
Cambridge Quantum

Quantum Computational Quantification of Protein-Ligand Interactions

by Josh J. M. Kirsopp, Cono Di Paola, David Zsolt Manrique, Michal Krompiec, Gabriel Greene-Diniz, Wolfgang Guba, Agnes Meyder, Detlef Wolf, Martin Strahm, David Muñoz Ramo

Cambridge Quantum
JOSH J. M. KIRSOPP
josh.kirsopp@cambridgequantum.com
CONO DI PAOLA
cono.dipaola@cambridgequantum.com
DAVID ZSOLT MANRIQUE
david.zsolt.manrique@cambridgequantum.com
MICHAL KROMPIEC
michal.krompiec@cambridgequantum.com
GABRIEL GREENE-DINIZ
gabriel.greene-diniz@cambridgequantum.com
DAVID MUÑOZ RAMO
david.ramo@cambridgequantum.com

Roche Pharmaceutical Research & Early Development
Roche Innovation Center Basel, F. Hoffmann-La Roche Ltd.
Basel 4070, Switzerland
WOLFGANG GUBA
wolfgang.guba@roche.com
AGNES MEYDER
DETLEF WOLF
MARTIN STRAHM

Cambridge Quantum
Terrington House
13-15 Hills Road
Cambridge CB2 1NL
United Kingdom

Published by Cambridge Quantum

15 October 2021

Quantum Computational Quantification of Protein-Ligand Interactions

Josh J. M. Kirsopp^{*1}, Cono Di Paola¹, David Zsolt Manrique¹, Michal Krompiec¹, Gabriel Greene-Diniz¹, Wolfgang Guba^{†2}, Agnes Meyder², Detlef Wolf², Martin Strahm², and David Muñoz Ramo¹

¹*Cambridge Quantum Computing Ltd, 13-15 Hills Road, CB2 1NL, Cambridge, UK*

²*Roche Pharmaceutical Research & Early Development, Roche Innovation Center Basel, F. Hoffmann-La Roche Ltd., Basel, 4070, Switzerland*

15 October 2021

Abstract

We have demonstrated a prototypical hybrid classical and quantum computational workflow for the quantification of protein-ligand interactions. The workflow combines the Density Matrix Embedding Theory (DMET) embedding procedure with the Variational Quantum Eigensolver (VQE) approach for finding molecular electronic ground states. A series of β -secretase (BACE1) inhibitors is rank-ordered using binding energy differences calculated on the latest superconducting transmon (IBM) and trapped-ion (Honeywell) Noisy Intermediate Scale Quantum (NISQ) devices. This is the first application of real quantum computers to the calculation of protein-ligand binding energies. The results shed light on hardware and software requirements which would enable the application of NISQ algorithms in drug design.

1 Introduction

The advent of quantum mechanics at the turn of the 20th century changed the way we look at the physical sciences. For chemistry, the implications were profound, as Dirac famously noted: “the fundamental laws ... for the whole of chemistry are thus completely known, and the difficulty lies only in the fact that application of these laws leads to equations that are too complex to be solved” [1]. Indeed, calculations of accurate solutions of the electronic Schrödinger equation, such as the Full Configuration Interaction method (FCI), of molecules scale exponentially with the number of atoms [2], rendering them applicable only to the smallest systems. Practical approximations to FCI, such as e.g. CCSD(T) [3], touted as computational chemistry’s gold standard, do exist, but their applicability is limited: single-reference methods like CCSD(T) fail for strongly-correlated (“multireference”) systems and the formal scaling of CCSD(T) with system size is $O(N^7)$, rendering it useful only for relatively small molecules. While the scaling limitation of coupled cluster methods (CC) can be somewhat overcome via localized methods such as DLPNO [4] at the expense of accuracy, methods for strongly-correlated systems scale exponentially with the number of correlated electrons and, even for systems small enough to be tractable, require expert knowledge to be applied. On the other hand, mean-field methods, such as Hartree-Fock (HF) and Density Functional Theory (DFT), which either dispense with electron correlation completely or treat it in an approximate, implicit manner, are routinely

applied to organic and inorganic systems of sizes up to few hundred atoms. This is possible due to their relatively low, $O(N^3)$ - $O(N^4)$, formal scaling, which could be reduced even to $O(N)$ in approximate implementations for systems comprised of thousands of atoms [5–7]. However, HF and DFT fail for multi-reference (strongly-correlated) systems and cannot describe dispersion interactions, which are a critical component of inter-molecular forces, otherwise than via an ad hoc correction [8]. Clearly, general-purpose, low-scaling and highly accurate electronic structure methods applicable to arbitrary molecular systems remain elusive. It is widely believed that the simulation of complex and strongly-correlated chemical systems, for which standard methods are either inaccurate or too expensive, is among the areas which stand to benefit most from the ongoing and rapid progress in quantum computation [9]. Indeed, the last half-decade has seen an eruption of research in quantum algorithms for materials simulations, including computation of ground and excited states of molecules, quantum dynamics and linear response, among a plethora of others [10–14]. Noisy Intermediate Scale Quantum (NISQ) devices, however, have so far limited the applicability of these algorithms to very small systems, such as H_2 , LiH , RbH and the like [15, 16]. Nevertheless, rapid progress in the capabilities of quantum hardware as well as intensive research on new quantum algorithms opens the possibility of utilising quantum computing in Computer Aided Drug Design (CADD) in the future.

The synthesis of new drug candidates requires substantial medicinal chemistry efforts. CADD workflows limit

^{*}josh.kirsopp@cambridgequantum.com

[†]wolfgang.guba@roche.com

the need for synthesis by screening candidate molecules using a variety of approaches in order to prioritise the most promising compounds [17]. In the context of protein-ligand binding, Free-Energy Perturbation (FEP) methods based on molecular mechanics force fields [18, 19] are most commonly used to assess relative free energies of binding. Yet, as shown by SAMPL blind predictive modelling challenges [20, 21], state-of-the-art force field methods often yield binding free energies only moderately correlated to experimental data, depending on the chemical space, e.g. R^2 ranging from 0.1 to 0.9 for the OA/TEMOA dataset vs. $R^2 < 0.4$ across the board for the CB8 benchmark [20]. Moreover, errors were found to be non-systematic and frequently exceeded 2 kcal/mol (ca. 3 mHa), thus potentially precluding meaningful ranking of molecules for the purposes of screening [20, 21]. Development of low-scaling quantum-mechanical methods capable of simulating proteins is ongoing, in DFT [22] and CC [23] frameworks, but clearly, there is a pressing need for accurate methods for large-scale simulations of binding energies, not only with accurate ranking of molecules in mind, but also for the purpose of benchmarking approximate methods.

Recently, Malone et al. presented an attempt to utilise quantum computing to compute protein-ligand binding energies, via a 1st order Symmetry Adapted Perturbation Theory (SAPT) procedure [24]. Variational Quantum Eigensolver (VQE) quantum simulations of 6 small molecules and a subsection of the KDM5A protein were performed on a state-vector emulator to estimate their 1- and 2-electron reduced density matrices, which were then employed in a classical SAPT⁽¹⁾ calculation. However computed interaction energies did not reproduce ligand rankings yielded by more accurate 2nd order SAPT calculations, due to the missing induction and dispersion components in the 1st order approximation [24].

In this paper we compute binding energies for a series of 12 BACE1 inhibitors complexed in the active site of the BACE1 enzyme [25] using quantum hardware in an exhibition of what is possible with machines on the market today. The BACE1 enzyme is part of the amyloidogenic pathway and therefore associated with the pathogenesis of Alzheimer’s disease [26]. The BACE1 enzyme itself is a much-studied system [27], and the motivations of the system selection are fourfold: availability of experimental activity data and high-resolution 3D structures of protein-ligand complexes for a series of inhibitors, limited induced fit and inhibitor modifications across the series being local to the central recognition motif [25]. Consequently, this system provides an ideal case for exploratory studies of protein-ligand interactions and algorithmic prototyping of fragmentation methods.

2 Methods

2.1 Exploratory Model of BACE1 - ligand complexes

State-of-the-art in-silico prediction of free energies of binding usually involves a Free Energy Perturbation (FEP) protocol based on long classical molecular dynamics sim-

ulations[18], which would be intractable if performed on a quantum-chemical level. A common approximation employed in quantum-mechanical studies of protein-ligand binding is to compute an average over single-point free energies of binding, calculated for selected configurations where the entropy and thermal corrections are obtained from a lower-level model [22].

In the present study, the (electronic) binding energy, E_{bind} , is chosen as a simple metric by which to measure the strength of the interaction between the protein and the ligand, and is calculated as

$$E_{\text{bind}} = E_{\text{complex (aq)}} - E_{\text{ligand (aq)}} - E_{\text{protein (aq)}}, \quad (1)$$

where the subscript (aq) indicates molecules in aqueous environment. We treat the protein as a set of distributed point charges and by invoking a “frozen protein” approximation, cancel protein-protein terms from E_{complex} and simplify the above equation to a single energy difference,

$$E_{\text{bind}} = E_{\text{ligand-in-protein (aq)}} - E_{\text{ligand-in-solvent (aq)}}. \quad (2)$$

The ligand is treated with a variety of electronic structure methods in the coming sections. For the treatment of solvent effects in the $E_{\text{ligand-in-protein}}$ quantity, an explicit-water solvation shell is added and the TIP3P water charges appended to the protein charges experienced by the ligand [28]. The solvation shell in $E_{\text{ligand-in-protein}}$ is not equilibrated. Solvent effects are accounted for in $E_{\text{ligand-in-solvent}}$ using the dd-COSMO implicit solvent model available in the PySCF package [29–31]. The assumption is also made that the crystal structures of the ligands in the active site of each protein are sufficiently close to the optimal geometries that they need not be optimized. An optimization of the ligand in the active site under the approximations made is not possible without the parameterisation of Lennard-Jones potentials, which is beyond the scope of this publication, as the lack of Pauli-repulsion would result in the ligand being drawn unphysically close to the point charges representing the active site. In contrast, the geometries of the ligands used in the evaluation of $E_{\text{ligand-in-solvent}}$ are optimised in the presence of the implicit solvent at Hartree-Fock level using the STO-3G basis together with the geometry optimiser, such that some amount of the geometric distortion needed for the ligand to occupy the active site is accounted for in E_{bind} [32]. The full series of oxazine ligands considered in this work is shown in figure 1. X-ray crystallographic structures of protein-ligand complexes are taken from prior work [25] and processed in MOE (Chemical Computing Group) [33]. The 3D Protonate routine [34, 35] was applied with standard settings. An additional 10Å shell of water molecules was added and the negative charge neutralised by adding sodium cations. Protein structures were aligned with respect to the 1b protein-ligand complex. Fixed-point charges were placed on the protein atoms using the AMBER10:EHT method, as implemented in MOE.

2.2 Ranking metric

The binding energy expression for ligand i (eq. 2) depends explicitly on the geometries as

$$E_b^{(i)}(\bar{P}_i, W_i) = E(\bar{P}_i, W_i, \bar{L}_i) - E(\bar{W}, \bar{L}_i)$$

where \bar{P}_i and \bar{L}_i are the X-ray protein and ligand structures for the ligand i , W_i are explicitly added water molecules, \bar{W} denotes the implicit dd-COSMO water environment, and \tilde{L}_i is the optimized ligand geometry with dd-COSMO implicit water. Ranking of the ligands depends on the accuracy and precision of the metric

$$\begin{aligned} m_{ij} &= E_b^{(i)}(\bar{P}_i, W_i) - E_b^{(j)}(\bar{P}_j, W_j) \\ &= E(\bar{W}, \tilde{L}_i) - E(\bar{W}, \tilde{L}_j) \\ &\quad + E(\bar{P}_i, W_i, \bar{L}_i) - E(\bar{P}_j, W_j, \bar{L}_j), \end{aligned} \quad (3)$$

where i and j index the system under consideration. A more precise metric could be obtained by sampling multiple geometries. To keep the number of quantum computations at a feasible level, but reduce uncertainties due to the explicit water we also used a modified metric with a single reference protein-water complex of the ligand r

$$\begin{aligned} m_{ij}^{(r)} &= E_b^{(i)}(\bar{P}_r, W_r) - E_b^{(j)}(\bar{P}_r, W_r) \\ &= E(\bar{W}, \tilde{L}_i) - E(\bar{W}, \tilde{L}_j) \\ &\quad + E(\bar{P}_r, W_r, \bar{L}_i) - E(\bar{P}_r, W_r, \bar{L}_j) \end{aligned} \quad (4)$$

and for the reference the protein-water complex of ligand 1b is chosen as it represents the unsubstituted prototype of the series. In other words, the energy of each protein-ligand complex is calculated using fixed point charges from the 1b-protein complex. This replacement is made possible by alignment of the protein residues, minimal induced fit and almost perfect alignment of the ligand molecules: all oxazines in the series can be seen super-imposed in the 1b protein cavity in Figure 2. We note that almost any other fixed protein configuration could have been chosen instead of that of the 1b complex. If, on the contrary, the metric uses ligand-specific protein fixed charges, the correlation to experimental data is relatively poor. A full sweep of the 12 protein structures performed at Hartree-Fock level using the STO-3G basis is shown on Figure 12 in the Appendix.

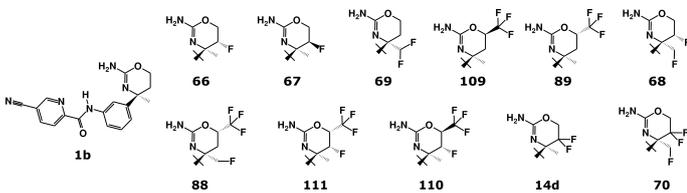


Figure 1: Series of oxazine BACE1 inhibitors investigated by Hilpert et al. [25]

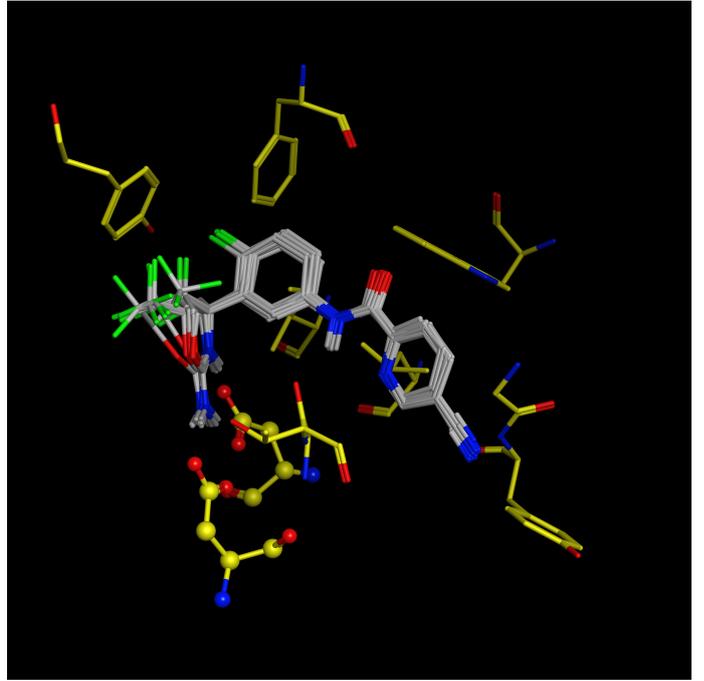


Figure 2: Experimental ligand geometries in the 1b cavity.

For the purposes of exposition we make the assumption that there exists a linear correlation [19] between the binding energy and the experimentally determined pIC_{50} data [25] shown in table 1.

Ligand	pIC_{50}
89	7.92
111	7.89
88	7.72
14d	7.64
66	7.55
110	7.40
70	7.31
1b	7.29
69	7.27
68	7.11
67	6.36
109	6.11

Table 1: Experimental pIC_{50} data for the investigated BACE1 inhibitors shown in Figure 1 [25].

2.3 Density Matrix Embedding Theory

To facilitate the simulation of such large systems on NISQ hardware, we reduce the quantum resources needed to evaluate the energy using density matrix embedding theory (DMET) [36, 37] implemented using VQE (see next section) as a fragment solver [38]. In DMET the atomic system is partitioned into fragments and the total energy of the system is the sum of the energies of fragments embedded in the environment comprising the rest of the system, represented by a quantum bath. The latter reproduces the entanglement of the embedded fragment with the environment through an interaction between N_{orb}^A fragment orbitals and $N_{\text{orb}}^B \leq N_{\text{orb}}^A$ bath orbitals [36].

The starting point of the DMET calculation is the

Hartree–Fock solution of the entire molecule. In order to spatially separate the fragment from its environment we apply the Löwdin [39] method to orthogonalise the atomic orbitals and use it as a localized basis for the DMET procedure.

The Hartree–Fock (HF) Reduced Density Matrix (RDM) in the localized basis is

$$\Gamma_{pq} = \sum_i \bar{C}_{pi} \bar{C}_{iq}^\dagger, \quad (5)$$

where the quantities \bar{C}_{pi} are the coefficients of the canonical molecular orbitals. The indices p and q refer to functions in the localised basis, and i indexes the molecular orbitals. Once the localized basis orbitals of a fragment have been identified by a fragmentation strategy, the fragment bath orbitals are computed from Γ_{pq} as described by Wouters et al. [37]. The fragment orbitals and the bath orbitals, together referred to as embedding orbitals, are then used to construct the fragment Hamiltonian operator

$$\begin{aligned} \hat{H}_{frag}(\mu) = & -\mu \sum_{p \in A} a_p^\dagger \hat{a}_p + \sum_{pqrs}^{N_{orb}^A} (pq|rs) \hat{a}_p^\dagger \hat{a}_r^\dagger \hat{a}_s \hat{a}_q \\ & + \sum_{pq}^{N_{orb}^A + N_{orb}^B} (h_{pq} + \sum_{rs}^{N_{orb}} ((pq|rs) - (ps|rq)) \Gamma_{rs}^{env}) \hat{a}_p^\dagger \hat{a}_q \end{aligned} \quad (6)$$

where N_{orb}^A is the number of orbitals on the fragment, N_{orb}^B is the number of bath orbitals, N_{orb} is the number of orbitals on the whole molecule, \hat{a}^\dagger and \hat{a} are creation and annihilation operators, respectively, Γ_{rs}^{env} is the RDM constructed from the fully occupied orbitals in the environment, and μ is the chemical potential. The quantities h_{pq} and $(pq|rs)$ are the one and two-electron integrals in chemists' notation, respectively. The fragment energy operator is

$$\begin{aligned} \hat{E}_{frag}(\mu) = & \sum_{p \in A} \left(\sum_{q}^{N_{orb}^A + N_{orb}^B} \left(h_{pq} + \frac{\sum_{N_{orb}} [(pq|rs) - (ps|rq)] \Gamma_{rs}^{env}}{2} \right) \hat{a}_p^\dagger \hat{a}_q \right. \\ & \left. + \frac{1}{2} \sum_{qrs}^{N_{orb}^A + N_{orb}^B} (pq|rs) \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s \right) \end{aligned} \quad (7)$$

This procedure is performed for each fragment set by the fragmentation strategy and by using an appropriate wavefunction method the ground state $\Psi_{frag}(\mu)$ of $\hat{H}_{frag}(\mu)$ Hamiltonian is computed. Given $\Psi_{frag}(\mu)$,

$$E_{frag}(\mu) = \sum_{frag} \langle \Psi_{frag}(\mu) | \hat{E}_{frag}(\mu) | \Psi_{frag}(\mu) \rangle \quad (8)$$

is computed for each fragment and for a particular global chemical potential, μ . The value of μ is determined by the constraint

$$\sum_{frag} \langle \Psi_{frag}(\mu) | \hat{N}_{frag} | \Psi_{frag}(\mu) \rangle = N_{tot} \quad (9)$$

where \hat{N}_{frag} is the particle number operator of the fragment $frag$ and \hat{N}_{tot} is the total number of electrons in the molecule. In practice, this is achieved by an iterative solver. Finally, the DMET total energy is computed as

$$E = E_{nuc} + \sum_{frag} E_{frag}(\mu), \quad (10)$$

where E_{nuc} is the nuclear-nuclear repulsion energy in the total molecule.

2.4 The Variational Quantum Eigensolver

The variational quantum eigensolver (VQE) [40] is a hybrid quantum/classical algorithm for the evaluation of molecular electronic ground or excited states [41]. We use the VQE as the quantum computational solver for the DMET fragment Hamiltonian.

A series of 4-qubit experiments are performed with the YXXX ansatz and a reduced active space of two electrons in four spin-orbitals. We compare the results in section 3 with their noiseless analogues: the same calculations performed with a state vector simulator to enable comparison with “perfect qubits”. The general VQE schema is shown in figure 3.

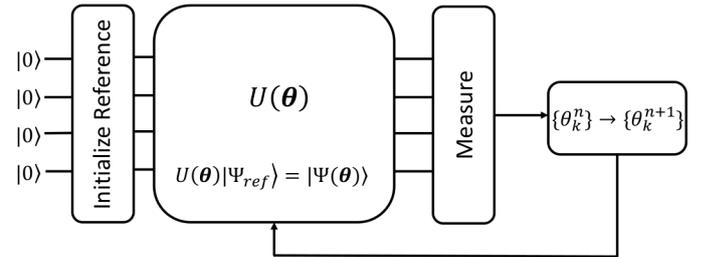


Figure 3: VQE loop. The set of parameters $\{\theta\}$ is updated iteratively by a classical optimiser, until the measured objective function is minimized.

The imaginary time evolution algorithm (ITE) was considered as an alternative to VQE for this work [42]. However, prohibitively large numbers of shots needed to reach the desired accuracy and difficulty converging to energetic minima suggested VQE, in this instance, represented a more robust approach.

2.5 DMET fragmentation

The chosen fragmentation pattern is based on known interactions between the headgroup and the binding pocket [25]. The primary inter-atomic interactions are shown in Figure 4. Each oxazine is subjected to the same fragmentation pattern, the tail group consisting of two aromatic rings and their substituents is treated as one large fragment and is denoted in this communication by the label [TG]. The headgroup, which is host to all substitutions throughout the ligand series is split into two fragments, one which contains the atoms responsible for the primary interactions with the catalytic aspartates, the other is made up of the remainder of the headgroup including any substitutions. These are referred to as $[\text{NH}_2\text{CNH}^+]$ and $[\text{R-HG}]$, respectively. The $[\text{NH}_2\text{CNH}^+]$ fragment is constant throughout the series of inhibitors being considered, but its energy is not, due to

dependence on varying bath orbitals. Moreover, this fragment is the only one treated with the VQE algorithm and the YXXX ansatz. The [R-HG] and [TG] fragments are treated at Hartree-Fock level. This choice of fragmentation ensures correlation energies and, crucially, binding energies remain comparable throughout the study. The fragmentation scheme is shown pictorially in Figure 5.

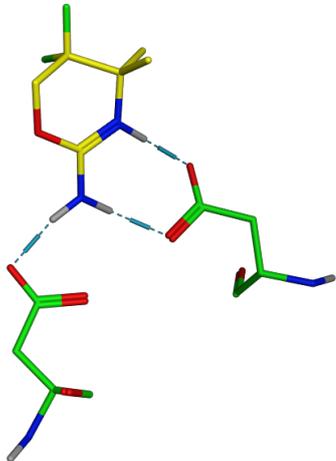


Figure 4: The primary interactions between the ligand and ASP residues in the active site are indicated by the blue dashed markers. The tail group is omitted for image clarity.

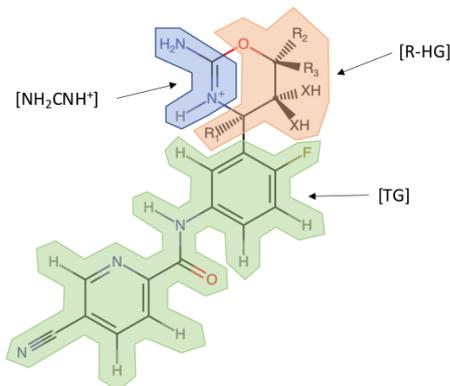


Figure 5: The chosen fragmentation pattern for the oxazine series.

2.6 Technical details of classical calculations and quantum simulations

All DMET calculations were performed using the EUMEN quantum chemistry software, utilizing the retargetable quantum compiler `t|ket` (via its Python interface, `pytket`) [43] for quantum circuit compilation and optimization for the target platforms and `PySCF` [29] for classical Hartree-Fock and CCSD calculations as well as for evaluation of molecular orbital integrals. `Qulacs` [44] and `Qiskit's AerSimulator` [45] emulators were used for the state-vector and shot-based simulations, respectively. Hardware experiments were performed on a superconducting transmon quantum computer, *ibmq_casablanca*, which is one of the IBM Quantum Falcon processors [46] and on a trapped-ion quantum computer, H1-S2, manufactured by Honeywell

Quantum Solutions [47]. For all calculations performed on quantum hardware we use the minimal STO-3G basis set [48] and limit the active space to two electrons in four spin-orbitals. A limited depth, custom Ansatz we shall refer to hereafter as “the YXXX Ansatz” [49] is used in all quantum computing simulations. The YXXX Ansatz contains only one two-electron excitation and provides a single-parameter circuit of lower depth than the analogous UCCD circuit [50]. The Jordan-Wigner transformation was used to encode the molecular Hamiltonian expressed in fermionic operators into the corresponding qubit operators [14]. For each energy calculation on the transmon machine, we first perform a variational calculation on the device to attain the optimised parameter for the YXXX ansatz, using 6000 shots in each iteration. The final energy value is evaluated with 60,000 shots (in 10 blocks of 6000) to reduce statistical error in the energy to that seen in the shot-based noiseless calculations. On the trapped-ion device, due to its slower speed and high qubit fidelity, we used statevector-optimised YXXX ansatz parameter and evaluated the VQE energy with 8000 shots. The error bars shown for the hardware and shot-based results are one standard deviation of the energy evaluated with the post-PMSV shot counts in 10 equally sized blocks with dimension $N_s/10$, where N_s is the total number of shots surviving the symmetry verification routine.

2.7 Error mitigation

In this study, we use two error mitigation techniques. The first is the SPAM correction [51–53] for noisy measurement and state preparation, implemented in `t|ket` [43]. The SPAM error mitigation technique works by profiling a device by first evaluating calibration circuits involving all pairs of used qubits and comparing it with the results expected in the absence of noise to obtain a transfer matrix, the inverse of which is then applied to post-process every measurement. The second technique applied in this paper is Partition Measurement Symmetry Verification (PMSV), a recently reported flavour of symmetry verification [54]. Symmetries such as mirror planes and particle number symmetries can be represented by single Pauli strings [55]. The symmetries and Hamiltonian elements are partitioned into sets of commuting terms [56, 57]. If each member of a group commutes with a given Pauli symmetry, that symmetry can be used for verification of the measurement result of the corresponding measurement circuit without the need for extra resources, i.e. any measurement which violates the known symmetry of the system must be due to noise, and hence is discarded during post-processing.

3 Results

We have constructed a simplified model for calculation of binding energies of ligands with the binding site of the BACE1 protein. We first briefly review two literature reports on theoretical calculations of ligand-BACE1 binding energies. We evaluate our model using classical mean-field methods and subsequently implement it within a quantum computing framework.

3.1 State-of-the-art in computation of ligand-BACE1 binding energies

Binding energies of 35 small molecules with monocyclic and bicyclic amidine or guanidine cores to a simplified model of the BACE1 protein have been calculated using B3LYP and M06-2X functionals in the 6-31+G** basis set and compared to experimentally determined potencies (pIC_{50}) [58]. To compute the binding energies, the authors used single-point electronic energy calculations of solvated ligand and ligand-protein complexes, neglecting entropy effects. A fair, solvation model-dependent correlation between predicted and experimental ligand affinities was observed, with the SM8 model yielding a R^2 of 0.65 and Mean Absolute Deviation (MAD) of 0.4 kcal/mol (0.6 mHa). However, the predicted relative binding energies were, according to the authors, exaggerated compared to experimental binding energies, with the slope of core potency to predicted affinity of about 0.4 (instead of the expected 1.0) [58]. The accuracy of this DFT-based model was similar to a state-of-the-art classical molecular dynamics-based Free Energy Perturbation (FEP) simulation of relative binding affinity of 32 BACE1 inhibitors to the full BACE1 protein (R^2 0.7, MAD 0.6 kcal/mol).[59] These two studies demonstrate that state-of-the-art methods offer only modest accuracy and that single-point energy calculations on a simplified system can yield comparable results to FEP simulations of the whole protein-ligand complex.

3.2 Evaluation of the Model with Classical Methods

To better understand the limitations of the present model, we compute the binding energy of all ligands in the series in the binding pocket of 1b as described in the Methods section, using a density functional of comparable or better quality [60], ω B97X [61], to the one used in the study by Roos *et al.* [58] cited in the previous section and a larger basis set, def2-TZVP, as well as at the HF/STO-3G level. The results of these calculations are shown in Figure 6 and demonstrate that application of a large basis set and a more sophisticated electronic structure method does not improve the agreement of binding energy ranking with the experimental ranking compared to HF/STO-3G. Hence, the use of a minimal basis set is not likely to be the main source of error in our model.

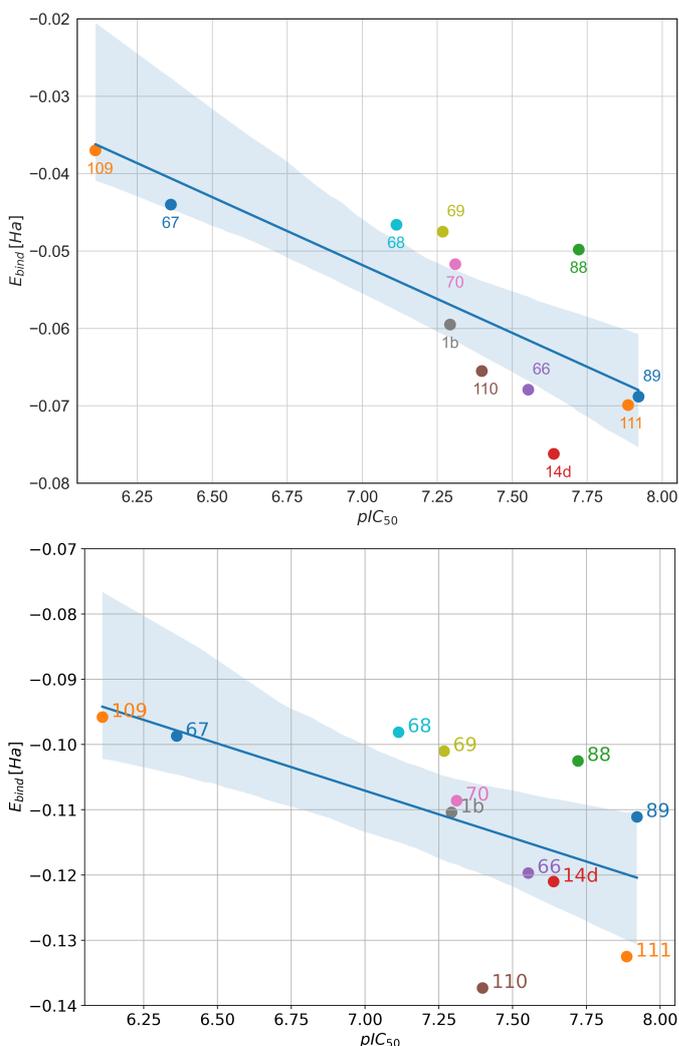


Figure 6: RHF/STO-3G (top, $R^2 = 0.61$) and ω 97X/def2-tzvp (bottom, $R^2 = 0.35$) binding energies against experimental pIC_{50} . Each point is labelled with the corresponding ligand ID.

3.3 Classical DMET Calculations and Noiseless Quantum Simulations

Adopting the fragmentation strategy outlined in the Methods section, we construct the full molecular energy as a sum of the fragment energies for the ligand both in solvent and in the active-site of the BACE1 protein and compute the binding energy according to equation 2. To establish a baseline for DMET experiments on quantum computers, we have performed a series of calculations where the solver for the $[NH_2CNH^+]$ fragment is either classical Hartree-Fock, classical CCSD (with or without an active space approximation) or state-vector VQE simulation on a quantum emulator. For the [TG] and [R-HG] fragments, the solver is always chosen to be classical Hartree-Fock.

We first consider a CCSD classical solver for the $[NH_2CNH^+]$ fragment both with and without an active space approximation. We show in figure 7a the binding energies obtained using CCSD with no active space restriction. Figure 7b shows the binding energies calculated using CCSD with a 2-electrons-in-4-spin-orbitals active space (to match the VQE simulations with the YXXX Ansatz) and

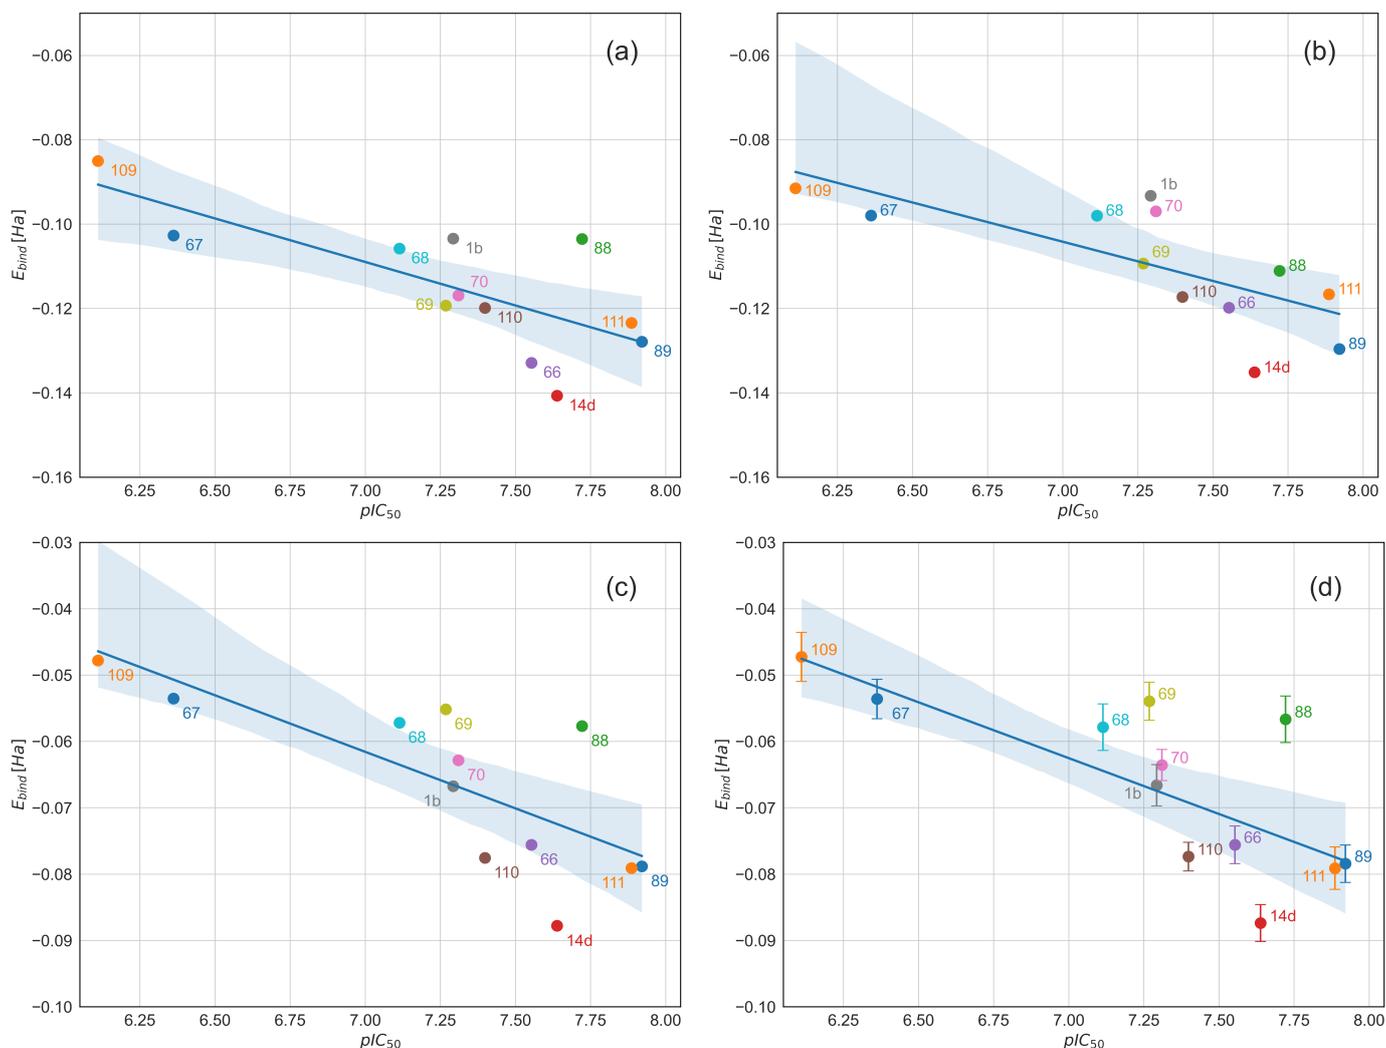


Figure 7: Binding energies obtained using (a) a classical CCSD solver without an active space, (b) a classical CCSD solver with a HOMO/LUMO active space, (c) a state vector YXXX-VQE solver with a HOMO/LUMO active space, and (d) a shot-based, noiseless YXXX-VQE solver with a HOMO/LUMO active space for the $[\text{NH}_2\text{CNH}^+]$ fragment. All energies are shown in Hartrees, the light blue shaded regions represent the 95% confidence intervals. In each graph the points are labelled with the corresponding ligand IDs.

demonstrates minor errors introduced by the active space approximation: the ordering of some strongly-bound ligands is reversed, but the overall trend is maintained.

Before appealing to hardware we first evaluate the series of binding energies using a state-vector simulator, which provides a noiseless benchmark for the quantum device calculations described in the coming subsection. The results of these simulations are shown in Figure 7c.

The outcome of running the VQE calculation is an approximation to the electron correlation energy, which, by definition, must be negative. With noisy quantum devices, it is not guaranteed to be the case: one could feasibly measure an energy above the mean-field result as a consequence of device errors. As such, to facilitate meaningful discussion in the later sections and benchmark the device errors we experience, we show in table 2 the correlation energies obtained from the state vector simulations. It is important to note that the correlation energies do not cancel out in the difference between energies of bound and solvated ligands. On the contrary, the interaction of the $[\text{NH}_2\text{CNH}^+]$ frag-

ment with the charge shell representing the protein results, in all cases, in a larger correlation energy. Since electron correlation depends on the separation between electrons, it is likely that a large portion of this correlation arises through the positively charged fragment drawing in electron density from the rest of the molecule.

Ligand	$E_{\text{ligand-in-protein (aq)}}^{\text{corr}}$	$E_{\text{ligand (aq)}}^{\text{corr}}$
89	-0.0345	-0.0245
111	-0.0354	-0.0262
88	-0.0324	-0.0246
14d	-0.0348	-0.0232
66	-0.0294	-0.0218
110	-0.0386	-0.0266
70	-0.0345	-0.0234
1b	-0.0273	-0.0201
69	-0.0282	-0.0205
68	-0.0326	-0.0220
67	-0.0313	-0.0218
109	-0.0358	-0.0252

Table 2: Correlation energies obtained using the YXXX ansatz in state-vector calculations of $E_{\text{ligand-in-protein}}(\text{aq})$ and $E_{\text{ligand}}(\text{aq})$ for each oxazine.

Figure 7d shows the results of shot-based noiseless calculations of all 12 oxazines binding energies. These quantities provide an estimate for what proportion of the uncertainty seen in the hardware simulations arises through the limited sampling of the distribution of states. In total we perform 60,000-80,000 shots per each energy evaluation and, on average, reduce the deviation from state vector simulations to less than 0.002 Ha.

3.4 Simulations on Quantum Computers

We perform all hardware experiments using IBM’s *ibmq_casablanca* superconducting transmon and Honeywell Quantum Solutions’ *H1-S2* trapped-ion NISQ devices. We find that application of error mitigation techniques is necessary for obtaining meaningful results of the simulation. Without error mitigation, we observe large variance of results and significant deviations from expected energies, in the worst case resulting in a shift of almost 0.77 Ha away from the state vector result, while in the best case we see a disagreement of only 0.012 Ha. Post-processing the device outputs with PMSV method, or with both PMSV and SPAM, results in suppression of errors to ± 0.02 Ha, see Figure 8. Comparison with Table 2 shows that using the raw (unmitigated) quantum device shot counts to construct the total energies leads to errors similar to, or in some cases larger than, the correlation energies for the oxazines in the series. This highlights the need for an effective error mitigation approach.

Figure 8 shows the impact of the post-processing of measurement counts using both regimes for both the fragment energy in the implicit solvent, and in the charge shell representing the protein, computed on the *ibmq_casablanca* device. We observe that the maximum deviation of the fragment energy from the state vector result is 0.019 Ha (an improvement over the unmitigated results by an order of magnitude), while the minimum deviation from the ideal result is just 0.0004 Ha. In all cases we observe that the deviation from the state vector result is less than the correlation energy of the respective fragment.

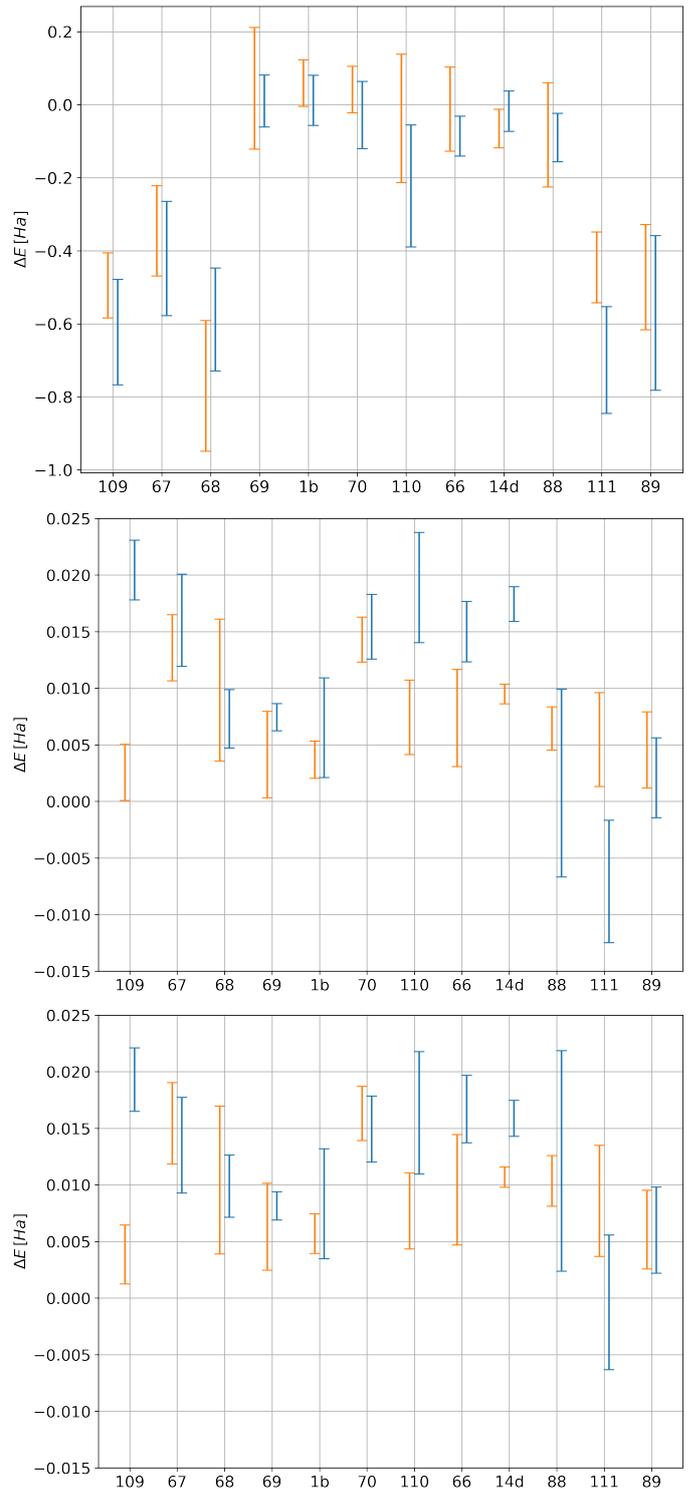


Figure 8: Deviations of fragment energies calculated on the *ibmq_casablanca* hardware (HW) from statevector (SV) simulations (in Hartree) of $E_{\text{ligand}}(\text{aq})$ (orange line) and $E_{\text{ligand-in-protein}}(\text{aq})$ (blue line), without error mitigation (top), after the application of PMSV in isolation (middle) and in conjunction with SPAM (bottom). The error bars represent standard deviations of data.

These results show that successive application of the SPAM and PMSV error mitigation techniques does not yield significant improvement over using PMSV on its own. Hence, we elect to only use PMSV for further experiments on the transmon and trapped-ion devices. In Table 3, we

compare deviations of binding energies obtained from hardware experiments on transmon (all 12 ligands) and trapped-ion devices (six selected ligands, in order to most effectively utilise limited computational resources) from ideal state-vector results and note similar accuracy delivered by both device architectures under our experimental conditions. Figures 9 and 10 show the binding energies plotted against pIC_{50} calculated on *ibmq-casablanca* and *H1-S2* after PMSV is applied.

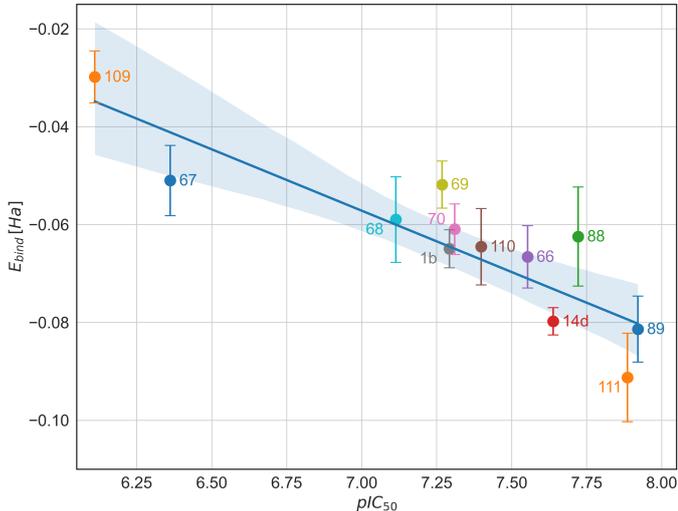


Figure 9: Binding energies E_{bind} (in Hartree) computed on *ibmq.casablanca* using DMET-VQE with the YXXX ansatz. The solid line is a line of best fit ($R^2 = 0.77$), the light-blue area represents 95% confidence interval. The error bars are standard deviations of the results. Each point is labelled with the corresponding ligand ID.

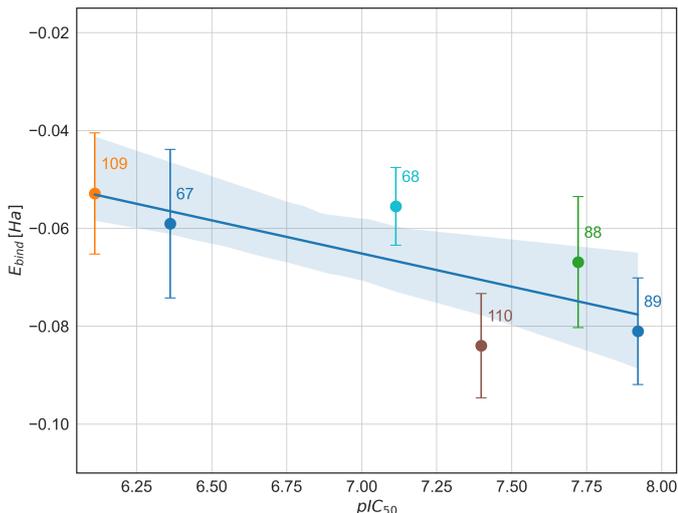


Figure 10: Binding energies E_{bind} (in Hartree) computed on *H1-S2* device using DMET-VQE with the YXXX ansatz. The solid line is a line of best fit ($R^2 = 0.56$), the light-blue area represents 95% confidence interval. The error bars are standard deviations of the results. Each point is labelled with the corresponding ligand ID.

Ligand	$E_{bind}^{Statevector}$	$\Delta E_{bind}^{Transmon}$	$\Delta E_{bind}^{Trapped-Ion}$
89	-0.07883	-0.00254	-0.00220
111	-0.07909	-0.01217	-
88	-0.05767	-0.00479	-0.00922
14d	-0.08779	0.00803	-
66	-0.07558	0.00899	-
110	-0.07755	0.01303	-0.00642
70	-0.06284	0.00189	-
1b	-0.06673	0.00179	-
69	-0.05515	0.00332	-
68	-0.05718	-0.00180	0.00168
67	-0.05353	0.00253	-0.00551
109	-0.04778	0.01795	-0.00510

Table 3: Deviations (in Hartree) of binding energies computed from hardware experiments performed on transmon (IBM) and trapped-ion (Honeywell) devices from statevector simulation results.

3.5 Ranking of BACE1 ligands

One of the main objectives of calculation of protein-ligand binding energies is ranking of ligands to support the prioritisation of design ideas [19]. Ideally, the simulated ranking should reproduce experimental trends in potencies, such as pIC_{50} values. We used coefficients of determination R^2 of pIC_{50} and computed binding energies to assess the quality of our model. For the ideal state-vector simulations, R^2 equals 0.55, indicating modest agreement with experimental data. Simulations obtained on *ibmq-casablanca* yield a R^2 of 0.77, indicating non-negligible impact of device noise on the ligand order and fortuitous error cancellation. Quantum simulations performed on the *H1-S2* device on half of the ligands result in a R^2 of 0.56, essentially identical to that of the state-vector simulation.

The ranking we obtained on the quantum computer is not improved with respect to the one calculated at HF/STO-3G level (see above), presumably due to the simplicity of our model of the protein-ligand interaction, which does not guarantee convergence to the experimental data. We note, however, that our results are comparable to the correlation of the data reported in the DFT study by Roos et al. [58] ($R^2 = 0.65$).

Aside from the coefficient of determination, one can assess the quality of our prediction by comparing the distributions of binding energies computed for the “weakly bound ligands” (67 and 109) to the “strongly bound” ones (the remaining molecules), see Figure 11: while the distributions are not perfectly separated, their means are shifted by over 0.02 Ha.

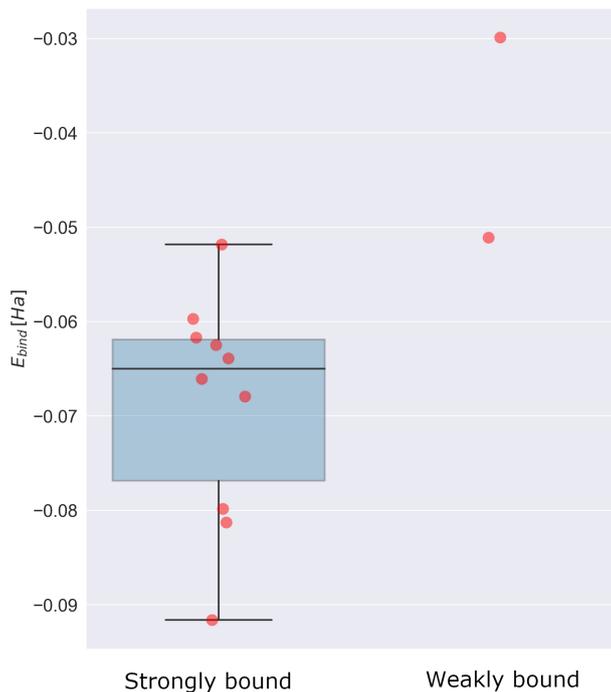


Figure 11: Comparison of distributions of binding energies, computed on *ibmq.casablanca*, of strongly bound and weakly bound ligands (see text).

4 Discussion

In this communication we provided a simple model and the first application of quantum simulation to pharmaceutically interesting chemical systems. We took a much-studied system for which benchmark data exists, and showed that by making simplifications and using embedding techniques, it is possible to make use of NISQ hardware to compute protein-ligand binding energies. While the ranking of ligands by the computed binding energies does not reproduce experimental ordering, it is comparable (in terms of the correlation coefficient) to results of state-of-the-art DFT and classical MD simulations on similar BACE1 inhibitors reported by others [58, 59].

Our approach differs significantly from that of Malone et al. [24], who also used quantum computation to estimate protein-ligand binding energies. Their simulations were, however, performed on a classical quantum emulator. They claim that their method, SAPT(VQE), is inherently more precise due to reliance only on a single VQE measurement, as opposed to calculating the difference between two VQE energies, because the latter results in accumulation of errors. However, the accuracy of Malone’s workflow is limited by truncation of the SAPT expansion at 1st order and it is not clear how their method can be effectively extended to higher orders.

The method presented in this publication, while hostage to many approximations and not production-ready yet, has a strong connection to established embedding and QM/MM

methods and can, therefore, be systematically improved. It would be possible to extend this model by taking many snapshot calculations as each protein-ligand complex is propagated classically, or to refine its QM part by including atoms of the protein’s binding site in the QM region, applying a correlated method to more fragments and including van der Waals interactions between the QM fragment and the rest of the system. The $[\text{NH}_2\text{CNH}^+]$ fragment is among the largest (in terms of number of atoms) chemical systems simulated on quantum hardware to date, the results obtained from these experiments show promisingly small amounts of noise after PMSV error mitigation is applied. It is worth emphasising that without PMSV error mitigation, deviations from ideal state-vector results are very large and far exceed the correlation energies themselves.

A major bottleneck in the application of quantum computation to the evaluation of molecular energies is found in the circuit depth required to construct the quantum state corresponding to chemical Ansätze, which very quickly becomes intractable on NISQ machines as active spaces expand to many qubits. Relatively high gate times (in comparison to transistor times in classical computers) lead to the requirement that the state of the qubit register remains coherent over time scales not currently possible. With parallel gate execution quickly becoming a reality, advances in quantum error correction and the continued improvement of QPUs in other aspects, we expect the size and complexity of models such as the one in this publication to increase rapidly, and the number of necessary approximations to facilitate quantum simulation to decrease. Specifically, we expect that quantitatively accurate results could be obtained if (1) a high-quality (e.g. triple- ζ) atomic basis set and a large (or even complete) active space are used, (2) a correlated calculation is performed on several (if not all) fragments and (3) at least the aminoacid residues closest to the ligand are included in the QM region. The first condition would warrant using up to 250 qubits for the $[\text{NH}_2\text{CNH}^+]$ fragment, necessitating significant improvements in capabilities of quantum hardware compared to current state-of-the-art, as well as novel software techniques to reduce quantum resource requirements. Such methods may comprise e.g. a combination of a compact qubit mapping [62] and PNO approximation [63], or application of a transcorrelated Hamiltonian [64]. Treatment of all fragments at a correlated level would be straightforward, as long as fragment size is kept limited (i.e. comparable to $[\text{NH}_2\text{CNH}^+]$). The requirement of having an explicit representation of the closest protein residues can be met by, for example, following the recipe for a minimal BACE1 active site model of Roos et al. [58]. Furthermore, while our workflow currently uses VQE executed on a NISQ device, there is nothing to prevent it from using a Quantum Phase Estimation solver running on a fault-tolerant quantum computer, once the latter is available. Precise evaluation of interaction energies, whether between small molecules or between proteins and ligands, is one of the central problems in computational chemistry and has countless practical applications. Hence, we expect the model presented herein to provide a platform for the development of more refined models as the capabilities of quantum computing hardware improve.

5 Conclusion

This exploratory study of protein-ligand binding energies evaluated using quantum hardware, as far as we are aware, is the first of its kind and provides an indicator of how rapidly the field of quantum computation has progressed in recent years. The purpose of this publication is to lay the groundwork for future development of a productive workflow for CADD and to evaluate the performance of hardware when the energies being measured are to the order of hundreds of Hartrees. We conclude that, if appropriate error mitigation techniques are applied, and the system simplified sufficiently, NISQ hardware is capable of recovering the trend obtained using classical computers. We suggest that improvement of the workflow described herein need not be limited to progress in hardware capabilities but also found in the refinement of the model, by expansion of the QM region to include atoms of the binding site as well as inclusion of dispersion interactions between the QM and MM regions.

Acknowledgements

We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team. We thank Brian Neyenhuis and all the Honeywell Quantum Solutions team for their availability and support with the H1-S2 device. The state-vector simulations in this work were performed on Microsoft Azure Virtual Machines provided by the “Microsoft for Startups” program. We would also like to thank Irfan Khan (CQ) for his helpful advice regarding hardware experiments. We thank Georgia Christopoulou and Andrew Tranter (CQ) for reviewing the manuscript of this work and for their valuable comments. We are grateful to Alyn Chad Edwards (CQ) for project management and providing feedback on the manuscript draft. We thank the members of the Roche Pharma Quantum Computing Task Force, specifically Mariëlle van der Pol, Stanislaw Adaszewski and Yvonna Li for insightful discussions and Bryn Roberts for initiating this Task Force at Roche.

References

- (1) P. A. M. Dirac and R. H. Fowler, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 1929, **123**, 714–733.
- (2) T. Helgaker, P. Jorgensen and J. Olsen, *Molecular electronic-structure theory*, John Wiley & Sons, 2014.
- (3) R. J. Bartlett and M. Musiał, *Rev. Mod. Phys.*, 2007, **79**, 291–352.
- (4) D. G. Liakos, Y. Guo and F. Neese, *The Journal of Physical Chemistry A*, 2020, **124**, 90–100.
- (5) J. C. A. Prentice et al., *The Journal of Chemical Physics*, 2020, **152**, 174111.
- (6) L. E. Ratcliff et al., *The Journal of Chemical Physics*, 2020, **152**, 194110.
- (7) J. Kusmann, M. Beer and C. Ochsenfeld, *WIREs Computational Molecular Science*, 2013, **3**, 614–636.
- (8) S. Grimme et al., *Chemical Reviews*, 2016, **116**, 5105–5154.
- (9) H. Liu et al., *arXiv preprint: arXiv:2102.10081*, 2021.
- (10) K. Bharti et al., *arXiv preprint arXiv:2101.08448*, 2021.
- (11) S. McArdle et al., *Reviews of Modern Physics*, 2020, **92**, 015003.
- (12) B. P. Lanyon et al., *Nature chemistry*, 2010, **2**, 106–111.
- (13) Y. Cao et al., *Chemical reviews*, 2019, **119**, 10856–10915.
- (14) B. Bauer et al., *Chemical Reviews*, 2020, **120**, 12685–12717.
- (15) K. Yeter-Aydeniz et al., *Advanced Quantum Technologies*, 2021, **4**, 2100012.
- (16) J. E. Rice et al., *The Journal of Chemical Physics*, 2021, **154**, 134115.
- (17) *Computer-Aided Drug Design*, ed. D. B. Singh, Springer Singapore, 2020.
- (18) W. Jaspers, J. Åqvist and H. Gutiérrez-de-Terán, in *Protein-Ligand Interactions and Drug Design*, ed. F. Ballante, Springer US, New York, NY, 2021, pp. 203–226.
- (19) Z. Cournia, B. Allen and W. Sherman, *Journal of Chemical Information and Modeling*, 2017, **57**, 2911–2937.
- (20) A. Rizzi et al., *Journal of Computer-Aided Molecular Design*, 2018, **32**, 937–963.
- (21) M. Amezcua, L. E. Khoury and D. L. Mobley, *Journal of Computer-Aided Molecular Design*, 2021, **35**, 1–35.
- (22) L. Gundelach et al., *Phys. Chem. Chem. Phys.*, 2021, **23**, 9381–9393.
- (23) M. E. Beck et al., *Journal of Computational Chemistry*, 2020, **42**, 293–302.
- (24) F. D. Malone et al., *arXiv preprint: arXiv:2110.01589*, 2021.
- (25) H. Hilpert et al., *Journal of medicinal chemistry*, 2013, **56**, 3980–3995.
- (26) F. Rombouts et al., *Expert Opinion on Therapeutic Patents*, 2020, **31**, 25–52.
- (27) A. K. Ghosh and H. L. Osswald, *Chemical Society Reviews*, 2014, **43**, 6765–6813.
- (28) W. L. Jorgensen et al., *The Journal of chemical physics*, 1983, **79**, 926–935.
- (29) Q. Sun et al., *The Journal of chemical physics*, 2020, **153**, 024109.
- (30) F. Lipparini et al., *Journal of chemical theory and computation*, 2013, **9**, 3637–3648.

- (31) F. Lipparini et al., *The Journal of chemical physics*, 2014, **141**, 184108.
- (32) L.-P. Wang and C. Song, *The Journal of chemical physics*, 2016, **144**, 214108.
- (33) *Molecular Operating Environment (MOE)*, version 2020.09, Chemical Computing Group ULC, 1010 Sherbrooke St. West, Suite #910, Montreal, QC, Canada, H3A 2R7, 2020.
- (34) P. Labute, *Proteins: Structure, Function, and Bioinformatics*, 2008, **75**, 187–205.
- (35) P. Labute, *Journal of Computational Chemistry*, 2008, **29**, 1693–1698.
- (36) G. Knizia and G. K.-L. Chan, *Journal of chemical theory and computation*, 2013, **9**, 1428–1432.
- (37) S. Wouters et al., *Journal of chemical theory and computation*, 2016, **12**, 2706–2719.
- (38) T. Yamazaki et al., *arXiv preprint: arXiv:1806.01305*, 2018.
- (39) P.-O. Löwdin, in ed. P.-O. Löwdin, Academic Press, 1970, vol. 5, pp. 185–199.
- (40) A. Peruzzo et al., *Nature communications*, 2014, **5**, 1–7.
- (41) O. Higgott, D. Wang and S. Brierley, *Quantum*, 2019, **3**, 156.
- (42) M. Motta et al., *Nature Physics*, 2019, **16**, 205–210.
- (43) S. Sivaraman et al., *Quantum Science and Technology*, 2020, **6**, 014003.
- (44) Y. Suzuki et al., *Qulacs: a fast and versatile quantum circuit simulator for research purpose*, 2021.
- (45) M. S. Anis et al., *Qiskit: An Open-source Framework for Quantum Computing*, 2021.
- (46) *IBM Quantum*, <https://quantum-computing.ibm.com/>, 2021.
- (47) *Honeywell and Cambridge Quantum reach new milestones*, <https://www.honeywell.com/us/en/press/2021/07/honeywell-and-cambridge-quantum-reach-new-milestones>, 2021.
- (48) W. J. Hehre, R. F. Stewart and J. A. Pople, *J. Chem. Phys.*, 1969, **51**, 2657–2664.
- (49) J. Lee et al., *Journal of chemical theory and computation*, 2018, **15**, 311–324.
- (50) Y. S. Yordanov et al., *arXiv preprint arXiv:2011.10540*, 2020.
- (51) C. Jackson and S. van Enk, *Physical Review A*, 2015, **92**, 042312.
- (52) C. Jackson and S. J. van Enk, *Phys. Rev. A*, 2015, **92**, 042312.
- (53) S. Bravyi et al., *Physical Review A*, 2021, **103**, DOI: 10.1103/physreva.103.042605.
- (54) K. Yamamoto et al., *arXiv preprint: arXiv:2109.08401*, 2021.
- (55) T.-C. Yen, R. A. Lang and A. F. Izmaylov, *The Journal of chemical physics*, 2019, **151**, 164111.
- (56) P. Gokhale et al., *arXiv preprint arXiv:1907.13623*, 2019.
- (57) A. Cowtan, W. Simmons and R. Duncan, *arXiv preprint arXiv:2007.10515*, 2020.
- (58) K. Roos et al., *Journal of Chemical Information and Modeling*, 2014, **54**, 818–825.
- (59) M. Ciordia et al., *Journal of Chemical Information and Modeling*, 2016, **56**, 1856–1871.
- (60) L. Goerigk et al., *Phys. Chem. Chem. Phys.*, 2017, **19**, 32184–32215.
- (61) J.-D. Chai and M. Head-Gordon, *The Journal of Chemical Physics*, 2008, **128**, 084106.
- (62) W. Kirby et al., *arXiv preprint arXiv:2109.14465*, 2021.
- (63) J. S. Kottmann et al., *The Journal of Physical Chemistry Letters*, 2021, **12**, 663–673.
- (64) M. Motta et al., *Phys. Chem. Chem. Phys.*, 2020, **22**, 24270–24281.

6 Appendix

6.1 Fixed or ligand-specific protein charges?

We have investigated the effect of protein charges used to compute $E_{\text{ligand-in-protein}}(\text{aq})$ on the ranking induced by computed binding energies and found that using ligand-specific charges (i.e. corresponding to the protein structure from each protein-ligand complex) yield much worse correlation than application of the same charges for all ligands, at both the RHF and DFT levels of theory, as seen on Figure 12. The relatively large R^2 value obtained with charges from the 1b-BACE1 complex, along with the 1b ligand being the unsubstituted base-molecule of the oxazine series, motivated the choice of these charges for the quantum simulations.

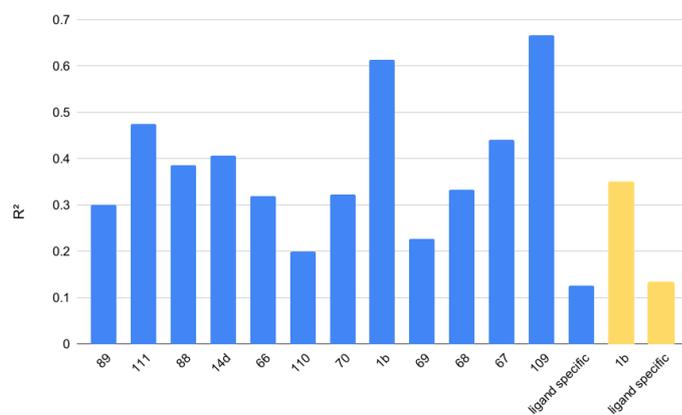


Figure 12: R^2 values for the full set of binding energies against pIC_{50} for each set of protein charges. The blue and gold bars indicate RHF/STO-3G and ω B97X/def2-TZVP levels of theory for the QM region, respectively.