In situ upgrade of quantum simulators to universal computers

Benjamin Dive¹, Alexander Pitchford², Florian Mintert¹, and Daniel Burgarth²

¹Department of Physics, Imperial College, SW7 2AZ London, UK

²Institute of Mathematics, Physics and Computer Science, Aberystwyth University, SY23 3FL Aberystwyth, UK
2018-07-23

Quantum simulators, machines that can replicate the dynamics of quantum systems, are being built as useful devices and are seen as a stepping stone to universal quantum computers. A key difference between the two is that computers have the ability to perform the logic gates that make up algorithms. We propose a method for learning how to construct these gates efficiently by using the simulator to perform optimal control on itself. This bypasses two major problems of purely classical approaches to the control problem: the need to have an accurate model of the system, and a classical computer more powerful than the quantum one to carry out the required simulations. Strong evidence that the scheme scales polynomially in the number of qubits, for systems of up to 9 qubits with Ising interactions, is presented from numerical simulations carried out in different topologies. This suggests that this *in situ* approach is a practical way of upgrading quantum simulators to computers.

Recent and ongoing work on building large quantum systems is leading to simulators that are able to model physical phenomena, allowing questions about the underlying science to be answered [1–3]. These machines contain a register of quantum particles, typically two level quantum systems (qubits) storing quantum information. The presence of interactions between these leads to dynamics that, by varying control parameters in the system Hamiltonian, can replicate the quantum behaviour of systems of interest. This, however, is less general than a quantum computer which is able to perform a universal set of logic gates on the qubits [4].

Provided some control parameters can be varied in time, it is in principle possible to do an

Figure 1: A classical computer finds a control pulse which enables a quantum simulator to perform logic gates. It does this in an iterative process by applying a control pulse to the simulator and then improving it based on the result of measurements.

arbitrary gate on a quantum many-body system such as a quantum simulator [5–7]. Finding the right time-dependency however relies almost exclusively on numerical methods, especially when physical constraints on the control fields are taken into account [8, 9]. These methods require a very precise knowledge of the parameters of a system, a daunting task for a machine with a huge number of degrees of freedom. Furthermore, they are intractable on a classical computer if the quantum simulator we want to solve the problem for is large enough to do something beyond the capabilities of classical computers. These two difficulties provide a major obstacle in using quantum simulators to perform arbitrary computation.

We circumvent these problems by showing how well-known existing numerical methods can be translated to run *in situ* on the quantum simulator itself, as illustrated in Fig.1. A mix of an-

alytical and numerical results point towards this being a scalable, bootstrapping scheme for performing a universal quantum computation, needing resources that grow only polynomially with the number of qubits.

The core principle is to use the system at hand to implement and improve control pulses until they perform the desired gate on it. Such an adaptive approach to finding controls is naturally used in laboratory work and has been studied as a fine-tuning tool in small systems [10, 11], a way of controlling quantum chemistry with light [12], a method for stochastic optimisation [13], and a way of correcting parameter drift [14]. Using simulators as oracles for reaching quantum states has also been explored [15], as has the potential for a quantum speed up in general optimisation problems [16]. In this paper we propose a way to transform this general approach into a large scale, constructive method that provides a new avenue to control many-body quantum systems. Independently, and concurrently to the writing of this paper, a similar approach as ours was developed and tested experimentally for the different problem of quantum state preparation [17].

The scheme

The model we consider is a quantum simulator that has the underlying ability to be a universal quantum computer (one that can run an arbitrary algorithm), and the task is to learn how to use it as such. For this reason, we take the simulator as consisting of n qubits with some interactions between them such that they form a fully connected graph where every qubit is (directly or indirectly) interacting with every other qubit. Furthermore we require that the timescale associated with this interaction is much shorter than the decoherence time in order for significant entanglement to be built up. In addition to this we need the ability to perform the following operations on each qubit individually: preparation in a complete basis set of states, fast rotations by applying strong Hamiltonians, and measurement in a complete basis set.

Such a system can be described by the Hamiltonian

$$H = \sum_i \left(f_x^i(t) \sigma_x^i + f_y^i(t) \sigma_y^i \right) + \sum_{\langle i,j \rangle} H^{i,j}, \quad (1)$$

where the first sum is over all qubits and the

time-dependent control functions $f(t)$ are to be determined. The second sum is over all connected qubits and $H^{i,j}$ is the interaction between qubit i and j . The choice of σ_x and σ_y for the controls is for convenience, any two Hamiltonians will work, and does not have to be the same for all qubits. As the controls are found *in situ*, it is not necessary to know beforehand the form of the interactions $H^{i,j}$.

These requirements are significant, but much easier than demanding direct control over two qubit operations, and correspond to the state-of-the-art in systems involving trapped ions [18–20], cold atoms [21, 22], NMR [23–25] or superconducting circuits [26, 27]. In these systems there already exist quantum simulators powerful enough to do simulations, and satisfy our requirements, but are not currently usable as computers as it is not known how to perform logic gates on them [2, 4]. Our numerical results show that the scheme developed here scales well for a range of systems where $H^{i,j}$ is of the Ising type ($\sigma_z \otimes \sigma_z$). As Ising machines are useful for a wide range of quantum simulations and can be built with many different technologies [19, 22, 28], this is a result with wide ranging applicability.

In the model we consider, the connectedness of the qubits and the ability to do fast single qubit operations guarantees that the two core requirements of our proposed optimisation scheme are satisfied: there exists a universal gate set that can be reached at short times [6], and process tomography can be performed [29]. While other systems satisfy these requirements and the approach detailed here would work, we focus on this model for clarity. As single qubit operations are assumed, the gates that controls are needed for are entangling ones, canonically the controlled-not (C-NOT) gate; these are vastly harder to perform using conventional methods and typically have much lower fidelities.

The steps for finding such a gate in the *in situ* scheme are outlined in Fig.2. These are very general and almost the same as in classical numerical optimisation: a guess for the optimal values is generated, these are fed into a function that computes their effect, the distance between this and the desired outcome is calculated, and this generates another guess for the optimal values. This is iterated until the values get close enough to the goal, or the process terminates unsuccessfully af-

Figure 2: Outline of the process used in optimal control, the two red steps in the middle are done *in situ* in our scheme while the others are classical processing. The starting point is an initial set of controls that parametrise the strength of the control Hamiltonians over the gate duration, in our examples these are generated randomly. The evolution of the system with these parameters is then calculated. On a classical computer this requires solving the time-dependent Schrödinger equation numerically for a model of the system, while in our scheme this is simply implementing the controls on the simulator. Evaluating the gate fidelity in the classical case is straightforward but, when done *in situ*, requires some form of tomography to measure it. We derived a tight bound for this gate fidelity in Eq.(2) that can be measured efficiently. If this fidelity is above a threshold, the process terminates successfully, otherwise the control parameters are updated based on the results of the latest and previous runs, and the process repeats. Such an approach can be used in a wide range of contexts, such as to perform quantum logic using random walks [30].

ter some timeout condition is reached. The difficulty with doing this for a quantum simulator of the type discussed above is in computing what unitary is produced by a given choice of control parameters; this requires both an accurate model of a high dimensional system and an exponentially large classical computer to solve it. Neither of those things can be done for a quantum system large enough to be an interesting quantum computer.

We eliminate these twin difficulties by using the quantum simulator to compute the effects of the control pulse on itself. This works because the simulator with a trial set of controls is guaranteed to be an accurate model of itself with those controls. The propagation step is therefore done *in situ*, but the method by which the control parameters are updated remains purely classical. This is because the information extracted from the quantum simulator (the gate fidelity) and the parametrisation of the control pulses are purely classical. An upshot of this is that the myriad of different methods to do numerical optimisation that have already been developed and work for quantum systems can be used in this protocol

directly.

In order to use the quantum simulator as a universal computer, this optimisation procedure needs to be repeated for a universal set of gates. As single qubit operations are assumed, it is sufficient to find a complete set of C-NOT gates. The minimum number of these gates, such that every quantum circuit can be implemented, is $n - 1$. In practice we expect it to be more efficient, and produce shorter circuits, to find the $\frac{1}{2}n(n - 1)$ C-NOT gates that act between every pair of qubits. As is shown in the numerical section, this is readily achieved even for pairs of qubits that are not directly interacting.

The question of whether the scheme works when the system Hamiltonian varies in time uncontrollably is important for an experimental implementation. If this change in time happens slowly compared to the time it takes to find a control pulse for a given gate, then it does not impede the ability of the scheme to find that pulse. However, it may mean that the pulse no longer produces the required dynamics at a later point in time when it is being used as part of an algorithm. In this case, it would be required to eventually rerun the optimisation scheme in order to correct for this drift.

On the other hand, if the stochastic fluctuations in the Hamiltonian are faster than the evolution time required for a gate, the problem is a different one. Repeatedly evolving the system with the same controls (necessary in order to measure the fidelity) will result in the system evolving with a different Hamiltonian on each repeat. This noise in the Hamiltonian thus translates into a lower fidelity being measured. Therefore, as long as the fast fluctuations in the Hamiltonian are small, they are not expected to prevent the scheme from finding a successful pulse but will limit the maximum possible fidelity.

Local gate fidelity

The measure used to gauge how close the system evolution is to the desired unitary is typically the gate fidelity [31]. This is a function between the dynamical map M which describes the evolution of the system under a set of controls (including potential decoherence which acts on the system), and the target unitary U . It is defined as $F(M, U) = \langle \psi | \rho_M | \psi \rangle$ where $|\psi\rangle = U \otimes \mathbb{1} |\Omega\rangle$ (with $|\Omega\rangle = \sum_k |kk\rangle$ being a maxi-

maximally entangled state) is the Choi state of U ; and $\rho_M = (M \otimes id) |\Omega\rangle \langle \Omega|$ is the Choi state of M . This distance measure is bounded between 0 and 1, with the upper limit being reached only when $M(\cdot) = U(\cdot)U^\dagger$. In the case of the system evolution being unitary simplifies down to $F(V, U) = |\frac{1}{d} \text{Tr}[V^\dagger U]|^2$. When the propagation step of Fig.2 is done classically, the whole unitary describing the evolution of the system is calculated as an exponentially large matrix from which the gate fidelity must be calculated.

In the *in situ* scheme this is no longer the case; the only thing which is accessible is a quantum state after it has been evolved by the quantum simulator under some control parameters. The standard method of extracting the gate fidelity is to perform a variant of process tomography, known as certification. This requires preparing the system in a specific state, evolving it, and then performing a set of measurements. The number of different preparation-measurement combinations, N_{meas} , is of order $O(d^2) = O(2^{2n})$, and thus scales exponentially [32].

However it is possible to do exponentially better for cases of interest where the target gate has a tensor product structure, $U = \bigotimes U_i$ where each U_i is a unitary which acts on a small number of qubits. This would typically be a single C-NOT on one pair of qubits and identity on the rest, $\text{C-NOT}_{1,2} \otimes \mathbb{1}_3 \otimes \mathbb{1}_4 \otimes \dots$, but it could be several simultaneous non-overlapping C-NOTs or even larger gates such as Toffoli. No matter what the exact form is, provided that the target can be decomposed into a tensor products of unitaries, the fidelity over the whole system is bounded by the *local estimator* F_{LE} according to

$$F(M, U) \geq F_{LE}(M, U) = 1 - \sum_i (1 - F(M_i, U_i)) \quad (2)$$

where $M_i(\rho_i) = M(\rho_i \bigotimes_{j \neq i} \frac{1}{d_j} \mathbb{1}_j)$. This is the reduced dynamical map acting on subsystem i where the other subsystems have been initialised in the maximally mixed state. This result is proved in the appendix below based on existing approaches [33].

The advantage of this local estimator to the fidelity is that it only requires certification to be performed over small dimensional subsystems of 1 or 2 qubits. As the size of these subsystems does not increase as the system is scaled up, the

cost of measuring the fidelity does not increase exponentially with the number of qubits. Applying existing results for certification to each term in Eq.(2) sequentially gives $N_{\text{meas}} = O(\sum_i d_i^2) = O(n)$. As is discussed in the methods section below, it is possible to remove this linear scaling by noting that each term in Eq.(2) can be recovered in parallel. This results in a constant cost, $N_{\text{meas}} = O(\max_i d_i^2) = O(1)$, a vast improvement over the previous exponential scaling, $O(2^{2n})$.

Beyond being efficiently recoverable, this estimator to the fidelity is useful for a number of reasons. It is a lower bound on the gate fidelity, so we are guaranteed that the true fidelity is at least as good. It converges to the exact fidelity in the limit that $F(M, U) \rightarrow 1$, this is important as we are most interested in having a measure of how good a gate is when it is close to the target. As can be seen in Fig.3, it is well behaved numerically and the initial convergence is fast. Increasing the number of qubits would increase the number of terms in Eq.(2) but not their structure, therefore we expect the qualitative features to remain the same as it is scaled up. The minimum value of the local fidelity is $1 - n$, while the true gate fidelity cannot go below 0, so it may be expected that the convergence to 1 is slower in larger systems as the local fidelity has a larger range to cover.

The scaling behaviour of the local fidelity was investigated by considering a target $U_T = \text{C-NOT} \otimes \mathbb{1} \otimes \mathbb{1} \dots$ and comparing the gate fidelity and local fidelity between it and the unitary $U = e^{-iH} U_T$, where H is a random Hamiltonian generated by a Gaussian distribution and normalised to $\|H\|_2 = 0.1$. As the number of qubits was increased from 3 to 14, the true gate fidelity averaged over different random H stayed the same (99.5%) while the local fidelity dropped linearly, by less than 0.2% per qubit. This is in accordance with our intuition that the local fidelity behaves similarly for different numbers of qubits, with the principle difference being the linearly increasing number of terms in the sum of Eq.(2).

Numerical investigation

The local fidelity detailed in the previous section shows that it is possible to estimate the fidelity of a quantum gate efficiently as the size of the system increases, removing one direct barrier

Figure 3: Comparison of the gate fidelity with the local estimator of the fidelity during an optimisation run. The gate fidelity and its local estimator, Eq.(2), are plotted as a function of iteration step for one complete run of the *in situ* optimisation scheme. The system is a five-qubit nearest-neighbour chain with the Hamiltonian of Eq.(1); Ising on the left where $H^{i,j} = \sigma_z \otimes \sigma_z$, and Heisenberg on the right where $H^{i,j} = \sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y + \sigma_z \otimes \sigma_z$. The target is a C-NOT gate on the first two qubits and identity on the others. The algorithm minimised the infidelity of the local estimator. The exact infidelity is plotted at each step for comparison. It is lower in both cases at all iteration steps, and highly correlated with the estimated infidelity, such that minimising the former also minimises the latter almost monotonically and the landscape remains trap free. Furthermore the difference between the two decreases rapidly as the infidelity approaches 0. In the Heisenberg case the true gate fidelity converges slower than for the Ising chain; this behaviour is closely mapped onto the local estimator. This demonstrates the validity of maximising the local estimator of the fidelity as a proxy for maximising the true gate fidelity.

from the scalability of the *in situ* optimisation scheme. There are, however, other factors that determine the time the protocol takes which need to be taken into account to assess its scalability. This requires an expression for the total time required to construct a control sequence for a gate in terms of the number of qubits in the system. As this is an optimisation problem that would be done ‘numerically’ on a hybrid classical-quantum computer, analytic expressions could not be obtained. In order to investigate this we conducted simulations of the protocol on a purely classical computer. We explored systems from 3 to 9 qubits; memory constraints on the cluster we used made larger systems unfeasible due to the difficulty of evolving (and doing gradient based optimisation of) operators.

The average time needed to find a control sequence for a gate can be expressed as:

$$T_{\text{total}} = T_{\text{run}} N_{\text{runs}} / p_{\text{succ}} \quad (3)$$

where T_{run} is the time it takes to do one run of a control sequence on the quantum simulator, N_{runs} is the number of sequences that are run on the simulator until the protocol halts, and p_{succ} is the probability that the protocol halts with a control pulse that reaches the desired fidelity. T_{run} can be decomposed as $T_{\text{run}} = T_{\text{init}} + T_{\text{gate}} + T_{\text{meas}}$

which is the time to initialise the system, evolve the system under the interaction and control Hamiltonians, and then measure it respectively. T_{init} and T_{meas} are determined by the type of system being used; we take them as fixed and independent of the number of qubits. The gate time, on the other hand, is a free parameter that must be decided before starting the *in situ* optimisation.

The total number of runs can be similarly decomposed as

$$N_{\text{runs}} = N_{\text{meas}} N_{\text{prec}} N_{\text{fids}} N_{\text{upds}}. \quad (4)$$

N_{meas} is the number of times the experiment with the same control pulse, but with different input states and measurement basis, must be repeated in order to measure the gate fidelity once. As the previous section showed, this is $O(1)$ for the local estimator to the fidelity, which is the measure used henceforth. N_{prec} is the number of times the fidelity must be measured to acquire sufficient statistics such that the fidelity is known to the desired precision. N_{fids} is the number of different fidelities that need to be measured for the optimisation algorithm to update the control sequence. It is 1 for gradient-free methods, while for steepest-ascent methods it is 1 plus the number of gradients (when they are measured by fi-

nite difference). N_{upds} is the number of the times the control sequences must be updated, corresponding to the number of times the scheme goes around the loop of Fig.2.

The scaling relation of the terms in Eq.(3) depends on the underlying classical algorithm. We used a steepest ascent gradient method similar to Gradient Ascent Pulse Engineering (GRAPE) algorithm [34] which is commonly used for optimising quantum control on classical computers with great success. In this approach, each of the independent Hamiltonians that can be controlled are taken as piecewise-constant with N_{ts} time-slots of equal widths that span the full gate time T_{gate} . We used analytical gradients in the simulations, for computational efficiency [8], which restricted us to piecewise constant and fixed T_{gate} .

The precision to which the fidelities need to be measured experimentally also needs to be specified. This could be done by either fixing N_{prec} itself, or by repeatedly measuring the fidelity until the error of the mean is below a specified value. We approximated the latter approach numerically by rounding each local fidelity measurement to some numerical accuracy A_{num} and calculated the equivalent N_{prec} . The *in situ* scheme therefore requires T_{gate} , N_{ts} and A_{num} to be chosen beforehand, as well as a target fidelity F_{targ} , and to know the number of different control Hamiltonians which, multiplied by N_{ts} , gives the total number of controls N_{ctrl} .

In order to simulate this completely numerically we also need to specify exactly what the control Hamiltonians and the constant interaction Hamiltonians are for a given system. This is not the case were this scheme done *in situ* experimentally. In an experiment T_{gate} minimisation could be included in optimisation objectives as the gradients would be calculated via a finite-difference method. Alternative pulse parametrisation to piecewise constant could also be used, as best suits the experimental setup. All the different parameters mentioned above are summarised in Fig.4.

We conducted a number simulations of this approach with a Hamiltonian of the type described in Eq.(1) for different number of qubits and interaction topologies. They were completed using the quantum optimal control modules in QuTiP [35–37]. These provide methods for optimising a control pulse to some fidelity measure. The

Parameter	Description
T_{gate}	Evolution time for the gate
N_{ts}	Number of timeslots for control pulse
A_{num}	Accuracy of fidelity measurements
F_{targ}	Target fidelity for the desired gate
N_{runs}	Number of (#) runs in total
N_{meas}	# different input-output pairs
N_{prec}	# repeats for required fidelity accuracy
N_{fids}	# different fidelities to update controls
N_{upds}	# control updates needed
p_{succ}	probability of success
N_{ctrl}	# parameters in control pulse

Figure 4: The different parameters defined in the text, summarised here for convenience. The top four are those which need to be fed into the classical optimiser in order for it to run GRAPE *in situ*; other classical protocols could be used, which would require different parameters. The bottom seven are used to quantify the efficiency of the scheme.

GRAPE implementation in QuTiP is described in the documentation, available at [37]. The code used to perform the numerical simulations is available in an open-source repository [38]. The *local Choi fidelity* measure customisation, and a method for automating locating the p_{succ} threshold, were developed for this study; they are fully described in the code documentation. The optimal T_{gate} and N_{ts} were determined by trialling a range of alternatives. As the result of the optimisation depends on the initial random control amplitudes (uniformly distributed in $[1, 1]$), each scenario was repeated multiple times to gain reliable statistics.

A high performance computing cluster was necessary for completing sufficient repetitions of the optimisation simulations of the larger systems in a reasonable time (the 9 qubit optimisations each required around four days of Intel Xeon CPU E5-2670 0 2.90GHz core processing time and were repeated hundreds of times). This is because the processing time required to optimise a pulse scales exponentially with system size due to the need to exponentiate the Hamiltonians in order to compute propagators. This difficulty precisely highlights the need to optimise pulses *in situ* for quantum systems of the size that would perform a useful quantum computation.

We found numerically that, for a range of examples, there exist values of T_{gate} and N_{ts} such

Topology	Coupling	T_{gate}	N_{ts}	N_{upds}
chain	Ising	π	12	60
star	Ising	π	12	214
fully connected	Ising	12π	160	295
chain	Heisenberg	16π	160	585
star	Heisenberg	12π	160	1043
fully connected	Heisenberg	12π	160	881

Figure 5: The cost of performing the *in situ* optimisation scheme is investigated for a range of different 5 qubit systems for the Hamiltonian of Eq.(1). The differences between the systems are their topology (a linear chain with nearest neighbour interactions, a star where all interact with a central qubit only, or fully connected where the interaction strengths are also randomised) and the interaction type. In each case the Hamiltonian used means that $N_{\text{ctrl}} = 10$, the target operation is a C-NOT gate on two qubits and identity on the rest, $F_{\text{target}} = 0.999$, and $p_{\text{succ}} > 0.98$. These simulations were done with full numerical precision. We see that, for five qubits, all six systems find the desired entangling gate, and do so at reasonable experimental cost. This indicates that the approach works for a range of possible quantum simulators.

that the *in situ* scheme converges. Fig.5 shows typical values of the most important parameters for a variety of topologies and interaction types. We consistently found that Ising systems were easier to find controls for than Heisenberg systems. In particular, our results suggest that in Heisenberg systems a GRAPE-based algorithm may require a T_{gate} that scales exponentially with the number of qubits in order for the optimisation to succeed. We found that this discrepancy also existed in purely classical optimisation techniques.

This suggests that Heisenberg systems are intrinsically harder to solve with optimal control methods than Ising ones, and that this does not depend on whether an *in situ* or classical approach is used. This is consistent with Fig.3 where we compare the local estimator to the fidelity with the true fidelity for a 5 qubit Heisenberg chain as a function of the iteration step as it is being optimised. The local estimator tracks the true fidelity steadily whether the underlying system is Heisenberg or Ising, the notable difference between the two plots is the plateau in the Heisenberg case. This shows that optimising the system is significantly harder and appears for both the local estimator and the true fidelity.

An exponential scaling in the required gate time would make such an approach infeasible for a quantum computer.

Regardless of these numerical difficulties, it can be shown [6] that it is possible to do fast entangling gates on such Heisenberg systems by using fast local unitaries and Trotter compositions to decouple the system into simple disconnected components. The problem is therefore with the choice of the particular optimisation algorithm that struggles to find the solution. It may be the case that using a different algorithm inside our *in situ* protocol, such as parametrising the control Hamiltonians as a Fourier series rather than as piecewise-constant, would find control pulses for shorter gate times.

For the case of an Ising chain we also varied the number of qubits in order to hypothesise a likely scaling for T_{total} in terms of the number of qubits in the system and found that a polynomial scaling matched very well. As the cost of doing a classical simulation of the *in situ* optimisation of the Ising chain is much lower than for the others, we also picked this system to investigate the impact of measurement noise by introducing a finite value of A_{num} , which parametrises the sensitivity of the system to measurement noise. The results are shown in Fig.6 and give us a good estimate of $N_{\text{upds}} = O(n)$ and $A_{\text{num}} = O(1/n)$. The latter implies that $N_{\text{prec}} = O(n^2)$, due to the central limit theorem that states the number of repetitions required scales quadratically with the desired accuracy which gives $N_{\text{prec}} \propto A_{\text{num}}^{-2} = O(n^2)$.

Putting this together with the previous results that $N_{\text{fids}} = O(n)$ for gradient based optimisation and $N_{\text{meas}} = O(1)$ for the local estimator fidelity gives $N_{\text{runs}} = O(n^4)$. As this is done with a constant gate time and with a constant success probability, this implies that the time required to find a control sequence that implements a C-NOT gate on an Ising chain using a steepest-gradient *in situ* scheme scales as $T_{\text{total}} = O(n^4)$.

The other system we investigated in depth was an Ising ring where the target C-NOT was between two next-nearest-neighbour or two randomly located qubits. Fig.7 shows the required N_{upds} for up to 9 qubits for this system. This graph is in agreement with the previous results of Fig.6 that the number of iterations required grows slowly with the number of qubits; in this

Figure 6: **Numerical simulation of the experimental cost of finding a C-NOT gate in an Ising chain using steepest ascent *in situ*.** The number of updates and the fidelity accuracy required for the optimisation protocol to succeed is plotted for a chain of qubits with the Hamiltonian of Eq.(1) with nearest-neighbour Ising interactions, a gate time of $T_{\text{gate}} = 4\pi$, and $N_{\text{ts}} = 48$ timeslots. The target in each case is a CNOT gate between two qubits in the middle of the chain, separated by one other qubit. The gate fidelity used is F_{LE} , therefore the true gate fidelity will be a little higher. Since the cost of the $n = 3$ case is significantly lower than for the other, it has been omitted from all fits.

Figure a) shows how N_{upds} scales with the number of qubits for different target gate fidelities (error bars are twice the standard error). For this plot, the accuracy to which the fidelity is measured, A_{num} , is picked to give a $p_{\text{succ}} = 50\%$ success rate. We see a strong linear relation in the number of iterations required as a function of the number of qubits giving $N_{\text{upds}} = O(n)$.

Figure b) shows how the accuracy to which the local fidelity needs to be measured, A_{num} , scales with the number of qubits for different target fidelities (error bars are 5 times the standard error). The data is expected to have an $O(1/n)$ scaling as, in order to reach a gate infidelity of ϵ , the fidelity ought to require a measurement accuracy $O(\epsilon)$. As this is calculated from the sum of the fidelities of the subsystem, we conjectured that they each need to be measured to an accuracy $O(\epsilon/n)$. The data points lie very close to a c/n curve, providing strong support for this argument. However, the constant c does not appear to have quite a linear relationship with ϵ ; we did not investigate this further as it does not affect scalability.

The fidelity accuracy for the $p_{\text{succ}} = 50\%$ is estimated using an interpolation of p_{succ} values for a range of A_{num} . Between 25 and 45 points are used in the interpolation. Each of these points are the average over a number of repetitions: 200 for $n = 3, 4$; 100 for $n = 5$; 50 for $n = 6, 7$. The method for selecting the A_{num} values for the simulations and the interpolations are described in more detail in the code repository [38].

case the results even show signs of being sub-linear. An unexpected feature shared by both sets of results is the low cost of the 3 qubit case. Our intuition is that it is due to the additional symmetries present and the ease with which the single qubit not part of the target C-NOT gate can be kept disentangled.

The reason for picking the ring topology and not-nearest-neighbour target gate was to check whether the Ising chain results were unique, to demonstrate the applicability of the *in situ* approach to different systems, and to test its ability to reach more complex gates. Specifically, it shows that the scalability of the scheme did not rely on boundary effects or on qubits being adjacent to each other. While a quantum computer could be built using only nearest-neighbour gates, being able to entangle two arbitrary qubits in the time of a single gate drastically reduces the potential run time of algorithms.

Discussion

This polynomial scaling observed numerically in these two different cases is encouraging evidence that the protocol may indeed be efficient. Some of the components that make up this exponential scaling come from numerical data, so several fits are possible. However the points lie so close to a linear fit in Fig.6 that a different fit, such as an exponential one, would diverge only slowly. Fig.7 suggests that corrections to the fit are more likely to make it sub-linear than more costly. Although it is clear that the results presented here do not form an absolute proof of the scalability of an *in situ* control scheme for all quantum simulators, they are at the very least strong evidence that this is a powerful approach to take for moderately large systems of a few tens of qubits.

Systems of such a size are interesting as they correspond to the state-of-the-art that can be realised experimentally. Using the *in situ* scheme

Figure 7: Number of control pulse updates needed to find a C-NOT gate in an Ising chain and ring. The number of updates required for the optimisation protocol to succeed is plotted for a chain (left) and a ring (right) of qubits with the Hamiltonian of Eq.(1) with nearest-neighbour Ising interactions $\sigma_z \otimes \sigma_z$, a gate time of $T_{\text{gate}} = 4\pi$, $N_{\text{ts}} = 48$ timeslots, and a gate fidelity of $F_{LE} = 0.999$. Full numerical accuracy was used in these simulations. Each data point represents repeated optimisations: 100 for $n < 8$; 96 for $n = 8$; 30 for $n = 9$. The number of successful optimisations is $p_{\text{succ}} > 90\%$ in all cases. The error bars are twice the standard error. In both cases there is no evidence that the scaling is more than linear, even disregarding the $n = 3$ data. As we did not have an obvious model to fit to these points, no best-fit is shown. The value of N_{upds} is slightly higher for the ring than the chain, but appears to have a smaller gradient in n . For both graphs the results for nearest-neighbour and random qubits are statistically indistinguishable. This highlights that the optimisation scheme operates equally well in both cases and works more efficiently than a naive dynamically decoupling protocol.

for such systems would likely find control sequences for entangling gates that are not currently known, and where purely classical numerical optimisation schemes would fail due to the enormous computational requirements. Furthermore, testing these predictions in such experiments would extend these results to numbers of qubits that are completely unattainable for a purely classical computer to model, and test this protocol closer to full-scale universal quantum computation.

One potential difficulty in optimal control is the existence of traps: local maxima of the fidelity that optimisation algorithms converge to which are not the global maxima. The question of whether traps exist in unitary control using the standard gate fidelity has been well studied [39–41], and the conclusion is that generic quantum control landscapes are almost always trap free. This may also apply to the local estimator of the fidelity; traps were not a problem for the numerical simulations we performed and found no evidence of any new traps in Fig.3 or elsewhere.

The numerical results presented here used GRAPE, which decomposed the control pulses into piece-wise constant functions. A potentially more powerful approach, and one which is harder

to do classically but may be easier to implement physically, would be to decompose them by frequency [42], such as in CRAB [43, 44] and GOAT [45]. Such algorithms are slow to run classically due to the difficulty of exponentiating the time-dependent Hamiltonian, a step which is bypassed in the *in situ* scheme. They typically require fewer parameters to describe a successful control pulse and thus may prove faster than GRAPE when done experimentally. A different variation would be to change from a gradient based algorithm to a geometric or genetic one, or even to use machine-learning algorithms to learn about the system [46]. Ideas from robust control [47, 48] may also be usable in an *in situ* framework, in order to make the approach more resilient to fluctuations in the system Hamiltonian.

A future direction to take this work is to apply it to another important aspect of quantum computation: error correction. The protocol detailed here can be used in much the same way for this by replacing the target operation from a C-NOT gate, to one protecting some logical qubits. Preliminary results show that with a tuneable interaction the system can discover decoherence free subspaces and simple error correcting codes this way. Work remains on what the most useful tasks to optimise for are, and on showing the scalability

of this approach.

Acknowledgements

This work was supported by EPSRC through the Quantum Controlled Dynamics Centre for Doctoral Training, the EPSRC Grant No. EP/M01634X/1, and the ERC Project ODYC-QUENT. We are grateful to HPC Wales for giving access to the cluster that was used to perform the numerical simulations. Many thanks to Stephen Glaser and David Leiner for discussions on possible implementations.

References

- [1] J. I. Cirac and Peter Zoller. Goals and opportunities in quantum simulation. *Nat. Phys.*, 8:264, apr 2012. DOI: [10.1038/nphys2275](https://doi.org/10.1038/nphys2275).
- [2] Tomi H. Johnson, Stephen R. Clark, and Dieter Jaksch. What is a quantum simulator? *EPJ Quantum Technol.*, 1:1, 2014. DOI: [10.1140/epjqt10](https://doi.org/10.1140/epjqt10).
- [3] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153, mar 2014. DOI: [10.1103/RevModPhys.86.153](https://doi.org/10.1103/RevModPhys.86.153).
- [4] David P. DiVincenzo and IBM. The Physical Implementation of Quantum Computation. *arXiv:quant-ph/0002077*, feb 2000. DOI: [10.1002/1521-3978\(200009\)48:9/11;1::AID-PROP771;3.0.CO;2-E](https://doi.org/10.1002/1521-3978(200009)48:9/11;1::AID-PROP771;3.0.CO;2-E).
- [5] Seth Lloyd. Almost Any Quantum Logic Gate is Universal. *Phys. Rev. Lett.*, 75:346, 1995. DOI: [10.1103/PhysRevLett.75.346](https://doi.org/10.1103/PhysRevLett.75.346).
- [6] Jennifer L. Dodd, Michael A. Nielsen, Michael J. Bremner, et al. Universal quantum computation and simulation using any entangling Hamiltonian and local unitaries. *Phys. Rev. A*, 65:040301, 2002. DOI: [10.1103/PhysRevA.65.040301](https://doi.org/10.1103/PhysRevA.65.040301).
- [7] Daniel Burgarth, Sougato Bose, Christoph Bruder, et al. Local controllability of quantum networks. *Phys. Rev. A*, 79:060305(R), 2009. DOI: [10.1103/PhysRevA.79.060305](https://doi.org/10.1103/PhysRevA.79.060305).
- [8] Shai Machnes, U. Sander, Steffen J. Glaser, et al. Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework. *Phys. Rev. A*, 84:022305, 2011. DOI: [10.1103/PhysRevA.84.022305](https://doi.org/10.1103/PhysRevA.84.022305).
- [9] Steffen J. Glaser, Ugo Boscain, Tommaso Calarco, et al. Training Schrodinger’s cat: Quantum optimal control: Strategic report on current status, visions and goals for research in Europe. *Eur. Phys. J. D*, 69:279, 2015. DOI: [10.1140/epjtd/e2015-60464-1](https://doi.org/10.1140/epjtd/e2015-60464-1).
- [10] D. J. Egger and Frank K. Wilhelm. Adaptive hybrid optimal quantum control for imprecisely characterized systems. *Phys. Rev. Lett.*, 112:240503, 2014. DOI: [10.1103/PhysRevLett.112.240503](https://doi.org/10.1103/PhysRevLett.112.240503).
- [11] J. Kelly, R. Barends, B. Campbell, et al. Optimal quantum control using randomized benchmarking. *Phys. Rev. Lett.*, 112:240504, 2014. DOI: [10.1103/PhysRevLett.112.240504](https://doi.org/10.1103/PhysRevLett.112.240504).
- [12] Herschel Rabitz, Regina de Vivie-Riedle, Marcus Motzkus, et al. Whither the Future of Controlling Quantum Phenomena? *Science*, 288:824, 2016. DOI: [10.1126/science.288.5467.824](https://doi.org/10.1126/science.288.5467.824).
- [13] Christopher Ferrie and Osama Moussa. Robust and efficient in situ quantum control. *Phys. Rev. A*, 91(5):052306, 2015. DOI: [10.1103/PhysRevA.91.052306](https://doi.org/10.1103/PhysRevA.91.052306).
- [14] J. Kelly, R. Barends, A. G. Fowler, et al. Scalable in-situ qubit calibration during repetitive error detection. *Phys. Rev. A*, 94:032321, 2016. DOI: [10.1103/PhysRevA.94.032321](https://doi.org/10.1103/PhysRevA.94.032321).
- [15] Jun Li, Xiaodong Yang, Xin Hua Peng, et al. Hybrid Quantum-Classical Approach to Quantum Optimal Control. *Phys. Rev. Lett.*, 118:150503, 2017. DOI: [10.1103/PhysRevLett.118.150503](https://doi.org/10.1103/PhysRevLett.118.150503).
- [16] Patrick Rebentrost, Maria Schuld, Leonard Wossnig, et al. Quantum gradient descent and Newton’s method for constrained polynomial optimization. *arxiv:1612.01789*, 2016. URL <http://arxiv.org/abs/1612.01789>.
- [17] Dawei Lu, Keren Li, Jun Li, et al. Enhancing quantum control by bootstrapping a quantum processor of 12 qubits. *njp Quantum Inf.*, 3:45, 2017. DOI: [10.1038/s41534-017-0045-z](https://doi.org/10.1038/s41534-017-0045-z).
- [18] Michael Johanning, Andrés F. Varón, and Christof Wunderlich. Quantum simulations with cold trapped ions. *J. Phys. B*, 42:154009, 2009. DOI: [10.1088/0953-4075/42/15/154009](https://doi.org/10.1088/0953-4075/42/15/154009).
- [19] B. P. Lanyon, C. Hempel, Daniel Nigg, et al. Universal Digital Quantum Simulation with Trapped Ions. *Science*, 334:57, 2011. DOI: [10.1126/science.1208001](https://doi.org/10.1126/science.1208001).
- [20] Rainer Blatt and Christian F. Roos. Quantum simulations with trapped ions. *Nat. Phys.*, 8:277, 2012. DOI: [10.1038/nphys2252](https://doi.org/10.1038/nphys2252).
- [21] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. Quantum simulations with ultracold quantum gases. *Nat. Phys.*, 8:267, 2012. DOI: [10.1038/nphys2259](https://doi.org/10.1038/nphys2259).
- [22] Henning Labuhn, Daniel Barredo, Sylvain Ravets, et al. Realizing quantum Ising models in tunable two-dimensional arrays of single Rydberg atoms. *Nature*, 534:667, 2016. DOI: [10.1038/nature18274](https://doi.org/10.1038/nature18274).
- [23] Xin Hua Peng and Dieter Suter. Spin qubits for quantum simulations. *Front. Phys. China*, 5:1, 2010. DOI: [10.1007/s11467-009-0067-x](https://doi.org/10.1007/s11467-009-0067-x).
- [24] J. Cai, A. Retzker, Fedor Jelezko, et al. A large-scale quantum simulator on a diamond surface

- at room temperature. *Nat. Phys.*, 9:168, 2013. DOI: [10.1038/nphys2519](https://doi.org/10.1038/nphys2519).
- [25] Isabela A. Silva, Alexandre M. Souza, Thomas R. Bromley, et al. Observation of time-invariant coherence in a nuclear magnetic resonance quantum simulator. *Phys. Rev. Lett.*, 117:160402, 2016. DOI: [10.1103/PhysRevLett.117.160402](https://doi.org/10.1103/PhysRevLett.117.160402).
- [26] Andrew A. Houck, Hakan E. Türeci, and Jens Koch. On-chip quantum simulation with superconducting circuits. *Nat. Phys.*, 8:292, 2012. DOI: [10.1038/nphys2251](https://doi.org/10.1038/nphys2251).
- [27] P. J. J. O’Malley, Ryan Babbush, I. D. Kivlichan, et al. Scalable Quantum Simulation of Molecular Energies. *Phys. Rev. X*, 6:031007, 2016. DOI: [10.1103/PhysRevX.6.031007](https://doi.org/10.1103/PhysRevX.6.031007).
- [28] P. L. McMahon, Alireza Marandi, Yoshitaka Haribara, et al. A fully programmable 100-spin coherent Ising machine with all-to-all connections. *Science*, 354:614, 2016. DOI: [10.1126/science.aah5178](https://doi.org/10.1126/science.aah5178).
- [29] J. F. Poyatos, J. I. Cirac, and Peter Zoller. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. *Phys. Rev. Lett.*, 78:390, 1997. DOI: [10.1103/PhysRevLett.78.390](https://doi.org/10.1103/PhysRevLett.78.390).
- [30] Yoav Lahini, Gregory R. Steinbrecher, Adam D. Bookatz, et al. Quantum logic using correlated one-dimensional quantum walks. *npj Quantum Inf.*, 4(1), 2018. DOI: [10.1038/s41534-017-0050-2](https://doi.org/10.1038/s41534-017-0050-2).
- [31] Alexei Gilchrist, Nathan K. Langford, and Michael A. Nielsen. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A*, 71:062310, 2005. DOI: [10.1103/PhysRevA.71.062310](https://doi.org/10.1103/PhysRevA.71.062310).
- [32] Marcus P. Da Silva, Olivier Landon-Cardinal, and David Poulin. Practical characterization of quantum devices without tomography. *Phys. Rev. Lett.*, 107:210404, 2011. DOI: [10.1103/PhysRevLett.107.210404](https://doi.org/10.1103/PhysRevLett.107.210404).
- [33] Marcus Cramer, Martin B. Plenio, Steven T. Flammia, et al. Efficient quantum state tomography. *Nat. Commun.*, 1:149, 2010. DOI: [10.1038/ncomms1147](https://doi.org/10.1038/ncomms1147).
- [34] Navin Khaneja, Timo Reiss, Cindie Kehlet, et al. Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms. *J. Magn. Reson.*, 172:296, 2005. DOI: [10.1016/j.jmr.2004.11.004](https://doi.org/10.1016/j.jmr.2004.11.004).
- [35] J R Johansson, P D Nation, and Franco Nori. QuTiP: An open-source Python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.*, 183(8), 2012. DOI: [10.1016/j.cpc.2012.11.019](https://doi.org/10.1016/j.cpc.2012.11.019).
- [36] J. R. Johansson, Paul D. Nation, and Franco Nori. QuTiP 2: A Python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.*, 184:1234, 2013. DOI: [10.1016/j.cpc.2012.11.019](https://doi.org/10.1016/j.cpc.2012.11.019).
- [37] Alexander Pitchford, Paul D. Nation, Robert Johansson, et al. [www.qutip.org](http://qutip.org/), 2018. URL <http://qutip.org/>.
- [38] Alexander Pitchford and Benjamin Dive. github.com/ajgpitch/quantum-in-situ-opt/, 2018. URL <https://github.com/ajgpitch/quantum-in-situ-opt/>.
- [39] Tak-San Ho, Jason Dominy, and Herschel Rabitz. Landscape of unitary transformations in controlled quantum dynamics. *Phys. Rev. A*, 79:013422, 2009. DOI: [10.1103/PhysRevA.79.013422](https://doi.org/10.1103/PhysRevA.79.013422).
- [40] Alexander N. Pechen and David J. Tannor. Are there traps in quantum control landscapes? *Phys. Rev. Lett.*, 106:120402, 2011. DOI: [10.1103/PhysRevLett.106.120402](https://doi.org/10.1103/PhysRevLett.106.120402).
- [41] Benjamin Russell, Herschel Rabitz, and Re-Bing Wu. Quantum Control Landscape Are Almost Always Trap Free. *arxiv:1608.06198*, 2016. URL <http://arxiv.org/abs/1608.06198>.
- [42] Björn Bartels and Florian Mintert. Smooth optimal control with Floquet theory. *Phys. Rev. A*, 88:052315, 2013. DOI: [10.1103/PhysRevA.88.052315](https://doi.org/10.1103/PhysRevA.88.052315).
- [43] Patrick Doria, Tommaso Calarco, and Simone Montangero. Optimal control technique for many-body quantum dynamics. *Phys. Rev. Lett.*, 106:190501, 2011. DOI: [10.1103/PhysRevLett.106.190501](https://doi.org/10.1103/PhysRevLett.106.190501).
- [44] Tommaso Caneva, Tommaso Calarco, and Simone Montangero. Chopped random-basis quantum optimization. *Phys. Rev. A*, 84:022326, 2011. DOI: [10.1103/PhysRevA.84.022326](https://doi.org/10.1103/PhysRevA.84.022326).
- [45] Shai Machnes, Elie Assémat, David Tannor, et al. Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits. *Phys. Rev. Lett.*, 120:150401, 2018. DOI: [10.1103/PhysRevLett.120.150401](https://doi.org/10.1103/PhysRevLett.120.150401).
- [46] Pantita Palittapongarnpim, Peter Wittek, Ehsan Zahedinejad, et al. Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics. *Neurocomputing*, 0:1, 2016. DOI: [10.1016/j.neucom.2016.12.087](https://doi.org/10.1016/j.neucom.2016.12.087).
- [47] D. Daems, A. Ruschhaupt, D. Sugny, et al. Robust quantum control by a single-shot shaped pulse. *Phys. Rev. Lett.*, 111(5):050404, 2013. DOI: [10.1103/PhysRevLett.111.050404](https://doi.org/10.1103/PhysRevLett.111.050404).
- [48] Edwin Barnes, Xin Wang, and S. Das Sarma. Robust quantum control using smooth pulses and topological winding. *Sci. Rep.*, 5:12685, 2015. DOI: [10.1038/srep12685](https://doi.org/10.1038/srep12685).
- [49] Benjamin Dive. *Controlling Open Quantum Systems*. Phd thesis, Imperial College London, 2017. URL <http://hdl.handle.net/10044/1/56626>.

Appendix

Local estimator to the fidelity

We derive a bound for the gate fidelity $F(M, U) = \langle \psi | \rho_M | \psi \rangle$ which is the fidelity of the Choi states of the target gate U with the realised channel M using methods derived from [33]. A more in depth analysis is available in [49]. The Choi state of U is defined as $|\psi\rangle = U \otimes \mathbb{1} |\Omega\rangle$, and the Choi state of M is $\rho = (M \otimes id) |\Omega\rangle \langle \Omega|$, where id is the identity map, and $|\Omega\rangle = \sum_k |kk\rangle$ is a maximally entangled state between the original Hilbert space and a copy of it. We consider the case where the target operation U is unitary and has a tensor product structure such that $U = \bigotimes_i U_i$. The Choi state of U inherits its tensor product structure and is given by $|\psi\rangle = \bigotimes_i |\psi_i\rangle$.

To find a bound for $F(M, U)$ we begin by constructing the projectors $h_i = \mathbb{1}_i - |\psi_i\rangle \langle \psi_i|$ for each U_i . These projectors have a very simple spectrum with a single 0 eigenvalue with corresponding eigenket $|\psi_i\rangle$, and a degenerate orthogonal space with eigenvalue 1. These projectors are summed together to form a Hamiltonian $H = \sum_i h_i \otimes \mathbb{1}_{\bar{i}}$ such that each h_i acts on its own part of the Hilbert space and is identity on the rest. This has a single $E_0 = 0$ eigenvalue with eigenstate $|E_0\rangle = |\psi\rangle$, while all its other eigenvalues are positive integers. By expanding this Hamiltonian in its eigenbasis $\{E_k, |E_k\rangle\}$ we have

$$\begin{aligned} \text{Tr}[H\rho] &= \sum_{k \geq 0} E_k \langle E_k | \rho | E_k \rangle \\ &\geq \sum_{k > 0} \langle E_k | \rho | E_k \rangle, \end{aligned}$$

as $E_0 = 0$ and all the other energies are one or greater. Next, using the identity $\text{Tr}[\rho] = \sum_{k \geq 0} \langle E_k | \rho | E_k \rangle = 1$ we find

$$\begin{aligned} \text{Tr}[H\rho] &\geq 1 - \langle E_0 | \rho | E_0 \rangle \\ &\geq 1 - \langle \psi | \rho | \psi \rangle. \end{aligned}$$

From the definition of the gate fidelity this give $F(M, U) \geq 1 - \text{Tr}[H\rho]$.

The expectation value of the Hamiltonian can

be evaluated according to

$$\begin{aligned} \text{Tr}[H\rho] &= \sum_i \text{Tr}[(\mathbb{1}_i - |\psi_i\rangle \langle \psi_i|) \otimes \mathbb{1}_{\bar{i}} \rho] \\ &= \sum_i \text{Tr}_i[(\mathbb{1}_i - |\psi_i\rangle \langle \psi_i|) \rho_i] \\ &= \sum_i (1 - \langle \psi_i | \rho_i | \psi_i \rangle), \end{aligned}$$

where $\rho_i = \text{Tr}_{\bar{i}}[\rho]$. This is also the Choi state of the map

$$M_i(\cdot) \equiv M(\cdot \bigotimes_{j \neq i} \frac{1}{d_j} \mathbb{1}_j),$$

which is the map M acting on subsystem i with the other subsystems in the maximally mixed state. This results in Eq.(2)

$$F(M, U) \geq 1 - \sum_i (1 - F(M_i, U_i)).$$

One way of measuring F_{LE} is to prepare the state of one subspace in a basis state and all the others in a maximally mixed state, perform the simulation, measure the initial subspace, and then repeat for a tomographically complete basis set and for each subsystem; giving a cost of $N_{\text{meas}} = \sum_i O((d_i)^2) = O(n)$. However by noting that a maximally mixed state is a random mixture of pure states, the fidelity of each subsystem can be measured at the same time by preparing each one in a random pure basis state. In this case there is no scaling with the number of qubits as the number of repetitions required depends only on the size of the largest subsystem. This gives $N_{\text{meas}} = O(1)$.

While this shows that local fidelity is efficient to extract from an experimental set up, it is no easier than the gate fidelity to compute in numerical simulations. This is because it must be calculated from a unitary (or CPT map) which represents the evolution of the whole system of dimension 2^n . Multiple partial traces of it are needed in order to compute the M_i from Eq.(2), which is typically a slow operation. Indeed, the need to calculate the local fidelity was a considerable strain on our numerical simulations and one of the compounding reasons why we could not simulate larger systems.