$, and a sharper bound $1.15J$ in a physically motivated symmetry sector. Although these bounds are far above existing numerical estimates—from gaps of order 10^{-2} to $10^{-1}J$ to possibly gapless scenarios [DMS12; He+17; YHW11]—they are, to our knowledge, the first rigorously certified upper bounds on this bulk spectral gap.

Our algorithm is built upon semidefinite programs (SDPs), which have proved powerful for certified results in quantum information theory [NPA08; PNA10] and, more recently, many-body physics [Ara+26; Ara25; FFS24; Mor+25; Wan+26; Wan+24]. These existing applications rely on standard SDP constraints that are linear in the state. By contrast, many important many-body properties, including spectral gaps, are naturally described by nonlinear constraints and are therefore not accessible to standard SDP methods in a direct way. The frustration-free case is special: the lower-bound problem admits positivity conditions that can be relaxed into a standard SDP hierarchy [Rai+26], although completeness remains open. Our work takes a different route: for general finite-range interacting systems, we develop a new method to translate a nonlinear formulation of the spectral gap into certified SDP relaxations, and we prove that this translation yields a complete characterization of the bulk spectral gap problem in the thermodynamic limit.

More broadly, our work opens practical and theoretical directions. On the practical side, stronger structural reductions, better exploitation of sparsity, and larger-scale computations could substantially sharpen the present algorithm and turn it into a competitive tool for difficult quantum many-body models. Even when numerically larger than the best variational estimates, such upper bounds are valuable because the framework provides rigorous thermodynamic-limit certificates, rather than numerical evidence that may change with improved Ansätze or larger simulations. On the theoretical side, our results raise a natural question about lower-bound certification for bulk spectral gaps. The known undecidability results concern finite-volume gap notions defined with fixed boundary conditions [Bau+20; CPW15]. We show that our method, formulated directly in the bulk, is compatible with those constructions because it addresses a different notion of spectral gaps. This leaves open whether the bulk spectral gap problem is itself undecidable, or whether it admits a complementary certification method for lower bounds.

SCOPE. For clarity, we present the main-text construction for a one-dimensional spin- $\frac{1}{2}$ chain. The framework extends to finite-range systems on arbitrary-dimensional, non-regular lattices, including the kagome lattice, and more general systems with appropriate operator-algebraic structures; see Supplementary Information Theorem 2.6.

FINITE-VOLUME SPECTRAL GAP WITH BOUNDARY CONDITIONS. Our work focuses on the bulk spectral gap, in




Figure 1: A portion of the kagome lattice. From “Kagome-lattice-bw.svg” by WilliamSix (Wikimedia Commons), CC BY-SA 4.0.

contrast to the conventional finite-volume notion with prescribed boundary conditions that we introduce first; see Supplementary Information Section 4.1. Consider a nearest-neighbour Hamiltonian

$$H = \sum_{i \in \mathbb{Z}} h_{i,i+1}, \quad (1)$$

where each local term $h_{i,i+1}$ acts on two neighbouring sites. Since this formal infinite sum is not a well-defined operator, one instead studies truncated Hamiltonians $H_B^{\Lambda(L)}$ on intervals $\Lambda(L) := \{-L, \dots, L\}$, with boundary condition B (e.g., open or periodic).

The finite-volume gap is then read from the spectrum of $H_B^{\Lambda(L)}$, and its behavior as $L \rightarrow \infty$ is taken as the thermodynamic-limit notion of spectral gaps. One (very strong) version, used for example in [CPW15], is a *non-degenerate uniform gap*: for all sufficiently large L , the finite-volume ground state is unique and the gap above it remains bounded below by a positive constant independent of L .

This definition is generally sensitive to boundary conditions: in the AKLT chain, for example, periodic boundary conditions give a unique finite-volume ground state, while four-fold degenerate for open boundary conditions [Aff+88]. Yet, the thermodynamic limit is meant to capture bulk properties of very large systems, so one would like a notion of spectral gaps that reflects intrinsic bulk excitations rather than artifacts of how the boundary is imposed. This motivates a formulation that captures the bulk behavior directly.

BULK SPECTRAL GAP IN THE THERMODYNAMIC LIMIT. In order to access the bulk property independent of chosen boundary conditions, we take the formulation rooted in the thermodynamic limit [BR81]; see Supplementary Information Section 1 for details.

Suppose that we have a state ρ on the infinite chain in the thermodynamic limit. On the infinite chain, physically meaningful excitations are local observables, supported on

finite intervals $\Lambda(L)$. Concretely, for a spin- $\frac{1}{2}$ chain, these are generated by finite products of local Pauli operators acting on $\Lambda(L)$. The set of all such local observables forms the local algebra \mathcal{A}_{loc} . A state is then described through its expectation values on these observables,

$$\omega(a) := \text{Tr}(\rho a), \quad (2)$$

which defines a positive normalized linear functional on \mathcal{A}_{loc} , and we also identify this functional ω with the state ρ . (More precisely, this expectation-value functional ω is defined on the norm closure of \mathcal{A}_{loc} .)

Although the formal Hamiltonian $H = \sum_i h_{i,i+1}$ is not itself a well-defined finitely supported operator, its commutator with any local observable a is well defined: only those interaction terms near the support of a fail to commute with it, so $[H, a]$ reduces to a finite sum. This gives a well-defined notion of energy change by local excitations and leads to the appropriate notion of infinite-volume ground state: ω is a *Kubo–Martin–Schwinger (KMS) ground state* when

$$\omega(a^*[H, a]) \geq 0 \quad (3)$$

for all local observables $a \in \mathcal{A}_{\text{loc}}$. Physically, this says that no finitely supported excitation a can lower the energy. In finite dimensions, this is equivalent to the usual spectral notion of ground states; see Supplementary Information Theorem 1.5. In the thermodynamic limit, it provides the corresponding bulk notion of ground states, independent of any particular choice of boundary conditions.

For a KMS ground state ω , we say that ω has a *locally non-degenerate bulk spectral gap of size at least $\gamma > 0$* if

$$\omega(a^*[H, a]) \geq \gamma(\omega(a^*a) - |\omega(a)|^2) \quad (4)$$

for all local observables $a \in \mathcal{A}_{\text{loc}}$ [NSY24]. (By the Cauchy–Schwarz inequality for states, Eq. (4) in particular implies Eq. (3).) For ω , the left-hand side of Eq. (4) is again the energy increase created by the local excitation a . Local degeneracy would mean that some local operator maps the chosen ground state ω to another orthogonal ground state. For such an operator, the energy increase is zero, while $\omega(a^*a) - |\omega(a)|^2$ is positive. Hence Eq. (4) could hold only with $\gamma = 0$. The local non-degeneracy condition excludes precisely this possibility: in that case, Eq. (4) says that every local excitation acting nontrivially on the chosen ω must pay at least γ units of energy after normalization. This should be distinguished from global degeneracy between distinct symmetry-broken or topological sectors, which need not be connected by local observables; see Supplementary Information Theorem 1.8.

The *bulk spectral gap of H* is then the supremum of all γ for which there exists a KMS ground state ω satisfying Eq. (4). This gap is defined directly in the thermodynamic limit through local excitations and is independent of any prescribed boundaries. For example, finite-volume AKLT

chains with open boundary conditions have boundary-induced ground-state degeneracy, while it is gapped in the bulk [Aff+88].

A COMPLETE HIERARCHY OF RELAXATIONS. The dynamical criterion Eq. (4) is the starting point of our complete SDP hierarchy; see Supplementary Information Section 2 for the full formulation and proof.

Suppose first that there exists a KMS ground state ω with a locally non-degenerate bulk spectral gap at least γ . Then Eq. (4) must hold for every local excitation a . In particular, it must hold for all excitations supported inside any fixed finite interval $\Lambda(L)$. Thus, every finite interval $\Lambda(L)$ gives a necessary consistency check for such a gapped KMS ground state.

This leads to a hierarchy of tests: for each L , we ask whether there exists a consistent assignment of expectation values to all excitations supported on $\Lambda(L)$ satisfying Eq. (4). If not for some L , then by contraposition, no KMS ground state with a locally non-degenerate bulk gap at least γ can exist. Thus, the hierarchy is a nested family of necessary conditions for the existence of a KMS ground state with a locally non-degenerate bulk gap of at least γ .

Although each test is formulated on a finite interval $\Lambda(L)$, it is *not* a finite-volume calculation. We do not diagonalize a truncated Hamiltonian on $\Lambda(L)$, nor do we extrapolate properties of a finite system with a prescribed boundary condition. Instead, each interval $\Lambda(L)$ is only a local window to test Eq. (4). The environment outside $\Lambda(L)$ is not fixed in advance but is left completely free, so the construction effectively ranges over *all* compatible boundary conditions at once. This is why infeasibility at some finite $\Lambda(L)$ gives a certified upper bound in the thermodynamic limit, rather than a mere finite-volume estimation.

For each fixed interval $\Lambda(L)$, this test becomes a finite-dimensional feasibility problem. That is, one asks whether there exists a consistent assignment of expectation values to all local excitations on $\Lambda(L)$ such that the bulk-gap inequality is satisfied. Existing works [Ara+26; FFS24; Wan+24] already use a similar local-consistency-check idea to, e.g., certify KMS ground states with Eq. (3). In those settings, the relevant constraints are linear in the functional ω and therefore the problem can be naturally cast into standard SDP formulations. Here, by contrast, Eq. (4) contains a nonlinear term $|\omega(a)|^2$, which is the major obstacle when one seeks to treat it with SDP methods.

To overcome this difficulty, we build on state polynomial optimization [Kle+24], further extending it so that Eq. (4) becomes an SDP constraint without changing the underlying thermodynamic-limit problem. In practice, for fixed $\Lambda(L)$, checking this inequality against all local excitations is not computationally viable. Thus, the algorithm is indexed by two parameters: L , which controls the size of the local test window, and d , which controls

the class of local excitations tested within that window; see Supplementary Information Theorems 2.1 and 2.4 for a precise definition.

Our central result is that the resulting algorithm, implemented through SDPs indexed by (L, d) , is not merely a sequence of necessary tests, but a *complete* certification procedure; see Supplementary Information Theorem 2.5. Namely, infeasibility at some finite level (L, d) rigorously proves that no KMS ground state can have a locally non-degenerate bulk gap above the chosen threshold γ . Conversely, if the algorithm remains feasible at every level (L, d) , then there exists a corresponding KMS ground state in the thermodynamic limit with a locally non-degenerate bulk spectral gap of at least γ . More generally, the algorithm also bounds expectation values of local observables over KMS ground states with locally non-degenerate bulk gap at least γ . Symmetry constraints can also be imposed: the same supremum is then restricted to symmetric KMS ground states, yielding a symmetry-restricted bulk gap; see Supplementary Information Theorem 2.7.

NUMERICAL SHOWCASING. From a numerical perspective, our algorithm rigorously certifies upper bounds on the bulk spectral gap of H , or on the corresponding symmetry-restricted bulk gap when symmetries are imposed. At each (L, d) , the largest feasible value of γ is such an upper bound: infeasibility above it certifies that no KMS ground state can satisfy Eq. (4) at that threshold. As the relaxation is tightened, these certified bounds decrease monotonically and become sharper at a higher computational cost. By completeness, arbitrary precision is achievable in principle.

As a proof of principle, we first test the hierarchy on the transverse-field Ising chain with spin-flip symmetry imposed, where the resulting behavior matches the known analytical results; see Fig. 2 or Supplementary Information Section 3.1 for more details.

We then apply the method to the spin- $\frac{1}{2}$ kagome lattice Heisenberg antiferromagnet, a canonical frustrated model in the study of quantum spin liquids. For this, the thermodynamic-limit gap remains under active debate: early density matrix renormalization group (DMRG) studies reported evidence compatible with a gapped ground state, with estimated singlet and spin gaps of about $0.04\text{--}0.05J$ and $0.13J$ [DMS12; YHW11], whereas later DMRG studies on infinitely long cylinders and tensor-network approaches argued for gapless behavior [He+17; Jia+19; Lia+17]. Exact diagonalization on clusters of up to 48 sites has not resolved the issue conclusively [LSM19]. These approaches provide valuable numerical evidence, but they do not by themselves constitute certified proofs for the thermodynamic-limit bulk gap.

Within our framework, we obtain certified upper bounds on the bulk spectral gap of the KLHM in the thermodynamic limit (Fig. 3). The tightest fully general upper bound is $2.18J$. We also test two symmetry sectors:

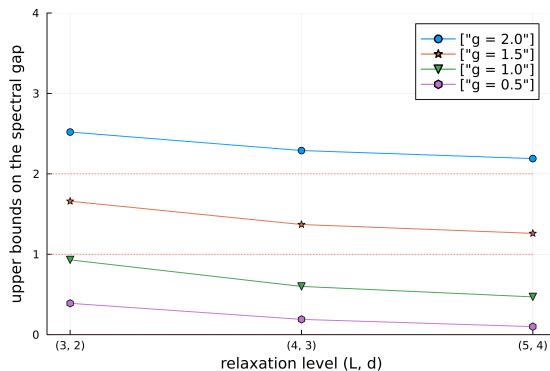


Figure 2: *Certified upper bounds for the transverse-field Ising chain with spin-flip symmetry imposed.* Largest feasible γ for $g = 0.5, 1.0, 1.5, 2.0$ along the relaxation levels $(L, d) = (n + 1, n)$, $n = 2, 3, 4$. As the relaxations are nested, these certified upper bounds decrease monotonically in L and d . In the disordered phase ($g = 1.5, 2.0$) the bounds approach the known gap value $2(g - 1)$ from above. At criticality ($g = 1.0$) the bound decays slowly, possibly reflecting divergent correlation length. For the ordered phase ($g = 0.5$), the sign-flip symmetry selects the symmetric mixed ground state whose local bulk gap vanishes, and our bounds approach to zero. Further calculations with different d are reported in Supplementary Information Table S1: the best bounds obtained for $g = 0.5, 1.0, 1.5, 2.0$ are $0.10, 0.45, 1.17, 2.11$, respectively.

one imposing only the three π -rotation spin symmetries, which allows chirality, and one additionally imposing finite spin-isotropy and time-reversal symmetry, where the tightest bound is $1.15J$. The algorithm therefore already provides rigorous thermodynamic-limit certificates beyond elementary analytic estimates and conventional finite-size or variational methods; see Supplementary Information Section 3.2 for details.

FURTHER DISCUSSIONS. Our current bounds on kagome model are proof-of-principle rather than an optimized kagome-specific computation. This points to a practical direction for future work: stronger structural reductions, better exploitation of symmetries (such as those in [Wan+26]), together with increased computational resources, may substantially sharpen the bounds and turn the method into a competitive tool for difficult models where traditional approaches remain inconclusive. Our bounds are computed using an SDP solver implemented in floating-point arithmetic and we did not take into account numerical errors. However, by extending the rigorous post-processing approach of [Nac+25] to our optimization problem, it would be possible to obtain exact rational bounds at a reasonable computational cost.

Moreover, improved certified upper bounds can have qualitative implications: if such a bound lies below the energy of a proposed excitation, then that excitation is rigorously excluded from being the lowest bulk excitation.

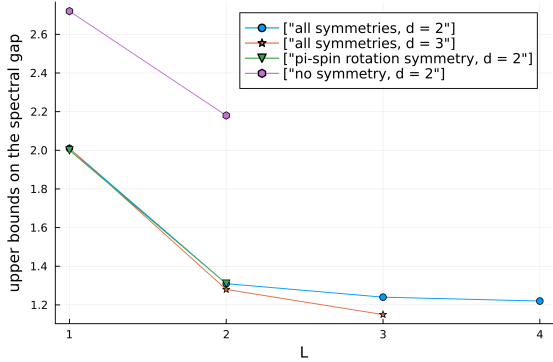


Figure 3: *Certified upper bounds for the spin- $\frac{1}{2}$ kagome lattice Heisenberg antiferromagnet.* Largest feasible gap parameter versus patch radius L , where L is the graph distance from a chosen reference site. We show four cases: the hierarchy with no symmetry restriction at $d = 2$ for $L = 1, 2$; the hierarchy with π -rotation spin symmetries at $d = 2$ for $L = 1, 2$; the same symmetries plus finite spin-isotropy and time-reversal symmetry at $d = 2$ for $L = 1, 2, 3, 4$; and the latter symmetry sector at $d = 3$ for $L = 1, 2, 3$. The corresponding kagome patches contain 5, 13, 27, and 45 sites for $L = 1, 2, 3, 4$, respectively. The tightest general certified upper bound on the bulk spectral gap is $2.18J$, while the tightest certified upper bound in the symmetry-restricted sectors is $1.15J$. Exact values and definitions are given in Supplementary Information Section 3.2 and Table S2.

Thus, even though the numerical value of a certified upper bound might be larger than traditional variational estimates, it provides rigorous information about the thermodynamic limit.

On the theoretical side, a natural complementary question is whether one can also certify *lower* bounds on bulk spectral gaps. At first sight, the known undecidability results for spectral gaps [Bau+20; CPW15] seem to obstruct such a possibility. However, they concern finite-volume gaps with prescribed boundary conditions, whereas our method addresses KMS bulk gappedness directly in the thermodynamic limit. This distinction is essential: finite-volume spectral data with fixed boundary conditions need not provide a reliable algorithmic route to deciding gappedness in the worst case scenarios. By contrast, our upper-bound hierarchy shows that the bulk formulation has a semi-decidable side: if Δ_{bulk} denotes the actual bulk gap, then every threshold strictly larger than Δ_{bulk} can be eventually ruled out by our algorithm. Thus, the two settings lead to different algorithmic questions rather than contradictory conclusions. In particular, our results imply that boundary excitations play an essential role in the construction of [CPW15]; see Supplementary Information Theorem 4.8. This leaves open whether the bulk spectral gap problem itself is undecidable, or instead admits a complementary lower-bound certification method. We refer to Supplementary Informa-

tion Sections 4.2 and 4.3 for more details. More generally, the distinction between these two gap notions illustrates a caution about mathematical idealizations of the physical world: different thermodynamic-limit idealizations can lead to genuinely different algorithmic questions.

Code availability

The code used to generate the numerical results and figures in this work is available at <https://github.com/wangjie212/SpectralGap>.

Acknowledgments

We thank Antonio Acín, Ignacio Cirac, Toby Cubitt, Hamza Fawzi, Tommaso Guaita, Lauritz van Luijk, Grégoire Misguich, James Purcell, Robert Salzmann, Lucas Tendick, Thomas Vidick, and James Watson for the helpful discussions. This work was performed within the project COMPUTE, which is funded within the QUANTERA II Programme that has received funding from the EU’s H2020 research and innovation programme under the GA No. 101017733 [1]. XX and MOR acknowledge funding by the ANR for the JCJC grants LINKS (No. ANR-23-CE47-0003) and the T-ERC QNET (No. ANR-24-ERCS-0008). JW was also supported by National Key R&D Program of China under grant No. 2023YFA1009401, Natural Science Foundation of China (NSFC) under grant No. 12571333. VM was also supported by the HORIZON-MSCA-2023-DN-JD of the European Commission under the Grant Agreement No. 101120296 (TENORS), the AI Interdisciplinary Institute ANITI funding through the French “France 2030” program under the Grant agreement No. ANR-23-IACL-0002, as well as the National Research Foundation, Prime Minister’s Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. IK also acknowledges support of the Slovenian Research Agency program P1-0222 and grants J1-50002, N1-0217, J1-60011, J1-50001, J1-3004 and J1-60025. Partially supported by the Fondation de l’École polytechnique as part of the Gaspard Monge Visiting Professor Program, IK and MS thank École polytechnique and Inria Paris Saclay for hospitality during the preparation of this manuscript. OF acknowledges funding by the European Research Council (ERC Grant AlgoQIP, Agreement No. 851716).

References

- [Aff+88] Ian Affleck, Tom Kennedy, Elliott H Lieb, and Hal Tasaki. “Valence bond ground states in isotropic quantum antiferromagnets”. In: *Communications in Mathematical Physics* 115.3 (1988), pp. 477–528.

- [Aff+87] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. “Rigorous results on valence-bond ground states in antiferromagnets”. In: *Phys. Rev. Lett.* 59 (7 Aug. 1987), pp. 799–802. DOI: 10.1103/PhysRevLett.59.799. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.59.799>.
- [AL18] Tameem Albash and Daniel A. Lidar. “Adiabatic quantum computation”. In: *Rev. Mod. Phys.* 90 (1 Jan. 2018), p. 015002. DOI: 10.1103/RevModPhys.90.015002. URL: <https://link.aps.org/doi/10.1103/RevModPhys.90.015002>.
- [AM85] Huzihiro Araki and Taku Matsui. “Ground states of the XY-model”. In: *Communications in mathematical physics* 101.2 (1985), pp. 213–245.
- [Ara+26] Mateus Araújo, Igor Klep, Andrew JP Garner, Tamás Vértesi, and Miguel Navascués. “First-order optimality conditions for non-commutative optimization problems”. In: *Foundations of Computational Mathematics* (2026). To appear.
- [Ara25] Mateus Araújo and Andrew JP Garner and Miguel Navascués. *Non-commutative optimization problems with differential constraints*. 2025. arXiv: 2408.02572 [quant-ph]. URL: <https://arxiv.org/abs/2408.02572>.
- [AAH88] Daniel P Arovas, Assa Auerbach, and FDM Haldane. “Extended Heisenberg models of antiferromagnetism: Analogies to the fractional quantum Hall effect”. In: *Physical review letters* 60.6 (1988), p. 531.
- [BDN16] Sven Bachmann, Wojciech Dybalski, and Pieter Naaijken. “Lieb–Robinson bounds, Arveson spectrum and Haag–Ruelle scattering theory for gapped quantum spin systems”. In: *Annales Henri Poincaré* 17 (2016), pp. 1737–1791.
- [BPR06] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. 2nd ed. Vol. 10. Algorithms and Computation in Mathematics. Berlin, Heidelberg: Springer, 2006. ISBN: 978-3-540-33098-1. DOI: 10.1007/3-540-33099-2. URL: <https://link.springer.com/book/10.1007/3-540-33099-2>.
- [Bau+20] Johannes Bausch, Toby S Cubitt, Angelo Lucia, and David Perez-Garcia. “Undecidability of the spectral gap in one dimension”. In: *Physical Review X* 10.3 (2020), p. 031038.
- [BS17] Federico Becca and Sandro Sorella. *Quantum Monte Carlo approaches for correlated systems*. Cambridge University Press, 2017.
- [Bla06] Bruce Blackadar. *Operator Algebras: Theory of C^* -Algebras and von Neumann Algebras*. Encyclopaedia of Mathematical Sciences. Springer Berlin Heidelberg, 2006. ISBN: 978-3-540-28517-5. DOI: 10.1007/978-3-540-28517-5.
- [BR81] Ola Bratteli and Derek W Robinson. *Operator Algebras and Quantum Statistical Mechanics II: Equilibrium States Models in Quantum Statistical Mechanics*. 1st ed. Vol. 2. Theoretical and Mathematical Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 1981. ISBN: 978-3-662-09089-3. DOI: 10.1007/978-3-662-09089-3.
- [BHM10] Sergey Bravyi, Matthew B Hastings, and Spyridon Michalakis. “Topological quantum order: stability under local perturbations”. In: *Journal of mathematical physics* 51.9 (2010).
- [CT17] Giuseppe Carleo and Matthias Troyer. “Solving the quantum many-body problem with artificial neural networks”. In: *Science* 355.6325 (2017), pp. 602–606.
- [CPW15] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. “Undecidability of the spectral gap”. In: *Nature* 528.7581 (Dec. 2015), pp. 207–211. DOI: 10.1038/nature16059. arXiv: 1502.04573 [quant-ph].
- [Den+02] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. “Topological quantum memory”. In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4452–4505.
- [DMS12] Stefan Depenbrock, Ian P McCulloch, and Ulrich Schollwöck. “Nature of the spin-liquid ground state of the $S=1/2$ Heisenberg model on the kagome lattice”. In: *Physical review letters* 109.6 (2012), p. 067201.
- [FNW92] Mark Fannes, Bruno Nachtergaele, and Reinhard F Werner. “Finitely correlated states on quantum spin chains”. In: *Communications in mathematical physics* 144 (1992), pp. 443–490.
- [FFS24] Hamza Fawzi, Omar Fawzi, and Samuel O Scalet. “Certified algorithms for equilibrium states of local quantum Hamiltonians”. In: *Nature Communications* 15.1 (2024), p. 7394.
- [Hal83] F Duncan M Haldane. “Nonlinear field theory of large-spin Heisenberg antiferromagnets: semiclassically quantized solitons of the one-dimensional easy-axis Néel state”. In: *Physical review letters* 50.15 (1983), p. 1153.
- [Has07] Matthew B Hastings. “An area law for one-dimensional quantum systems”. In: *Journal of statistical mechanics: theory and experiment* 2007.08 (2007), P08024–P08024.

- [Has04] Matthew B Hastings. “Lieb-Schultz-Mattis in higher dimensions”. In: *Physical review b* 69.10 (2004), p. 104431.
- [He+17] Yin-Chen He, Michael P. Zaletel, Masaki Oshikawa, and Frank Pollmann. “Signatures of Dirac Cones in a DMRG Study of the Kagome Heisenberg Model”. In: *Phys. Rev. X* 7 (3 July 2017), p. 031020. DOI: 10.1103/PhysRevX.7.031020. URL: <https://link.aps.org/doi/10.1103/PhysRevX.7.031020>.
- [Jia+19] Shenghan Jiang, Panjin Kim, Jung Hoon Han, and Ying Ran. “Competing Spin Liquid Phases in the $S=1/2$ Heisenberg Model on the Kagome Lattice”. In: *SciPost Physics* 7.1 (2019), p. 006.
- [Kle+24] Igor Klep, Victor Magron, Jurij Volčič, and Jie Wang. “State polynomials: positivity, optimization and nonlinear Bell inequalities”. In: *Mathematical Programming* 207.1 (2024), pp. 645–691.
- [Kna88] Stefan Knabe. “Energy gaps and elementary excitations for certain VBS-quantum antiferromagnets”. In: *Journal of statistical physics* 52.3 (1988), pp. 627–638.
- [LSM19] Andreas M. Läuchli, Julien Sudan, and Roderich Moessner. “ $S=1/2$ kagome Heisenberg antiferromagnet revisited”. In: *Phys. Rev. B* 100 (15 Oct. 2019), p. 155142. DOI: 10.1103/PhysRevB.100.155142. URL: <https://link.aps.org/doi/10.1103/PhysRevB.100.155142>.
- [Lem19] Marius Lemm. *Finite-size criteria for spectral gaps in D -dimensional quantum spin systems*. 2019. arXiv: 1902.07141 [quant-ph]. URL: <https://arxiv.org/abs/1902.07141>.
- [Lia+17] Hai-Jun Liao, Zhi-Yuan Xie, Jing Chen, Zhi-Yuan Liu, Hai-Dong Xie, Rui-Zhen Huang, Bruce Normand, and Tao Xiang. “Gapless spin-liquid ground state in the $S=1/2$ kagome antiferromagnet”. In: *Physical review letters* 118.13 (2017), p. 137202.
- [Mar08] Murray Marshall. *Positive polynomials and sums of squares*. Mathematical surveys and monographs 146. American Mathematical Society, 2008. ISBN: 978-0-8218-4402-1. DOI: 10.1007/978-0-8218-4402-1.
- [MRS24] Glen Bigan Mbeng, Angelo Russomanno, and Giuseppe E Santoro. “The quantum Ising chain for beginners”. In: *SciPost Physics Lecture Notes* (2024), p. 082.
- [Mor+25] Luke Mortimer, Donato Farina, Grazia Di Bello, David Jansen, Andreas Leitherer, Pere Mujal, and Antonio Acín. “Certifying steady-state properties of open quantum systems”. In: *Physical Review Research* 7.3 (2025), p. 033237.
- [Nac+25] Younes Naceur, Jie Wang, Victor Magron, and Antonio Acín. *Certified bounds on optimization problems in quantum theory*. 2025. arXiv: 2512.17713 [quant-ph]. URL: <https://arxiv.org/abs/2512.17713>.
- [Nac06] Bruno Nachtergaele. *Quantum Spin Systems after DLS1978*. 2006. arXiv: math-ph/0603017 [math-ph]. URL: <https://arxiv.org/abs/math-ph/0603017>.
- [Nac96] Bruno Nachtergaele. “The spectral gap for some spin chains with discrete symmetry breaking”. In: *Communications in mathematical physics* 175.3 (1996), pp. 565–606.
- [NSY24] Bruno Nachtergaele, Robert Sims, and Amanda Young. “Stability of the bulk gap for frustration-free topologically ordered quantum lattice systems”. In: *Letters in Mathematical Physics* 114.1 (2024), p. 24.
- [NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. “A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations”. In: *New Journal of Physics* 10.7 (2008), p. 073013.
- [Oga23] Yoshiko Ogata. “Classification of gapped ground state phases in quantum spin systems”. In: *International Congress of Mathematicians*. European Mathematical Society-EMS-Publishing House GmbH. 2023, pp. 4142–4161.
- [Orú14] Román Orús. “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”. In: *Annals of Physics* 349 (2014), pp. 117–158. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2014.06.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491614001596>.
- [PNA10] Stefano Pironio, Miguel Navascués, and Antonio Acín. “Convergent relaxations of polynomial optimization problems with noncommuting variables”. In: *SIAM Journal on Optimization* 20.5 (2010), pp. 2157–2180.
- [Rai+26] Kshiti Sneh Rai, Ilya Kull, Patrick Emonts, Jordi Tura, Norbert Schuch, and Flavio Bacchari. “A Hierarchy of Spectral Gap Certificates for Frustration-Free Spin Systems”. In: *Quantum* 10 (Apr. 2026), p. 2065. ISSN: 2521-327X. DOI: 10.22331/q-2026-04-13-2065.

URL: <https://doi.org/10.22331/q-2026-04-13-2065>.

- [Sac99] Subir Sachdev. “Quantum phase transitions”. In: *Physics world* 12.4 (1999), pp. 33–38.
- [Sch11] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. In: *Annals of Physics* 326.1 (2011). January 2011 Special Issue, pp. 96–192. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2010.09.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491610001752>.
- [Tak02] Masamichi Takesaki. *Theory of operator algebras. I*. Vol. 124. Encyclopaedia of Mathematical Sciences. Springer-Verlag, Berlin, 2002, pp. xx+415. ISBN: 3-540-42248-X.
- [Wan+26] Jie Wang, David Jansen, Irénée Frerot, Marc-Olivier Renou, Victor Magron, and Antonio Acín. “Scalable Ground-State Certification of Quantum Spin Systems via Structured Noncommutative Polynomial Optimization”. In: *arXiv preprint arXiv:2604.01555* (2026).
- [Wan+24] Jie Wang, Jacopo Surace, Irénée Frérot, Benoît Legat, Marc-Olivier Renou, Victor Magron, and Antonio Acín. “Certifying ground-state properties of many-body systems”. In: *Physical Review X* 14.3 (2024), p. 031006.
- [YHW11] Simeng Yan, David A Huse, and Steven R White. “Spin-liquid ground state of the $S=1/2$ kagome Heisenberg antiferromagnet”. In: *Science* 332.6034 (2011), pp. 1173–1176.
- [You23] Amanda Young. *Quantum Spin Systems*. 2023. arXiv: 2308.07848 [math-ph]. URL: <https://arxiv.org/abs/2308.07848>.
- [ZGS25] W Zhu, Shou-Shu Gong, and DN Sheng. “Quantum spin liquids in frustrated Kagome Heisenberg model”. In: *Quantum Frontiers* 4.1 (2025), p. 11.

Supplementary Information

We give an outline of this supplementary material, emphasizing where the main definitions and technical results enter. In Section 1 we introduce the operator-algebraic framework for many-body systems in the thermodynamic limit, including dynamical systems (Section 1.1) and KMS ground states (Section 1.2). We then define the state-dependent locally non-degenerate bulk gap and the corresponding system-level bulk spectral gap (Section 1.3), and record several consequences used later. In Section 2 we construct our algorithm via a SDP hierarchy: first at a high level (Section 2.1), then in the state-polynomial framework, where we prove completeness (Section 2.2). In Section 3 we use the hierarchy to compute certified upper bounds and report numerical case studies for the transverse-field Ising model (Section 3.1) and the kagome lattice Heisenberg model (Section 3.2). Finally, in Section 4 we compare bulk spectral gaps with finite-volume gap notions under fixed boundary conditions (Section 4.1), revisit the construction of [CPW15] and show that boundary excitations play an essential role there (Section 4.2), and explain why bulk-gap undecidability remains open despite [Bau+20] (Section 4.3).

1 Operator algebraic formalism for spectral gaps in the thermodynamic limit

The theory of operator algebras provides a natural framework for studying bulk behavior directly in the thermodynamic limit [BR81]; see also [You23] for a high-level introduction. In this section we recall the minimal formalism needed for our construction. After defining the quasi-local algebra and its dynamics (Section 1.1), we introduce KMS ground states and their GNS bulk Hamiltonians (Section 1.2). We then define the state-dependent locally non-degenerate bulk gap, give its equivalent dynamical criterion, identify its relationship with purity/extremity, and define the associated system-level bulk spectral gap (Section 1.3).

1.1 Thermodynamic limit and dynamical systems with operator algebras

We begin by introducing an algebraic formulation for thermodynamic limits, following [BR81, Chapter 6.2].

Let us denote by $M_q(\mathbb{C})$ the set of $q \times q$ matrices with complex entries. Let $\Lambda(L) := \{-L, \dots, L\}^D$ be the set of vertices of a D -dimensional lattice of side length $2L+1 \in \mathbb{N}$. To each site $i \in \Lambda(L)$ we associate a Hilbert space $\mathcal{H}^{(i)} \simeq \mathbb{C}^q$ and let $\mathcal{A}_i = B(\mathcal{H}^{(i)}) \simeq M_q(\mathbb{C})$ be the C^* -algebra of (bounded) observables acting on $\mathcal{H}^{(i)}$. Thanks to finite dimensionality, we may assume that \mathcal{A}_i is generated by a finite set of self-adjoint operators $\{x_j^{(i)}\} \subset \mathcal{A}_i$ and is bounded such that $M_j^{(i)} - (x_j^{(i)})^2 \succeq 0$ for some fixed constant $M_j^{(i)} > 0$. Similarly, for any finite subset $S \subset \Lambda(L)$, let $\mathcal{A}_S = \bigotimes_{i \in S} \mathcal{A}_i$ with the generating set $\{x_j^{(i)}\}_{i \in S}$. Then, the thermodynamic limit $\bigcup_L^\infty \Lambda(L)$ can be modeled by an approximately finite-dimensional C^* -algebra of quasi-local observables

$$\mathcal{A} := \overline{\mathcal{A}_{\text{loc}}}, \quad \text{where } \mathcal{A}_{\text{loc}} := \bigcup_{L \in \mathbb{N}} \mathcal{A}_{\Lambda(L)}, \quad (\text{S1})$$

and the closure is taken with respect to the operator norm.

With the observables of the system in the thermodynamic limit defined, we now introduce the dynamics. For any finite subset $S \subset \Lambda$, the interaction Φ among the sites of S is given by a self-adjoint matrix $\Phi(S) = \Phi(S)^* \in \mathcal{A}_S$. Fixing an interaction Φ , the local Hamiltonian on the sublattice $\Lambda(L)$ is given by

$$H^{\Lambda(L)} = \sum_{S \subset \Lambda(L)} \Phi(S) \in \mathcal{A}_{\Lambda(L)}, \quad (\text{S2})$$

which is well-defined as a finite sum. In our manuscript, we consider only *finite-range interactions*: there exists an $l \geq 0$ such that $\Phi(S) \neq 0$ only if $\text{diam}(S) \leq l$, where $\text{diam}(S) := \sup_{x,y \in S} d(x,y)$ denotes the diameter of S in the lattice metric.

Example 1.1. *For example, with a translationally invariant nearest neighbour interaction, the local Hamiltonian of the transverse-field Ising model on the sublattice $\{-L, \dots, L\} \subset \mathbb{Z}$ at some constant $g \geq 0$ [MRS24]*

is given by

$$H_{\text{TFIM},g}^{\Lambda(L)} = - \sum_{i=-L}^{L-1} Z_i Z_{i+1} + g \sum_{i=-L}^L X_i. \quad (\text{S3})$$

where $Z_i, X_i \in \mathcal{A}_i$ are the Pauli matrices acting on the site i .

While the thermodynamic limit of the observables of Λ is straightforward, for the family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$, taking the norm limit $L \rightarrow \infty$ inside \mathcal{A} is, however, *not* meaningful due to unboundedness. Instead, we consider the *Heisenberg dynamics* generated by the family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. For every local observable $a \in \mathcal{A}_{\text{loc}}$, the limit

$$\tau_{\mathcal{F}}^t(a) := \lim_{L \rightarrow \infty} e^{itH^{\Lambda(L)}} a e^{-itH^{\Lambda(L)}}, \quad t \in \mathbb{R}, \quad (\text{S4})$$

exists and defines a strongly-continuous one-parameter automorphism group on \mathcal{A} (thanks to the finite range interaction). Furthermore, let $\delta_{\mathcal{F}}$ be the densely defined *derivation* (or *generator*) of the dynamics $\tau_{\mathcal{F}}$, which is a closed derivation whose dense domain contains \mathcal{A}_{loc} as a core. For any local observable $a \in \mathcal{A}_{\Lambda(L)}$, it satisfies

$$\delta_{\mathcal{F}}(a) = i[H^{\Lambda(L)}, a]. \quad (\text{S5})$$

Definition 1.2. *Given the family of the local Hamiltonians $\mathcal{F} = \{H^{\Lambda(L)}\}_L$ on the lattice $\Lambda = \bigcup_{L=1}^{\infty} \Lambda(L)$, let $\tau_{\mathcal{F}}$ be the associated Heisenberg dynamics and $\delta_{\mathcal{F}}$ be its derivation. We call the tuple $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ a dynamical system induced by \mathcal{F} .*

Crucially, with the operator algebraic formalism of the thermodynamic limit of many-body systems, we directly access the bulk properties of the systems without the need to fix any boundary conditions on the finite-size approximants.

1.2 KMS ground states

We are ready to introduce a dynamical formulation of ground states, the Kubo–Martin–Schwinger (KMS) state [BR81, Chapter 5.3] at temperature 0.

Definition 1.3 (KMS ground state). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$ and let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a state. Then, ω is a KMS ground state for the family \mathcal{F} if and only if*

$$-i\omega(a^* \delta_{\mathcal{F}}(a)) \geq 0 \quad (\text{S6})$$

for all a in the domain of $\delta_{\mathcal{F}}$. In particular, for $a \in \mathcal{A}_{\text{loc}}$, Eq. (S6) simplifies to $\omega(a^*[H^{\Lambda(L)}, a]) \geq 0$ for some L . It also follows that ω is $\tau_{\mathcal{F}}$ -invariant, i.e., $\omega(\tau_{\mathcal{F}}(a)) = \omega(a)$ for all $a \in \mathcal{A}$ and $t \in \mathbb{R}$.

Note that Eq. (S6) implicitly requires $-i\omega(a^* \delta_{\mathcal{F}}(a)) \in \mathbb{R}$ for every a , so $\omega(a^* \delta_{\mathcal{F}}(a))$ is purely imaginary. Hence

$$\begin{aligned} 0 &= \omega(a^* \delta_{\mathcal{F}}(a)) + \overline{\omega(a^* \delta_{\mathcal{F}}(a))} \\ &= \omega(a^* \delta_{\mathcal{F}}(a)) + \omega(\delta_{\mathcal{F}}(a^*)a) = \omega(\delta_{\mathcal{F}}(a^*a)). \end{aligned} \quad (\text{S7})$$

Since every element of a unital $*$ -algebra is a linear combination of squares, it follows that every KMS ground state satisfies the stationary condition

$$\omega(\delta_{\mathcal{F}}(a)) = 0 \quad (\text{S8})$$

for all a in the domain of $\delta_{\mathcal{F}}$. Equivalently, ω is stationary under the dynamics, $\omega \circ \tau_{\mathcal{F}}^t = \omega$, in agreement with [BR81, Lemma 5.3.16]. Moreover,

$$-i\omega(a^* \delta_{\mathcal{F}}(a)) = -\frac{i}{2}\omega(a^* \delta_{\mathcal{F}}(a) - \delta_{\mathcal{F}}(a^*)a). \quad (\text{S9})$$

This rewriting, where the right-hand side is hermitian, will be useful later in Section 2.2.

Eq. (S6) captures the intuition that any perturbation of ω by an observable must increase the energy of the system. We note that the KMS ground state is generally not unique. In fact, the set of all KMS ground states is weak-* compact and convex, and its extremal points are also pure states on \mathcal{A} [BR81, Chapter 5.3.3]. The following proposition/definition of *bulk Hamiltonians* connects the dynamical formulation with the usual spectral definition of ground states. It relies on the Gelfand-Naimark-Segal (GNS) representation [Tak02, Theorem 9.14].

Proposition 1.4 (Bulk Hamiltonian). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a KMS ground state for \mathcal{F} and let $(\mathcal{H}_{\omega}, \pi_{\omega}, |\Omega_{\omega}\rangle)$ be its GNS representation.*

Then there exists a unique positive self-adjoint (possibly unbounded) operator H_{ω} with a dense domain in \mathcal{H}_{ω} , called the bulk Hamiltonian associated with ω , such that

$$\begin{aligned} \pi_{\omega}(\tau_{\mathcal{F}}^t(a)) &= e^{itH_{\omega}} \pi_{\omega}(a) e^{-itH_{\omega}}, \\ e^{itH_{\omega}} \pi_{\omega}(a) |\Omega_{\omega}\rangle &= \pi_{\omega}(\tau_{\mathcal{F}}^t(a)) |\Omega_{\omega}\rangle, \\ H_{\omega} |\Omega_{\omega}\rangle &= 0, \\ H_{\omega} \pi_{\omega}(a) |\Omega_{\omega}\rangle &= -i\pi_{\omega}(\delta_{\mathcal{F}}(a)) |\Omega_{\omega}\rangle, \end{aligned} \tag{S10}$$

for all $t \in \mathbb{R}$ and $a \in \mathcal{A}$. Consequently, $\text{Spec}(H_{\omega}) \subset [0, \infty)$, and $|\Omega_{\omega}\rangle$ is a ground state for H_{ω} in the usual spectral sense.

Proof. See [BR81, Proposition 5.3.19]. □

In other words, for any KMS ground state ω of $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$, there exists a unique bulk Hamiltonian H_{ω} in the GNS representation that implements the dynamics induced by the family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Moreover, the GNS cyclic state $|\Omega_{\omega}\rangle$ is the zero-energy ground state of this bulk Hamiltonian H_{ω} in the usual sense.

Remark 1.5 (Connection to finite-dimensional system). *For a finite-dimensional quantum system $M_q(\mathbb{C})$, the KMS ground state in Theorem 1.3 for a Hamiltonian $H \in M_q(\mathbb{C})$ coincides with the usual spectral ground state for H . Indeed, every state ω on $M_d(\mathbb{C})$ satisfies $\omega(\cdot) = \text{Tr}(\rho \cdot)$ for some density operator ρ . It is straightforward to check that every ground state for H (as the eigenvector corresponding to the lowest eigenvalue) satisfies Eq. (S6). Conversely, diagonalizing H with its eigenvectors $\{|E_i\rangle\}$ and choosing $a_{ij} = |E_i\rangle\langle E_j|$ forces a KMS ground state to be supported only on the minimum eigenspace of H .*

It follows that the KMS ground state can be regarded as a natural generalization of the usual definition of ground states to infinite systems.

1.3 Spectral gaps in the bulk

We now discuss the algebraic formalism of spectral gaps in the thermodynamic limit, which characterizes bulk excitations in many-body systems [Oga23, Definition 3.2].

Definition 1.6 (Locally non-degenerate bulk gap on states). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a KMS ground state for \mathcal{F} and let $(\mathcal{H}_{\omega}, \pi_{\omega}, |\Omega_{\omega}\rangle)$ be its GNS representation with the bulk Hamiltonian H_{ω} .*

We say that the system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with the state ω is locally non-degenerate bulk-gapped if there exists a constant $\gamma > 0$ such that the bulk Hamiltonian H_{ω} :

1. *is locally non-degenerate in the sense that*

$$\ker(H_{\omega}) = \text{Span}\{|\Omega_{\omega}\rangle\};$$

2. *satisfies the bulk spectral gap condition*

$$\text{Spec}(H_{\omega}) \cap (0, \gamma) = \emptyset.$$

In this case, we say the system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with the state ω has a locally non-degenerate bulk spectral gap of at least γ .

There is a dynamical characterization of locally non-degenerate bulk-gapped systems (by substituting $A = a - \omega(a)\mathbb{1}$ in [NSY24, Eq. (2.22)]), which plays a central role in Section 2.

Proposition 1.7 (Dynamical criterion for locally non-degenerate bulk gap on ground states). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a KMS ground state for \mathcal{F} and let $(\mathcal{H}_{\omega}, \pi_{\omega}, |\Omega_{\omega}\rangle)$ be its GNS representation with the bulk Hamiltonian H_{ω} .*

Then, the system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with the state ω has a locally non-degenerate bulk spectral gap of at least γ if and only if

$$-i\omega(a^*\delta_{\mathcal{F}}(a)) \geq \gamma(\omega(a^*a) - |\omega(a)|^2). \quad (\text{S11})$$

for all a in the domain of $\delta_{\mathcal{F}}$.

Note that $\omega(a^*a) - |\omega(a)|^2 \geq 0$ for any state ω and any a by the Cauchy–Schwarz inequality for states [Bla06, p. II.6.2.6] so Eq. (S11) automatically implies Eq. (S6). Thus, by the discussion following Theorem 1.3, Eq. (S11) is equivalent to the stationary condition $\omega(\delta_{\mathcal{F}}(a)) = 0$ and

$$-\frac{i}{2}\omega(a^*\delta_{\mathcal{F}}(a) - \delta_{\mathcal{F}}(a^*)a) \geq \gamma(\omega(a^*a) - |\omega(a)|^2), \quad (\text{S12})$$

for all a in the domain of $\delta_{\mathcal{F}}$. This symmetrized expression of Eq. (S12) is used later in Section 2.

Eq. (S11), equivalently $\omega(\delta_{\mathcal{F}}(a)) = 0$ and Eq. (S12), corresponds to the following physical intuition:

- (a) The left-hand side $-i\omega(a^*\delta_{\mathcal{F}}(a))$ measures the energy increase due to the excitation a . It is “energy increase” because ω satisfies Eq. (S6).
- (b) The term $\omega(a^*a) - |\omega(a)|^2$ measures the variance of the operator a . Geometrically, it represents the squared norm of the component of the vector $\pi_{\omega}(a)|\Omega_{\omega}\rangle$ that lies in the orthogonal complement $\{|\Omega_{\omega}\rangle\}^{\perp}$.
- (c) The parameter γ represents a lower bound on the “conversion rate” between distance (the right-hand-side) and energy (the left-hand-side). The inequality implies that the energy increase is proportional to how far $\pi_{\omega}(a)$ pushes the ground state $|\Omega_{\omega}\rangle$ away from 0-eigenspace. Specifically, every unit that a drives $|\Omega_{\omega}\rangle$ out of the 0-eigenspace costs at least γ units of energy.

Remark 1.8. *Note that “non-degeneracy” of a given KMS ground state ω does not refer to the same thing in the conventional physics literature. Local non-degeneracy does not exclude the global degeneracy coming from several distinct pure KMS ground states with inequivalent GNS representations and only excludes zero-energy excitations that can be reached by local operators from the chosen KMS ground state.*

A simple example is the ordered phase of the transverse-field Ising chain; see Section 3.1. In the usual physics language this phase is two-fold degenerate, corresponding to the two pure KMS ground states induced by $|\uparrow \cdots\rangle$ and $|\downarrow \cdots\rangle$. Each of these two states admits a locally non-degenerate bulk spectral gap, yet is separated by a global spin flip and is not connected by local observables. Thus, this global degeneracy should not be confused with the local non-degeneracy tested by Theorems 1.6 and 1.7.

Recall that the set of all KMS ground states is convex. We make the following observation about local non-degeneracy and pure states (i.e., the extremal points of the KMS ground state set).

Proposition 1.9. *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a KMS ground state for \mathcal{F} .*

If ω is not pure, then its bulk Hamiltonian H_{ω} is locally degenerate. Subsequently, Eq. (S11) holds for all a only when $\gamma = 0$.

Proof. Let $(\mathcal{H}_{\omega}, \pi_{\omega}, |\Omega_{\omega}\rangle)$ be the GNS representation of ω , and let H_{ω} be the corresponding bulk Hamiltonian. Since ω is not pure, there exist two distinct KMS ground state ω_1 and ω_2 , $\omega_1 \neq \omega_2$, and $s \in (0, 1)$ such that $\omega = s\omega_1 + (1 - s)\omega_2$.

Since $s\omega_1 \leq \omega$ and $(1-s)\omega_2 \leq \omega$, the Arveson's Radon–Nikodym theorem [Bla06, p. II.6.4.6] gives unique operators $T_1, T_2 \in \pi_\omega(\mathcal{A})'$, $0 \leq T_1, T_2 \leq \mathbb{1}$, such that, for all $a \in \mathcal{A}$,

$$\begin{aligned}\langle \Omega_\omega | T_1 \pi_\omega(a) | \Omega_\omega \rangle &= s\omega_1(a), \\ \langle \Omega_\omega | T_2 \pi_\omega(a) | \Omega_\omega \rangle &= (1-s)\omega_2(a).\end{aligned}\tag{S13}$$

Moreover, uniqueness gives $T_1 + T_2 = \mathbb{1}$, and neither T_1 nor T_2 is a scalar multiple of the identity since the convex decomposition is assumed to be nontrivial. Both $T_1 |\Omega_\omega\rangle, T_2 |\Omega_\omega\rangle \neq 0$ since $\omega_1, \omega_2 \neq 0$.

We now show that $T_1 |\Omega_\omega\rangle$ lies in $\ker(H_\omega)$. Recall from Theorem 1.4 that $\pi_\omega(\tau_{\mathcal{F}}^t(a)) = e^{itH_\omega} \pi_\omega(a) e^{-itH_\omega}$ and $e^{itH_\omega} |\Omega_\omega\rangle = |\Omega_\omega\rangle$. For any $a, b \in \mathcal{A}$, since ω_1 is a $\tau_{\mathcal{F}}$ -invariant KMS ground state, we compute

$$\begin{aligned}\langle \Omega_\omega | \pi_\omega(a) T_1 \pi_\omega(b) | \Omega_\omega \rangle &= s\omega_1(ab) \\ &= s\omega_1(\tau_{\mathcal{F}}^t(ab)) \\ &= s\omega_1(\tau_{\mathcal{F}}^t(a) \tau_{\mathcal{F}}^t(b)) \\ &= \langle \Omega_\omega | T_1 \pi_\omega(\tau_{\mathcal{F}}^t(a)) \pi_\omega(\tau_{\mathcal{F}}^t(b)) | \Omega_\omega \rangle \\ &= \langle \Omega_\omega | \pi_\omega(\tau_{\mathcal{F}}^t(a)) T_1 \pi_\omega(\tau_{\mathcal{F}}^t(b)) | \Omega_\omega \rangle \\ &= \langle \Omega_\omega | (e^{-itH_\omega} \pi_\omega(a) e^{itH_\omega}) T_1 (e^{-itH_\omega} \pi_\omega(b) e^{itH_\omega}) | \Omega_\omega \rangle \\ &= \langle \Omega_\omega | \pi_\omega(a) (e^{itH_\omega} T_1 e^{-itH_\omega}) \pi_\omega(b) | \Omega_\omega \rangle,\end{aligned}\tag{S14}$$

where we also use the fact that $\tau_{\mathcal{F}}$ is a $*$ -automorphism on \mathcal{A} and $T_1 \in \pi_\omega(\mathcal{A})'$. It follows from the cyclicity of $|\Omega_\omega\rangle$ that $T_1 = e^{itH_\omega} T_1 e^{-itH_\omega}$, i.e., T_1 commutes with e^{itH_ω} . Thus,

$$H_\omega T_1 |\Omega_\omega\rangle = -i \lim_{t \rightarrow 0} \frac{e^{itH_\omega} T_1 |\Omega_\omega\rangle - T_1 |\Omega_\omega\rangle}{t} = 0.\tag{S15}$$

The same argument shows that $T_2 |\Omega_\omega\rangle \in \ker(H_\omega)$ as well.

The vector $T_1 |\Omega_\omega\rangle$ is not colinear with $|\Omega_\omega\rangle$. Indeed, if $T_1 |\Omega_\omega\rangle = \lambda |\Omega_\omega\rangle$, then $s\omega_1 = \lambda\omega$. Evaluating $a = 1$ gives $\lambda = s$, and hence $\omega_1 = \omega$, contradicting the nontriviality of the convex decomposition. It follows that $\ker(H_\omega)$ is of at least dimension 2 and H_ω is locally degenerate.

For clarity, we directly prove the ‘‘subsequent’’ statement without invoking Theorem 1.7. As $\langle \Omega_\omega | T_1 | \Omega_\omega \rangle = s\omega_1(1) = s$, set

$$|\eta\rangle := \frac{T_1 |\Omega_\omega\rangle - s |\Omega_\omega\rangle}{\sqrt{\langle \Omega_\omega | T_1^2 | \Omega_\omega \rangle - s^2}}.\tag{S16}$$

Then $|\eta\rangle \perp |\Omega_\omega\rangle$, $\langle \eta | \eta \rangle = 1$, and $|\eta\rangle \in \ker(H_\omega)$ since both $T_1 |\Omega_\omega\rangle$ and $|\Omega_\omega\rangle$ are.

Therefore, since $\pi_\omega(\mathcal{A}_{\text{loc}}) |\Omega_\omega\rangle$ is a core for H_ω , there exists a sequence $\{a_k\}_k \subset \mathcal{A}_{\text{loc}}$ such that $\pi_\omega(a_k) |\Omega_\omega\rangle \rightarrow |\eta\rangle$ and $H_\omega \pi_\omega(a_k) |\Omega_\omega\rangle \rightarrow H_\omega |\eta\rangle = 0$ as $k \rightarrow \infty$. It follows from Theorem 1.4 that

$$-i\omega(a_k^* \delta_{\mathcal{F}}(a_k)) = \langle \Omega_\omega | \pi_\omega(a_k)^* H_\omega \pi_\omega(a_k) | \Omega_\omega \rangle \rightarrow \langle \eta | H_\omega | \eta \rangle = 0.\tag{S17}$$

On the other hand,

$$\omega(a_k^* a_k) - |\omega(a_k)|^2 \rightarrow \langle \eta | \eta \rangle - |\langle \Omega_\omega | \eta \rangle|^2 = 1.\tag{S18}$$

Thus, Eq. (S11) can hold only when $\gamma = 0$. \square

Let \mathcal{G}_γ denote the set of KMS ground states of $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with locally non-degenerate bulk gap at least γ . By Theorem 1.9, \mathcal{G}_γ need not be convex, even though the full set of KMS ground states is convex. On the other hand, \mathcal{G}_γ is weak- $*$ -closed. Indeed, for every fixed a , both sides of Eq. (S11) depend weak- $*$ -continuously on ω , and thus define a weak- $*$ -closed set; intersecting over all a preserves closedness. In fact, \mathcal{G}_γ is even weak- $*$ -compact as a closed subset of the weak- $*$ -compact space of all states.

So far, the discussion has been state-dependent, it concerns whether a given KMS ground state is locally non-degenerate bulk-gapped. We now define bulk spectral gaps at the level of dynamical systems.

Definition 1.10 (Bulk spectral gap). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. The bulk spectral gap Δ_{bulk} of $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ is the supremum of all $\gamma \geq 0$ for which there exists a KMS ground state ω satisfying Eq. (S11).*

To finish up the discussion of the algebraic spectral gap, we introduce the notion of bulk gaplessness on KMS ground states.

Definition 1.11 (Bulk-gapless KMS ground state). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by the local Hamiltonian family $\mathcal{F} = \{H^{\Lambda(L)}\}_L$. Let $\omega : \mathcal{A} \rightarrow \mathbb{C}$ be a KMS ground state for \mathcal{F} and let $(\mathcal{H}_{\omega}, \pi_{\omega}, |\Omega_{\omega}\rangle)$ be its GNS representation with the bulk Hamiltonian H_{ω} .*

We say that the system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with the state ω is bulk-gapless if, for every ε ,

$$\text{Spec}(H_{\omega}) \cap (0, \varepsilon) \neq \emptyset.$$

That is, the eigenvalue 0 of H_{ω} is not an isolated point of the spectrum.

2 A complete SDP hierarchy for certifying bulk spectral gaps

We now construct the SDP hierarchy used to certify upper bounds on the bulk spectral gap. The hierarchy is based on state polynomial optimization [Kle+24] and is designed to test, for a fixed threshold γ , whether there exists a KMS ground state satisfying the dynamical bulk-gap criterion of Theorem 1.7. Its central property is completeness: feasibility at every level is equivalent to the existence of such a KMS ground state, while infeasibility at some finite level rules out that threshold and hence gives a certified upper bound on the bulk spectral gap of Theorem 1.10. We first present the hierarchy at a high level in Section 2.1, then give the state-polynomial formulation and proof of completeness in Section 2.2.

2.1 The hierarchy of relaxations on a high level

Let the interaction Φ be finite-range with the interaction length l . Our goal is to turn the dynamical characterization in Theorem 1.7 into a tractable algorithm. Rather than working with the full quasi-local C^* -algebra \mathcal{A} , we restrict, for each n , to the local algebra $\mathcal{A}_{\Lambda(n)}$ of observables supported on $\Lambda(n)$, equipped with the corresponding finite-volume Hamiltonian $H^{\Lambda(n)}$. Note that we use n (rather than L as in Section 1) to label the finite regions $\Lambda(n)$, since n will also denote the corresponding level of the relaxation hierarchy.

Within $\mathcal{A}_{\Lambda(n)}$ we introduce the notion of operator complexity. Any $a \in \mathcal{A}_{\Lambda(n)}$ can be written as a noncommutative polynomial in the local generators $\{x_j^{(i)}\}_{i \in \Lambda(n)}$. We define the degree of a , denoted by $\deg(a)$, to be the minimum total degree among all such polynomial representations. Finally, let $W_{\Lambda(n)}^d \subset \mathcal{A}_{\Lambda(n)}$ denote the set of all monomials in the generators of total degree $\leq d$ and $\mathcal{A}_{\Lambda(n)}^d = \text{Span}(W_{\Lambda(n)}^d)$ the polynomials in $\mathcal{A}_{\Lambda(n)}$ of total degree $\leq d$.

In the following we assume that the degree of the local Hamiltonians is uniformly bounded, i.e., there exists a (minimal) integer $\deg(H)$ independent of the region $\Lambda(n)$ such that $H^{\Lambda(n)} \in \mathcal{A}_{\Lambda(n)}^{\deg(H)}$ for all n . This assumption is fulfilled for many physical models; for example, for all translation-invariant Hamiltonians with finite-range interactions.

Definition 2.1 (The hierarchy of relaxations). *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by a family of local Hamiltonians $\mathcal{F} = \{H^{\Lambda(n)}\}_n$ with the interaction length l . Fix a gap parameter $\gamma \geq 0$.*

We first define the feasibility hierarchy for certifying the existence of a KMS ground state with a locally non-degenerate bulk spectral gap at least γ . For $n, d \in \mathbb{N}$, with n, d sufficiently large, the relaxation on $\mathcal{A}_{\Lambda(n)}^{2d}$

asks whether there exists a linear functional $\omega_n^d : \mathcal{A}_{\Lambda(n)}^{2d} \rightarrow \mathbb{C}$ satisfying

$$\begin{aligned}
\omega_n^d(1) &= 1, & (\text{normalization}) \\
\omega_n^d(a^*a) &\geq 0, \quad \forall a \in \mathcal{A}_{\Lambda(n)}^d, & (\text{state positivity}) \\
\omega_n^d([H^{\Lambda(n)}, a]) &= 0, \quad \forall a \in \mathcal{A}_{\Lambda(n-l)}^{2d-\deg(H)}, & (\text{stationarity}) \\
\frac{1}{2} \omega_n^d(a^*[H^{\Lambda(n)}, a] - [H^{\Lambda(n)}, a^*]a) &\geq \gamma(\omega_n^d(a^*a) - |\omega_n^d(a)|^2), \quad \forall a \in \mathcal{A}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil} & (\text{bulk spectral gap})
\end{aligned} \tag{S19}$$

More generally, let $O = O^* \in \mathcal{A}$ be a local observable. We are interested in the extremal expectation values of O over \mathcal{G}_γ , the weak- $*$ -closed set of KMS ground states of $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with locally non-degenerate bulk gap at least γ :

$$\begin{aligned}
\langle O \rangle_{\min} &:= \min_{\omega \in \mathcal{G}_\gamma} \omega(O), \\
\langle O \rangle_{\max} &:= \max_{\omega \in \mathcal{G}_\gamma} \omega(O).
\end{aligned} \tag{S20}$$

Whenever n, d are sufficiently large such that $O, H^{\Lambda(n)} \in \mathcal{A}_{\Lambda(n)}^{2d}$, we define the corresponding upper and lower relaxations by optimizing O over the same feasible set:

$$\begin{aligned}
\alpha_{n,d}^{\max/\min} &= \underset{\omega_n^d : \mathcal{A}_{\Lambda(n)}^{2d} \rightarrow \mathbb{C} \text{ linear}}{\text{maximize/minimize}} \omega_n^d(O), \\
\text{s.t. } \omega_n^d(1) &= 1, & (\text{normalization}) \\
\omega_n^d(a^*a) &\geq 0, \quad \forall a \in \mathcal{A}_{\Lambda(n)}^d, & (\text{state positivity}) \\
\omega_n^d([H^{\Lambda(n)}, a]) &= 0, \quad \forall a \in \mathcal{A}_{\Lambda(n-l)}^{2d-\deg(H)}, & (\text{stationarity}) \\
\frac{1}{2} \omega_n^d(a^*[H^{\Lambda(n)}, a] - [H^{\Lambda(n)}, a^*]a) &\geq \gamma(\omega_n^d(a^*a) - |\omega_n^d(a)|^2), \quad \forall a \in \mathcal{A}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil}. & (\text{bulk spectral gap})
\end{aligned} \tag{S21}$$

By setting $O = 1$, the trivial optimization reduces to the feasibility hierarchy of Eq. (S19).

The normalization condition and state positivity constraint correspond to the definition of states on \mathcal{A} , and the stationarity constraint and bulk spectral gap condition are meant to recover Eqs. (S8) and (S12), which are equivalent to Eq. (S11). Note that while the term $|\omega_n^d(a)|^2$ in the constraints is nonlinear, this optimization problem can be cast as an SDP using the framework of state polynomial optimization [Kle+24]; see Theorem 2.4 for the full formulation.

Remark 2.2. The subscript $\Lambda(n-l)$ in $\mathcal{A}_{\Lambda(n-l)}^{2d-\deg(H)}$ in stationarity constraint and $\mathcal{A}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil}$ in the bulk spectral gap constraint is essential. Since the interaction has a range l , any such excitation a is separated from the boundary of $\Lambda(n)$ by enough room for the commutator $[H^{\Lambda(n)}, a]$ to coincide with the infinite-volume dynamic derivation $\delta_{\mathcal{F}}(a)$. Thus, the finite relaxation tests the bulk gap condition in a local window without imposing any boundary condition outside $\Lambda(n)$. Without this buffer, the constraint would involve boundary terms of the chosen finite-volume Hamiltonian, reducing the test to a finite-volume gap calculation with that boundary condition.

Remark 2.3. For fixed (n, d) , the relaxation Eq. (S19) is a finite-dimensional feasibility problem (over finitely many parameters specifying the linear functional ω_n^d on $\mathcal{A}_{\Lambda(n)}^{2d}$). Consequently, it can in principle be solved using quantifier elimination over the reals (the Tarski–Seidenberg theorem), for instance via cylindrical algebraic decomposition (CAD). However, quantifier elimination/CAD has doubly exponential worst-case complexity, and is therefore impractical except for very small instances [BPR06]. In contrast, the state polynomial optimization yields an SDP hierarchy, which is capable of handling optimization problems of Eq. (S21) beyond feasibility, and is far more effective in practice; see Section 3.

This hierarchy of relaxations admits a clear physical interpretation. If the system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ admits a

KMS ground state ω with a locally non-degenerate bulk gap of at least γ , then ω restricted to the subspace $\mathcal{A}_{\Lambda(n)}^{2d}$, denoted by ω_n^{2d} , must be a valid solution to Eqs. (S19) and (S21) for every (n, d) . Furthermore, it is straightforward to see that the optimal values of the hierarchy form a monotonic sequence as n (or d) increases.

Consequently, if the problem is found to be infeasible at any level (n, d) , it provides a rigorous certificate that the system cannot admit a bulk spectral gap of at least γ . On the other hand, feasibility at a single level (n, d) does not certify a bulk spectral gap; it only confirms the absence of gap-violating excitations due to operators in $\mathcal{A}_{\Lambda(n)}^{2d}$ rather than the full \mathcal{A} .

We will show in Theorem 2.5 that the state polynomial optimization formulation of Theorem 2.1 is complete. Furthermore, the hierarchy estimates the expectation values in the sense that $\alpha_{n,d}^{\min} \nearrow \langle O \rangle_{\min}$ and $\alpha_{n,d}^{\max} \searrow \langle O \rangle_{\max}$.

2.2 The state polynomial formulation and its completeness

We first formalize the state polynomial optimization problem corresponding to Theorem 2.1 and then prove that the resulting hierarchy is complete: feasibility at all relaxation levels is equivalent to the existence of a KMS ground state with a locally non-degenerate bulk spectral gap at least γ , and the corresponding optimization values converge to the exact extremal expectation values over such states.

Recall that $\mathcal{A}_{\text{loc}} = \bigcup_{L \in \mathbb{N}} \mathcal{A}_{\Lambda(L)}$ and $\mathcal{A} := \overline{\mathcal{A}_{\text{loc}}}$, where each matrix algebra $\mathcal{A}_i \simeq M_q(\mathbb{C})$, $i \in \Lambda(L)$, is generated by $\{x_j^{(i)}\}_j$. As \mathcal{A}_i is just matrix algebra, all elements of \mathcal{A}_i are bounded; there exists a constant $M_j^{(i)}$ such that $M_j^{(i)} - (x_j^{(i)})^2 \succeq 0$, and in particular, the square root $s_j^i := (M_j^{(i)} - (x_j^{(i)})^2)^{1/2} \in \mathcal{A}_i$ exists. (For example, with spin-1/2 chain and Pauli matrices, we have $M_j^{(i)} = 1$ and $1 - (X^{(i)})^2 = 1 - (Y^{(i)})^2 = 1 - (Z^{(i)})^2 = 0$.)

To begin with, we write the explicit state polynomial optimization formulation of Eq. (S21), following the notation of [Kle+24]. For every (noncommutative) monomial w in the local generators $\{x_j^{(i)}\}$, we introduce a formal *state symbol* $\varsigma(w)$. Heuristically, $\varsigma(w)$ represents the scalar expectation of w in the state ς .

We extend $\varsigma(\cdot)$ by linearity and involution, and impose the compatibility rules for all monomials w, v and complex numbers c

$$\varsigma(1) = 1, \quad \varsigma(w^*) = \varsigma(w)^*, \quad \varsigma(w + cv) = \varsigma(w) + c\varsigma(v), \quad (\text{S22})$$

together with commutativity of all state symbols and the convention that $\varsigma(\cdot)$ treats state symbols as scalars:

$$[\varsigma(w), \varsigma(v)] = 0, \quad \varsigma(w\varsigma(v)) = \varsigma(w)\varsigma(v). \quad (\text{S23})$$

These rules give rise to the ring of state polynomials

$$\mathcal{S} := \mathbb{C}[\varsigma(w) : w \neq 1 \text{ is a monomial in the generators}] \quad (\text{S24})$$

with an antilinear involution $*$. We then define the ring of *noncommutative state polynomials* by

$$\mathcal{P} := \mathcal{S} \otimes \mathcal{A}_{\text{loc}}, \quad (\text{S25})$$

whose elements are finite sums $\sum_r p_r a_r$ with $p_r \in \mathcal{S}$ and $a_r \in \mathcal{A}_{\text{loc}}$.

We set $\deg(\varsigma(w)) := \deg(w)$ and extend $\deg(\cdot)$ multiplicatively, so that $\deg(\varsigma(w)\varsigma(v)) = \deg(w) + \deg(v)$ and $\deg(\varsigma(w)u) = \deg(w) + \deg(u)$ for $u \in \mathcal{P}$. For each (n, d) we write $\mathcal{S}_{\Lambda(n)}^d$ and $\mathcal{P}_{\Lambda(n)}^d$ for the truncations to (noncommutative) state polynomials supported on $\Lambda(n)$ and of total degree $\leq d$. By a *monomial in \mathcal{S}* (or \mathcal{P}) we mean a product of factors $\varsigma(w)$ (and one factor w') with w (and w') monomials in \mathcal{A}_{loc} .

Definition 2.4 (State polynomial optimization for bulk spectral gap). *Fix a gap parameter $\gamma \geq 0$ and assume there is a finite maximal interaction length l . Let $n, d \in \mathbb{N}$ be large enough such that $O, H^{\Lambda(n)} \in \mathcal{A}_{\Lambda(n)}^{2d}$.*

Let $\mathcal{L}_n^d : \mathcal{S}_{\Lambda(n)}^{2d} \rightarrow \mathbb{C}$ be a linear functional. The associated moment matrix $\mathbf{M}_d(\mathcal{L}_n^d)$ is indexed by the monomials $s, t \in \mathcal{P}_{\Lambda(n)}^d$ and defined by

$$[\mathbf{M}_d(\mathcal{L}_n^d)]_{s,t} := \mathcal{L}_n^d(\varsigma(s^*t)). \quad (\text{S26})$$

The moment matrix associated with the bulk spectral gap, $\mathbf{M}_{d-\lceil \deg(H)/2 \rceil}^{\text{gap}, \gamma}(\mathcal{L}_n^d)$, is indexed by the monomials $s, t \in \mathcal{S}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil}$ and defined by

$$\left[\mathbf{M}_{d-\lceil \deg(H)/2 \rceil}^{\text{gap}, \gamma}(\mathcal{L}_n^d) \right]_{s,t} := \mathcal{L}_n^d \left(\frac{1}{2} \left(\varsigma(s^*[H^{\Lambda(n)}, t] - [H^{\Lambda(n)}, s^*]t) \right) - \gamma \left(\varsigma(s^*t) - \varsigma(s^*)\varsigma(t) \right) \right). \quad (\text{S27})$$

The state polynomial relaxation at level (n, d) for the local observable $O \in \mathcal{A}_{\Lambda(n)}^{2d}$ is defined as

$$\begin{aligned} \alpha_{n,d}^{\max/\min} &:= \underset{\mathcal{L}_d: \mathcal{S}_{\Lambda(n)}^{2d} \rightarrow \mathbb{C} \text{ linear functional}}{\text{maximize/minimize}} \quad \mathcal{L}_n^d(\varsigma(O)), \\ \text{s.t.} \quad \mathcal{L}_n^d(1) &= 1, & (\text{normalization}) \\ \mathbf{M}_d(\mathcal{L}_n^d) &\succeq 0, & (\text{positivity}) \\ \mathcal{L}_n^d(\varsigma([H^{\Lambda(n)}, w])) &= 0, \quad \forall w \in \mathcal{S}_{\Lambda(n-l)}^{2d-\deg(H)}, & (\text{stationarity}) \\ \mathbf{M}_{d-\lceil \deg(H)/2 \rceil}^{\text{gap}, \gamma}(\mathcal{L}_n^d) &\succeq 0. & (\text{bulk spectral gap}) \end{aligned} \quad (\text{S28})$$

By setting $O = 1$, the trivial optimization reduces to the hierarchy for Eq. (S19).

Let us spell out the meaning of the constraints in Eq. (S28). The positivity constraint $\mathbf{M}_d(\mathcal{L}_n^d) \succeq 0$ is equivalent to

$$\mathcal{L}_n^d(\varsigma(q^*q)) \geq 0, \quad \forall q \in \mathcal{S}_{\Lambda(n)}^d \quad (\text{S29})$$

as in [Kle+24, Lemma 6.4]. In particular, it contains conditions $\mathcal{L}_n^d(\varsigma(a^*a)) \geq 0$ for $a \in \mathcal{A}_{\Lambda(n)}^d$ as special cases, because the moment matrix $\mathbf{M}_d(\mathcal{L}_n^d)$ is indexed by $\mathcal{S}_{\Lambda(n)}^d$ rather than only by $\mathcal{S}_{\Lambda(n)}^d$. The stationarity constraint is equivalent to

$$\mathcal{L}_n^d(\varsigma([H^{\Lambda(n)}, q])) = 0, \quad \forall q \in \mathcal{S}_{\Lambda(n-l)}^{2d-\deg(H)}. \quad (\text{S30})$$

The bulk spectral gap constraint $\mathbf{M}_{d-\lceil \deg(H)/2 \rceil}^{\text{gap}, \gamma}(\mathcal{L}_n^d) \succeq 0$ is equivalent to

$$\mathcal{L}_n^d \left(\frac{1}{2} \left(\varsigma(q^*[H^{\Lambda(n)}, q] - [H^{\Lambda(n)}, q^*]q) \right) - \gamma \left(\varsigma(q^*q) - |\varsigma(q)|^2 \right) \right) \geq 0, \quad \forall q \in \mathcal{S}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil} \quad (\text{S31})$$

which is a state-polynomial formulation of the symmetrized bulk-gap criterion Eq. (S12).

Let $\mathcal{Q}_n^{d,\text{pos}}$ and $\mathcal{Q}_n^{d,\text{gap}}$ be the convex cones in $\mathcal{S}_{\Lambda(n)}^{2d}$ generated by the hermitian elements appearing in Eqs. (S29) and (S31):

$$\begin{aligned} \mathcal{Q}_n^{d,\text{pos}} &:= \text{cone} \left\{ \varsigma(q^*q) : q \in \mathcal{S}_{\Lambda(n)}^d \right\}, \\ \mathcal{Q}_n^{d,\text{gap}} &:= \text{cone} \left\{ \frac{1}{2} \left(\varsigma(q^*[H^{\Lambda(n)}, q] - [H^{\Lambda(n)}, q^*]q) \right) - \gamma \left(\varsigma(q^*q) - |\varsigma(q)|^2 \right) : q \in \mathcal{S}_{\Lambda(n-l)}^{d-\lceil \deg(H)/2 \rceil} \right\}. \end{aligned} \quad (\text{S32})$$

Similarly, let \mathcal{I}_n^d be the hermitian part of the linear subspace appearing in (S30), i.e.,

$$\mathcal{I}_n^d := \left\{ \varsigma([H^{\Lambda(n)}, q - q^*]) : q \in \mathcal{S}_{\Lambda(n-l)}^{2d-\deg(H)} \right\} \quad (\text{S33})$$

where we use that $[H^{\Lambda(n)}, a]$ is hermitian when $a^* = -a$. Then the positivity, stationarity, and bulk spectral gap constraints are equivalent to the requirement that \mathcal{L}_n^d is nonnegative on the convex cone

$$\mathcal{Q}_n^d := \mathcal{Q}_n^{d,\text{pos}} + \mathcal{I}_n^d + \mathcal{Q}_n^{d,\text{gap}}. \quad (\text{S34})$$

By construction, $\mathcal{Q}_n^d \subseteq \mathcal{Q}_{n'}^{d'}$ whenever $n \leq n'$ and $d \leq d'$.

Theorem 2.5. *Let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be a dynamical system induced by a family of local Hamiltonians $\mathcal{F} = \{H^{\Lambda(n)}\}_n$ with interaction length l , and let $O \in \mathcal{A}$ be a local observable. Fix a gap parameter $\gamma \geq 0$.*

Then, the state polynomial reformulation (Theorem 2.4) of the relaxation hierarchy (Theorem 2.1) is feasible at every level (n, d) with the parameter γ if and only if there exists a KMS ground state ω on the dynamical system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ with a locally non-degenerate bulk spectral gap of at least γ .

Furthermore, the hierarchy values $\alpha_{n,d}^{\min}$ and $\alpha_{n,d}^{\max}$ converge monotonically to the extremal expectation values from Eq. (S20):

$$\alpha_{n,d}^{\min} \nearrow \langle O \rangle_{\min}, \quad \alpha_{n,d}^{\max} \searrow \langle O \rangle_{\max} \quad (\text{S35})$$

as $n, d \rightarrow \infty$.

For the proof of this theorem, we need to construct an archimedean quadratic module $\mathcal{Q} \subset \mathcal{S}$. As in [Kle+24], let \mathcal{Q} be the convex cone generated by $\varsigma(q^*q)$, the hermitian bulk-gap polynomials $\frac{1}{2}(\varsigma(q^*\delta_{\mathcal{F}}(q) - \delta_{\mathcal{F}}(q^*)q)) - \gamma(\varsigma(q^*q) - \varsigma(q^*)\varsigma(q))$, and the hermitian polynomials encoding the stationarity condition, $\varsigma(\delta_{\mathcal{F}}(q - q^*))$, for all $q \in \mathcal{S}$. Then \mathcal{Q} is the union of the cones \mathcal{Q}_n^d on which the \mathcal{L}_n^d are required to be positive. Moreover, \mathcal{Q} is a quadratic module of \mathcal{S} , i.e., $p^*\mathcal{Q}p \subseteq \mathcal{Q}$ for all $p \in \mathcal{S}$, because the maps ς and $\delta_{\mathcal{F}}$ treat elements of \mathcal{S} as scalars. Next, recall that for each $x_j^{(i)}$ there is an $s_j^{(i)} = (M_j^{(i)} - (x_j^{(i)})^2)^{1/2} \in \mathcal{A}_i$. Hence, for every $q \in \mathcal{S}$,

$$\varsigma\left(q^*(M_j^{(i)} - (x_j^{(i)})^2)q\right) = \varsigma\left((s_j^{(i)}q)^*(s_j^{(i)}q)\right) \in \mathcal{Q}. \quad (\text{S36})$$

Thus, [Kle+24, Lemma 5.4] applies and therefore the quadratic module \mathcal{Q} is archimedean.

Proof. The necessity of the hierarchy is immediate from the discussion at the end of Section 2.1 and by letting both ς and \mathcal{L}_n^d be induced by the KMS ground state ω . For the sufficiency of the hierarchy, assuming the feasibility of Theorem 2.4 for every (n, d) , we need to construct a state $\omega : \mathcal{A} \rightarrow \mathbb{C}$ that satisfies Eq. (S11) for the dynamical system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ induced by $\mathcal{F} = \{H^{\Lambda(n)}\}_n$. To this end, we first extract a limit \mathcal{L} on \mathcal{S} from the feasible solutions \mathcal{L}_n^d using a standard argument, and then we use the representation theorem for archimedean quadratic modules [Mar08, Theorem 5.4.4] (also known as the Kadison-Dubois Theorem) to obtain an extremal solution \mathcal{K} to construct the desired ω .

First, given the functionals $\mathcal{L}_n^d : \mathcal{S}_{\Lambda(n)}^{2d} \rightarrow \mathbb{C}$, we extract a state \mathcal{L} on the full \mathcal{S} . To this end, note that \mathcal{S} is the inductive limit of $\mathcal{S}_{\Lambda(n)}^{2d}$ as $n, d \rightarrow \infty$. As a vector space of countable dimension, we can enumerate a basis $\{p_1, p_2, \dots\}$ for \mathcal{S} . We can further assume that these basis elements are hermitian, i.e., $p_i^* = p_i$.

For each basis element $p_i \in \mathcal{S}$, since $\mathcal{Q} \subset \mathcal{S}$ is archimedean, there exists a constant $m_i > 0$ such that $m_i \pm p_i \in \mathcal{Q}$. As \mathcal{Q} is the union of the cones \mathcal{Q}_n^d on which \mathcal{L}_n^d is nonnegative, for all sufficiently large (n, d) , we have $m_i \pm p_i \in \mathcal{Q}_n^d$ and hence $|\mathcal{L}_n^d(p_i)| \leq m_i$. That is, for each fixed p_i , the sequence $\{\mathcal{L}_n^d(p_i)\}_{n,d}$ is uniformly bounded and has a convergent subsequence by compactness. A standard diagonal argument over the countable basis $\{p_1, p_2, \dots\}$ gives a convergent subsequence of functionals $\mathcal{L}_{n_k}^{d_k}$ such that $\mathcal{L}_{n_k}^{d_k}(p_i)$ is convergent for all i as $k \rightarrow \infty$. We define $\mathcal{L} : \mathcal{S} \rightarrow \mathbb{C}$ on the basis $\{p_1, p_2, \dots\}$ by

$$\mathcal{L}(p_i) := \lim_{k \rightarrow \infty} \mathcal{L}_{n_k}^{d_k}(p_i), \quad (\text{S37})$$

and extend linearly to all of \mathcal{S} . By construction, $\mathcal{L}(1) = 1$ and $\mathcal{L}(\mathcal{Q}) \subset \mathbb{R}_{\geq 0}$.

Since $\mathcal{Q} \subset \mathcal{S}$ is archimedean and $\mathcal{L}(\mathcal{Q}) \subset \mathbb{R}_{\geq 0}$, by the representation theorem for archimedean quadratic modules [Mar08, Theorem 5.4.4], there exists a probability measure μ , supported on the set of characters \mathcal{K} on \mathcal{S} satisfying $\mathcal{K}(1) = 1$ and $\mathcal{K}(\mathcal{Q}) \subset \mathbb{R}_{\geq 0}$, such that

$$\mathcal{L}(a) = \int \mathcal{K}(a) d\mu(\mathcal{K}). \quad (\text{S38})$$

Therefore, every $\mathcal{K} \in \text{supp}(\mu)$ gives rise to a feasible solution to Eq. (S28) for all n, d .

Moreover, since $\mathcal{L}(\varsigma(O)) = \int \mathcal{K}(\varsigma(O)) d\mu(\mathcal{K})$, there exists $\mathcal{K} \in \text{supp}(\mu)$ such that $\mathcal{K}(\varsigma(O)) \geq \mathcal{L}(\varsigma(O))$ (and similarly there exists $\mathcal{K} \in \text{supp}(\mu)$ with $\mathcal{K}(\varsigma(O)) \leq \mathcal{L}(\varsigma(O))$). In case of maximization with $\alpha_{n,d}^{\max}$, choose \mathcal{K} such that $\mathcal{K}(\varsigma(O)) \geq \mathcal{L}(\varsigma(O))$; otherwise, for minimization with $\alpha_{n,d}^{\min}$, fix \mathcal{K} such that $\mathcal{K}(\varsigma(O)) \leq \mathcal{L}(\varsigma(O))$.

Define the functional $\omega : \mathcal{A}_{\text{loc}} \rightarrow \mathbb{C}$ by

$$\omega(a) := \mathcal{K}(\zeta(a)). \quad (\text{S39})$$

It is straightforward to check that $\omega(1) = 1$, and $\omega(a^*a) \geq 0$ for all $a \in \mathcal{A}_{\text{loc}}$ because $\zeta(a^*a) \in Q$ for all $a \in \mathcal{A}_{\text{loc}}$. Hence, ω is a state on \mathcal{A}_{loc} . Moreover, from $\pm\zeta(\delta_{\mathcal{F}}(a - a^*)) \in Q$ for all $a \in \mathcal{A}_{\text{loc}}$ it follows that $\omega(\delta_{\mathcal{F}}(a - a^*)) = 0$ for all $a \in \mathcal{A}_{\text{loc}}$, hence $\omega(\delta_{\mathcal{F}}(a)) = 0$ for all $a \in \mathcal{A}_{\text{loc}}$ because $a = \frac{1}{2}((a - a^*) - i((ia) - (ia)^*))$. This shows that ω fulfils the stationarity condition Eq. (S8) on \mathcal{A}_{loc} . Similarly, as $\frac{1}{2}(\zeta(a^*\delta_{\mathcal{F}}(a) - \delta_{\mathcal{F}}(a^*)a)) - \gamma(\zeta(a^*a) - \zeta(a^*)\zeta(a)) \in Q$ for all $a \in \mathcal{A}_{\text{loc}}$ and as characters are *-homomorphic,

$$\frac{1}{2}\mathcal{K}(\zeta(a^*\delta_{\mathcal{F}}(a) - \delta_{\mathcal{F}}(a^*)a)) \geq \gamma(\mathcal{K}(\zeta(a^*a)) - |\mathcal{K}(\zeta(a))|^2) \quad (\text{S40})$$

for all $a \in \mathcal{A}_{\text{loc}}$, so ω also satisfies Eq. (S12) on \mathcal{A}_{loc} . Equivalently, ω satisfies

$$\omega(a^*[H^{\Lambda(n)}, a]) \geq \gamma(\omega(a^*a) - |\omega(a)|^2), \quad (\text{S41})$$

for all $a \in \mathcal{A}_{\text{loc}}$, that is, it fulfils Eq. (S11) on \mathcal{A}_{loc} .

Now, with $\mathcal{A}_{\text{loc}} = \bigcup_{n \in \mathbb{N}} \mathcal{A}_{\Lambda(n)}$ and $\mathcal{F} = \{H^{\Lambda(n)}\}_n$, following the same steps as in Section 1.1, we recover the desired dynamical system $(\mathcal{A} = \overline{\mathcal{A}_{\text{loc}}}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$. As \mathcal{A}_{loc} inherits the C^* -norm from \mathcal{A} , the positivity and normalization of ω imply that ω is norm-bounded on \mathcal{A}_{loc} , and hence ω extends by continuity to a state on the quasi-local algebra \mathcal{A} , which we again denote by ω . Since \mathcal{A}_{loc} is a core for $\delta_{\mathcal{F}}$, the extension ω satisfies Eq. (S11) for all a in the domain of $\delta_{\mathcal{F}}$.

To sum up the proof of the convergence statement, we recover a KMS ground state ω on the dynamical system $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ admitting a locally non-degenerate bulk spectral gap at least γ , proving the “if and only if” statement of the theorem.

Finally, we prove the “furthermore” statement. Since level- (n, d) feasible sets are nested as n, d increase, $\alpha_{n,d}^{\min}$ is monotonically nondecreasing and $\alpha_{n,d}^{\max}$ is monotonically nonincreasing, hence both converge.

Every γ -gapped KMS ground state of $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ restricts to a feasible solution at every level. So, by optimality and taking $n, d \rightarrow \infty$, one has

$$\lim_{n,d \rightarrow \infty} \alpha_{n,d}^{\min} \leq \langle O \rangle_{\min} \leq \langle O \rangle_{\max} \leq \lim_{n,d \rightarrow \infty} \alpha_{n,d}^{\max}. \quad (\text{S42})$$

It remains to show $\lim_{n,d \rightarrow \infty} \alpha_{n,d}^{\max} \leq \langle O \rangle_{\max}$. By the convergence proof, the hierarchy admits a limiting state ω on $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$, and therefore

$$\lim_{n,d \rightarrow \infty} \alpha_{n,d}^{\max} = \mathcal{L}(\zeta(O)) \leq \mathcal{K}(\zeta(O)) = \omega(O) \leq \langle O \rangle_{\max}. \quad (\text{S43})$$

The first equality is by construction, the first inequality follows from the choice of \mathcal{K} (see the paragraph below Eq. (S38)), the equality $\mathcal{K}(\zeta(O)) = \omega(O)$ follows from the definition of ω , and the final inequality follows from the definition of $\langle O \rangle_{\max}$. The case of minimization is analogous. \square

Consequently, if the hierarchy with $O := 1$ is infeasible at some level for a given γ , then γ is larger than the bulk spectral gap of Theorem 1.10.

Remark 2.6 (Scope of Construction). *The proof does not rely on the specific lattice realization of the quasi-local algebra beyond the existence of (i) a distinguished dense *-subalgebra $\mathcal{A}_{\text{loc}} \subset \mathcal{A}$ (typically an increasing union of finite-dimensional C^* -subalgebras), (ii) a strongly continuous dynamics τ on \mathcal{A} with a generator δ whose dense domain contains \mathcal{A}_{loc} as a core, and (iii) an archimedean constraint on the local generators. Whenever these structural assumptions hold, the corresponding variant of Theorem 2.4 leads to an analogous Theorem 2.5. In particular, the locally non-degenerate bulk spectral gap inequality can be replaced by any family of constraints that is expressible in the framework of state polynomial optimization, in order to tackle a specific class of problems or achieve better convergence behavior.*

As an example beyond square lattices, we adapt our SDP hierarchy to the kagome lattice in Section 3.2.

Remark 2.7 (Hierarchy with local symmetry constraints). *Let α be a $*$ -automorphism of \mathcal{A} preserving \mathcal{A}_{loc} (mapping local observables to local observables) and commuting with the dynamics $\tau_{\mathcal{F}}$. One may restrict the hierarchy to α -symmetric KMS ground states by adding linear invariance constraints. At the level of Eq. (S21), this means imposing*

$$\omega_n^d(\alpha(a)) = \omega_n^d(a) \quad (\text{S44})$$

for all local observables a for which both a and $\alpha(a)$ lie in the truncation $\mathcal{A}_{\Lambda(n)}^{2d}$. In the state-polynomial formulation Eq. (S28), the same symmetry is imposed by extending α to \mathcal{S} and requiring

$$\mathcal{L}_n^d(\alpha(p)) = \mathcal{L}_n^d(p) \quad (\text{S45})$$

for the corresponding state polynomials p and $\alpha(p)$ in the truncation $\mathcal{S}_{\Lambda(n)}^{2d}$. The proof of Theorem 2.5 extends directly, with the feasible set replaced by the α -symmetric feasible set. Hence feasibility for every (n, d) is equivalent to the existence of an α -symmetric KMS ground state with locally non-degenerate bulk spectral gap of at least γ .

Imposing physically motivated symmetries therefore reduces the certification problem to a specified subclass of KMS ground states. This restriction can substantially improve numerical performance, and is used in the numerical examples of Section 3.

3 Numerical showcases

For each finite level (n, d) , feasibility of Theorem 2.4 at a given γ is a necessary condition for the existence of a KMS ground state satisfying Eq. (S11). Hence, if the relaxation is infeasible, then γ is ruled out as a possible bulk spectral gap in the sense of Theorem 1.10. Equivalently, maximizing γ subject to feasibility gives a certified upper bound on the bulk spectral gap. As (n, d) increases, the relaxation becomes stricter, so these certified upper bounds form a monotonically nonincreasing sequence, yielding progressively tighter bounds in the thermodynamic limit. This shows that bulk spectral gap is semi-decidable: one can always exclude a proposed bulk gap in finite time.

We test the SDP hierarchy on the transverse-field Ising model (Section 3.1) and the kagome lattice Heisenberg model (Section 3.2). In practice, the computational cost can be reduced by exploiting structural constraints, reduced density matrix positivity, and symmetries specific to the model under study, as in [Wan+26]. For the numerical implementations reported below, we use a selected subset of state monomials for each degree d to index the moment matrix in order to better balance tightness of SDPs and computational costs. The precise monomial sets used in each computation, together with reproducibility details, are provided at <https://github.com/wangjie212/SpectralGap>. All numerical experiments were performed on a desktop computer with Intel(R) Core(TM) i9-10900 CPU@2.80GHz and 64G RAM except that the computation on the Kagome lattice Heisenberg model with $L = 3, 4$ was conducted on a workstation with Intel(R) Xeon(R) Gold 6348 CPU@2.60GHz and 773G RAM.

3.1 Transverse-field Ising model

A standard analytically solvable example is the transverse-field Ising model on the infinite chain \mathbb{Z} . Here each site i carries a qubit algebra $\mathcal{A}_i \simeq M_2(\mathbb{C})$ with Pauli operators $X_i, Z_i \in \mathcal{A}_i$. For $g \geq 0$, the finite-volume Hamiltonian on $\Lambda(L) = \{-L, \dots, L\} \subset \mathbb{Z}$ is

$$H_{\text{TFIM},g}^{\Lambda(L)} = - \sum_{i=-L}^{L-1} Z_i Z_{i+1} + g \sum_{i=-L}^L X_i. \quad (\text{S46})$$

The exact solution indicates three regimes [MRS24]:

1. *Disordered phase* $g > 1$: unique ground state with a finite-volume spectral gap that converges to $2(g - 1)$ in the thermodynamic limit.
2. *Critical point* $g = 1$: unique ground state; the finite-volume gap closes as $1/L$.

Table S1: Largest feasible gap parameter returned by the SDP hierarchy for the transverse-field Ising chain. Rows correspond to the system size parameter L , and columns to the monomial degree d . All computations impose the spin-flip \mathbb{Z}_2 -symmetry. Missing entries indicate values not computed within the available computational budget. Mirror symmetry $i \mapsto -i$ is also imposed to further improve the numerical tightness.

$g = 0.5$				$g = 1.0$			
L	$d = 2$	$d = 3$	$d = 4$	L	$d = 2$	$d = 3$	$d = 4$
2	0.52	0.44	0.44	2	1.18	1.14	1.14
3	0.39	0.25	0.24	3	0.93	0.78	0.78
4	0.36	0.19	0.15	4	0.85	0.60	0.59
5	0.35	0.17	0.10	5	0.81	0.52	0.47
6	0.34	0.16	–	6	0.80	0.47	–
7	0.34	0.15	–	7	0.79	0.45	–

$g = 1.5$				$g = 2.0$			
L	$d = 2$	$d = 3$	$d = 4$	L	$d = 2$	$d = 3$	$d = 4$
2	2.00	1.98	1.98	2	2.89	2.88	2.88
3	1.66	1.57	1.57	3	2.52	2.47	2.47
4	1.52	1.37	1.37	4	2.36	2.29	2.29
5	1.46	1.27	1.26	5	2.28	2.19	2.19
6	1.42	1.20	–	6	2.24	2.14	–
7	1.40	1.17	–	7	2.21	2.11	–

3. *Ordered phase* $0 \leq g < 1$: two degenerate ground states in the thermodynamic limit aligned along the z -axis, each with an energy gap $2(1 - g)$.

The set of KMS ground states for transverse-field Ising model is fully characterized by [AM85, Theorem 1]. For $0 \leq g < 1$, there exist exactly two pure KMS ground states whose equal classical mixture is the unique KMS ground state that is \mathbb{Z}_2 -symmetric under the mapping $Y_i \mapsto -Y_i, Z_i \mapsto -Z_i$. By Theorem 1.9, this \mathbb{Z}_2 -symmetric equal classical mixture has a locally non-degenerate bulk gap of 0. For $g \geq 1$, there exists a unique pure KMS ground state and it is also \mathbb{Z}_2 -symmetric.

Recall that our SDP hierarchy tests the necessary condition for a locally non-degenerate bulk gap γ over all KMS ground states of the given dynamical system. In the numerical tests here, we impose the spin-flip \mathbb{Z}_2 symmetry,

$$Y_i \mapsto -Y_i, Z_i \mapsto -Z_i, \tag{S47}$$

for every site i , using the symmetry-restricted hierarchy of Theorem 2.7. The resulting hierarchy therefore certifies only the absence of a gapped \mathbb{Z}_2 -symmetric KMS ground state. For $g \geq 1$, this restriction does not change the result, since the unique pure KMS ground state is \mathbb{Z}_2 -symmetric. For $0 \leq g < 1$, the symmetry constraint selects the unique \mathbb{Z}_2 -symmetric equal classical mixture which has a bulk gap of 0. We impose mirror symmetry $i \mapsto -i$ to further improve the numerical tightness.

In Fig. 2 of the main text, we plot the largest feasible gap parameter at relaxation level $(L, d) = (n + 1, n)$, for $n \in \{2, 3, 4\}$ and $g \in \{0.5, 1.0, 1.5, 2.0\}$, with the spin-flip \mathbb{Z}_2 -symmetry imposed. The corresponding values, together with additional computations at other relaxation levels, are reported in Table S1. As expected from the nested relaxations, the maximal feasible values decrease monotonically as either L or d is increased.

For $g = 1.5$ and $g = 2.0$, the upper bounds approach the known value $2(g - 1)$ from above. For $g = 1$, the decrease is slower, likely due to critical finite-size scaling. For $g = 0.5$, the symmetry-restricted hierarchy returns bounds decreasing towards 0, consistent with testing the \mathbb{Z}_2 -symmetric equal classical mixture, whose locally non-degenerate bulk gap is 0.

Table S1 also illustrates the different numerical roles of L and d . For the ordered case $g = 0.5$, increasing

the monomial degree d is more effective than increasing the system size L within our computational budget; the best bound obtained is 0.10 at $L = 5$ and $d = 4$. For the other values of g considered here, increasing L gives the more significant improvement at comparable computational cost. Lastly, if the spin-flip symmetry is not imposed, the computations become substantially more expensive; within the same budget we could only reach $L = 3$ and $d = 3$, leading to much looser upper bounds in all cases.

3.2 Kagome lattice Heisenberg model

The kagome lattice is a two-dimensional periodic graph (V, E) built from corner-sharing triangles, where V is the set of vertices and E is the set of nearest-neighbour edges; see Fig. 1 of the main text for illustration. It is a canonical setting for geometric frustration and has become a central playground for the study of quantum spin liquids.

We express the kagome lattice as the limit of finite sublattices. To this end, we fix a distinguished vertex $v_0 \in V$ and let $d(\cdot, \cdot)$ denote the graph distance on (V, E) (i.e., the length of a shortest edge-path). For each $L \in \mathbb{N}$ we define the finite region

$$\Lambda_{\text{kagome}}(L) := \{v \in V : d(v, v_0) \leq L\}. \quad (\text{S48})$$

Then $(\Lambda_{\text{kagome}}(L))_{L \in \mathbb{N}}$ is an increasing exhaustion of V by finite subsets.

Consider a spin- $\frac{1}{2}$ degree of freedom at each site i , with local algebra $\mathcal{A}_i \simeq M_2(\mathbb{C})$ and Pauli matrices $X_i, Y_i, Z_i \in \mathcal{A}_i$. Clearly, $\mathcal{A}_{\Lambda_{\text{kagome}}(L)} := \bigotimes_{i \in \Lambda_{\text{kagome}}(L)} \mathcal{A}_i$ is finite-dimensional and the algebra of quasi-local operators is $\mathcal{A} := \bigcup_L \mathcal{A}_{\Lambda_{\text{kagome}}(L)}$. On $\Lambda_{\text{kagome}}(L)$, the spin- $\frac{1}{2}$ kagome lattice Heisenberg model (KLHM) has the local Hamiltonian

$$H_{\text{KLHM}}^{\Lambda_{\text{kagome}}(L)} = J \sum_{\langle i, j \rangle \in E: i, j \in \Lambda_{\text{kagome}}(L)} \frac{1}{4} (X_i X_j + Y_i Y_j + Z_i Z_j), \quad (\text{S49})$$

for a fixed $J > 0$.

The numerical consensus suggests that the ground state of KLHM is consistent with quantum spin liquids. However, whether the thermodynamic-limit spectrum is gapped or gapless remains debated. DMRG studies on long kagome cylinders report evidence compatible with a gapped ground state [DMS12; YHW11], with an estimated singlet gap of 0.04–0.05 J [YHW11] and an estimated spin gap of $\approx 0.13J$ from cylinder extrapolations [DMS12]. (Recall that spin and singlet gaps are defined by restricting the allowed excitations to fixed spin sectors; they are not upper bounds on the bulk spectral gap which is due to all excitations. Furthermore, the values cited are estimates and should not be identified directly with the bulk spectral gap.) Subsequent DMRG work on infinitely long cylinders, however, argued that the spin gap may be significantly smaller and favors a gapless Dirac quantum spin liquid [He+17]. Tensor-network approaches based on projected entangled pair/simplex states have also argued for a gapless ground state [Jia+19; Lia+17]. Exact diagonalization [LSM19] is restricted to comparatively small clusters (e.g., 48 sites) and has not resolved the issue conclusively. We refer to the recent review [ZGS25] for a modern overview and further references.

In addition to the hierarchy with no symmetry restriction, to get sharper numerical bounds within physically motivated symmetry sector we also use the symmetry-restricted variants described in Theorem 2.7. First, consider the three π -rotations

$$\begin{aligned} R_x &: (X_i, Y_i, Z_i) \mapsto (X_i, -Y_i, -Z_i), \\ R_y &: (X_i, Y_i, Z_i) \mapsto (-X_i, Y_i, -Z_i), \\ R_z &: (X_i, Y_i, Z_i) \mapsto (-X_i, -Y_i, Z_i), \end{aligned} \quad (\text{S50})$$

for every site $i \in \Lambda_{\text{kagome}}(L)$. Invariance under the three π -rotations enforces $\omega(X_i) = \omega(Y_i) = \omega(Z_i) = 0$, and is therefore compatible with the expected absence of one-site magnetization in a spin liquid. Second, consider the even permutations A_3 of the spin components,

$$(X_i, Y_i, Z_i) \mapsto (Y_i, Z_i, X_i), \quad (X_i, Y_i, Z_i) \mapsto (Z_i, X_i, Y_i). \quad (\text{S51})$$

Together with the three π -rotations R_x, R_y, R_z , these generate the tetrahedral group $T \subset SO(3)$, a finite

subgroup of the global spin rotation symmetry, which is motivated by the $SU(2)$ -symmetric spin liquids considered in [DMS12]. Third, we consider the time-reversal symmetry,

$$\Theta : (X_i, Y_i, Z_i) \mapsto (-X_i, -Y_i, -Z_i). \quad (\text{S52})$$

This additional condition rules out states with nonzero scalar chirality defined for sites $i, j, k \in \Lambda_{\text{kagome}}(L)$ on a triangle:

$$\chi_{ijk} = \frac{1}{8} \left(X_i Y_j Z_k + Y_i Z_j X_k + Z_i X_j Y_k - X_i Z_j Y_k - Z_i Y_j X_k - Y_i X_j Z_k \right), \quad (\text{S53})$$

since $\Theta(\chi_{ijk}) = -\chi_{ijk}$. Thus bounds obtained with the time-reversal symmetry Θ apply only to nonchiral ground states, whereas bounds obtained without imposing Θ , for instance with only the three π -rotation symmetry and the permutation symmetries, do not exclude chiral spin-liquid states by symmetry. These symmetry restrictions are motivated by the properties of quantum spin liquids commonly discussed in the KLHM literature; they should not be interpreted as rigorously proven properties of all KMS ground states.

Remark 3.1. *Before presenting the numerical bounds, we record a crude analytical upper bound obtained directly from the bulk-gap inequality Eq. (S11) and Theorem 1.10, based only on the locality and the operator-norm bound. Indeed, for any local observable a with $\omega(a^*a) - |\omega(a)|^2 > 0$, we have*

$$\Delta_{\text{bulk}} \leq \frac{\omega(a^* \delta_{\mathcal{F}_{\text{KLHM}}}(a))}{\omega(a^*a) - |\omega(a)|^2}, \quad (\text{S54})$$

where $\delta_{\mathcal{F}_{\text{KLHM}}}$ is the dynamic generator induced by the Hamiltonian family $\{H_{\text{KLHM}}^{\Lambda_{\text{kagome}}(L)}\}$.

Fix a KMS ground state ω and a site 0 of the kagome lattice. For a unit vector $\vec{n} \in \mathbb{R}^3$ and $\vec{\sigma}_0 = (X_0, Y_0, Z_0)^T$, define the 0-site projection

$$P_{\vec{n},+} := \frac{1 + \vec{n} \cdot \vec{\sigma}_0}{2}. \quad (\text{S55})$$

We choose \vec{n} orthogonal to the one-site magnetization $\vec{m} := (\omega(X_0), \omega(Y_0), \omega(Z_0))$, so that $\omega(P_{\vec{n},+}) = 1/2$. Since $P_{\vec{n},+}$ is a projection, this gives

$$\omega(P_{\vec{n},+}^* P_{\vec{n},+}) - |\omega(P_{\vec{n},+})|^2 = \omega(P_{\vec{n},+}) - \omega(P_{\vec{n},+})^2 = \frac{1}{4}. \quad (\text{S56})$$

For the numerator, only the four nearest-neighbour bonds incident to the site 0 contribute to $\delta_{\mathcal{F}_{\text{KLHM}}}(P_{\vec{n},+})$. For each such bond,

$$H_{0j} = \frac{J}{4} (X_0 X_j + Y_0 Y_j + Z_0 Z_j). \quad (\text{S57})$$

Using the rotational invariance of this interaction, we may estimate in a basis where $P_{\vec{n},+} = (1 + Z_0)/2$. Then

$$\|[H_{0j}, P_{\vec{n},+}]\| \leq \frac{J}{4} \left(\|[X_0, P_{\vec{n},+}]\| + \|[Y_0, P_{\vec{n},+}]\| \right) \leq \frac{J}{2}. \quad (\text{S58})$$

Hence, using that the kagome lattice has a coordination number 4,

$$\omega(P_{\vec{n},+} \delta_{\mathcal{F}_{\text{KLHM}}}(P_{\vec{n},+})) \leq \|P_{\vec{n},+} \delta_{\mathcal{F}_{\text{KLHM}}}(P_{\vec{n},+})\| \leq \sum_j \|[H_{0j}, P_{\vec{n},+}]\| \leq 2J. \quad (\text{S59})$$

It then follows that Eq. (S11) cannot hold for any $\gamma > 8J$, therefore $\Delta_{\text{bulk}} \leq 8J$.

Furthermore, suppose that ω is invariant under the three π -rotation R_x, R_y, R_z symmetry. Then $\omega(X_0) = \omega(Y_0) = \omega(Z_0) = 0$. By setting $a = X_0$, direct calculation shows that $\|[H_{0j}, a]\| \leq J$ and $\omega(a^*a) - |\omega(a)|^2 = 1$. This implies that the bulk spectral gap for all ω that is R_x, R_y, R_z -symmetric is upper bounded by $4J$.

By Theorem 2.6, our algorithm generalizes to KLHM using $\Lambda_{\text{kagome}}(L)$. At monomial degree $d = 2, 3$,

Table S2: Certified upper bounds for the KLHM under different symmetry restrictions. Missing entries indicate values not computed within the available computational budget. Note that $L = 1, 2, 3, 4$ correspond to kagome patches of 5, 13, 27, 45 sites, respectively.

	No symmetry	R_x, R_y, R_z		$R_x, R_y, R_z, A_3, \Theta$	
L	$d = 2$	$d = 2$	$d = 3$	$d = 2$	$d = 3$
1	2.72J	2.00J	–	2.01J	2.01J
2	2.18J	1.31J	–	1.31J	1.28J
3	–	–	–	1.24J	1.15J
4	–	–	–	1.22J	–

the relaxation hierarchy already gives certified bounds substantially sharper than the crude estimate of Theorem 3.1; see Table S2 and also Fig. 3 in the main text.

We first test the hierarchy without imposing any symmetry restriction on the KMS ground state. At $d = 2$, we obtain the certified upper bound $2.72J$ at $L=1$, and $2.18J$ at $L = 2$. Thus, we certify that no KMS ground state of the KLHM can have a locally non-degenerate bulk spectral gap larger than $2.18J$. Equivalently, the bulk spectral gap of KLHM is less than $2.18J$.

Next, we impose the three π -spin rotation symmetries R_x, R_y, R_z . These constraints restrict the optimization to KMS ground states with vanishing one-site magnetization. The tightest upper bound we obtain in this symmetry class is $1.31J$ at $(L, d) = (2, 2)$. Equivalently, within this symmetry class, we certify that no KMS ground state of the KLHM can have a locally non-degenerate bulk spectral gap larger than $1.31J$. Without imposing the time-reversal symmetry Θ , the relaxation may still allow chiral ground states. To see this, we maximize the scalar chirality χ_{ijk} within the same SDP feasibility problem at fixed gap parameter γ . At $(L, d) = (1, 2)$ with $\gamma = 1.99J$, we obtain

$$\omega(\chi_{ijk}) \leq 0.0484, \quad (\text{S60})$$

and at $(L, d) = (2, 2)$ with $\gamma = 1.30J$, we obtain

$$\omega(\chi_{ijk}) \leq 0.0558. \quad (\text{S61})$$

Namely, this means that any hypothetical KMS ground state in this symmetry section with locally non-degenerate bulk gap $\geq 1.99J$ has scalar chirality at most 0.0484, and any hypothetical KMS ground state in this symmetry sector with gap $\geq 1.30J$ has scalar chirality at most 0.0558. (The later value 0.0558, even though tested at a higher level of the hierarchy, is larger than 0.0484, because now we also consider more ground states with possibly smaller gaps.) These values should not be interpreted as evidence for or against chirality of the physical KLHM ground state; they only diagnose what is allowed by the corresponding SDP relaxation.

Finally, we impose all aforementioned symmetries: the three π -spin rotation symmetry, the cyclic permutation symmetry A_3 , and the time-reversal symmetry Θ . These constraints restrict the search to nonchiral, spin-isotropic KMS ground states. In practice, note that composing time reversal with the π -rotations gives the single-axis sign flips,

$$(X_i, Y_i, Z_i) \mapsto (-X_i, Y_i, Z_i), \quad (X_i, Y_i, Z_i) \mapsto (X_i, -Y_i, Z_i), \quad (X_i, Y_i, Z_i) \mapsto (X_i, Y_i, -Z_i), \quad (\text{S62})$$

which is what we imposed in the calculation. This enforces additional conditions, for example, $\omega(X_i Y_j) = \omega(X_i Z_j) = \omega(Y_i Z_j) = 0$. In this sector, the best bound we obtain is $1.15J$ at $(L, d) = (3, 3)$. Thus, within this symmetry sector, we certify that no KMS ground state of the KLHM can have a locally non-degenerate bulk spectral gap larger than $1.15J$. (For the entries reported as $2.01J$ at $L = 1$, the corresponding SDPs are actually infeasible for a γ very close to $2J$ from above, e.g., $2.00000001J$.)

The bounds in Table S2 are far above the best numerical estimates and extrapolations from variational methods, but they are fully certified upper bounds for the corresponding classes of KMS ground states in the thermodynamic limit. The calculation with no symmetry restriction gives bounds to the bulk spectral gap

of the KLHM, while the symmetry-restricted calculations give sharper bounds within physically motivated sectors. The present kagome calculation is not optimized and should be viewed as a proof of principle. Sharper certified bounds could become physically informative: for instance, driving the bound below the proposed singlet gap of [YHW11] or the proposed spin gap of [DMS12] would rule out those values within the corresponding class of states, while driving the bound to zero would certify the absence of a bulk spectral gap in that class. We leave a more optimized numerical implementation to future work.

4 Undecidability and notions of spectral gaps

Undecidability results [Bau+20; CPW15] show that, for certain translationally invariant lattice Hamiltonians with bounded local dimensions and coefficients, the existence of a spectral gap is algorithmically undecidable when the gap is defined via a thermodynamic limit of finite systems under fixed boundary conditions. On the other hand, as in the leading paragraph of Section 3, our SDP hierarchy shows that the spectral gap problem is semi-decidable.

Since our SDP hierarchy addresses the *bulk* (algebraic) notion of spectral gap, there is no contradiction. This section recalls the results of [Bau+20; CPW15] and explains why the undecidability for the algebraic bulk gap remains open. More broadly, this illustrates a caution about mathematical idealizations of the physical world: distinct infinite-volume limits and gap notions can lead to genuinely different algorithmic questions.

In Section 4.1 we recall the more conventional finite-volume definition of the spectral gap under fixed boundary conditions. In Section 4.2 we summarize the 2D undecidability construction of [CPW15], explain why it does not contradict our bulk-gap SDP hierarchy (Theorem 4.7), and show that the construction is necessarily sensitive to (generalized) boundary conditions (Theorem 4.8). Finally, in Section 4.3 we discuss the 1D result of [Bau+20] and clarify what additional ingredients would be needed to extend the undecidability to the algebraic bulk definition.

4.1 Finite-volume spectral gaps and boundary conditions

We recall the finite-volume notion of spectral gap used in [CPW15], which is formulated in terms of the spectra of finite-volume Hamiltonians along a fixed choice of boundary condition. This dependence will be central to the comparison with algebraic bulk notions.

We start by formalizing the boundary conditions.

Definition 4.1. *Let $\{H^{\Lambda(L)}\}$ be a family of local Hamiltonians on \mathbb{Z}^D associated with a finite-range interaction Φ . A boundary condition on $\Lambda(L)$ is a choice of a self-adjoint operator $B^{\Lambda(L)} \in \mathcal{A}$ whose support is contained in an R -neighbourhood of the boundary $\partial\Lambda(L)$, for some fixed R independent of L . We set the family of Hamiltonians with the boundary condition $\{B^{\Lambda(L)}\}$ to be*

$$H_B^{\Lambda(L)} := H^{\Lambda(L)} + B^{\Lambda(L)}, \tag{S63}$$

which is again finitely supported for every L .

The open boundary condition (OBC) corresponds to $B_{\text{OBC}}^{\Lambda(L)} := 0$, that is, $H_{\text{OBC}}^{\Lambda(L)} = H^{\Lambda(L)}$. For the periodic boundary condition (PBC), we replace $\Lambda(L)$ with the discrete torus $\mathbb{T}_L^D := (\mathbb{Z}/L\mathbb{Z})^D$ and define $H_{\text{PBC}}^{\Lambda(L)}$ by the same interaction Φ , with sites identified modulo L . (In the nearest-neighbour case, this is equivalent to adding wrap-around couplings across opposite faces.)

Fix a boundary condition $\{B^{\Lambda(L)}\}_L$. Let $\text{Spec}(H_B^{\Lambda(L)}) = \{\lambda_0(H_B^{\Lambda(L)}), \lambda_1(H_B^{\Lambda(L)}), \dots\}$ denote the spectrum with non-decreasing ordering $\lambda_0(H_B^{\Lambda(L)}) \leq \lambda_1(H_B^{\Lambda(L)}) \leq \dots$. Define the ground state energy as the smallest eigenvalue $\lambda_0(H_B^{\Lambda(L)})$ and call the associated normalized eigenvector a ground state. The ground state energy is said to be non-degenerate if $\lambda_0(H_B^{\Lambda(L)}) < \lambda_1(H_B^{\Lambda(L)})$, that is, the dimension of the eigenspace corresponding to the lowest eigenvalue is 1. There are two important classes of behavior of the spectrum $\text{Spec}(H_B^{\Lambda(L)})$ at the thermodynamic limit as $L \rightarrow \infty$.

Definition 4.2 (Non-degenerate uniform gap with a boundary condition). Consider the family $\{H^{\Lambda(L)}\}_L$ of Hamiltonians with the boundary condition $\{B^{\Lambda(L)}\}$. We say that the system under the boundary condition $\{B^{\Lambda(L)}\}$ is non-degenerate uniformly gapped if $\exists \gamma > 0, \exists L_0, \forall L > L_0$ the ground state energy $\lambda_0(H_B^{\Lambda(L)})$ is non-degenerate, and

$$\text{Spec}\left(H_B^{\Lambda(L)}\right) \cap \left(\lambda_0(H_B^{\Lambda(L)}), \lambda_0(H_B^{\Lambda(L)}) + \gamma\right) = \emptyset. \quad (\text{S64})$$

In this case, we say that the system under the boundary condition $\{B^{\Lambda(L)}\}$ has a non-degenerate uniform spectral gap of at least γ .

Remark 4.3. Theorem 4.2 is stricter than the conventional definition in the usual physics literature for gapped systems. In particular, this notion is stronger than the locally non-degenerate bulk spectral gap on KMS ground states (Theorem 1.6) since it requires a uniform bound for excitations in finite systems. Indeed, a finite-volume uniform gap for some chosen boundary condition is a lower bound for the bulk gap [BDN16, Proposition 5.4].

The converse implication is false, that is, Theorem 1.6 does not imply Theorem 4.2. First, the local non-degeneracy of Theorem 1.6 does not necessarily imply the global non-degeneracy of Theorem 4.2. In addition, finite-volume gaps may close due to boundary (edge) excitations even when the bulk is gapped (e.g., the AKLT model with OBC [Aff+87]).

Moreover, as in Theorem 1.8, we remark again that “non-degeneracy” in Theorem 4.2 refers to the global uniqueness of ground states, whereas the algebraic Theorem 1.6 refers to non-degeneracy in the GNS representation of KMS ground states.

We next state the conventional gapless condition.

Definition 4.4 (Gapless with a boundary condition). Consider the family $\{H^{\Lambda(L)}\}_L$ of Hamiltonians with the boundary condition $\{B^{\Lambda(L)}\}$. We say that the system under the boundary condition $\{B^{\Lambda(L)}\}$ is gapless if $\forall \varepsilon > 0, \exists L_0, \forall L > L_0$

$$\Delta(H_B^{\Lambda(L)}) = \lambda_1(H_B^{\Lambda(L)}) - \lambda_0(H_B^{\Lambda(L)}) < \varepsilon. \quad (\text{S65})$$

Importantly, being gapless as in Theorem 4.4 does not imply that the system is still bulk-gapless (Theorem 1.11), due to the potential existence of edge states associated with the given boundary condition; again, see the AKLT model with open boundary condition. We remark that the gapless Theorem 4.4 is not as stringent as the gapless condition considered in [CPW15], where they instead adopt the continuum/uniform condition near the ground state energy.

Definition 4.5 (Uniformly gapless with a boundary condition). Consider the family $\{H^{\Lambda(L)}\}_L$ of Hamiltonians with the boundary condition $\{B^{\Lambda(L)}\}$. We say that the system under the boundary condition $\{B^{\Lambda(L)}\}$ is uniformly gapless if $\exists c > 0, \forall \varepsilon > 0, \exists L_0, \forall L > L_0$ the spectrum $\text{Spec}(H_B^{\Lambda(L)})$ is ε -dense in the interval

$$[\lambda_0(H_B^{\Lambda(L)}), \lambda_0(H_B^{\Lambda(L)}) + c]. \quad (\text{S66})$$

That is, every point in the interval lies within a distance ε of some eigenvalue of $H_B^{\Lambda(L)}$.

The uniformly gapless condition is strictly stronger than Theorem 4.4: For example, $H_B^{\Lambda(L)}$ with eigenvalues $\lambda_0, \lambda_k = \lambda_0 + 1/(1+k)$ for $k = 1, \dots, L$ is gapless but not uniformly gapless, since the spectrum near λ_0 is too discrete to be ε -dense in any nontrivial interval. Another example is a system with a degenerate ground space but a spectral gap above it (e.g., $\lambda_0 = \lambda_1 < \lambda_2$): it is gapless in the sense of Theorem 4.4 but not uniformly gapless. The AKLT chain with open boundary condition is an example of this type. On the other hand, $H_B^{\Lambda(L)}$ with $\lambda_0, \lambda_k = \lambda_0 + k/L$ is both gapless and uniformly gapless.

Finally, we recall the relationship between weak-* limit points of finite-volume ground states (as the volume grows and boundary conditions vary) and KMS ground states of the associated infinite-volume dynamics.

Remark 4.6. Consider a family of finite-volume Hamiltonians $\mathcal{F} = \{H^{\Lambda(L)}\}_L$ induced by a finite-range interaction Φ , and let $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$ be the associated dynamical system. For any boundary condition $\{B^{\Lambda(L)}\}_L$ as in Theorem 4.1, let $|\psi_B^{\Lambda(L)}\rangle$ be a ground state of $\lambda_0(H_B^{\Lambda(L)})$ and let $\omega_B^{\Lambda(L)}$ denote the corresponding (finite-volume) vector state defined as

$$\omega_B^{\Lambda(L)}(\cdot) := \langle \psi_B^{\Lambda(L)} | \cdot | \psi_B^{\Lambda(L)} \rangle. \quad (\text{S67})$$

Clearly, any weak-* limit point of $\{\omega_B^{\Lambda(L)}\}_L$ (for a fixed boundary condition) is a KMS ground state ω_B on $(\mathcal{A}, \tau_{\mathcal{F}}, \delta_{\mathcal{F}})$.

The converse direction requires a more general notion of boundary conditions beyond Theorem 4.1 from a variational perspective. Consider the conditional Hamiltonian on $\Lambda(L)$

$$\tilde{H}^{\Lambda(L)} := H^{\Lambda(L)} + W^{\Lambda(L)}, \quad W^{\Lambda(L)} := \sum_{\substack{S \cap \Lambda(L) \neq \emptyset \\ S \cap \Lambda(L)^c \neq \emptyset}} \Phi(S), \quad (\text{S68})$$

where the surface term $W^{\Lambda(L)}$ is well-defined since Φ is assumed to be of finite-range. By [BR81, Theorem 6.2.52], a state ω is a KMS ground state if and only if, for every finite region $\Lambda(L)$,

$$\omega(\tilde{H}^{\Lambda(L)}) = \inf_{\omega' \in C_{\Lambda(L)}^{\omega}} \omega'(\tilde{H}^{\Lambda(L)}), \quad (\text{S69})$$

where

$$C_{\Lambda(L)}^{\omega} := \{\omega' \text{ is a state on } \mathcal{A} \mid \omega'|_{\mathcal{A}_{\Lambda(L)^c}} = \omega|_{\mathcal{A}_{\Lambda(L)^c}}\}. \quad (\text{S70})$$

In this sense, the exterior restriction $\omega|_{\mathcal{A}_{\Lambda(L)^c}}$ can be regarded as a boundary condition of the finite subsystem $\Lambda(L)$. Thus, every KMS ground state is locally a finite-volume ground state in this conditional sense, with boundary condition given by its own exterior restriction. Consequently, every KMS ground state arises as a weak-* limit point of such finite-volume ground states for a suitable choice of these generalized boundary conditions. This notion is more general than the usual boundary conditions of Theorem 4.1.

4.2 Revisiting 2D spectral gap undecidability through bulk-gap SDP hierarchy

The authors of [CPW15] construct, for each input m , a family of Hamiltonians $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ on the two-dimensional lattice \mathbb{Z}^2 with the following properties: the interaction is translationally invariant, nearest-neighbour, and has uniformly bounded coefficients, and the Hilbert space dimension on each site is uniformly bounded. For open boundary conditions, each such family falls into one of two promised cases: it is either non-degenerate uniformly gapped in the sense of Theorem 4.2 with the gap $\gamma_{\text{CPW}}(m) \geq 1$, or uniformly gapless in the sense of Theorem 4.5.

Crucially, the spectral behavior of $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ under open boundary conditions encodes the halting problem for a fixed universal Turing machine (UTM):

$$\left\{ \begin{array}{l} \text{YES: UTM halts on } m \implies \{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L \text{ is non-degenerate uniformly gapped} \\ \quad \text{under OBC (Theorem 4.2) with } \gamma_{\text{CPW}}(m) \geq 1, \\ \text{NO: UTM does not halt on } m \implies \{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L \text{ is uniformly gapless, and thus gapless under OBC} \\ \quad \text{(Theorem 4.5, and thus Theorem 4.4).} \end{array} \right. \quad (\text{S71})$$

In particular, deciding whether $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ is gapped or gapless under open boundary conditions is equivalent to deciding whether the UTM halts on m , and is therefore undecidable.

Now, let us apply our SDP hierarchy for bulk spectral gaps (Theorems 2.1 and 2.4) to the decision problem Eq. (S71) of $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$. Specifically, consider the following algorithm:

Algorithm 1 Spectral gap semi-decision via SDP hierarchy (Theorems 2.1 and 2.4)

Require: Dynamical system $(\mathcal{A}, \tau_{\mathcal{F}_{\text{CPW}}(m)}, \delta_{\mathcal{F}_{\text{CPW}}(m)})$ for $\mathcal{F}_{\text{CPW}}(m) = \{H_{\text{CPW}}^{\Lambda(L)}(m)\}$

Ensure: Fixed $\gamma_0 \in (0, 1)$, returns **NO** if no KMS ground state with locally non-degenerate bulk gap $\geq \gamma_0$ exists (equivalently, if the bulk spectral gap of the dynamical system is $< \gamma_0$); otherwise loops forever.

- 1: **loop**
 - 2: Increment both n and d by 1 (starting at $n = d = 1$)
 - 3: Solve the level- (n, d) feasibility SDP of Theorems 2.1 and 2.4 with γ_0
 - 4: **if** SDP is infeasible **then**
 - 5: **return NO**
 - 6: **end if**
 - 7: **end loop**
-

The completeness of our SDP hierarchy (Theorem 2.5) has the following implications:

1. If the input $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ corresponds to a **YES**-instance of the UTM, then with the open boundary condition, the system admits a non-degenerate uniform gap ≥ 1 (in the sense of Theorem 4.2). By Theorem 4.3, there exists a KMS ground state with a local non-degenerate bulk spectral gap ≥ 1 (in the sense of Theorem 1.6) due to [BDN16, Proposition 5.4]. Consequently, Algorithm 1, which certifies the existence of a KMS ground state with a gap of at least $\gamma_0 \in (0, 1)$, runs indefinitely.
2. If Algorithm 1 returns **NO** at some (n_0, d_0) , then we have a proof that no KMS ground state with a gap ≥ 1 exists, including the limiting state ω_{OBC} from the ground states due to the open boundary condition (Theorem 4.6). By the spectral property of $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$, it follows that the system must be gapless with the open boundary condition (Theorem 4.4).

That is, Algorithm 1 can semi-decide if the given $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ is gapless with the open boundary condition at some finite level (n_0, d_0) and thus in finite time.

This appears to be in tension with the well-known undecidability results of [CPW15] by the following statement:

(A) Run the UTM and Algorithm 1 simultaneously. In the **YES**-instance of Eq. (S71), the UTM halts in finite time and Algorithm 1 runs indefinitely. In the **NO**-instance of Eq. (S71), the UTM runs indefinitely. But, since the system is gapless, Algorithm 1 halts in finite time. Therefore, one can fully decide the halting problem in finite time with Algorithm 1, which is a contradiction.

The Statement (A) is, however, *false* due to a mix-up of various definitions modeling spectral gaps in the thermodynamic limit and dependence on boundary conditions. We explain this in the following remark.

Remark 4.7 (Compatibility with [CPW15]). *Let us analyze two cases of the given UTM on input m .*

- (i) If the UTM halts: *This is already discussed by bullet point 1. Our SDP hierarchy is feasible at every level and Algorithm 1 runs indefinitely.*
- (ii) If the UTM does not halt: *The system is proven to be gapless with open boundary condition in the sense of Theorem 4.4. Let ω_{OBC} be a weak-* limit point of finite-volume ground states of $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}$ with open boundary conditions. However, it is not immediate that this finite-volume gaplessness implies that ω_{OBC} is bulk-gapless in the GNS sense. Moreover, this does not exclude the possibility that there exists another KMS ground state which admits a locally non-degenerate bulk spectral gap of at least 1. (By the equivalence of Theorem 4.6, there might exist a different boundary condition B such that $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}$ admits a limit state with a locally non-degenerate spectral gap.) If this is the case, Algorithm 1 does not terminate in finite time necessarily, and therefore falsifies the Statement (A).*

To sum up, the **NO**-instance of the decision problem Eq. (S71) does not mean that Algorithm 1 necessarily halts in finite time, due to, e.g., the possibility of an alternative boundary condition for which the spectral gap is gapped. Nonetheless, such dependence on the boundary conditions is expected since they serve as the initial states to the UTM in the construction of [CPW15].

We conclude this subsection by emphasizing that, for the Hamiltonians constructed by [CPW15], spectral properties in the thermodynamic limit can depend on the choice of generalized boundary conditions.

Proposition 4.8. Consider the family $\{H_{\text{CPW}}^{\Lambda(L)}(m)\}_L$ associated with the fixed UTM in [CPW15]. Fix any threshold $\gamma_0 \in (0, 1)$. Then there exist a non-halting input m' and a KMS ground state $\tilde{\omega}$ with a locally non-degenerate bulk spectral gap at least γ_0 .

It follows from Theorem 4.6 that this $\tilde{\omega}$ is a weak-* limit of finite-volume ground states with some generalized boundary condition.

Proof. Suppose not. Then for every non-halting input m' , every KMS ground state of the dynamical system induced by $\{H_{\text{CPW}}^{\Lambda(L)}(m')\}_L$ has locally non-degenerate bulk gap strictly smaller than 1. By the completeness of the SDP hierarchy (Theorem 2.5), it follows that running Algorithm 1 with γ_0 on the instance $\{H_{\text{CPW}}^{\Lambda(L)}(m')\}_L$ must terminate in finite time for every such m' . Consequently, Statement (A) can be used to contradict the undecidability of the halting problem. Finally, we invoke the discussion of Theorem 4.6. \square

4.3 The 1D construction and why bulk gap undecidability remains open

A natural question is whether the *algebraic bulk* spectral-gap problem is undecidable, in analogy with the undecidability results for finite-volume gaps under fixed boundary conditions. As discussed in Section 4.2, our SDP hierarchy (Algorithm 1) semi-decides the absence of KMS ground states with locally non-degenerate bulk spectral gaps in finite time. Equivalently, it semi-decides when a proposed γ lies above the bulk spectral gap of Theorem 1.10. Thus, to establish undecidability in the bulk sense, one would need a fixed UTM and a computable map $m \mapsto \{H_{\text{bulk}}^{\Lambda(L)}(m)\}_L$ such that

$$\left\{ \begin{array}{l} \mathbf{YES}: \text{UTM halts on } m \implies \text{every KMS ground state is bulk-gapless} \\ \hspace{10em} \text{in the sense of Theorem 1.11,} \\ \mathbf{NO}: \text{UTM does not halt on } m \implies \text{there exists a locally non-degenerate bulk-gapped KMS} \\ \hspace{10em} \text{ground state in the sense of Theorem 1.6 with } \gamma_{\text{bulk}}(m) \geq 1. \end{array} \right. \quad (\text{S72})$$

Note that this has the opposite halting/gap direction compared to Eq. (S71).

A natural candidate to obtain such a reversal is the one-dimensional construction of [Bau+20]. There, for each input m , the authors construct a family of translationally invariant nearest-neighbour Hamiltonians $\{H_{\text{BCLP}}^{\Lambda(L)}(m)\}_L$ on the one-dimensional lattice \mathbb{Z} . The local interaction terms are uniformly bounded and the Hilbert space dimension of each site is uniformly bounded. Under open boundary conditions, the family is promised to fall into one of two cases: it is either non-degenerate uniformly gapped in the sense of Theorem 4.2 with the gap $\gamma_{\text{BCLP}}(m) \geq 1$, or uniformly gapless in the sense of Theorem 4.5.

As in the 2D case, undecidability follows because the spectral behavior under open boundary conditions encodes the halting problem for a fixed UTM; however, the correspondence is reversed relative to [CPW15]:

$$\left\{ \begin{array}{l} \mathbf{YES}: \text{UTM halts on } m \implies \{H_{\text{BCLP}}^{\Lambda(L)}(m)\}_L \text{ is uniformly gapless, and thus gapless under OBC} \\ \hspace{10em} \text{(Theorem 4.5, and thus Theorem 4.4).} \\ \mathbf{NO}: \text{UTM does not halt on } m \implies \{H_{\text{BCLP}}^{\Lambda(L)}(m)\}_L \text{ is non-degenerate uniformly gapped} \\ \hspace{10em} \text{under OBC (Theorem 4.2) with } \gamma_{\text{BCLP}}(m) \geq 1, \end{array} \right. \quad (\text{S73})$$

Moreover, they obtain a partial extension to periodic boundary conditions for chain lengths promised to be coprime to a fixed integer P , at the cost of a local dimension growing linearly with P .

However, the construction of [Bau+20] does not directly yield Eq. (S72). Indeed:

1. If the UTM does not halt on m , then the KMS ground state induced by the weak-* limit of open boundary ground states (denoted by ω_{OBC}) is a locally non-degenerate bulk-gapped KMS ground state. This matches the **NO**-instance of Eq. (S72).
2. If the UTM halts on m , then Eq. (S73) does not immediately imply that the limit KMS ground state ω_{OBC} is bulk-gapless. Furthermore, this does not exclude the existence of other KMS ground states (arising from different boundary conditions) that remain locally non-degenerate bulk-gapped. Hence, the **YES**-instance of Eq. (S72) is not established.

In summary, while it is plausible that the bulk spectral-gap problem is undecidable, a proof would require additional ingredients ensuring that *all* KMS ground states are bulk-gapless in the halting case, rather than only the state selected by a particular boundary condition. For this reason, undecidability in the algebraic bulk sense remains open.