Chemically Motivated Simulation Problems are Efficiently Solvable by a Quantum Computer

Philipp Schleich,[1,2,*] Lasse Bjørn Kristensen,[1,3,†] Jorge A. Campos Gonzalez Angulo,[3] Davide Avagliano,[3] Mohsen Bagherimehrab,[1] Abdulrahman Aldossary,[3] Christoph Gorgulla,[4,5] Joe Fitzsimons,[6] and Alán Aspuru-Guzik[1,2,3,7,8,9,10,11,‡]

[1]*Department of Computer Science, University of Toronto, Toronto, ON M5S 3G4, Canada.*
[2]*Vector Institute of Artificial Intelligence, Toronto, ON M5G 1M1, Canada.*
[3]*Department of Chemistry, University of Toronto, Toronto, ON M5S 3H6, Canada.*
[4]*Department of Structural Biology, St. Jude Children's Research Hospital, Memphis, TN 38105, USA.*
[5]*Department of Physics, Harvard University, Cambridge, MA 02138, USA.*
[6]*Horizon Quantum Computing, Singapore 138565.*
[7]*Department of Chemical Engineering & Applied Chemistry, University of Toronto, Toronto, ON M5S 3E5, Canada.*
[8]*Department of Materials Science & Engineering, University of Toronto, Toronto, ON M5S 3E4, Canada.*
[9]*Acceleration Consortium, Toronto, ON M5S 3H6, Canada.*
[10]*Canadian Institute for Advanced Research, Toronto, ON M5G 1Z8, Canada.*
[11]*Lebovic Fellow, Canadian Institute for Advanced Research, Toronto, ON M5S 1M1, Canada.*

Simulating chemical systems is highly sought after and computationally challenging, as the simulation cost exponentially increases with the system size. Quantum computers have been proposed as a computational means to overcome this bottleneck. Most efforts recently have been spent on determining the ground states of chemical systems. Hardness results and the lack of efficient heuristics for initial-state generation sheds doubt on the feasibility. Here we propose an inherently efficient approach for solving chemical simulation problems, meaning it requires quantum circuits of size scaling polynomially in relevant system parameters. If a set of assumptions can be satisfied, our approach finds good initial states by assembling initial states for dynamical simulation in a scattering tree. We discuss a variety of quantities of chemical interest that can be measured based on quantum simulation, e.g. of a reaction, succeeding the initial state preparation.

I. INTRODUCTION

The idea of using quantum computers for the simulation of quantum systems reaches back to the first proposals of quantum computing by Benioff, Feynman and Manin [1–3]. This idea has generated substantial effort towards the application of quantum computing to chemical problems, as quantum many-body simulations for chemistry are inherently limited by the curse of dimensionality and constitute a significant portion of current supercomputing usage. This effort has largely focused on two problems: quantum simulation of dynamics (also known simply as Hamiltonian simulation or quantum simulation in the field of quantum computing) and the ground-state problem. Quantum Simulation [4–8] describes the time evolution according to the Schrödinger equation under a Hamiltonian. The relevant Hamiltonians for chemical and most physical processes are efficiently computable and the simulation of their time evolution according to the Schrödinger equation on a quantum computer is provably within the computational complexity class BQP, the class of decision problems solvable on a quantum computer with polynomial time complexity and bounded error probability [6, 9]. Conversely, the ground-state problem – phrased as a decision problem known as the local Hamiltonian problem in quantum complexity theory – is complete for the class Quantum Merlin Arthur (QMA) [9–11], and is hence believed to not be efficiently solvable even on a quantum computer.

Ground-state search and Hamiltonian simulation are often used in an intertwined manner. Simulation algorithms serve as a subroutine to find the ground-state energy (e.g., via phase estimation [10, 12–14] or in the elucidation of reaction mechanisms described in [15], while ground states can be input states for performing quantum dynamics. State-of-the art implementations of Hamiltonian simulation fall into the classes of algorithms based on Trotter-Suzuki

* philipps@cs.toronto.edu
† l.kristensen@mail.utoronto.ca
‡ aspuru@utoronto.ca

formulas [9, 16], which split up an exponential of a sum into products of exponentials, algorithms based on "qubitization" [17, 18], which makes use of quantum signal processing [19] to encode the simulation in a quantum walk and randomized algorithms such as QDRIFT [20] and extensions [21, 22]. Each of these approaches has their advantages in different scenarios [23, 24]. However, as they all lead to polynomially-sized circuits for Hamiltonians of constant sparsity, we refrain from going into details here and refer to relevant literature. A good overview can be found in [25].

We propose a departure from how a majority of the quantum computing community approaches chemistry at this time – moving from a 1950's computational chemist's way of thinking [26, 27] that has been shaped by the limitations of classical computing to a new era: With fault-tolerant quantum computers, dynamical simulation of quantities that a practicing chemist might directly care about is within reach – most relevant quantum chemistry problems are inherently addressable through dynamical evolution alone, leading to efficient quantum algorithms for these problems.

We propose a framework that makes use of a limited set of attainable atomic initial states and then prepares input states for a reaction of interest through a scattering process – dynamically. Then, time evolution embodies the reaction, and a wide set of relevant quantities can be measured. This yields an algorithm that is not limited by the QMA-hardness of preparing ground states and thermal states. Similarly, we can circumvent an orthogonality catastrophe, namely a vanishing success probability to retrieve a ground state. As it has been recently shown [28], this effect is why exponential speedups for ground-state energy estimation of molecular Hamiltonians on quantum computers may be hard to attain as the (state-of-the-art) methods considered for state preparation fail to produce reliable overlaps for molecules of increasing size. Our framework stems from two observations. First, the simulation of k-local Hamiltonians is achievable by polynomial-size quantum circuits. Hamiltonians that stem from chemical problems are 2-local due to the nature of the Coulomb interaction and, thus, also have finite locality when represented as quantum operations on qubits. The inclusion of photons to simulate light-matter interactions in the computation increases the maximum support of operations but not beyond a constant factor. The second observation is that we aim to simulate processes corresponding to experiments that can be performed in a lab in finite time, i.e. problems that are, in some sense, experimentally verifiable. Molecular ground states as viewed from the perspective of computational chemistry, which imply a frozen molecular geometry at a temperature of zero Kelvin, typically do not represent a system's thermal equilibrium state and hence do not belong to this class. We propose to simulate the process of producing reactants with a hierarchical scattering process. We first prepare the ground states of atoms by quantum phase estimation and then use a simple scattering process to produce molecular reactants. Our method integrates artificial potentials and photonic fields to induce success of scatterings, leading to a lower-bounded probability of success. This means that a fixed number of repetitions of a weak measurement scheme to herald success will suffice to ensure success of intermediary scattering processes. This serves as efficient state preparation of a meaningful, molecular input state, which does not need to be a ground state. Our framework facilitates the modeling of complex chemical reactions by hierarchically operating the scattering with N atoms to create M reactants which then can undergo a quantum simulation corresponding to a reaction, followed by measurements of reaction rates and time correlation functions.

II. HAMILTONIAN SIMULATION AND COMPLEXITY CONSIDERATIONS

We continue by discussing relevant complexity classes for the problems we consider and argue about their efficiency. The complexity class BQP (bounded-error quantum polynomial-time) is oftentimes considered the quantum generalization of the complexity class P (polynomial time), or more precisely its probabilistic extension BPP (bounded-error probabilistic polynomial-time). Polynomial complexity is usually considered efficient as the increase in cost when scaling relevant parameters is somewhat moderate. Problems from QMA (Quantum Merlin-Arthur), in contrast, can be considered "hard"; QMA is the quantum analog of NP or, more precisely, the probabilistic class MA (Merlin-Arthur). QMA describes promise problems that can be verified in polynomial time with bounded error probability. However, there is no guarantee on efficiency in their solution. [29–31] We claim that a large class of chemically relevant phenomena can be addressed by efficiently solvable problems – more precisely, the intuitive argument we make is that phenomena that occur in finite time in a chemical laboratory can be simulated in finite time. The gap between the time in nature and on a quantum computer is only polynomially-sized, which is why we relate it to BQP. Systems in nature typically do not cool down to the ground state, thus it would not be part of the chemically relevant quantities we consider; in fact, it was recently shown that local minima are more favourable to reach in quantum systems that ground states [32].

Here, we can draw a parallel between polynomially-sized problems, such those from BQP, and QMA-hard problems such as finding the ground state. In this way, we push for a fundamental change in the way chemistry problems on quantum computers are typically approached and we display this in Fig. 1. We denote the set of computational problems of chemical relevance – including "unnatural" phenomena like ground states – Chem. Note that these are not decision problems; our argumentation is based on circuit fragments that stem from known complexity results rather than decision problems and the language of complexity theory. The focal set of problems within this work is what we call ChemPoly, namely problems of relevance to chemistry that have polynomial complexity. Due to the BQP-completeness of Hamiltonian simulation [33], one might even argue that the set of polynomially-sized problems are contained within the set of problems with chemical relevance. Irrespective of this point, our present study highlights a wide set of chemically interesting problems that involve simulation as a subroutine yet avoid subroutines and decision problems that are known to be QMA-hard. Many approaches in chemistry that involve search for the ground state are either QMA- or NP-hard [34–36]; examples beyond the local Hamiltonian problem are finding the universal functional in density functional theory or the Hartree-Fock problem. Significant progress was made on solving the latter numerically, to the point that it is mostly seen as "solved" despite the hardness results. This gives rise to hope regarding the ground-state problem as well. Yet still at this point we seem to be far from similar success on chemistry ground-state search, and the biggest hope for solving useful chemical problems lies in dynamics.

To represent chemical systems, we choose molecular Hamiltonians, not necessarily restricted to the Born-Oppenheimer (BO) approximation. On quantum computers, non-BO was previously explored in the split-operator formalism, where kinetic and potential dynamics are factored [37, 38]. While it was argued in [37] that the non-BO approach is more efficient than implementing the BO procedure, there have been recent advances that render BO more interesting again thanks to a fully coherent implementation [39]. Hamiltonians occurring in chemistry are composed by the operators outlaid in Table I from which locality is an obvious consequence, as k-locality from the 2-body nature of the interactions.

Charged particle kinetic energy	$p^2/2m$
Photon energy	$\omega(a^\dagger a + 1/2)$
Interparticle potential	$\frac{qq'}{\lvert r-r'\rvert}$
Photon-particle interaction	$\frac{q}{m}\mathbf{A}(r)[\frac{q}{2m}\mathbf{A}(r) - p]$

TABLE I. Hamiltonian components for molecular systems interacting with photons. $\mathbf{A}(r) = k[c(k,r)a + c(k,r)a^\dagger]$ is the vector potential corresponding to the photon field.

For a polynomially-sized circuit complexity, it is further necessary that the operator norm of the simulated Hamiltonian, or rather its individual terms that act on at most k qubits, grow at most polynomially in the number of qubits used to represent them [7–9]. The fulfillment of this requirement can be guaranteed because the system energy does not increase by refinement of the numerical representation. Further, because the energy is an extensive thermodynamic property, it grows roughly linearly by increasing the system size (i.e. the number of particles). Therefore, the thermodynamic relation of quantity of matter to internal energy is, in fact, linear. Our framework is independent of the choice of basis used to numerically represent the Hamiltonian as well as of the choice between first and second quantization. Asymptotically, a first-quantized representation tends to be more efficient for fault-tolerant quantum algorithms with an abundant number of logical qubits, as the number of required qubits grows linearly with the number of particles and logarithmically with the number of basis functions (since we store the basis information for every particle), as opposed to the linear dependence in the number of basis functions for second quantization, where we store occupation for each basis function.

III. COMPUTATIONAL FRAMEWORK

In what follows, we describe the computational framework in more detail. We discuss how to represent systems, their evolution, and chemically interesting observables that can be framed as measurements. Our overall idea is based on the experiment in [40], where molecules were "built" using two atoms with finite success probability. The chemical

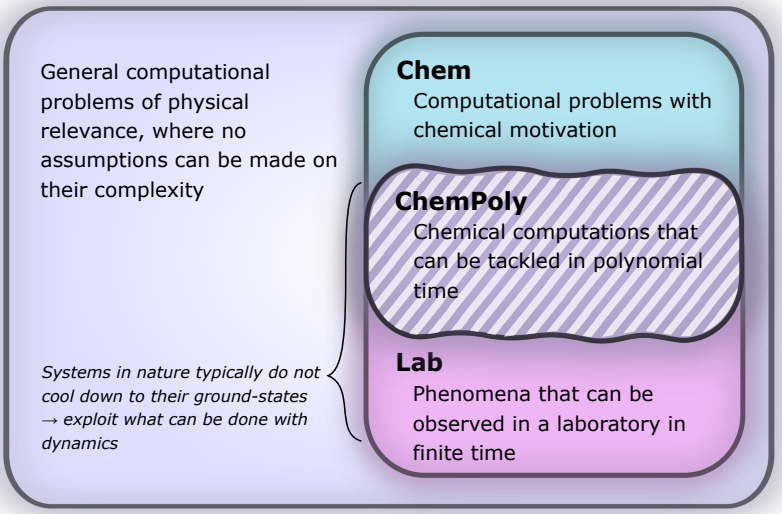

FIG. 1. Complexity of simulating chemical problems. We target the set of computational problems, ChemPoly, within chemistry that have polynomial complexity when simulated on a quantum computer and are observable in a lab; the sets Lab and ChemPoly may coincide.

dynamics we aim to simulate requires a sensible initial state for the time evolution. This needs to be a state that reproduces a natural state with high fidelity to successfully perform a desired experiment. A ground state would be an obvious candidate. However, as previously argued, general molecular ground states are hard to access. Nevertheless, preparing the ground states of atoms is feasible – the lighter the atom, the easier – and can be done efficiently as constant overlap for heuristic input states is expected [28]. This is why we propose to follow a hierarchical approach as outlined in Fig. 2. All processes generating the reactants and products are carried out through Hamiltonian simulation with local Hamiltonians, which include both the molecular and external field components. Since the simulation of k-local Hamiltonians is efficient, its overall computational cost is polynomial in the evolution time and system size; therefore, the simulation is feasible using quantum circuits with size that grows polynomially in relevant parameters. After simulation, we aim to e.g. measure reaction rates according to the scheme in [37, 41] or auto-correlation functions.

We start with a state representing N atoms, all of which are assumed to be in their ground states or another interesting state $\rho_i^{(\mathrm{atom})}$, such as a thermal state. The preparation of these states can be achieved by existing algorithms for ground- [10, 42, 43] and thermal-state preparation [44] followed by amplitude amplification to enhance the probability of obtaining the desired state. Since the atoms are all finitely sized and we can prepare an initial state with a significant overlap with the ground/thermal state for each atom, each of their amplitude amplification costs are constant. The overall cost to prepare the atomic initial states is $O(N\ \mathrm{poly}(\frac{1}{\varepsilon}))$, with ε being the accuracy in preparing the atomic states. Using these atoms encoded as a state $\rho_1^{(\mathrm{atom})}\rho_2^{(\mathrm{atom})}\cdots\rho_N^{(\mathrm{atom})}$, we seek to create M molecular reactants $\rho_1^{(\mathrm{mol})}\rho_2^{(\mathrm{mol})}\cdots\rho_M^{(\mathrm{mol})}$.

Based on these initial states, we next prepare the reactants by a scattering process mimicking a physical experiment. This means that we jointly evolve a set of atoms meant to form a reactant molecule, given appropriate positioning in space and initial momenta. We discuss modeling and treatment of a bath that allows exchange of energy with an environment further below.

Consider one of the reactant molecules to be prepared. We combine the constituent atoms by scattering in a tree-like manner, see Fig. 2. This is done successively up until the desired molecular state is attained. If each one of the molecules is, in the worst-case, generated by combining only two participants at a time, the overall complexity is quasi-linear in the number of input atoms . It is essential to make sure that at each level in this procedure, we have high overlap with the desired intermediate state. Otherwise, the overlap would decrease at each combination step – if we have initial overlap of $(1-\delta)$, then after n, steps, the overlap would drop exponentially to $(1-\delta)^n$. An obvious choice for avoiding this

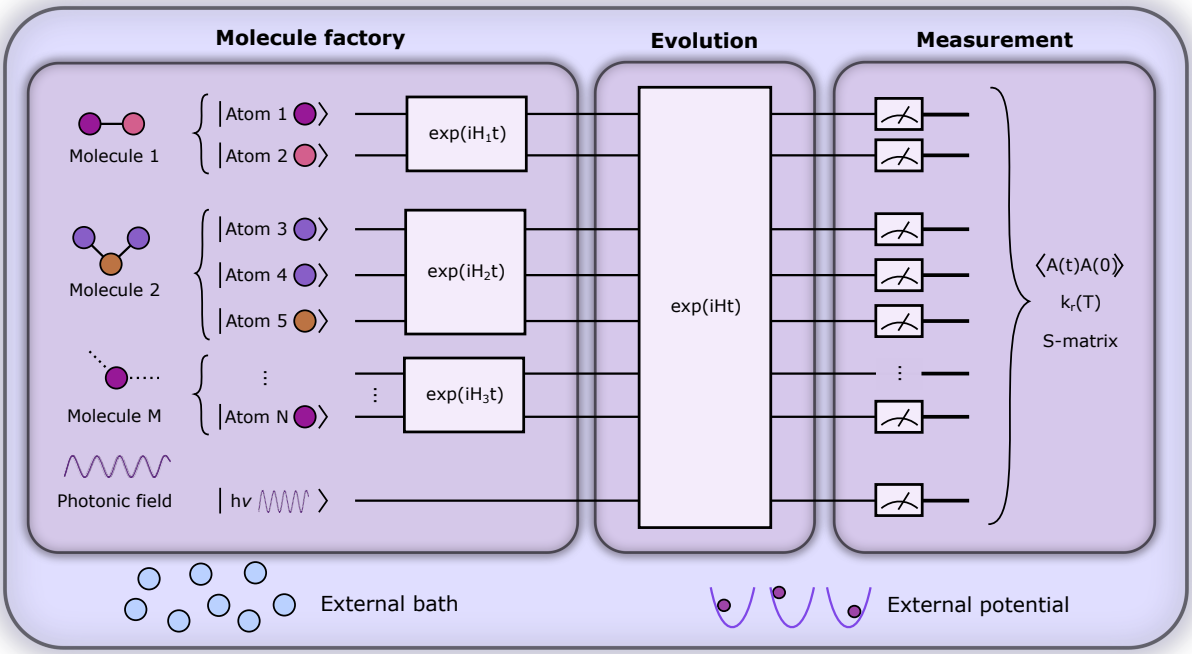

FIG. 2. Computational framework. Our simulation framework can be separated into a state preparation procedure ("molecule factory"), the evolution of the reaction of interest and a measurement step that extracts usable information. The molecule factory prepares a set of molecular input states for the reaction, which, e.g., resemble thermal or ground states. These states are produced in a tree-like fashion together with a weak measurement scheme to ensure that the target states are prepared with sufficient probability. A photonic field serves as a source of energy for reactions and an external bath, either explicit or implicit, serves as an energy sink. Furthermore, we utilize external potentials in the spirit of optical tweezers tailored to any Hamiltonians along the procedure to ensure sufficient success probability and as a means to control positions in space.

problem would be to use oblivious amplitude amplification [45]. However, in our case, the open-system character of the simulation and the fact that both the state we want to reflect about and its probability depend on the (unknown) input state are major roadblocks for this approach, which we therefore leave up to further research. Instead, we propose the following approach to ensure bounded success probability. Following the nanoreactor approach in [46], we introduce artificial potentials which confine the products to be combined in each scattering step, as well as an additional photonic field as energy source, and a bath as energy sink. This allows to chemically enable the formation of reactants so that there is a constant lower bound P on the probability of success. Envisioning the scattering as an at least binary tree, for each node in the tree, there is a simulation channel $\mathcal{E}(\rho)$, parametrized by subsystem size, dissipation model (bath), and confinement properties (artificial potentials). Then we may assume that this channel produces a state of the kind

$$\mathcal{E}(\rho_0) = p_0\rho_{\text{suc}} + (1 - p_0)\rho_{\neg\text{suc}} + C, \tag{1}$$

given input ρ_0 and probability of success p_0; C describes any potential coherence between the subspace of successful and unsuccessful scatterings. The following procedure is also displayed in Fig. 3. We may similarly assume we have an observable O_{suc} , which allows us to distinguish between the subspaces spanned by ρ_{suc} and $\rho_{\neg suc}$ by weak measurements. This serves the purpose of heralding the success of the scattering and to project the state into the successful subspace. The presence of a coherence C does not influence measurement outcomes of $\text{tr}[O_{\text{suc}}\mathcal{E}(\rho_0)]$. One thing to note is that said measurement of O_{suc}, even if carried out in a weak fashion, could disturb the state in a sense that it does not describe a realistic state encountered in nature anymore after the measurement. However, we can assume

that our simulation channel $\mathcal{E}(\cdot)$ does resemble a process in nature. Thus, it tends to make states decay towards such physically meaningful states, and another application of the channel will therefore map us back into a state resembling those in nature,

$$\mathcal{E}(\rho_{\text{suc}}) = p_{\text{suc}}\rho_{\text{suc}}^{\text{nat}} + p_{\neg\text{suc}}\rho_{\neg\text{suc}}^{\text{nat}} + C^{\text{nat}}. \tag{2}$$

If the molecule is unstable but the reacted state is desired, we may simply skip this reapplication and proceed. Hence, for the case when measuring O_{suc} yields a positive outcome, generation of meaningful states with high success rate seems plausible. Further, we can use the following strategy to ensure success even if O_{suc} shows that scattering was unsuccessful. Although we are projected into the unsuccessful subspace, the failure is heralded. This means we can apply another simulation channel $\mathcal{E}'(\cdot)$, which may be slightly modified from $\mathcal{E}(\cdot)$ (e.g., stronger confinement or dissipation), leading to

$$\mathcal{E}(\rho_{\neg\text{suc}}) = p_1\rho'_{\text{suc}} + (1-p_1)\rho'_{\neg\text{suc}} + C'. \tag{3}$$

This is a similar assumption as the case in [43], in the sense that we expect a redistribution of probability towards the state of interest. We can then redo the measurement and iterate between measurements of O_{suc} and (progressively more) modified simulation channels until success occurs. To summarize, we can use the above strategy to create a tree-like sequence of scatterings, so that at each step, the success probability at each of the nodes in the tree is bounded by say $1 \geq P > 0$. As we can herald success, at most $O(1/P)$ repetitions are required per node. Since P is promised to be fixed, the number of repetitions per node is therefore $O(1)$. Furthermore, since failure does not require a restart from the leafs of the tree, the complexities of individual nodes straightforwardly add. Hence, the number of repetitions grows linearly with the number of nodes in the tree, which with N initial constituents goes as at most $O(N\log(N))$.

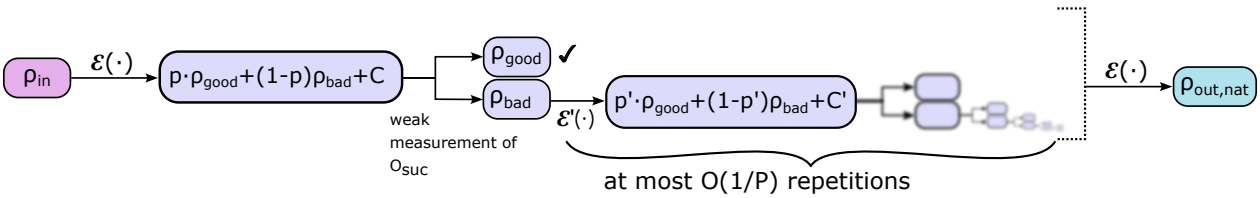

FIG. 3. Procedure of a single scattering step. Evolution through the simulation channel $\mathcal{E}(\cdot)$ produces states with overlap in successful and not successful spaces. In the case of success, proceed, and potentially apply another step of time evolution to make sure the state represents a natural state. If not successful, keep repeating additional, possibly modified, time evolutions which again produce overlap in the successful subspace. Repeat this until success, the expected number of repetitions scaling inversely in the lower bound P on the success probability per step. Success can be quantified by a weak measurment of an observable O_{suc} that, e.g., signifies the success of forming a bond by capturing spatial proximity.

We motivate parallels to optical tweezers [47, 48] or molecular beams [49, 50] that are used in physical experiments in order to engineer the aforementioned artificial potentials to boost the success probability of scatterings. We could use an abundance of particles, requiring an abundance of available qubits, to use well-established success rates of molecular beams. On the other hand, if we trap particles with an artificial potential – e.g., adding to the Hamiltonian a harmonic potential term confined to a specific region – we can considerably boost success without high additional cost in necessary space. Many of the applications we will discuss benefit from reactant states in the form of nuclear wave packets. To that end, consider preparation of the input states for the scattering process to be of this form. Then, said process supported by artificial potentials would in most physical cases not significantly harm the locality in configuration space and preserve its wave packet nature. To induce the interaction of two reactants, we conceptually follow the experimental realization in [40]. Photons are then brought into play to excite the reactants, so that the reaction can occur [51]. A bath can then absorb excess energy once the reaction has occurred, allowing the products to relax.

One can also think of avoiding the scattering approach by using a molecular Hartree-Fock (HF) state as the initial state, with subsequent mapping of the single-particle electronic wavefunctions to plane waves by a Fourier transform, and the nuclei can be initialized as wave packets in nuclear configuration space. Even though it is known that preparing the HF state is not efficient in general [35], practical implementations throughout the Roothaan-Hall equations usually

scale polynomially with system size (e.g. if the number of iterations is bounded by a constant, which can be chosen to be arbitrarily large) [52]. Then, we can expect that a Hartree-Fock state as input to a simulation channel $\mathcal{E}(\cdot)$ would be close enough to the manifold of physically meaningful states so that convergence to a desired state as input to a reaction happens in controllable time. One advantage of the Hartree-Fock approach will likely be that this has considerably lower constant factors than the scattering approach. Among the advantages of the scattering approach are that it should work for general states (including those where Hartree-Fock does not provide a sufficiently good approximation), and that it provides initial states that are closer to nature.

As mentioned above, the embedding of processes into a larger environment plays an important role in the framework during the molecular preparation stage. The ability to dissipate excess energy is essential for both the success probability of forming stable bound states and for the ability of the dynamics to emulate the open-system evolution of chemical experiments. The framework, therefore, employs recently proposed methods for simulating the weak coupling to a large, memory-less (i.e., Markovian) environment [44]. This simulation is efficient and shown to converge to a thermal state of the system. The only parameter of this procedure that does not scale polynomially with system size is the thermalization time, which is difficult to predict and can, in principle, grow exponentially with system size. However, our observation is that slowly thermalizing systems in our simulation corresponds to systems that also thermalize slowly in nature. Thus, for physically motivated open-system models, we would expect to produce either thermalized, metastable, or slowly thermalizing states, depending on which ones are prevalent in nature. A good example for these types of systems would be a glassy molecular mixture. From this perspective, we conjecture that polynomial-time simulations are sufficient to reach all chemically relevant states.

One thing to note is that the system on which readout is to be performed may need to include some degrees of freedom surrounding the molecule, e.g., if solvent-effects, photon or phonon couplings, or non-Markovian dynamics are important. The preparation of this more explicit bath follows the same framework as the main molecular degrees of freedom.

IV. MEASURING DYNAMICAL QUANTITIES OF INTEREST AND EXEMPLARY APPLICATIONS

Simulating the dynamics of molecular processes provides access to meaningful information about the rearrangements of atomic nuclei and electrons, and their interactions with electromagnetic fields, as they unfold. Already with classical resources, time-dependent approaches offer numerous advantages over time-independent ones, as the former are more amenable to handling the continuum portion of the spectrum, and they grant access to the relevant elements of the scattering matrix (S-matrix) over a meaningful range of energies. We provide an overview of chemical problems, with time-dependent versus stationary quantities, solved on different hardware, in Fig. 4.

Measuring observables in a dynamical picture requires considering the time evolution of a wave packet, i.e., a superposition of solutions of the time-dependent Schrödinger equation (TDSE). Most measurements of dynamical quantities can be phrased in terms of wave packet correlation functions, whose calculation fits perfectly into our framework. Transition amplitudes are measured, e.g., using a Hadamard test [37, 41, 53]. We can follow the scheme introduced in [54] capable of obtaining two-point correlation functions, or the cumulative correlation function method in [55]. This scheme is extendable to the S-matrix, and to n-point correlators by additional time-evolutions and block-encodings, similarly to the linear response framework in [53]. If results at more than two times are required, we can use a history state encoding following the conditional time evolutions of the quantum circuit of quantum phase estimation, similar to the construction in [56] who also consider observables measured across various time-steps. To that end, we add a clock register $\sum_{\ell=0}^{n_t} |\ell\Delta t\rangle$ so that $n_t\Delta t = t$. Then, instead of a direct time evolution of the overall system after the molecule factory in Fig. 2 in the state $\rho_1^{\text{mol}}\rho_2^{\text{mol}}\cdots\rho_M^{\text{mol}}$, we split the evolution into chunks of Δt and condition on the clock register to produce a superposition of the state at different evolution time steps. Furthermore, if needed, the application of a Quantum Fourier Transform before measurement is straightforward and does not compromise the overall efficiency of the algorithm. While this framework could be interpreted as purely theoretical, we can easily show that this approach covers a vast range of context extending from nature to chemical laboratories.

Reaction rates

Chemists are commonly interested in observed kinetics, i.e., reaction rates, a topic inherently suitable for simulation of quantum dynamics For instance, dynamical simulation for nonadiabatic processes such as charge transfer reactions fall into the class of quantum circuits of polynomial complexity [57]. Our framework allows use of the measurement schemes proposed in [37] in which the reaction rate can be computed from the degree of localization of the wave packet in nuclear configuration space corresponding to the formed products. In addition to making simulation itself achievable, the quantum approach reduces the measurement to a binary search as opposed to the more complicated evaluation of a time correlation function [58]. Nonetheless, if needed, the latter could, in principle, be computed from the dynamics so that the rate constant is calculated as a function of the scattering cross-sections.

Photochemistry and photophysics

Photochemical processes are triggered when a molecule absorbs a photon. In these transformations, electronically excited states become populated, giving access to reactive channels that are thermally unachievable [59]. The interaction between a molecule and an external photonic field alters the potential energy surface on which wave packets propagate. Therefore, the nuclear and electronic degrees of freedom cannot be decoupled and the Born-Oppenheimer approximation breaks down [60]. The dynamics simulation of these systems on classical computers becomes prohibitive after including very few degrees of freedom. Our framework can seamlessly include photonic fields in the "molecular factory" and the subsequent reaction simulation from Fig. 2 and is amenable to coupling the system with external potentials. Other interesting molecular photophysical processes are also within the scope of our framework. For instance, spectroscopic measurements rarely go to higher orders beyond second order and, instead, ultrafast spectroscopy is preferred, where the system is tracked during a transformation initiated by light. Light-harvesting complexes, photovoltaics, and molecular machines are examples of photophysical applications that our framework can simulate. Our approach would make it possible to simulate quantum systems with quantum computers, e.g., like optoelectronic devices, where the quantum dynamics are far more important than the eigenstates of the Hamiltonian. In these devices, the relaxation pathways of the excitons are exploited for light generation and harvesting. Classical simulations suffer from the large space of excitons and from phonon coupling, making current simulations beyond hopeless. Conversely, simulating the time evolution of thermal quantum states inherently captures all the necessary behavior.

Linear and non-linear molecular spectroscopy

Our proposed framework enables computations connected with the in-laboratory characterization of molecular structures and properties. The absorption spectrum of a molecule is given by the Fourier transform of the wave packet autocorrelation function [61]. By judiciously selecting an initial state, measuring the autocorrelation function can give us different spectra, including electronic, vibrational, and rotational spectra [62]. Through computation of emission and absorption spectra, our approach accommodates the simulation of fluorescent systems, as those used in biomolecule marking, or thermally activated delayed fluorescence, where forbidden relaxations are assisted by thermal coupling to the environment. Additionally, n-time correlation functions allow the exploration of linear-response molecular spectroscopy beyond UV-Vis and fluorescence, such as rotational or vibrational spectroscopy, in which contributions from excited states are typically significant at room temperature. Simulating spectroscopic measurements can be used not only to reproduce experiments but also to probe the simulated quantum system, e.g., the presence of an IR signal may indicate the formation of a molecule, as in the molecular factory in Fig. 2. Two-dimensional spectroscopy could enormously benefit from extracting n-time correlation functions from dynamics through following our approach too. For instance, two-dimensional infrared (2D IR) spectroscopy reveals second-order vibrational couplings, which characterize molecular interactions [63, 64]. Classical simulations of this state-of-the-art technique typically exclude anharmonicity and produce errors associated with the BO approximation and vibrational population transfer. In this example and many others, finding the ground state of the system is far removed from calculating the experimental spectrum.

Free Energy Simulations

Free energies play a role in many naturally occurring physical processes. The free energy determines whether a process occurs spontaneously in a closed system, such as whether a ligand binds to a protein, whether a material such as salt dissolves in water, or into what shape a protein folds. The type of free energy relevant to a specific system depends on the nature of the system. In the case of an isothermal, isobaric system that only allows volumetric work, the relevant free energy is the Gibbs free energy. In applications, one is mostly interested in free energy differences between two different states of the system, which can be characterized by two Hamiltonians, H_1 and H_2. There exist multiple methods on how one can calculate the free energy difference between these two states. One class of methods uses fluctuation relations, such as the Jarzinsky equation, that require ensembles of dynamical simulation systems of interest [65]. To estimate the free energy difference, these methods require evaluating the total energy of each trajectory for both the initial and final states. A time-dependent Hamiltonian $H(\lambda(t))$ is used to alter the system from state 1 to state 2, where λ is an externally controlled parameter such that $H(0) = H_1$ and $H(1) = H_2$ (the transition is generally non-adiabatic, and states do not have to be or remain in the ground state such as in adiabatic quantum computing). The number of energy values required depends on the type of problem considered, the speed of how quickly the Hamiltonian is transformed, and the desired accuracy. Within our computational framework, there are two straightforward ways how one can obtain the required energies. First, one can create a sufficiently large number of identical systems and perform separate quantum simulations of them before measuring the total energies of the initial and final states. Alternatively, one can use the clock-register approach from above and simulate a single trajectory. To obtain a sufficiently large number of statistically independent energy values of both initial and final states, one can simulate each state long enough so that measurements can be taken with sufficiently large time intervals in between. Once the energies are obtained, the averages and final calculation of the free energy difference can be carried out on a classical computer.

Quantum Machine Learning

Instead of directly measuring interesting quantities, it can be envisioned that the output of the dynamical quantum simulation is processed using machine learning. This can be understood as machine learning with quantum input data, with possible classical or quantum outputs [66, 67]. This is particularly promising in light of recent results which indicate that there is a provable advantage in the efficiency of extracting information when given access to multiple copies of a state in a format that can be manipulated by a quantum computer compared to having access to only measurements performed on the state [68]. Thus, certain chemically relevant properties of states and quantum evolutions may be learnable more efficiently in our quantum computing framework than in a conventional experiment. Alternatively, it is also possible to measure spectroscopic quantities and replace physical experiments in machine learning pipelines that operate on spectroscopic results, potentially aiding further developments in molecular design [69, 70].

V. CONCLUSION AND OUTLOOK

We have provided an algorithmic framework that, in principle, can solve a wide set of chemically relevant problems using inherently efficient building blocks. To that end, we considered that while general ground states are hard to obtain, we may assume to be able to prepare atomic ones as a single-cost effort that can enable a building block library. A scattering process, implemented by simulating dynamics and boosted by artificial potentials, is able to produce a molecular input state for a subsequent dynamical simulation, which then is followed by the measurement of dynamical quantities. We provide examples of applications from, e.g., spectroscopy, photochemistry and beyond.

It remains for future work now to build upon this approach and provide a more concrete algorithm, e.g., to gauge feasibility of the procedure with respect to constant factors and more precise scaling coming from choosing specific problem instances and methodologies. For instance, the detailed computational cost of using artificial potentials to boost success probabilities of the scattering process falls into this category. Another interesting extension would be the inclusion of other modeling tools used in classical simulation, such as the Nosé-Hoover thermostat [71]. Much ado has been made in quantum chemistry about the finding of ground state energies, a problem known to be QMA-complete for

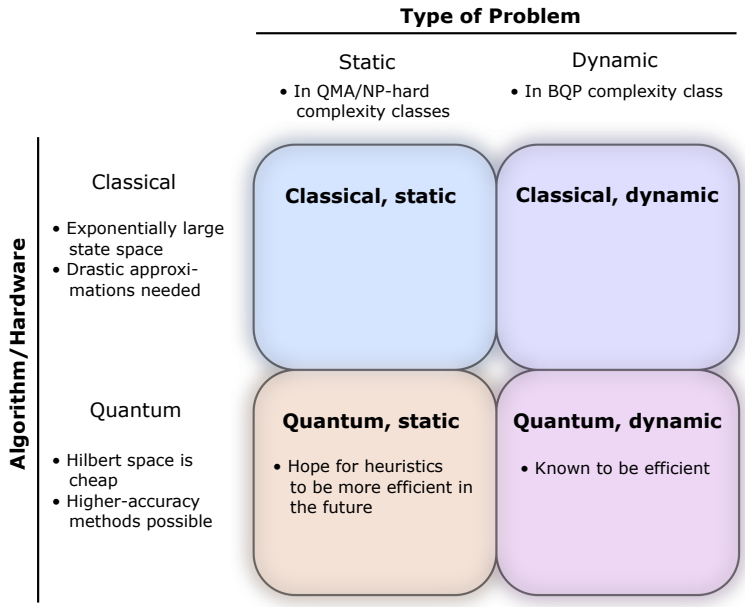

FIG. 4. Classifying chemical problems related to their hardness and space complexity. Dynamical properties are quantumly efficient whereas static properties are generally hard. As quantum computers do not suffer from the curse of dimensionality, one can expect the sweet spot of quantum simulations, up to constant factors in the cost, to lie in the evaluation of dynamic properties.

local Hamiltonians like those seen in molecular systems. We should cease to hold ourselves as captives of this artificial perspective which focuses on a problem known to be hard, and allow ourselves to seek dynamics directly. Additionally, for situations when the ground state is of interest, it may prove useful to give up the search for exact solutions to a hard problem and look at heuristics. The Hartree-Fock problem is known to be NP-hard yet practically efficient thanks to approximations [35]. Attempts based on open systems dynamics such as in [43] may be a promising path towards quantum heuristics for ground states. Nevertheless, we call upon what dynamical quantum simulation has to offer for chemistry.

ACKNOWLEDGEMENTS

The authors thank Kouhei Nakaji, Mohammad Ghazi Vakili, FuTe Wong and Sumner Alperin-Lea for helpful discussions and comments.

AA gratefully acknowledges King Abdullah University of Science and Technology (KAUST) for the KAUST Ibn Rushd Postdoctoral Fellowship. LBK acknowledges support from the Carlsberg Foundation. This research was developed with funding from the Defence Advanced Research Projects Agency (DARPA) and the Department of the Navy, Office of Naval Research. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. AAG acknowledges support from the NSERC Industrial Research Chairs Program and the Canada 150 Research Chairs.

REFERENCES

[1] P. Benioff, Journal of Statistical Physics **22**, 563 (1980).
[2] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).
[3] Y. Manin, Sovetskoye Radio, Moscow **128**, 28 (1980).

[4] C. Zalka, Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **454**, 313 (1998).
[5] S. Wiesner, arXiv preprint quant-ph/9603028 (1996).
[6] D. S. Abrams and S. Lloyd, Physical Review Letters **79**, 2586 (1997).
[7] D. Aharonov and A. Ta-Shma, in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, STOC '03 (Association for Computing Machinery, New York, NY, USA, 2003) pp. 20–29.
[8] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Communications in Mathematical Physics **270**, 359 (2007).
[9] S. Lloyd, Science **273**, 1073 (1996).
[10] D. S. Abrams and S. Lloyd, Physical Review Letters **83**, 5162 (1999).
[11] J. Kempe, A. Kitaev, and O. Regev, in *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science*, Lecture Notes in Computer Science, edited by K. Lodaya and M. Mahajan (Springer, Berlin, Heidelberg, 2005) pp. 372–383.
[12] A. Y. Kitaev, arXiv preprint quant-ph/9511026 (1995).
[13] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309**, 1704 (2005).
[14] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Molecular Physics **109**, 735 (2011).
[15] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Proceedings of the National Academy of Sciences **114**, 7555 (2017).
[16] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Physical Review X **11**, 011020 (2021).
[17] G. H. Low and I. L. Chuang, Quantum **3**, 163 (2019).
[18] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Quantum **3**, 208 (2019).
[19] D. Motlagh and N. Wiebe, arXiv preprint arXiv:2308.01501 (2023).
[20] E. Campbell, Physical Review Letters **123**, 070503 (2019).
[21] K. Nakaji, M. Bagherimehrab, and A. Aspuru-Guzik, arXiv preprint arXiv:2302.14811 (2023).
[22] M. Pocrnic, M. Hagan, J. Carrasquilla, D. Segal, and N. Wiebe, Physical Review Research **6**, 013224 (2024).
[23] A. Rajput, A. Roggero, and N. Wiebe, Quantum **6**, 780 (2022).
[24] M. Hagan and N. Wiebe, Quantum **7**, 1181 (2023).
[25] Y. Su, D. W. Berry, N. Wiebe, N. Rubin, and R. Babbush, PRX Quantum **2**, 040332 (2021).
[26] P.-O. Löwdin, in *Advances in Chemical Physics* (John Wiley & Sons, Ltd, 1958) pp. 207–322.
[27] P.-O. Löwdin, International Journal of Quantum Chemistry **55**, 77 (1995), _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.560550203.
[28] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell, E. F. Valeev, L. Lin, and G. K.-L. Chan, Nature Communications **14**, 1952 (2023).
[29] A. Ambainis, in *2014 IEEE 29th Conference on Computational Complexity (CCC)* (2014) pp. 32–43, iSSN: 1093-0159.
[30] J. D. Watson and J. Bausch, arXiv preprint arXiv:2105.13350 (2021).
[31] J. Bausch, T. S. Cubitt, and J. D. Watson, Nature Communications **12**, 452 (2021).
[32] C.-F. Chen, H.-Y. Huang, J. Preskill, and L. Zhou, arXiv preprint arXiv:2309.16596 (2023).
[33] R. P. Feynman, Optics News **11**, 11 (1985).
[34] N. Schuch and F. Verstraete, Nature Physics **5**, 732 (2009), number: 10 Publisher: Nature Publishing Group.
[35] J. D. Whitfield, P. J. Love, and A. Aspuru-Guzik, Physical Chemistry Chemical Physics **15**, 397 (2012).
[36] B. O'Gorman, S. Irani, J. Whitfield, and B. Fefferman, PRX Quantum **3**, 020322 (2022).
[37] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, Proceedings of the National Academy of Sciences **105**, 18681 (2008).
[38] H. H. S. Chan, R. Meister, T. Jones, D. P. Tew, and S. C. Benjamin, Science Advances **9**, eabo7484 (2023).
[39] S. Simon, R. Santagati, M. Degroote, N. Moll, M. Streif, and N. Wiebe, arXiv preprint arXiv:2307.13033 (2023).
[40] L. R. Liu, J. D. Hood, Y. Yu, J. T. Zhang, N. R. Hutzler, T. Rosenband, and K.-K. Ni, Science **360**, 900 (2018).
[41] J. S. Pedernales, R. Di Candia, I. L. Egusquiza, J. Casanova, and E. Solano, Physical Review Letters **113**, 020505 (2014).
[42] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309**, 1704 (2005).
[43] Z. Ding, C.-F. Chen, and L. Lin, arXiv preprint arXiv:2308.15676 (2023).
[44] M. J. Kastoryano, F. G. Brandão, A. Gilyén, and others, arXiv preprint arXiv:2303.18224 (2023).
[45] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing - STOC '14* (ACM Press, New York, New York, USA, 2014) pp. 283–292.
[46] L.-P. Wang, A. Titov, R. McGibbon, F. Liu, V. S. Pande, and T. J. Martínez, Nature Chemistry **6**, 1044 (2014).
[47] P. A. M. Neto and H. M. Nussenzveig, Europhysics Letters **50**, 702 (2000).
[48] J. E. Molloy and M. J. Padgett, Contemporary Physics **43**, 241 (2002).
[49] P. Casavecchia, Reports on Progress in Physics **63**, 355 (2000).
[50] S. Y. T. van de Meerakker, H. L. Bethlem, and G. Meijer, Nature Physics **4**, 595 (2008), number: 8 Publisher: Nature Publishing Group.
[51] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Photons and Atoms: Introduction to Quantum Electrodynamics*, 1st ed.

(Wiley, 1997).
[52] S. Gulania and J. D. Whitfield, The Journal of Chemical Physics **154**, 044112 (2021).
[53] E. Kökcü, H. A. Labib, J. K. Freericks, and A. F. Kemper, arXiv preprint arXiv:2302.10219 (2023).
[54] F. Tacchino, A. Chiesa, S. Carretta, and D. Gerace, Advanced Quantum Technologies **3**, 1900052 (2020).
[55] L. Lin and Y. Tong, PRX Quantum **3**, 010318 (2022).
[56] N. Diaz, P. Braccia, M. Larocca, J. Matera, R. Rossignoli, and M. Cerezo, arXiv preprint arXiv:2308.12944 (2023).
[57] P. J. Ollitrault, G. Mazzola, and I. Tavernelli, Physical Review Letters **125**, 260511 (2020).
[58] H. Liu, G. H. Low, D. S. Steiger, T. Häner, M. Reiher, and M. Troyer, Materials Theory **6**, 11 (2022).
[59] R. Bhuyan, J. Mony, O. Kotov, G. W. Castellanos, J. Gómez Rivas, T. O. Shegai, and K. Börjesson, Chemical Reviews **123**, 10877 (2023).
[60] B. F. E. Curchod and T. J. Martínez, Chemical Reviews **118**, 3305 (2018).
[61] D. J. Tannor, *Introduction to Quantum Mechanics* (University Science Books, 2007) google-Books-ID: t7m08j3Wi9YC.
[62] V. Vitale, J. Dziedzic, S. M.-M. Dubois, H. Fangohr, and C.-K. Skylaris, Journal of Chemical Theory and Computation **11**, 3321 (2015).
[63] N. T. Hunt, Chemical Society Reviews **38**, 1837 (2009).
[64] I. Noda, Applied Spectroscopy **44**, 550 (1990).
[65] L. Bassman Oftelie, K. Klymko, D. Liu, N. M. Tubman, and W. A. de Jong, Physical Review Letters **129**, 130603 (2022).
[66] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Nature **549**, 195 (2017).
[67] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Nature Computational Science **2**, 567 (2022).
[68] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Science **376**, 1182 (2022).
[69] A. Angulo, L. Yang, E. S. Aydil, and M. A. Modestino, Digital Discovery **1**, 35 (2022).
[70] J. F. Joung, M. Han, J. Hwang, M. Jeong, D. H. Choi, and S. Park, JACS Au **1**, 427 (2021).
[71] W. G. Hoover, *Computational Statistical Mechanics*, 1st ed. (Elsevier, 1991).