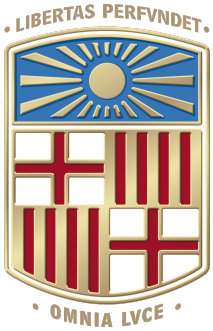

Towards Efficient Spatial Variational 2-RDM via Measurement Constraints

Annika Nel

Supervised by: Antonio Acín, Leonardo Zambrano

A thesis presented for the Master in Quantum Science and Technology
at Universitat de Barcelona

Quantum Information Theory Group
Institute of Photonic Sciences
Barcelona, Catalunya

August 2025

Master in Quantum Science and Technology Barcelona

ICFO

Abstract

Reduced density matrices (RDMs) offer a more scalable alternative to full wavefunctions when performing chemical calculations. The variational two-electron RDM (v2RDM) method exploits the efficiency of RDMs, employing semidefinite programming (SDP) to enable polynomial scaling of ground state simulations. Recent work by Avdic & Mazziotti seeks to improve the performance of the v2RDM by incorporating classical shadow constraints, simultaneously reducing the number of measurements required for tomography. Drawing from this work, we introduce a spatial orbital variant of the v2RDM with measurement constraints (m-v2RDM). The proposed method achieves comparable accuracy for small to medium-sized molecules such as H_2, H_4, and HF, while substantially reducing memory and runtime costs. Its comparatively simple implementation also allows for the approximation of larger systems like N_2, which are otherwise intractable on modest computational resources using standard v2RDM. As a pedagogical resource, the spatial variant more closely resembles the underlying theory, making it an accessible introduction to RDMs. The spatial m-v2RDM further highlights the complementary nature of measurement constraints and N-representability conditions, framing the RDM as a potential tool for noise mitigation in quantum information processing.

Keywords: Reduced Density Matrices, Semidefinite Programming, N-representability, Quantum Chemistry, Quantum Tomography, Classical Shadows

Acknowledgements

I would like to acknowledge my supervisors, Antonio Acín and Leonardo Zambrano, for making this work possible. To Toni, for his expert guidance, and for giving me the opportunity to explore my interests in tomography and quantum simulation at a world-class institution like ICFO. To Leo, whose kindness, constant support, and readiness to help lightened every challenge. My sincere thanks to you both.

I am grateful to the coordinators and lecturers of the Master in Quantum Science and Technology, Barcelona, for their dedication to their students. You play an invaluable role in shaping future quantum scientists.

A special thank you to Teo and Júlia for their collaboration throughout this work. To Teo, for his patience and careful explanation during our many conversations on quantum chemistry. To Júlia, for her sharp insight, creativity, and for graciously lending me her laptop when mine was damaged, allowing me to finish this thesis in time.

I also wish to thank the members of the Quantum Information Theory Group at ICFO, who were always welcoming and willing to share ideas.

My deepest gratitude goes to my mom and sister, who have always supported my dreams even in the hardest times.

To the friends I made in the Master, especially Marco, Lucía, and Julie - thank you for making this year unforgettable. I feel lucky to have met such wonderful people.

To Salma and Victoria, you were my home away from home.

Finally, I express my sincere appreciation to ICFO for financially supporting me through both the Catalonia Quantum Academy Scholarship and the María Yzuel Fellowship Award. Without this generosity, I would not have been able to pursue this Master's program. Thank you for helping me start my career in quantum sciences.

This thesis is dedicated to my dad, Gideon, whose example instilled in me an enduring curiosity and love for nature. *Toe word jy reën...*

Contents

1 Introduction

Accurate simulation of molecular systems remains a central challenge in applications of quantum chemistry. Many observables, such as the ground-state energy, can be computed using only p-local interactions, which are fully described by the corresponding p-particle reduced density matrix (p-RDM). Unlike the exponential scaling of wavefunction approaches, RDMs typically enable polynomial scaling with system size. As such, RDM theory is an active area of research that has influenced the development of modern computational chemistry tools [Maz11b, Lev13, MOC+24].

RDMs have recently been recognized in quantum information processing for their potential to mitigate exponential resource requirements. Characterizing a subset of the system using RDMs presents a more efficient alternative to full state tomography, which scales exponentially with system size. Indeed, RDMs have shown promise as tools for error correction in quantum algorithms [RBM18, WDPT25].

Variational methods have emerged as a notable technique for calculating RDMs of molecular ground states. The variational 2-RDM (v2RDM) approach optimizes the 2-RDM subject to N-representability conditions, which ensure that the resulting matrix corresponds to a valid many-electron wavefunction [Maz07]. Since the N-representability conditions can be expressed as a semidefinite program (SDP), the v2RDM benefits from convex optimization theory, yielding highly accurate results. The v2RDM has been successfully applied to molecules as large as 64 electrons [MMK+19]. However, such calculations remain computationally demanding and must compete with modern configuration interaction (CI) methods [ED24].

Recent work has sought to reduce resource costs by combining the v2RDM with techniques from quantum tomography. In particular, the framework of classical shadows allows for the estimation of many observables by storing an efficient classical representation of a quantum state [HKP20]. In their 2024 paper entitled "*Fewer measurements from shadow tomography with N-representability conditions*", Avdic & Mazziotti integrate classical shadows with the v2RDM to accelerate convergence and improve performance [AM24b]. Their results demonstrate that shadow constraints improve the accuracy of the v2RDM, while N-representability conditions in turn enhance shadow prediction and reduce the number of measurements required for a given accuracy.

Building on this work, we propose a spatial orbital variant of v2RDM with additional measurement constraints (m-v2RDM). This new approach offers a simpler implementation than existing algorithms, and substantially reduced memory and runtime demands. It achieves a comparable accuracy for small molecules and approximates larger systems with modest resources, albeit at reduced accuracy. Furthermore, its clarity and ease of implementation make it easily extendable and valuable as a pedagogical resource.

The remainder of this thesis is organized as follows. Section 2 reviews the theoretical background of SDPs, N-representability, quantum tomography, and classical shadows. It concludes by unifying these concepts in the v2RDM and examining the shadow constraints of Avdic & Mazziotti. In Section 3, we address some essential practical considerations for N-representability and review existing v2RDM implementations. Subsequently, we motivate the spatial m-v2RDM and detail its implementation. Section 4 presents the performance of the proposed method and compares it with Avdic and Mazziotti's shadow v2RDM, followed by conclusions in Section 5. Finally, Section 6 discusses directions for future work and potential improvements.

2 Theoretical Background

This section begins with an overview of semidefinite programming, followed by an introduction to the N-representability problem and the role of reduced density matrices in quantum chemistry. Thereafter, we review quantum state tomography and classical shadows. We conclude with a synthesis of these concepts in the variational 2-RDM with shadows.

2.1 Semidefinite Programming

In *semidefinite programming*, we aim to optimize a linear function of some matrix variable over a convex set, subject to linear constraints. It is typically expressed using both primal and dual formulations. The standard primal formulation of an SDP reads,

$$\begin{aligned} & \text{maximize} && \operatorname{Tr}(AX) \\ \text{subject to} \quad & && \Phi_i(X) = B_i, \quad i = 1, \ldots, m \\ & && \Gamma_j(X) \leq C_j, \quad j = 1, \ldots, n \quad \text{i.e.} \quad C_j - \Gamma_j(X) \geq 0 \end{aligned} \tag{1}$$

where X is a Hermitian operator that constitutes the optimization variable. Additionally, the matrices A, B_i, and C_j are Hermitian operators that define the linear objective function, equality constraints, and inequality constraints, respectively. Note that any linear function of X can be expressed as $\operatorname{Tr}(AX)$. Finally, the linear maps $\Phi_i(\cdot)$ and $\Gamma_j(\cdot)$ are required to be Hermiticity-preserving. The need for Hermiticity becomes clear when we consider that the notion of positive semidefiniteness is only well-defined for matrices with real eigenvalues, i.e. Hermitian matrices. As such, both the variables and the constraint functions must preserve Hermiticity to ensure the well-definedness of the solution space.

Every primal SDP possesses a complementary dual formulation. While the primal problem seeks to maximize its objective function, the dual seeks to minimize, thus providing an upper bound on the optimal solution to the primal problem. Similarly, solutions to the primal problem provide a lower bound on the optimal solution to the dual problem.

The dual problem can be derived by introducing Lagrange multipliers to the primal problem. After some simplification, we arrive at the following formulation of the dual:

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^{m} \operatorname{Tr}(Y_i B_i) + \sum_{j=1}^{n} \operatorname{Tr}(Z_j C_i) \\ \text{subject to} \quad & A - \sum_{i=1}^{m} \Phi_i^\dagger(Y_i) - \sum_{j=1}^{n} \Gamma_j^\dagger(Z_j) = 0, \\ & Z_j \succeq 0 \quad j = 1, \ldots, n. \end{aligned} \tag{2}$$

where Y_i and Z_j are the Lagrange multipliers associated with the equality and inequality constraints on the primal problem, respectively. Unsurprisingly, both Y_i and Z_j should be Hermitian. The full derivation can be found in Chapter 2 of [SC23].

Having both the primal and dual solutions to an optimization problem is useful for several reasons. Firstly, under mild assumptions, SDPs exhibit *strong duality*. That is, the optimal solutions to the primal and dual coincide. Hence, the self-consistency of SDPs can be used to validate the optimal solution. Additionally, when strong duality does not hold, the primal and dual can provide lower and upper bounds the the optimal solution. Finally, the dual is at times a more efficient or practical alternative to the primal.

SDPs have become ubiquitous owing to the fairly recent development of efficient solvers like the interior-point [Kar84, AA00] and first-order algorithms [OCPB16]. Their convexity guarantees a global optimum as opposed to several local optima, making them considerably more tractable than non-convex optimization problems. Furthermore, many quantum

information tasks are naturally framed as SDPs, as they deal primarily with Hermitian operators.

2.2 The N-representability Problem

2.2.1 Two-particle Reduced Density Matrices

In quantum chemistry, it is of particular interest to describe the correlations between pairs of electrons. Since electrons are indistinguishable and interact under Coulomb repulsion, many properties of an N-electron system can be computed from at most pairwise interactions [May55, L̈55]. These are fully characterized by the two-electron reduced density matrix or 2-RDM, defined as

$$ {}^2D^{ij}_{kl} = \langle\psi|\, a^\dagger_i a^\dagger_j a_l a_k \,|\psi\rangle \tag{3} $$

where $|\psi\rangle$ is typically the fermionic ground-state wavefunction [Maz07]. Accordingly, $a^\dagger_i$ and a_j denote the fermionic creation and annihilation operators in second quantization. In the context of molecules, $a^\dagger_i$ creates an electron in the ith orbital, while a_j annihilates an electron in the jth orbital.

Consider the Hamiltonian for an N-electron system in Born-Oppenheimer approximation:

$$ \hat{H} = \sum_{ij} h^i_j a^\dagger_i a_j + \frac{1}{2}\sum_{ijkl} h^{ij}_{kl} a^\dagger_i a^\dagger_j a_l a_k \tag{4} $$

where h^i_j and h^{ij}_{kl} represent the one-body kinetic energy operator and two-body interactions, respectively. The one-body terms can be incorporated into the two-body terms to give a more compact form of the Hamiltonian,

$$ \hat{H} = \sum_{ijkl} {}^2K^{ij}_{kl} a^\dagger_i a^\dagger_j a_l a_k \tag{5} $$

where ${}^2K^{ij}_{kl}$ is the reduced Hamiltonian operator, given by

$$ {}^2K^{ij}_{kl} = \frac{1}{N-1}(h^i_k \delta^j_l + h^j_l \delta^i_k) + h^{ij}_{kl}. \tag{6} $$

Taking the expectation value of Eq.(5), we see that the electronic energy can be expressed entirely in terms of the 2-RDM as defined in Eq.(3),

$$ E = \sum_{ijkl} {}^2K^{ij}_{kl}\, {}^2D^{ij}_{kl} = \mathrm{Tr}({}^2K^2D) \tag{7} $$

This convenient representation allows one to calculate the energy of an atom or molecule using *only* the 2-RDM. Indeed, any observable involving at most two-body interactions can be calculated using the 2-RDM [May55, Col63, Maz07]. The potential of the 2-RDM was first emphasized by Coulson in 1955, when he questioned whether it could ultimately replace the wavefunction in chemical calculations entirely - a remark that became known as *Coulson's Challenge* [Jud01].

By the variational principle, the ground-state energy is always less than or equal to the expectation value of the Hamiltonian calculated with some trial wavefunction. Therefore, one can treat ${}^2D$ as a variable over which the ground-state energy in Eq. (7) can be minimized to obtain the 2-RDM, effectively solving the many-body Schrödinger equation [L̈55]. This approach has the especially attractive advantage that the 2-RDM scales polynomially with the number of orbitals, unlike the exponential scaling of configuration interaction methods [LTWHG17].

However, early variational energy calculations by Mayer and Lowdin [May55, L̈55] demonstrated that 2-RDMs constrained solely by wavefunction normalization produced energies that were too low. Rather, additional constraints were required to guarantee that the 2-RDM obtained via minimization could be obtained from the integration of an N-electron density matrix. These became known as the *N-representability* conditions and would be vital for realizing Coulson's vision.

2.2.2 N-representability

The 2-RDM is obtained by integrating over all but two particles of the full N-electron density matrix,

$$^2D = \mathrm{Tr}_{3\ldots N}(^N D) \tag{8}$$

where

$$^N D = |\psi\rangle\langle\psi| = \sum {}^N D^{i_1,\ldots,i_N}_{j_1,\ldots,j_N} \, |i_N \cdots i_1\rangle\langle j_N \cdots j_1| \, . \tag{9}$$

It is indeed possible to define a p-RDM by tracing out all but p of the N electrons:

$$^p D = \mathrm{Tr}_{p+1\ldots N}(^N D) \tag{10}$$

resulting in a set of $\binom{N}{p}$ p-RDMs. Similarly to the 2-RDM, any p-local observable can be calculated using only the set of p-RDMs [RBM18]. It is immediately evident that the set of p-RDMs needs to be consistent. That is, each element should be contractable from a single global state. For general quantum systems, ensuring consistency is referred to as the *quantum marginal problem* and is known to be QMA-complete [Liu06].

The fermionic version of the quantum marginal problem is known as the *N-representability problem* [Kly06]. In addition to consistency and general requirements for density matrices, fermionic p-RDMs are further subject to *N-representability* conditions. First conceptualized by Coleman in 1963 [Col63], these constitute the necessary and sufficient conditions to ensure that a given p-RDM can be derived from an antisymmetric N-particle wavefunction.

Finding the full set of conditions proved to be a decades-long endeavour [LTC13, Maz07]. To date, most progress has been made with respect to 2-RDMs, with the complete N-representability for 2-RDMs being formulated by Mazziotti as recently as 2012 [Maz12]. Arguably the most significant advancement came in the form of Mazziotti and Erdahl's *p-positivity* conditions [ME01], which enabled the N-representability problem to be cast as a semidefinite program. This is expanded upon for the 2-RDM in Section 2.4.

2.2.3 The N-representable 2-RDM

The basic conditions for N-representability of the 2-RDM follow from the general requirements for density matrices established by Von Neumann. A valid fermionic 2-RDM is required to be at least,

(i) Positive semidefinite i.e. all its eigenvalues should be non-negative (to ensure non-negative probabilities),

$$^2D \succeq 0 \tag{11}$$

(ii) Antisymmetric (due to fermionic anticommutation relations),

$$^2D^{ij}_{kl} = -{}^2D^{ij}_{lk} = {}^2D^{ji}_{lk} = -{}^2D^{ji}_{kl} \tag{12}$$

(iii) Hermitian

$$\left({}^2D^{ij}_{kl}\right)^* = \left(\langle\psi| \, a^\dagger_i a^\dagger_j a_l a_k \, |\psi\rangle\right)^* = \langle\psi| \, a^\dagger_k a^\dagger_l a_j a_i \, |\psi\rangle = {}^2D^{kl}_{ij} = {}^2D^{ij}_{kl} \tag{13}$$

(iv) Normalized (to conserve particle number)

$$\mathrm{Tr}({}^2D) = \tfrac{1}{2}N(N-1) \tag{14}$$

However, the above conditions are not sufficient to ensure that a 2-RDM is N-representable, that is, traceable from an N-electron density matrix. While not complete, an important subset of the sufficient conditions are known as the *p-positivity* conditions [ME01]. We derive them as follows:

For a p-particle system, consider that the overlap matrices, M_J^I, must be positive semidefinite:

$$M_J^I = \langle \Phi_I | \Phi_J \rangle = \langle \Psi | \hat{C}_I \hat{C}_J^\dagger | \Psi \rangle \succeq 0 \tag{15}$$

where $\hat{C}_I$ is a product of p creation and annihilation operators, $|\Psi\rangle$ is the groundstate, and the set of basis functions $\langle \Phi_I|$ are defined as

$$\langle \Phi_I | = \langle \Psi | \, \hat{C}_I. \tag{16}$$

When constructing $\hat{C}_I$, permuting $a_i^\dagger$ and a_i admits $p+1$ different overlap matrices, each of which are required to be positive semidefinite. The one- and two-electron cases yield the overlap matrices ${}^1D$, ${}^1Q$, ${}^2D$, ${}^2Q$ and ${}^2G$ given in Table 1.

$\hat{C}_I$	RDM	Description
$a_i^\dagger$	${}^1D_j^i = \langle \psi \vert \, a_i^\dagger a_j \, \vert \psi \rangle$	one-electron
a_i	${}^1Q_j^i = \langle \psi \vert \, a_i a_j^\dagger \, \vert \psi \rangle$	one-hole
$a_i^\dagger a_j^\dagger$	${}^2D_{kl}^{ij} = \langle \psi \vert \, a_i^\dagger a_j^\dagger a_l a_k \, \vert \psi \rangle$	two-electron
$a_i a_j$	${}^2Q_{kl}^{ij} = \langle \psi \vert \, a_i a_j a_l^\dagger a_k^\dagger \, \vert \psi \rangle$	two-hole
$a_i^\dagger a_j$	${}^2G_{kl}^{ij} = \langle \psi \vert \, a_i^\dagger a_j a_l^\dagger a_k \, \vert \psi \rangle$	electron-hole

Table 1: Overlap matrices generated for the one and two-electron case.

Hence, we obtain the *2-positivity* conditions,

$${}^2D \succeq 0 \tag{17}$$

$${}^2Q \succeq 0 \tag{18}$$

$${}^2G \succeq 0 \tag{19}$$

Furthermore, rearranging the creation and annihilation operators allow ${}^2Q$ and ${}^2G$ to be expressed as linear functions of ${}^2D$:

$${}^2Q_{kl}^{ij} = 2\,{}^2I_{kl}^{ij} - 4\,{}^1D_k^i \wedge {}^1I_l^j + {}^2D_{kl}^{ij} \tag{20}$$

$${}^2G_{kl}^{ij} = {}^1I_l^j\,{}^1D_k^i - {}^2D_{kj}^{il} \tag{21}$$

as first derived by Garrod and Percus in 1964 [GP64]. Here, $\wedge$ denotes the Grassmann wedge product[1] and ${}^1D$ is the 1-RDM which is related to ${}^2D$ by a partial trace,

$${}^1D_k^i = \frac{1}{N-1}\sum_j {}^2D_{kj}^{ij}. \tag{22}$$

Note that the 2-positivity conditions for ${}^2D$ and ${}^2Q$ imply the 1-positivity conditions for ${}^1D$ and ${}^1Q$, which are necessary and sufficient for N-representability of the 1-RDM [Col63].

[1] Antisymmetric tensor product given in appendix A.

2.3 Quantum Tomography

2.3.1 Quantum State Tomography

Quantum state tomography refers to the estimation of an unknown quantum state, ρ, by measuring a finite number of copies of that state. If the appropriate measurement scheme is chosen, ρ can be entirely characterized by the probability vector, $\boldsymbol{p}$, which results from repeated measurement of the state. The relationship between ρ and $\mathbf{p}$ is defined by a *tomographic map*,

$$\mathcal{T} : \rho \mapsto \boldsymbol{p} \tag{23}$$

which is required to be linear in ρ and injective, i.e. $\rho \neq \rho' \implies \mathcal{T}(\rho) \neq \mathcal{T}(\rho')$. More simply, the chosen measurement scheme should map distinct quantum states to distinct probability distributions - a property referred to as *tomographical completeness* [BCMTS24]. The tomographical completeness of $\mathcal{T}$ guarantees that we can unambiguously reconstruct ρ from $\boldsymbol{p}$ by inverting the map:

$$\rho = \mathcal{T}^{-1}(\boldsymbol{p}). \tag{24}$$

Quantum measurement is typically performed using positive operator-valued measures (POVMs). Let $\{E_i\}$ be a set of positive semi-definite operators $E_i \in \mathcal{L}(\mathcal{H}_d)$ such that $\sum_i E_i = \mathbb{1}$, where $\mathcal{L}(\mathcal{H}_d)$ is the set of Hermitian operators on d-dimensional Hilbert space $\mathcal{H}_d$. The set $\{E_i\}$ constitutes a POVM. As per the Born rule, the probability of measuring outcome i is given by

$$p_i = \mathrm{Tr}(E_i \rho). \tag{25}$$

To ensure tomographical completeness, we can select a POVM that is *informationally complete*. An informationally complete POVM is one for which the set of operators $\{E_i\}$ spans $\mathcal{L}(\mathcal{H}_d)$ [Cze21].

In practice, we cannot obtain the exact probability vector $\boldsymbol{p}$ since we are limited to a finite number of copies of ρ. Instead, we can estimate $\boldsymbol{p}$ using empirical frequencies, $\boldsymbol{f}$ with

$$f_i = \frac{n_i}{N} \tag{26}$$

where n_i is the number of times outcome i was observed when measuring N copies of ρ. By the law of large numbers,

$$\boldsymbol{f} \to \boldsymbol{p} \quad \text{as} \quad N \to \infty\,, \qquad \text{and} \qquad \langle \boldsymbol{f} \rangle = \boldsymbol{p} \tag{27}$$

where $\langle \boldsymbol{f} \rangle$ is the expectation value of $\boldsymbol{f}$. We can thus construct a map, $\boldsymbol{f} \mapsto \hat{\rho}$ called an *estimator*, such that we estimate the state ρ using:

$$\hat{\rho} = \mathcal{T}^{-1}(\boldsymbol{f}) \tag{28}$$

For example, consider the *least squares estimator* (LSE), where

$$\hat{\rho} = \underset{\rho \in \mathcal{L}(\mathcal{H}_d)}{\arg\min} ||\boldsymbol{f} - \mathcal{T}(\rho)||_2^2 \tag{29}$$

such that ρ is a valid density matrix, i.e. $\rho \succeq 0$ and $\mathrm{Tr}(\rho) = 1$. This estimator aims to minimize the square difference between the empirical frequency, $\boldsymbol{f}$, and the probability distribution generated by ρ.

2.3.2 Classical Shadows

As quantum devices grow in size, conventional tomographical techniques have become less practical since they suffer from the *curse of dimensionality*. That is, the number of copies required to accurately reconstruct a quantum system scales exponentially with the size of that system, consequently demanding exponential classical memory and computing power. In the ongoing effort to develop more efficient techniques, recent breakthroughs include *shadow tomography*.

The central idea of shadow tomography is to estimate selected target functions of a quantum state without full state reconstruction. Many interesting properties of quantum states are linear functions their density matrix. For example, the probability distribution, fidelity with a pure target state, and entanglement witnesses are all examples of functions that take the form of expectation values:

$$o_i(\rho) = \mathrm{Tr}(O_i\rho) \qquad 1 \leq i \leq M \tag{30}$$

where o_i is the expectation value associated with observable O_i. Limiting the prediction task to only these M target functions makes it possible to predict an exponential number of target functions from only a polynomial number of samples. However, shadow tomography as originally proposed by Aaronson is difficult to implement, since it requires exponentially long quantum circuits and all copies of the state to be stored in quantum memory [Aar18].

Classical shadows is one fairly recent technique designed to address this, in which an efficient classical representation of a quantum state suffices to predict *any* linear function of that state in expectation [HKP20]. Crucially, both the memory and number of copies required do not depend on the size of the system (number of qubits) [HKP20]. Rather, the resource requirements scale with the choice of measurement and the number of linear target functions to be estimated. The procedure is described for an n-qubit state as follows:

Randomly select a unitary U from a tomographically complete fixed ensemble $\mathcal{U}$ and apply it to the state, $\rho \mapsto U\rho U^\dagger$. Measure the rotated state in the computational basis to obtain an n-bit measurement outcome, $|\hat{b}\rangle \in \{0,1\}^n$. Thereafter, store an efficient classical description of $U^\dagger|\hat{b}\rangle\langle\hat{b}|U$ in classical memory. In expectation, the above process can be viewed as a quantum channel, $\rho \mapsto \mathcal{M}(\rho)$:

$$\mathbb{E}\left[U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right] = \mathcal{M}(\rho) \tag{31}$$

The tomographical completeness of $\mathcal{U}$ guarantees that $\mathcal{M}$ is invertible, resulting in an estimator,

$$\hat{\rho} = \mathcal{M}^{-1}\left(U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right) \tag{32}$$

called a *classical shadow*. In expectation, we recover the original state exactly: $\rho = \mathbb{E}[\hat{\rho}]$. The inverse mapping $\mathcal{M}^{-1}$ depends on the measurement ensemble, $\mathcal{U}$, and generally involves a non-trivial derivation. Fortunately, Huang et al. provide the Clifford and Pauli measurements as examples.

In order to estimate a target function, we repeat the process to generate a set of N independent shadows, $\{\hat{\rho}_1, ..., \hat{\rho}_N\}$ and average over the estimators to predict the target function:

$$\hat{o}_i(N,1) = \frac{1}{N}\sum_{j=1}^{N} \mathrm{Tr}\left(O_i\hat{\rho}_j\right) \tag{33}$$

For greater accuracy, Huang et al. recommend using *median of means* to estimate the target function instead of just the sample mean: construct K independent sample means to form the set

$$\hat{o}_i(N,K) = \text{median}\left\{\hat{o}_i^{(1)}(N,1),\ldots,\hat{o}_i^{(K)}(N,1)\right\} \quad \text{where} \quad \hat{o}_i^{(k)} = \frac{1}{N}\sum_{j=N(k-1)+1}^{Nk} \text{tr}\left(O_i\hat{\rho}_j\right) \tag{34}$$

with $1 \leq k \leq K$. Theorem 1 provides the corresponding accuracy guarantee,

Theorem 1. *Fix a measurement primitive $\mathcal{U}$, a collection $O_1,...,O_M$ of $2^n \times 2^n$ Hermitian matrices and accuracy parameters $\epsilon, \delta \in [0,1]$. Set*

$$K = 2\log(2M/\delta) \quad \text{and} \quad N = \frac{34}{\epsilon^2}\max_{1\leq i\leq M}\left\|O_i - \tfrac{\text{Tr}(O_i)}{2^n}\mathbb{I}\right\|_{\text{shadow}}^2 \tag{35}$$

where $\|\cdot\|_{\text{shadow}}$ denotes the shadow norm. Then, a collection of NK independent classical shadows allows for accurately predicting all features via median of means prediction:

$$|\hat{o}_i(N,K) - \text{Tr}(O_i\rho)| \leq \epsilon \quad \textit{for all} \quad 1 \leq i \leq M \tag{36}$$

with probability at least $1-\delta$.

The shadow norm is defined as

$$\|O\|_{\text{shadow}} = \max_{\sigma:\,\text{state}}\left(\mathbb{E}_{U\sim\mathcal{U}}\sum_{b\in\{0,1\}^n}\langle b|U\sigma U^\dagger|b\rangle\langle b|U\mathcal{M}^{-1}(O)U^\dagger|b\rangle^2\right)^{1/2}. \tag{37}$$

The above theorem implies a sample complexity of

$$N_{\text{tot}} = \mathcal{O}\left(\frac{\log M}{\epsilon^2}\max_{1\leq i\leq M}\left\|O_i - \tfrac{\text{Tr}(O_i)}{2^n}\mathbb{I}\right\|_{\text{shadow}}^2\right). \tag{38}$$

2.4 Variational 2-RDM with Shadows

It is easy to show that the set of N-representable p-RDMs, ${}_p^N P$, is convex [Maz12]. When restricted to only the p-positivity conditions, we obtain a set of approximately N-representable p-RDMs, ${}_p^N \tilde{P}$ which is also convex, with ${}_p^N \tilde{P} \subseteq {}_p^N P$. Thus, the minimization of the ground state energy subject to the p-positivity conditions constitutes a semidefinite program. Since this is a relaxation of the full N-representability problem, its solution will lower bound the true energy. Employing the 2-positivity conditions yields the following SDP, called the variational 2-RDM (v2RDM) [Maz11a]:

$$\begin{aligned} &\min_{{}^2D} E[{}^2D] \\ \text{such that} \quad &{}^2D \succeq 0 \\ &{}^2Q \succeq 0 \\ &{}^2G \succeq 0 \\ &\text{Tr}({}^2D) = \tfrac{1}{2}N(N-1) \\ &{}^2Q = f_Q({}^2D) \\ &{}^2G = f_G({}^2D) \end{aligned} \tag{39}$$

where $E[{}^2D] = \text{Tr}({}^2K^2D)$ and f_Q and f_G are the linear mappings defined in Eq. (20) and Eq. (21), respectively.

2.4.1 Classical Shadow Constraints

In their paper, Avdic & Mazziotti propose the use of classical shadows to improve the performance of the v2RDM, and in turn reduce the number of measurements required [AM24b]. Classical shadows are characterized as unitary transformations of the 2-RDM, and constructed as follows:

$$S_n^{pq} = \langle \Psi | \hat{U}_n^\dagger \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_q \hat{a}_p \hat{U}_n | \Psi \rangle, \tag{40}$$

where

$$\hat{U}_n = \exp\left(\sum_{uv} A_n^{uv} \hat{a}_u^\dagger \hat{a}_v \right) \tag{41}$$

are unitaries sampled using the Haar measure. A_n is a one-body anti-Hermitian matrix, and n is the shadow index. Thereafter, constraints are added to the v2RDM SDP:

$$S_n^{pq} = \left((U_n \otimes U_n)\ {}^2D\ (U_n \otimes U_n)^T \right)_{pq}^{pq} \tag{42}$$

where

$$U_n = \exp(A_n). \tag{43}$$

Note that the above implies the equivalence of Equations (40) and (42).

There are a few points in this formulation that merit clarification. Firstly, the transpose in Eq. (42) appears to be intended as a *conjugate* transpose; otherwise, S_n^{pq} could contain non-zero imaginary components despite being expectation values. The authors also provide a correction to the original derivation of Eq. (42) in Appendix B of this work.

Additionally, the definition of $\hat{U}_n$ is somewhat ambiguous. Appendix A of [AM24a] describes generating $\hat{U}_n$ using normal sampling and Gram–Schmidt decomposition, which does not involve an anti-Hermitian matrix such as A_n. Moreover, the definition of $\hat{U}_n$ given in Eq. (41) is presented without motivation, even though this construction does not necessarily yield Haar-distributed unitaries.

Most importantly, the authors note that the use of the term "classical shadows" may not be entirely accurate. The method described above does not incorporate the defining features of classical shadows, namely,

(i) storing an efficient classical representation of the measurements,

(ii) reconstructing the state with an inverse map of the unitary channel,

(iii) and estimating an observable using an aggregation of the reconstructed state.

Alternatively, we suggest that these more closely resemble typical measurements with unitaries, and hence adopt this terminology going forward. That being said, considering the *measurements* in isolation constitutes another semidefinite program:

$$\begin{aligned} &\min_{{}^2D} \ E[{}^2D] \\ \text{such that} \quad &S_n^{pq} = \left((U_n \otimes U_n)\ {}^2D\ (U_n \otimes U_n)^\dagger \right)_{pq}^{pq} \end{aligned} \tag{44}$$

When combined with Eq. (39), we obtain the main contribution of Avdic & Mazziotti's paper, namely, a variational 2-RDM SDP with additional measurement constraints. They refer to this as the *shadow v2RDM* (sv2RDM) method.

3 Implementation

Thus far, we have provided an elegant formulation of the v2RDM. However, most literature does not explicitly mention that its implementation necessitates very specific structure for the ${}^2D$, ${}^2Q$ and ${}^2G$ matrices. Additionally, the conventional notation is somewhat vague as it indexes by 'orbitals', typically without specifying whether these refer to spatial or spin-orbitals - a crucial distinction. As such, this section begins with a brief overview of molecular orbitals. Thereafter, we review existing implementations of the v2RDM. We conclude by proposing a new, spatial version of the v2RDM with additional measurement constraints.

3.1 Spatial vs Spin-orbitals

A *spatial* (or molecular) orbital is a one-electron spatial wavefunction that depends on either cartesian (x, y, z) or spherical (r, θ, ϕ) coordinates. These are visually represented by regions around the nucleus where an electron may occur, as shown in Figure 1. A *spin-orbital* adds a fourth coordinate, namely the electron spin, taking the value of either α (spin-up) or β (spin-down). Formally, a spin-orbital is the product of a one-electron spatial orbital and a one-electron spin function. Typically, they are approximated using *Slater determinants* - a set of basis functions that enforces the antisymmetry required by the Pauli exclusion principle [Lev13]. It thus possible for two electrons to occupy the same spatial orbital, provided they have different spins. Figure 2 compares the electron configuration of spatial orbitals with that of spin-orbitals.

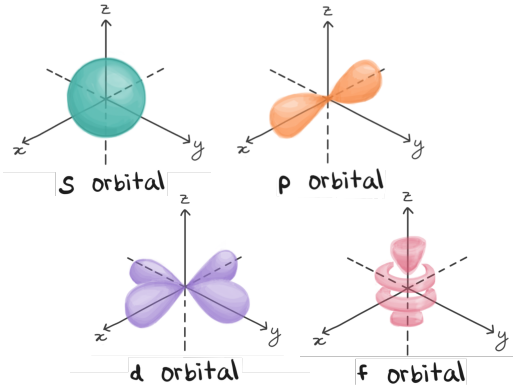

Figure 1: Spatial orbitals are functions of spatial coordinates and are designated the letters s, p, d, and f based on their angular momentum quantum number.

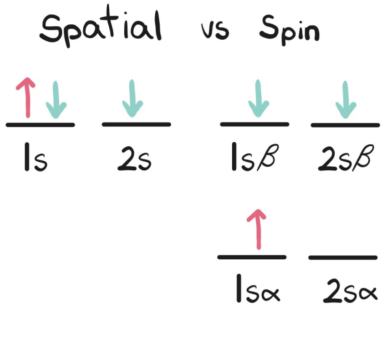

Figure 2: One possible configuration for the ground state of Li. Notice that spin is handled explicitly when dealing with spin-orbitals.

3.2 Existing v2RDMs

Very few implementations of the v2RDM are publicly available. The authors are aware of only three open-source versions[2] and two proprietary implementations[3] of the v2RDM. Furthermore, these employ the Boundary Point SDP algorithm (BPSDP) [Maz11a], which entirely reformulates the v2RDM SDP provided in the literature. Fortunately, recent papers by DePrince [ED24] and Rubin [RBM18] shed some light on the topic. These are

[2] By Mazziotti, DePrince, and Rubin.

[3] Maplesoft and Q-Chem

among the rare instances that explicitly detail the spin-block structure of the ${}^2D$, ${}^2Q$ and ${}^2G$ matrices.

In practice, it is more apt to include spin components when defining 2-RDMs:

$$ {}^2D^{i_\sigma j_\tau}_{k_\sigma l_\tau} = \langle\psi|\, a^\dagger_{i_\sigma} a^\dagger_{j_\tau} a_{l_\tau} a_{k_\sigma}\, |\psi\rangle \tag{45} $$

$$ {}^2Q^{i_\sigma j_\tau}_{k_\sigma l_\tau} = \langle\psi|\, a_{i_\sigma} a_{j_\tau} a^\dagger_{l_\tau} a^\dagger_{k_\sigma}\, |\psi\rangle \tag{46} $$

$$ {}^2G^{i_\sigma j_\tau}_{k_\kappa l_\lambda} = \langle\psi|\, a^\dagger_{i_\sigma} a_{j_\tau} a^\dagger_{l_\lambda} a_{k_\kappa}\, |\psi\rangle \tag{47} $$

where $\sigma, \tau, \kappa, \lambda \in \{\alpha, \beta\}$ are the spin components and $i, j, k, l \in \{1, ..., r\}$ are the spatial components of the orbitals. The second quantization operator $a^\dagger_{i_\sigma}$ (a_{i_σ}) thus corresponds to creating (destroying) an electron in orbital i with spin σ. The resulting matrices have the following spin-block structure:

$$ {}^2D = \begin{pmatrix} {}^2D^{\alpha\alpha}_{\alpha\alpha} & 0 & 0 & 0 \\ 0 & {}^2D^{\beta\beta}_{\beta\beta} & 0 & 0 \\ 0 & 0 & {}^2D^{\alpha\beta}_{\alpha\beta} & 0 \\ 0 & 0 & 0 & {}^2D^{\beta\alpha}_{\beta\alpha} \end{pmatrix}, \; {}^2G = \begin{pmatrix} {}^2G^{\alpha\alpha}_{\alpha\alpha} & {}^2G^{\alpha\alpha}_{\beta\beta} & 0 & 0 \\ {}^2G^{\beta\beta}_{\alpha\alpha} & {}^2G^{\beta\beta}_{\beta\beta} & 0 & 0 \\ 0 & 0 & {}^2G^{\alpha\beta}_{\alpha\beta} & 0 \\ 0 & 0 & 0 & {}^2G^{\beta\alpha}_{\beta\alpha} \end{pmatrix}. \tag{48} $$

Each spin-block has dimensions (r^2, r^2) for a total dimension of $(4r^2, 4r^2)$ per matrix. Notice that the only non-zero entries are those that preserve spin. For example, a non-zero value for ${}^2D^{\alpha\alpha}_{\beta\beta}$ would imply that we can create two α electrons and destroy two β, but this would change the total spin of the system. The v2RDM assumes that the system has zero spin i.e. is not subject to an external magnetic field. Spin-adapted versions of the algorithm have also been proposed [Maz05]. Additionally, note that ${}^2D$ and ${}^2Q$ have a similar structure as they represent analogous operations - creating two electrons is analogous to creating two holes elsewhere. However, ${}^2G$ has a slightly different structure as the electron-hole 2-RDM. The spatial 2-RDM can be obtained by summing over the spin components of the spin-orbital 2-RDM [Mar19],

$$ {}^2D^{ij}_{kl} = \tfrac{1}{2}(D^{i_\alpha j_\alpha}_{k_\alpha l_\alpha} + D^{i_\beta j_\beta}_{k_\beta l_\beta} + D^{i_\alpha j_\beta}_{k_\alpha l_\beta} + D^{i_\beta j_\alpha}_{k_\beta l_\alpha}) \tag{49} $$

The BPSDP approach flattens each spin-block component into a large vector $\mathbf{x}$. To calculate the energy, it takes the dot product of $\mathbf{x}$ with another vector, $\mathbf{c}$, which contains the one and two-body integrals of the Hamiltonian. The resulting objective function reads,

$$ \begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^T\mathbf{x} \\ \text{such that} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & M(\mathbf{x}) \succeq 0. \end{aligned} \tag{50} $$

Constraints are built into $M(\mathbf{x})$ by constructing the ${}^2Q$ and ${}^2G$ matrices in terms of each spin-block basis using 0s, 1s and -1s. Overall, the solution is not straightforward and difficult to relate to the theoretical formulation of the v2RDM. Hence, extending the available implementations to include shadows proved to be unworkable.

3.3 Proposed Spatial v2RDM

While very helpful, DePrince and others do not explain why the spin-block structure is strictly necessary for constraining the v2RDM, nor why a spatial version would not suffice. As a sanity test, we examined the 2-RDM obtained from FCI (full configuration interaction)

using OpenFermion [MRS+20]. Importantly, this ${}^2D_{\text{FCI}}$ is given in spatial orbital basis and is obtained via diagonalization of the full Hamiltonian. We also computed the reduced Hamiltonian, ${}^2K$, using the one- and two-body integrals accessible via OpenFermion. The relevant code can be found in Appendix D. Initial calculations with the ${}^2D_{\text{FCI}}$ consistently produced energies that matched FCI, indicating that a much simpler implementation of the v2RDM may be possible. The full investigation is detailed in Section 4.

A simpler, *spatial* version of the v2RDM is desirable for a number of reasons. Firstly, the number of optimization variables scales as r^4 instead of $(2r)^4$, as for the spin-orbital case. For example, the spatial treatment of H_4, which has 4 orbitals, would optimize 256 variables. In comparison, the spin-orbital treatment would require a substantial 4096 variables. Even Mazziotti's open source Python version of the v2RDM, which employs the more efficient BPSDP, requires 2256 variables. A spatial version would thus necessitate significantly fewer computational resources and less time. Secondly, a spatial implementation aligns more naturally with the underlying theory, making it easier to extend or modify and better suited for pedagogical purposes.

The spatial v2RDM was implemented in Python. A short excerpt demonstrating the simplicity of the code is provided in Listing 1. The full code is given in Appendix D. The GitHub repository for this project can be found here. The CVXPY library was used with the MOSEK solver to implement the SDP. Molecules were simulated using the OpenFermion library [MRS+20]. The mappings from ${}^2D$ to ${}^2Q$ and ${}^2G$ were implemented as CVXPY functions and used to constrain the SDP. The mappings were validated by manually constructing ${}^2D$, ${}^2Q$ and ${}^2G$ from the ground state using the definitions in Table 1. These were subsequently compared with the result of applying each mapping to ${}^2D$. This, along with other validation steps, can be found in the `mappings-demo.ipynb` notebook in the GitHub repository for this project.

3.4 Measurement Constraints

The measurement constraints were implemented similarly to those given in Avdic & Mazziotti's paper. The measurements were obtained by applying a random unitary to the ${}^2D_{\text{FCI}}$ obtained from OpenFermion, as per Eq. (42). Each unitary was generated by sampling a normal distribution, then performing Gram-Schmidt decomposition. Such unitaries are known to be Haar distributed [AM24a]. Thereafter, the unitaries and resulting S_n matrices were used to constrain the SDP, so that any feasible ${}^2D$ satisfies the condition in Eq. (42).

To simulate noisy measurement, a random matrix was added to the S_n matrix. The matrix was sampled from a Gaussian distribution with a mean of zero and standard deviation ϵ, where ϵ is the desired noise level. This is akin to the noise introduced by Avdic & Mazziotti. To accommodate the noisy measurements, the constraints in Eq. (42) were relaxed similarly to the reference paper:

$$S_n^{pq} - 3\epsilon^{pq} \leq X_n^{pq} \leq S_n^{pq} + 3\epsilon^{pq} \tag{51}$$

where

$$X_n^{pq} = \left((U_n \otimes U_n)\, {}^2D\, (U_n \otimes U_n)^\dagger\right)_n^{pq} \tag{52}$$

The term ϵ^{pq} is a matrix with the same shape as S_n, with each entry equal to the desired noise level, ϵ. The factor of 3 was added to account for noise samples that exceed the standard deviation of ϵ, since omitting the prefactor would inevitably result in an infeasible SDP. The associated Python code can be found in Appendix D.

```python
import cvxpy as cp

### assume 'molecule' has been defined using OpenFermion
r, N = molecule.n_orbitals, molecule.n_electrons

# define optimization variable
D2 = cp.Variable((r**2, r**2))

# define objective function
K2 = getK2(molecule)
E = lambda D : cp.trace(K2 @ D) + molecule.nuclear_repulsion
objective = cp.Minimize(E(D2))

# define constraints
constraints = [
    D2 >> 0,
    Q2_cvxpy(D2, r, N) >> 0,
    G2_cvxpy(D2, r, N) >> 0,
    cp.trace(D2) == 0.5 * N * (N-1)
]

# solve
problem = cp.Problem(objective, constraints)
problem.solve(solver=cp.MOSEK, eps=1e-8)
```

Listing 1: Basic implementation of the spatial v2RDM. The code can be directly related to the theory, unlike the BPSDP used in most implementations.

4 Results

Presented below are the results of the spatial v2RDM with measurements (m-v2RDM). We begin with a closer examination of OpenFermion's spatial ${}^2D_{\text{FCI}}$ so as to contextualize subsequent findings. Thereafter, we consider the spatial v2RDM in isolation, followed by the effect of adding the measurement constraints. We do so for both noiseless and noisy measurements. Finally, we compare the performance of the spatial m-v2RDM to Avdic & Mazziotti's sv2RDM implementation. All results were averaged over ten runs and produced on a personal computer with a 2.3 GHz 8-core Intel processor and 16GB of RAM.

4.1 Examining the Spatial ${}^2D_{\text{FCI}}$

The 2-RDM obtained from OpenFermion provides insight into the limitations of the spatial v2RDM. Recall that the energy of a molecule can be calculated using $E = \text{Tr}({}^2K\,{}^2D)$. When accounting for normalization and indexing, we can obtain a ${}^2K$ and ${}^2D_{\text{FCI}}$ from OpenFermion that align closely with theoretical expectations. Figure 3 shows that the energy calculated using the ${}^2D_{\text{FCI}}$ matches the theoretical prediction. The spatial ${}^2D_{\text{FCI}}$ also exhibits the expected normalization. The results seem to indicate that the spatial 2-RDM could be used in variational calculations instead of the spin-orbital version.

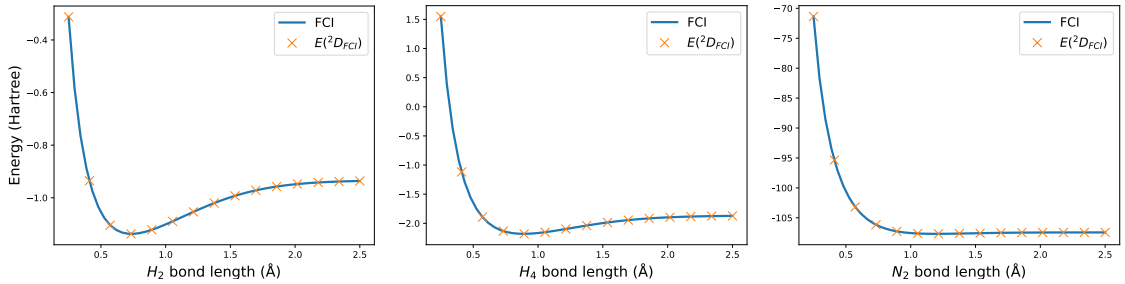

Figure 3: Calculating the groundstate energy of a molecule using the reduced Hamiltonian ${}^2K$ and the spatial ${}^2D_{\text{FCI}}$ consistently matches the energy obtained by FCI, i.e. full diagonalization of the Hamiltonian. This is shown for the hydrogen molecule (H_2), the hydrogen chain with evenly spaced atoms (H_4), and diatomic nitrogen (N_2).

The spatial ${}^2D_{\text{FCI}}$ is consistently positive semidefinite, and so is the ${}^2Q$ matrix when acquired using Eq. (20). However, Table 2 demonstrates that the spatial ${}^2D_{\text{FCI}}$ does not typically meet all requirements for a valid fermionic density matrix. For a variety of molecules, the ${}^2G$ matrix obtained from Eq. (21) fails to be positive. This contradicts the G-positivity constraint on the v2RDM. Finally, the antisymmetry requirement in Eq. (12) is consistently violated.

Molecule	Orbitals	Electrons	${}^2D_{\text{FCI}} \succeq 0$	${}^2Q \succeq 0$	${}^2G \succeq 0$	Antisymm.	Hermitian
H_2	2	2	True	True	True	False	True
H_4	4	4	True	True	False	False	True
LiH	6	4	True	True	False	False	True
HF	6	10	True	True	False	False	True
H_2O	7	10	True	True	False	False	True
NH_3	8	10	True	True	False	False	True
N_2	10	14	True	True	False	False	True

Table 2: OpenFermion's spatial 2-RDM consistently fails to meet all requirements for a valid fermionic density matrix. All molecules were constructed with the STO-3G basis using evenly spaced atoms, each separated by 1.0Å.

The failure of the spatial ${}^2D_{\text{FCI}}$ to exhibit the correct antisymmetry and G-positivity are likely closely related. The spatial ${}^2D$ is obtained by summing over the spin components of the spin-orbital version, as per Eq. (49). Since the spin-block structure of ${}^2D$ and ${}^2Q$ are the same, the spin-information appears to be preserved when mapping between them. However, the off-diagonal spin-block elements of ${}^2G$ contain information about the spin that is clearly lost when converting to the spatial representation. This raises the question of whether the G-positivity condition is applicable to the spatial v2RDM, or should be reformulated to better account for spin.

4.2 Spatial v2RDM Performance

The performance of the spatial v2RDM without measurements affirms the findings of our investigation above. Figure 4 shows the energy curves obtained by the spatial v2RDM compared with the FCI energy for the hydrogen, hydrogen chain, and hydrogen fluoride molecules. Given alongside are the minimum eigenvalues of the ${}^2D$, ${}^2Q$ and ${}^2G$ matrices at the corresponding bond length.

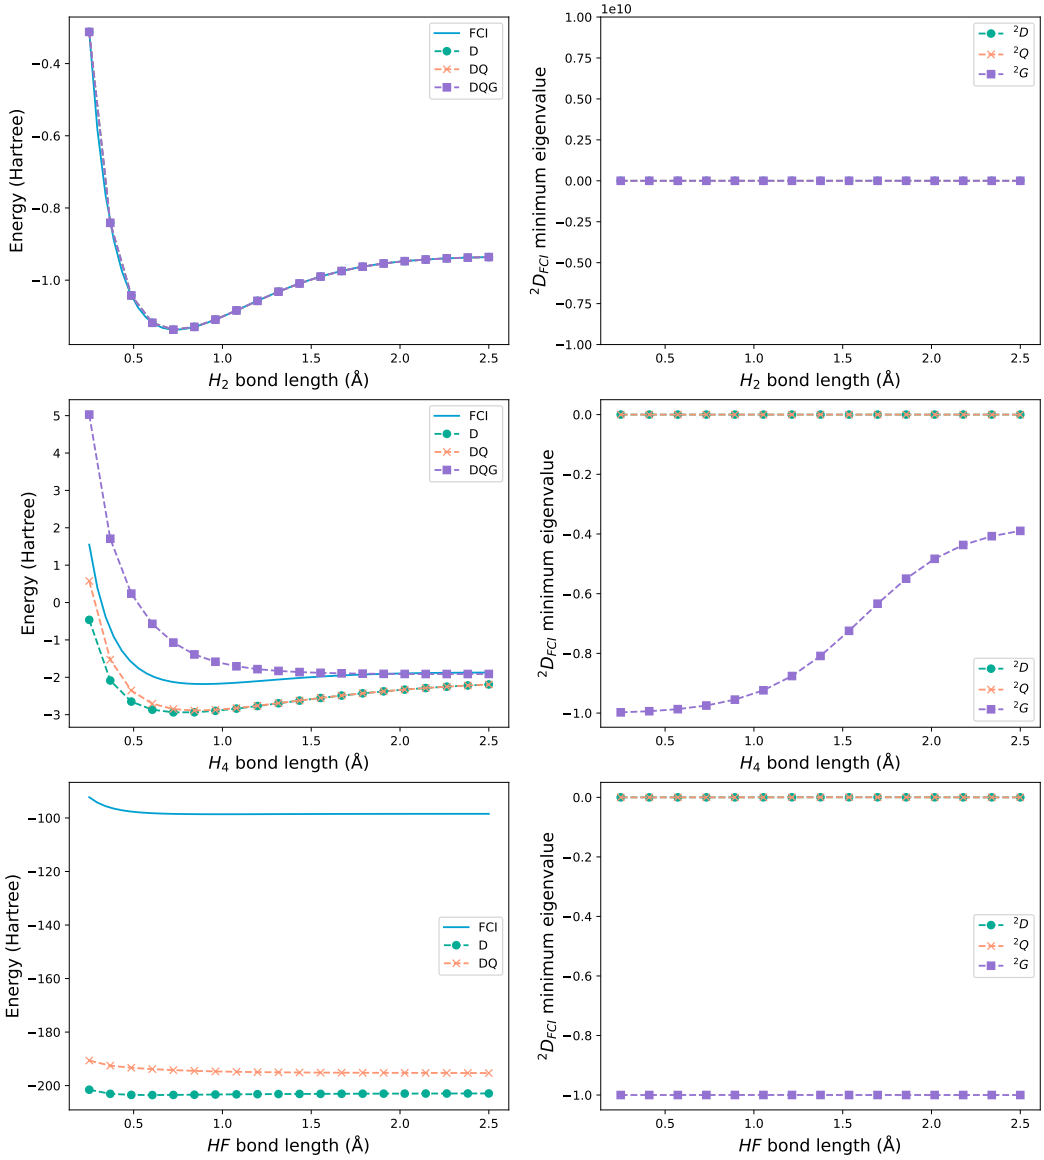

Figure 4: The positivity of the ${}^2G$ matrix affects the feasibility and accuracy of the spatial v2RDM. The full DQG conditions are accurate for H_4 only in regions where is ${}^2G$ is more positive, while accurate everywhere for H_2. For HF, adding the G condition results in an infeasible SDP. Hence, only the D and Q conditions can be applied.

For a simple molecule like H_2, we observe the FCI energy matched exactly using conditions D, DQ and DQG. All eigenvalues of the positivity matrices are non-negative up to machine precision. In contrast, consider HF in the bottom row. The ${}^2G$ matrix is negative throughout, with a constant value for the smallest eigenvalue. Correspondingly, adding the G condition makes the spatial v2RDM infeasible for all bond lengths. Similar behaviour was observed for other larger molecules like water and ammonia.

H_4 provides an interesting case study on the effect of the G condition. At shorter bond lengths, the energy is overestimated to compensate for the negativity of ${}^2G$. Although negative throughout, the minimum eigenvalue of ${}^2G$ increases at longer bond lengths, resulting in greater accuracy at those points. For all molecules tested, the D and DQ conditions consistently provide a lower bound for the FCI energy, as expected for a relaxation of the full N-representability conditions.

4.3 Introducing Measurements

Constraining the SDP using only measurements was found to be largely infeasible, except for small molecules or many measurements. When feasible, the measurements-only SDP is highly accurate without noise. In the noisy case, more measurements are required to obtain the same accuracy as without noise.

Simply adding the D-positivity condition to the measurements-only SDP makes many infeasible problems become feasible. For example, adding the D condition renders H_4 feasible with only 1 measurement, while it otherwise requires at least 11 measurements to be feasible. The D and DQ conditions have the same effect on noisy measurements, albeit with slightly less accuracy. Figure 5 compares the D and DQ conditions with and without noise for H_4 at 1.0Å over a range of measurements. Additionally, it shows that the accuracy of the spatial m-v2RDM is proportional to the level of noise.

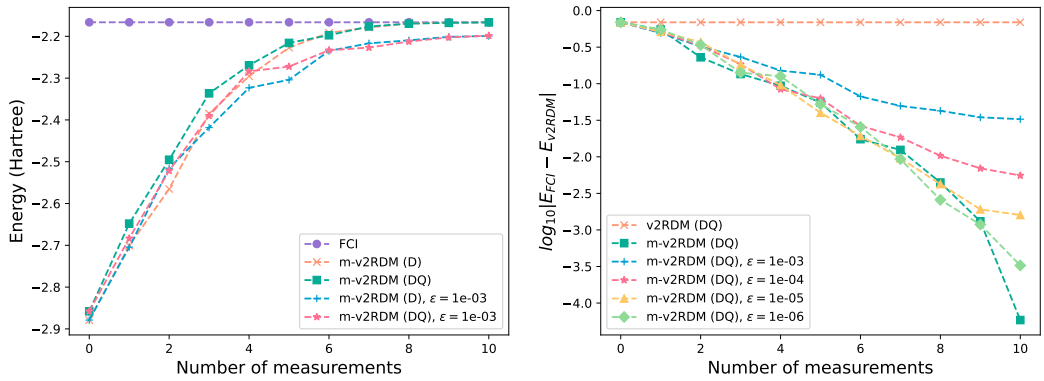

Figure 5: Adding the D condition makes H_4 feasible for any amount of either noiseless or noisy measurements. Moreover, the DQ conditions exhibit an improved accuracy over D in both cases. As expected, increased noise results in reduced accuracy.

The spatial v2RDM and measurement constraints are highly complementary. Figure 6a demonstrates the remarkable effect of combining these two approaches. Without measurements, the spatial v2RDM performs poorly on the hydrogen fluoride molecule. Moreover, at least 22 measurements are required for the measurements-only SDP to be feasible. However, combining the constraints garners an improvement of ~90 Hartree (93%) on the spatial v2RDM for just 4 measurements. The full effect of the measurement constraints can be seen at 22 measurements, where the difference with the FCI energy sharply decreases.

In general, adding the Q condition yields a marginal increase in accuracy. However, the time complexity of the problem scales notably slower with the Q condition than without, as shown in Figure 6b. If it is possible to formulate a spatial version of the G condition, we expect that adding G should improve the accuracy and, potentially, the time complexity of the SDP.

Lastly, the spatial m-v2RDM allows us to approximate large molecules. In Figure 7, 40 measurements achieve a ~1.8 Hartree energy difference with FCI. Although this does not reflect the standard desired chemical accuracy of approximately 1.6×10^{-3} Hartrees [Lev13], there is potential for improvement with more measurements.

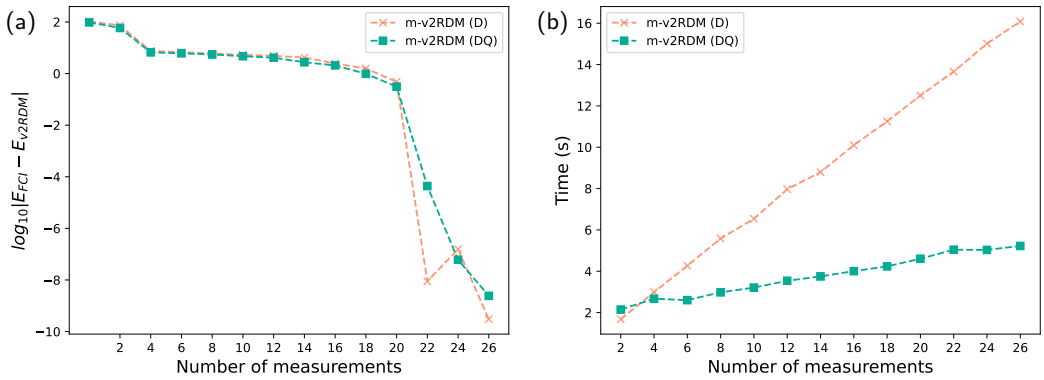

Figure 6: Adding measurements vastly improves the spatial v2RDM. Without at least the D condition, the HF molecule is infeasible with fewer than 22 measurements. The DQ conditions significantly improve the time scaling of HF compared to D only.

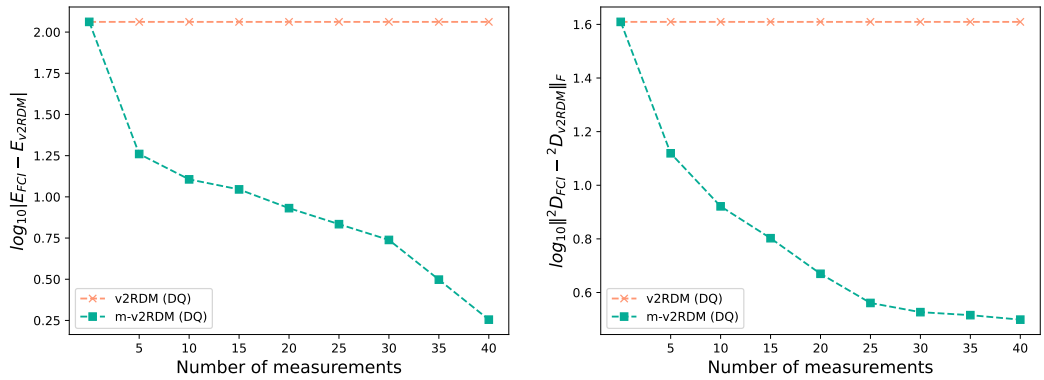

Figure 7: Even large molecules like N_2 (shown here at 1.75Å) can be approximated using the spatial m-v2RDM.

4.4 Comparison with Avdic & Mazziotti

Here we compare the results in [AM24b] with the spatial m-v2RDM above. We also make use of Mazziotti's open source v2RDM Python implementation to examine memory and time requirements for the v2RDM without shadows.

Mazziotti's v2RDM implementation achieves high accuracy even without shadow constraints. Nonetheless, algorithms such as the BPSDP still have substantial time and memory requirements. Despite its shortcomings, the spatial m-v2RDM offers a practical alternative when computational resources are limited.

Consider Figure 3a in Avdic & Mazziotti's paper. For H_4, the sv2RDM achieves an accuracy on the order of 10^{-4} within 8 shadows. The spatial version presented above obtains a comparable accuracy within 10 measurements (Figure 5), while optimizing only 256 variables (16×16 matrix) with a single run taking on average ~0.45s. In comparison, Mazziotti's v2RDM requires 2 256 variables and an average of ~25s per run. It is not known whether the addition of shadows decreases the time requirement, but it will certainly require at least as much memory.

For hydrogen fluoride at 1.0Å, Mazziotti's v2RDM yields energies within $10^{-3.5}$ of FCI, but a single run takes approximately ~1 100s and optimizes 11 412 variables. The spatial version achieves a similar accuracy within 22 measurements, while taking only ~3.6s per run and optimizing 1 296 variables (36×36 matrix).

The spatial m-v2RDM can approximate the ground state energy of nitrogen on a standard laptop. Consider Figure 7 above and Figure 2a in [AM24b]. With 30 measurements, we obtain an accuracy comparable to the sv2RDM (D) with 2 shadows. With 40 measurements, the spatial m-v2RDM (DQ) achieves similar accuracy to Mazziotti's v2RDM (DQ) without shadows. On average, a single run of the spatial m-v2RDM with 30 shadows takes 225s, while 40 measurements takes 305s. The number of optimization variables is 10 000 (100 x 100 matrix). In comparison, Mazziotti's v2RDM requires 88 500 optimization variables and failed to run on the machine used in this paper due to insufficient memory.

5 Conclusion

Based on recent work by Avdic & Mazziotti [AM24b], we have presented a novel spatial orbital variant of the variational 2-RDM (v2RDM) with additional measurement constraints (m-v2RDM). The proposed method offers a simpler implementation than existing algorithms, such as the boundary point semidefinite programming (BPSDP) algorithm, which are often complex and resource intensive.

Computationally, the spatial m-v2RDM drastically reduces memory and runtime requirements. Its simplified structure allows for the approximation of larger molecules like N_2, which are costly or intractable with traditional implementations. For small to medium molecules, including H_2, H_4, and HF, the spatial m-v2RDM achieves accuracy comparable to standard v2RDM methods while using far fewer variables and shorter runtimes. These results demonstrate the method's potential to make reduced density matrix approaches more practical, particularly when combined with measurement constraints.

Although chemical accuracy is not consistently achieved, performance improves substantially with more measurements, suggesting clear opportunities for refinement. In particular, a key limitation arises from the spatial 2-RDM's lack of antisymmetry and G-positivity. As a result, many molecules could be treated only with the DQ conditions as opposed to the full DQG conditions.

This work also illustrates the complementary nature of N-representability conditions and measurement constraints. Adding the D-positivity condition to a measurements-only SDP makes many otherwise infeasible problems feasible. Adding Q-positivity conditions further improves accuracy and runtime complexity. Conversely, even a small number of measurements significantly enhances the spatial v2RDM's performance.

Finally, the spatial approach aligns naturally with the underlying theory, making it easier to implement, extend, and modify. Its simplicity makes the spatial m-v2RDM a valuable pedagogical tool for researchers and students exploring reduced density matrix methods.

6 Outlook

Future work may include applying the spatial m-v2RDM to larger molecules to test its computational limits and accuracy. A key theoretical direction is the development of a spatial analogue of the G-positivity condition. Results by Avdic & Mazziotti [AM24b] indicate that the DQG conditions provide a substantial improvement over DQ compared to that of DQ over D, suggesting that a spatial G condition could significantly improve performance.

Implementing the method with true classical shadow constraints may be another valuable line of investigation. This could further reduce measurement requirements and improve scalability. Moreover, it could allow the estimation of excited states, as in Avdic

& Mazziotti's follow-up paper [AM24a]. A proposal for true classical shadows is given in Appendix C.

Finally, it is worth exploring the potential of spatial N-representability conditions for error correction in quantum algorithms or noisy quantum communication. These types of applications may provide the most promising avenue for impact, since v2RDM methods struggle to contend with the efficiency of modern configuration interaction methods.

Bibliography

[AA00] Erling D. Andersen and Knud D. Andersen. *The Mosek Interior Point Optimizer for Linear Programming: An Implementation of the Homogeneous Algorithm*, page 197–232. Springer US, 2000.

[Aar18] Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC'18, page 325–338. ACM, June 2018.

[AM24a] Irma Avdic and David A. Mazziotti. Enhanced shadow tomography of molecular excited states via the enforcement of N-representability conditions by semidefinite programming. *Physical Review A*, 110(5), November 2024.

[AM24b] Irma Avdic and David A. Mazziotti. Fewer measurements from shadow tomography with n-representability conditions. *Phys. Rev. Lett.*, 132:220802, May 2024.

[BCMTS24] Emilio Bagan, John Calsamiglia, Ramon Munoz-Tapia, and Gael Sentís. Lecture notes. Quantum Statistical Inference, Universitat Autònoma Barcelona, October 2024.

[Col63] A. J. Coleman. Structure of fermion density matrices. *Rev. Mod. Phys.*, 35:668–686, Jul 1963.

[Cze21] Artur Czerwinski. Quantum state tomography with informationally complete povms generated in the time domain. *Quantum Information Processing*, 20(3), March 2021.

[ED24] A. Eugene DePrince. Variational determination of the two-electron reduced density matrix: A tutorial review. *WIREs Computational Molecular Science*, 14(1), January 2024.

[GP64] Claude Garrod and Jerome K. Percus. Reduction of the N-particle variational problem. *Journal of Mathematical Physics*, 5(12):1756–1776, December 1964.

[HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, June 2020.

[Jud01] Brian R. Judd. Reduced Density Matrices: Coulson's Challenge. *Physics Today*, 54(2):58–59, February 2001.

[JW28] P. Jordan and E. Wigner. Ber das Paulische Äquivalenzverbot. *Zeitschrift Physik*, 47(9-10):631–651, 1928.

[Kar84] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing - STOC '84*, STOC '84, page 302–311. ACM Press, 1984.

[Kly06] Alexander A Klyachko. Quantum marginal problem and N-representability. *Journal of Physics: Conference Series*, 36:72–86, April 2006.

[L$\ddot{\text{o}}$55] Per-Olov Löwdin. Quantum theory of many-particle systems. I. Physical interpretations by means of density matrices, natural spin-orbitals, and convergence problems in the method of configurational interaction. *Phys. Rev.*, 97:1474–1489, Mar 1955.

[Lev13] Ira N. Levine. *Quantum Chemistry*. Pearson, Upper Saddle River, NJ, 7 edition, February 2013.

[Liu06] Yi-Kai Liu. *Consistency of Local Density Matrices Is QMA-Complete*, page 438–449. Springer Berlin Heidelberg, 2006.

[LTC13] Eduardo V. Ludeña, F. Javier Torres, and Cesar Costa. Functional N-

representability in 2-matrix, 1-matrix, and density functional theories. *Journal of Modern Physics*, 04(03):391–400, 2013.

[LTWHG17] Susi Lehtola, Norm M. Tubman, K. Birgitta Whaley, and Martin Head-Gordon. Cluster decomposition of full configuration interaction wave functions: A tool for chemical interpretation of systems with strong correlation. *The Journal of Chemical Physics*, 147(15), October 2017.

[Mar19] Elvis Maradzike. *Development and Application of The Variational Two-Electron Reduced Density Matrix Complete Active Space Self-Consistent Field Method to Address The Electron Correlation Problem in Quantum Chemistry.* PhD thesis, The Florida State University, 2019.

[May55] Joseph E. Mayer. Electron correlation. *Phys. Rev.*, 100:1579–1586, Dec 1955.

[Maz05] David A. Mazziotti. Variational two-electron reduced density matrix theory for many-electron atoms and molecules: Implementation of the spin- and symmetry-adapted T_2 condition through first-order semidefinite programming. *Physical Review A*, 72(3), September 2005.

[Maz07] David A. Mazziotti, editor. *Reduced-Density-Matrix Mechanics: With Application to Many-Electron Atoms and Molecules.* Wiley, March 2007.

[Maz11a] David A. Mazziotti. Large-scale semidefinite programming for many-electron quantum mechanics. *Physical Review Letters*, 106(8), February 2011.

[Maz11b] David A. Mazziotti. Two-electron reduced density matrix as the basic variable in many-electron quantum chemistry and physics. *Chemical Reviews*, 112(1):244–262, August 2011.

[Maz12] David A. Mazziotti. Structure of fermionic density matrices: Complete N-representability conditions. *Physical Review Letters*, 108(26), June 2012.

[ME01] David A. Mazziotti and Robert M. Erdahl. Uncertainty relations and reduced density matrices: Mapping many-body quantum mechanics onto four particles. *Physical Review A*, 63(4), March 2001.

[MMK+19] J. Wayne Mullinax, Elvis Maradzike, Lauren N. Koulias, Mohammad Mostafanejad, Evgeny Epifanovsky, Gergely Gidofalvi, and A. Eugene DePrince. Heterogeneous CPU + GPU algorithm for variational two-electron reduced-density matrix-driven complete active-space self-consistent field theory. *Journal of Chemical Theory and Computation*, 15(11):6164–6178, 2019.

[MOC+24] Gustavo E. Massaccesi, Ofelia B. Oña, Pablo Capuzzi, Juan I. Melo, Luis Lain, Alicia Torre, Juan E. Peralta, Diego R. Alcoba, and Gustavo E. Scuseria. Determining the N-representability of a reduced density matrix via unitary evolution and stochastic sampling. *Journal of Chemical Theory and Computation*, 20(22):9968–9976, November 2024.

[MRS+20] Jarrod R McClean, Nicholas C Rubin, Kevin J Sung, Ian D Kivlichan, Xavier Bonet-Monroig, Yudong Cao, Chengyu Dai, E Schuyler Fried, Craig Gidney, Brendan Gimby, Pranav Gokhale, Thomas Häner, Tarini Hardikar, Vojtěch Havlíček, Oscar Higgott, Cupjin Huang, Josh Izaac, Zhang Jiang, Xinle Liu, Sam McArdle, Matthew Neeley, Thomas O'Brien, Bryan O'Gorman, Isil Ozfidan, Maxwell D Radin, Jhonathan Romero, Nicolas P D Sawaya, Bruno Senjean, Kanav Setia, Sukin Sim, Damian S Steiger, Mark Steudtner, Qiming Sun, Wei Sun, Daochen Wang, Fang Zhang, and Ryan Babbush. Openfermion: The electronic structure package for quantum computers. *Quantum Science and Technology*, 5(3):034014, June 2020.

[OCPB16] Brendan O'Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding.

Journal of Optimization Theory and Applications, 169(3):1042–1068, February 2016.

[RBM18] Nicholas C Rubin, Ryan Babbush, and Jarrod McClean. Application of fermionic marginal constraints to hybrid quantum algorithms. *New Journal of Physics*, 20(5):053020, May 2018.

[SC23] Paul Skrzypczyk and Daniel Cavalcanti. *Semidefinite Programming in Quantum Information Science*. 2053-2563. IOP Publishing, 2023.

[WBAG11] James D. Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure Hamiltonians using quantum computers. *Molecular Physics*, 109(5):735–750, March 2011.

[WDPT25] Zherui Jerry Wang, David Dechant, Yash J. Patel, and Jordi Tura. Mitigating shot noise in local overlapping quantum tomography with semidefinite programming. *Physical Review A*, 111(5), May 2025.

A Grassmann Wedge Product

The *wedge* product or *exterior* product, denoted by $\wedge$, is an operation in exterior algebra named after Hermann Grassmann. Its defining property is anti-commutativity. Let $a, b \in V$ where V is a vector space. Then,

$$a \wedge b = -b \wedge a \tag{53}$$

and, consequently,

$$a \wedge a = 0 \quad \forall\, a \in V. \tag{54}$$

For example, the wedge product of two rank-2 tensors (one-particle matrices) yields a rank-4 tensor (two-particle matrix) [Maz07]:

$$c^{i,j}_{k,l} = a^i_k \wedge b^j_l = \frac{1}{4}\left(a^i_k b^j_l - a^j_k b^i_l - a^i_l b^j_k + a^j_l b^i_k\right) \tag{55}$$

For higher-dimensional tensors, the wedge product can be written as

$$a^{j_1 j_2 \ldots j_p}_{i_1 i_2 \ldots i_p} \wedge b^{i_{p+1} \ldots i_N}_{j_{p+1} \ldots j_N} = \left(\frac{1}{N!}\right)^2 \sum_{\pi,\sigma} \epsilon(\pi)\epsilon(\sigma)\pi a^{\sigma(i_{p+1})\ldots\sigma(i_N)}_{i_1 i_2 \ldots i_p} \pi b^{\sigma(j_{p+1})\ldots\sigma(j_N)}_{j_1 j_2 \ldots j_p} \tag{56}$$

where π and σ represent all permutations of the upper and lower indices, respectively. The function ϵ determines the sign of each term, returning $+1$ for even permutations and -1 for odd permutations. Geometrically, the wedge product can be thought of as a generalization of area and volume to higher dimensions. Indeed, in $\mathbb{R}^3$, it is closely related to the cross product.

B Correction to Constraint Derivation

To derive the measurement constraints, Avdic & Mazziotti provide Eq. (3) of [AM24b] as follows:

$$S^{pq}_n = \sum_{ijkl} U^{pi}_n U^{pj}_n \, {}^2D^{ij}_{kl} U^{ql}_n U^{qk}_n \tag{57}$$

However, this does not correspond to the eventual expression obtained in Eq. (6) of their work,

$$S^{pq}_n = \left(\left(U_n \otimes U_n\right) {}^2D \left(U_n \otimes U_n\right)^T\right)^{pq}_{pq}. \tag{58}$$

We provide the following correction. First, rewrite U_n in terms of its spectral decomposition. For convenience, we omit the subscript n:

$$U = \sum U_{uv} \left|u\right\rangle \left\langle v\right| \tag{59}$$

The tensor product of U with itself can therefore be written

$$U \otimes U = \left(\sum_{uv} U_{uv} \left|u\right\rangle \left\langle v\right|\right) \otimes \left(\sum_{st} U_{st} \left|s\right\rangle \left\langle t\right|\right) = \sum_{uvst} U_{uv} U_{st} \left|us\right\rangle \left\langle vt\right| \tag{60}$$

$$\implies (U \otimes U)^T = \sum_{uvst} U_{uv} U_{st} \left|vt\right\rangle \left\langle us\right| \tag{61}$$

We can retrieve the elements of the above using,

$$(U \otimes U)_{ab,ij} = U_{ai} U_{bj} \qquad \text{and} \qquad (U \otimes U)^T_{kl,cd} = (U \otimes U)_{cd,kl} = U_{ck} U_{dl}. \tag{62}$$

Next, let X be the result of sandwiching ${}^2D$ between the unitary products, that is,

$$X_{ab,cd} = (U \otimes U)_{ab,ij} \, {}^2D^{ij}_{kl} \, (U \otimes U)^T_{kl,cd} \tag{63}$$

The diagonal entries of X are thus given by

$$X_{pq,pq} = U_{pi} U_{qj} D^{ij}_{kl} U_{pk} U_{ql}. \tag{64}$$

Hence, the corrected indexing for Eq. (3) reads

$$S^{pq}_n = \sum_{ijkl} U^{pi}_n U^{qj}_n \, {}^2D^{ij}_{kl} U^{pk}_n U^{ql}_n \tag{65}$$

In order, the correct unitary indices are pi, qj, pk, and ql, not pi, pj, ql, and qk as in the original version. This was also numerically validated.

C True Classical Shadow Constraints

Here, we outline a possible implementation of classical shadow constraints on the spatial v2RDM. Let r be the number of spatial orbitals for the molecule under consideration. This corresponds to $2r$ spin-orbitals or qubits. Note that the conversion between qubit and fermionic spin-orbital representation needs to be handled carefully using an appropriate transformation, such as the Jordan-Wigner transform [JW28, WBAG11]. The correct treatment will depend on the quantum chemistry library used. As the classical shadows technique is only defined for qubits, the ground state below, ρ, should be taken as being in qubit representation.

Sample the unitary U from the Clifford group. Apply it to the ground state and measure in the computational basis to obtain the $2r$-bit measurement outcome, $|\hat{b}\rangle \in \{0,1\}^{2r}$. Huang et al. provide the inverted channel for the Clifford group in [HKP20]. Apply the inverted channel to construct a classical shadow:

$$\hat{\rho} = (2^{2r} + 1) U^\dagger |\hat{b}\rangle\langle\hat{b}| U - \mathbb{I} \tag{66}$$

Repeat this N times to obtain a collection of classical shadows. Since we wish to estimate the 2-RDM, we construct the p^{th} sample mean using

$${}^2\hat{D}^{ij\,(p)}_{kl}(N,1) = \frac{1}{N} \sum_{n=1}^{N} \mathrm{Tr}(\mathrm{JW}(a^\dagger_i a^\dagger_j a_l a_k) \hat{\rho}_n) \tag{67}$$

where i, j, k, and l are indexed over the number of qubits, $2r$. JW is the Jordan-Wigner transform, which transforms the fermionic operator $a^\dagger_i a^\dagger_j a_l a_k$ to qubit representation before applying it to the qubit state estimator $\hat{\rho}_n$. Repeat this P times to obtain the median of means estimator,

$${}^2\hat{D}^{ij}_{kl}(N,P) = \mathrm{median}\left\{ {}^2\hat{D}^{ij\,(1)}_{kl}(N,1), ..., {}^2\hat{D}^{ij\,(P)}_{kl}(N,1) \right\} \tag{68}$$

Now, let the p^{th} shadow be given by

$$\hat{S}_p = \sum_{ijkl} {}^2\hat{D}^{ij}_{kl}(N,P) \, |ij\rangle \, \langle lk| \tag{69}$$

This can be used to constrain the spin-orbital version of the v2RDM. To apply it to the spatial version, we first need to sum over the spin components to convert in to spatial

representation, as per Eq. (49). Let $\hat{S}'_p$ be the spatial shadow. Theorem 1 conveniently provides the error bounds for our constraints on the SDP:

$$\min_{^2D} \; E[^2D]$$
$$\text{such that} \quad \hat{S}'_p - \epsilon\,\mathbb{I} \leq {}^2D \leq \hat{S}'_p + \epsilon\,\mathbb{I} \quad \text{for all} \quad 1 \leq p \leq P \tag{70}$$

Where ϵ is determined by the desired sample complexity. The true value for ^2D should lie within the bounds above with probability $1-\delta$, where $\epsilon, \delta \in [0,1]$. Here, the number of target functions is equal to the dimensions of $\hat{S}_p$, i.e. $M = (2r)^2 \times (2r)^2 = 16r^4$. The sample complexity is thus,

$$N_{\text{tot}} = \mathcal{O}\left(\frac{\log 16r^4}{\epsilon^2} \max_{1\leq t\leq 16r^4} \left\| O_t - \tfrac{\text{Tr}(O_t)}{2^{2r}}\mathbb{I} \right\|^2_{\text{shadow}}\right). \tag{71}$$

where $O_t = \text{JW}(a_i^\dagger a_j^\dagger a_l a_k)$ for each unique permutation of i, j, k and l. An initial attempt to code this can be found in the `classical_shadows.ipynb` notebook in the GitHub repository.

D Spatial m-v2RDM Code

The full spatial v2RDM code with measurements is given below. An example usage is provided in Listing 2. Listing 3 contains the main SDP function. The code for handling measurement constraints is given in Listing 4. Listing 5 contains the functions that interface with OpenFermion to calculate the reduced Hamiltonian and $^2D_{\text{FCI}}$. The code for the ^2Q and ^2G mappings can be found in the `dqg.py` file in on GitHub.

```python
from openfermionpyscf import run_pyscf
from openfermion import MolecularData
from helpers import get_spatial_D2, generate_measurement

# Create molecule
geom = [('H', (0.0, 0.0, x)) for x in range(4)]
molecule = MolecularData(geometry=geom, basis='sto-3G', multiplicity=1, description='H4')
molecule = run_pyscf(molecule, run_fci=True)

# Gather measurements
r = molecule.n_orbitals
n_measurements = 10
D2fci = get_spatial_D2(molecule)
measurements = [generate_measurement(D2fci, r) for _ in range(n_measurements)]

# Run SDP
result = run_sdp(molecule, conditions='DQ', measurements=measurements)

# Output
print('FCI energy:', molecule.fci_energy)
print('SDP Result:', result['primal'])
```

Listing 2: How to run the spatial m-v2RDM for the hydrogen chain at 1.0Å with 10 measurements.

```python
import cvxpy as cp
from molecule_helper import getK2
from dqg import Q2_cvxpy, G2_cvxpy
from measurements import make_measurement_constraint

def run_sdp(molecule, conditions='DQG', measurements=[], noisy=False, epsilon=1e-8):
    r = molecule.n_orbitals
    N = molecule.n_electrons

    # Define optmization variable
    D2 = cp.Variable((r**2, r**2))

    # Define objective function
    K2 = getK2(molecule)
    E = lambda D : cp.trace(K2 @ D) + molecule.nuclear_repulsion
    objective = cp.Minimize(E(D2))

    # Generate measurement constraints
    measurement_constraints = []
    for (Un, Sn) in measurements:
        measurement_constraints += make_measurement_constraint(D2, Un, Sn, r, noisy, epsilon)

    # Define DQG constraints
    D_constraint = [D2 >> 0]
    Q_constraint = [Q2_cvxpy(D2, r, N) >> 0]
    G_constraint = [G2_cvxpy(D2, r, N) >> 0]
    trace_constraint = [cp.trace(D2) == 0.5 * N * (N-1)]

    # Build full constraints
    constraints = measurement_constraints + \
        (D_constraint if 'D' in conditions else []) + \
        (Q_constraint if 'Q' in conditions else []) + \
        (G_constraint if 'G' in conditions else []) + \
        (trace_constraint if len(conditions) > 0 else [])

    # Solve SDP
    problem = cp.Problem(objective, constraints)
    problem.solve(solver=cp.MOSEK, eps=1e-8)

    # Get results
    primal = E(D2).value
    dual = -constraints[-1].dual_value * 0.5 * N*(N-1) # dual obtained from trace

    return {'primal': primal, 'dual': dual, 'D2': D2.value}
```

Listing 3: The main function for the spatial m-v2RDM. Note that in quantum chemistry packages like OpenFermion, the nuclear repulsion constant is typically separated from the one-body and two-body integrals. Hence, it is added back during the energy calculation.

```python
import numpy as np
import cvxpy as cp
from scipy.stats import unitary_group

def generate_measurement(D2_fci, r, noisy=False, epsilon=1e-8):
    Un = unitary_group.rvs(r) # returns Haar distributed unitary
    UxU = np.kron(Un, Un)
    Sn = np.diag(UxU @ D2_fci @ UxU.conj().T).reshape((r, r), order='C')
    if noisy:
        gaussian_noise = np.random.normal(loc=0.0, scale=epsilon, size=Sn.shape)
        Sn = Sn + gaussian_noise
    return Un, Sn.real # .real discards imaginary part in case of tiny errors

def make_measurement_constraint(Dvar, Un, Sn, r, noisy=False, epsilon=1e-8):
    UxU = cp.kron(Un, Un)
    X = cp.real(cp.diag(UxU @ Dvar @ UxU.conj().T)).reshape((r, r), order='C')
    if noisy:
        return [Sn - 3 * epsilon <= X, X <= Sn + 3 * epsilon]
    else:
        return [Sn == X]

```

Listing 4: Functions to generate measurements and corresponding constraints.

```python
import numpy as np

def getK2(molecule):
    r, N = molecule.n_orbitals, molecule.n_electrons
    K2 = np.zeros((r, r, r, r))
    h1 = molecule.one_body_integrals
    h2 = molecule.two_body_integrals.transpose(0, 1, 3, 2) # chemist -> physicist notation

    for i, j, k, l in product(range(r), repeat=4):
        # Embed one-body terms into two-body form
        term1 = h1[i, k] * (1 if j == l else 0)
        term1 += h1[j, l] * (1 if i == k else 0)
        K2[i, j, k, l] = term1/(N-1) + h2[i, j, k, l] # Add two-body integrals
    return K2.reshape((r**2, r**2), order='C')

def get_spatial_D2(molecule):
    r = molecule.n_orbitals
    D2 = molecule.fci_two_rdm # returns spatial orbital version
    D2 = 0.5 * D2.transpose(0, 1, 3, 2) # adjust for normalization and chemist's notation
    return D2.reshape((r**2, r**2), order='C')

```

Listing 5: Functions used to obtain the reduced Hamiltonian and ${}^2D_{\mathrm{FCI}}$ from OpenFermion. Note that the two-body integrals and 2-RDM returned by OpenFermion are indexed using chemist's notation. Thus, transposition is required to convert to physicist's notation for energy calculations.