Strategies for braiding and ground state preparation in digital quantum hardware

Herasymenko, Y.

7 Measurement-driven navigation in many-body Hilbert space

7.1 Introduction

Quantum state preparation is a prominent routine in quantum information processing toolbox [185–199]. Such procedure often implies steering a quantum system from a "simple" towards a more complex, pre-designated resourceful state, e.g. a many-body entangled state. A steering protocol is characterized by an as short as possible runtime and high resulting overlap with the target state. Constructing such protocols can be done in multiple distinct ways. One is to design the Hamiltonian of the system, such that its unitary evolution leads to a designated state. This paradigm is represented by methods like digital computation or analog simulation [187, 189–192]. Such protocols require exact knowledge of the starting state, as well as the precise timing of the unitary evolution, to be accurate. Another strategy is to make use of the environment, adding a dissipative element to the evolution. Combined with the Hamiltonian evolution, this results in methods such as drive-and-dissipation [195, 196]. Finally, one can design a sequence of generalized measurements, which brings the system towards the target state via measurement back-action alone [200–202]. The relevant part of the evolution is then completely governed by the system-detector coupling (see also Ref. [203]). Unlike protocols involving pre-defined unitary evolution, such measurement-driven state preparation may not require knowledge of the starting state and fine-tuning of the system Hamiltonian [202].

The above types of state-preparation strategies can be referred to as *passive*, meaning that these protocols are pre-determined and pursued regardless of how the system evolves. Given this perspective, it appears beneficial to go beyond the forms of control described above, and introduce the concept of *active* decision making. This type of steering exploits information extracted during the system's evolution to decide on the

operations that follow. This is also referred to as closed-loop quantum control and is typically used to improve the Hamiltonian-based state preparation [69–73, 205]. In that case, extracting the necessary information requires the introduction of measurements into the protocol, which may result in an undesired back-action. Nevertheless, in certain cases, closed-loop control of Hamiltonian evolution does yield an improvement to the speed and the fidelity of the protocol.

Another possibility, which is a subject of increasing interest, is to employ active decision in measurement-driven protocols (implying no Hamiltonian drive) [74, 75, 206]. In such protocols, the necessary information about the system is naturally available from the employed measurements. The active decision is then being made about possible changes in the subsequent generalized measurements, such that the target state is prepared as rapidly and accurately as possible [74, 75, 201, 206]. Some general theorems have been stated concerning such active measurement-driven state preparation [206], along with some specific protocols designed to reach single-qubit target states [74, 75]. However, it remains unclear how an active measurement-driven protocol can be effectively harnessed to engineer resourceful many-body states. In this case, the large size of the Hilbert space makes it challenging to actively steer the system evolution in the desired direction.

In this chapter, we establish a general framework for measurement-driven active navigation in Hilbert space and construct active-decision protocols for measurement-only steering of many-body states. In particular, we focus on states manifesting genuine multipartite entanglement. When attempting to address the problem, one is naturally constrained by a few factors. One is that only reasonably local system-detector couplings are to be used in the protocol. Moreover, we require that the number of distinct system-detectors couplings available for steering does not scale up faster than the system's size (this number should not be super extensive). This natural requirement restricts the capabilities of the protocol and results in the need of correlating different system-detector couplings. Another prerequisite is that applying one type of coupling generally leads to an update in the expected benefits from other couplings. This phenomenon, which we refer to as "coupling frustration", calls for nontrivial coordination between different coupling applications. Finally, there is a problem of orienteering: it is relatively easy to "get lost" in the many-body Hilbert space when exploring it with the set of tools limited by locality and extensivity (cf. Ref. [205]).

We note that upon the availability of indefinite computational power, one can always find an optimum sequence of measurements through dynamic

programming techniques (cf. Ref. [206]). Roughly speaking, this can be done by considering all possible future quantum trajectories of the system. However, in a large Hilbert space, it is practically impossible to realize the theoretically optimal feedback policy. This is because the extensive consideration of outcome scenarios is too complex for a many-body Hilbert space of an already not very large system (it increases at least exponentially with the system's size and the duration of the protocol). Instead, we aim at designing heuristic strategies for active decision-making, which would allow for a significant – but not necessarily optimal – speedup of the protocol.

To meet these challenges, we introduce *Hilbert-space navigation techniques*. The first technique, which we term *greedy orienteering policy*, is based on the notion of a cost function. A simple example of such a cost function could be the target state infidelity. Minimizing it in a greedy protocol may already yield a reasonable advantage compared to the passive policy. To test this approach, we study numerically the preparation of a ground state of Affleck-Lieb-Kennedy-Tasaki (AKLT) spin-1 model [207]. The numerical study shows the speedup factor that increases with system size, reaching factor 9.5 for $N = 6$. Looking ahead, we discuss the fundamental challenge of landscape flatness that may arise for some target states when using simple infidelity as the cost function. Although this issue did not arise in the example we considered, we propose a possible modification to the cost function which should remedy this problem if it occurs.

The second technique is to map the Hilbert space onto a colored multigraph, referred to as the Quantum State Machine. The vertices of such a graph correspond to the basis states, and the edges represent the actions of generalized measurements. Upon an appropriate choice of basis states, such Quantum State Machine representations allow for improved navigation in Hilbert space. This can be done by heuristically representing it as quantum wayfinding on the graph. To substantiate this heuristic, we introduce the notion of semiclassical coarse-graining of a Quantum State Machine graph. Optimizing the exploration of these graphs by choosing the most appropriate system-detector couplings results in advantageous active-decision protocols. To exemplify this navigation paradigm we consider the preparation of the 3-qubit W-state, with a numerical study demonstrating a 12.5-fold improvement in protocol runtime.

Throughout the chapter, we assume that we know the initial state of the system. This can be a "cheap" (say, product) and robust quantum state that does not require many resources for its preparation. However, one can directly generalize the above approaches to the case where the initial state is unknown and is therefore represented by a density matrix. In the

more intricate case of a Quantum State Machine-based policy, one would then need to take a weighted combination of graph navigation protocols with different initial states.

The remainder of the chapter is organized as follows. In Sec. 7.2, we introduce the basics of measurement-induced steering. Specifically, in Sec. 7.2.1, we define the steering protocols and their elements, as well as the quantitative measure of the protocol's success. Then, in Sec. 7.2.2, we illustrate these definitions as applied to passive steering of a single qubit. The general selection criteria, including locality and extensivity, for the system-detector couplings, which are to be used for the active steering, are addressed in Sec. 7.2.3. In Sec. 7.3, we introduce the notion of frustration of steering and discuss the possibilities of protocols' speed-up for mutually commuting (Sec. 7.3.1) and non-commuting (Sec. 7.3.2) couplings. In the latter case, we develop a "parent-Hamiltonian" approach. A "quantum-compass" approach to active-decision steering, based on the greedy cost-function accumulation policy, is developed in Sec. 7.4, where we also employ this scheme to the preparation of the AKLT state. In Sec. 7.5, we develop an alternative active-steering framework – a "Quantum State Machine." In Sec. 7.5.1, we introduce the generalities of this approach based on the underlying representation of the steering protocol in terms of a quantum graph. Next, we discuss the quantum parts of this graph (Sec. 7.5.2), as well as the coarse-graining procedure, with the resulting coarse-grained graph being semiclassical (Sec. 7.5.3). This type of Hilbert-space orienteering is illustrated in Sec. 7.5.4, where an active-decision steering protocol for preparation of a three-qubit W-state is presented. Our findings are summarized and discussed in Sec. 7.6.

7.2 Measurement-driven state preparation

7.2.1 Generalities

Measurement-driven state steering is a specific class of state-preparation protocols. Its basic building blocks are coupling the quantum system (s) to quantum detectors (ancillary systems) utilizing engineered interactions, followed by strong measurement on the detectors (d). The goal of designing a measurement-based steering protocol is to generate a process that prepares the desired system state by utilizing a sequence of measurement back-actions.

Here, we will additionally assume that the internal evolution of the system and the detector are trivial (their Hamiltonians are kept null:

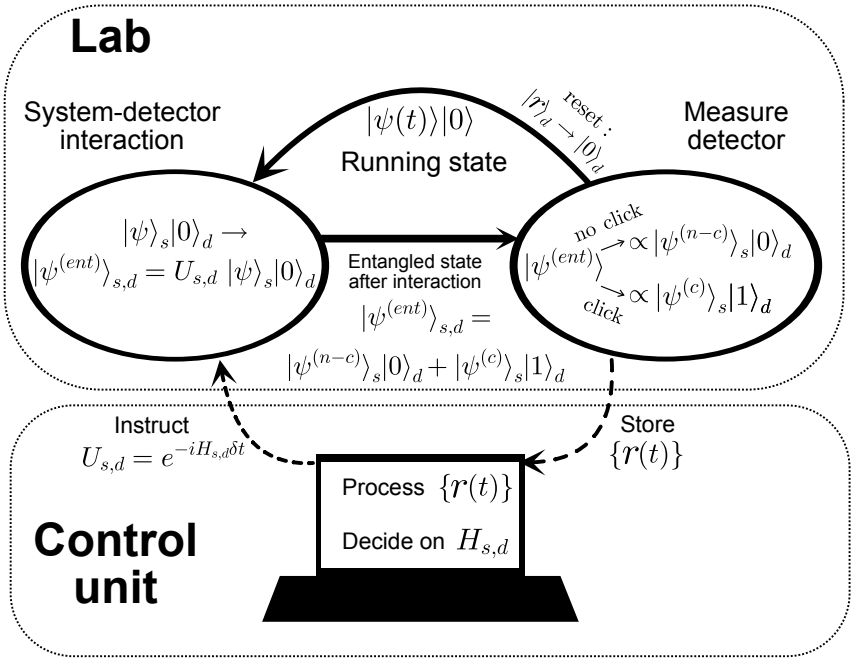

Figure 7.1: Basic design of the measurement-driven state preparation. The procedure starts with a given initial state $\rho_s^{(in)}$ and proceeds with a protocol, as described in Def. 32, until a good accuracy of the target state $|\psi_0\rangle$ is achieved. The control unit decides on the system-detector interaction unitary $U_{s,d}$ based on the stored record of detector readouts. We focus on constructing an optimized policy for decision-making, such that the target state is simulated as efficiently as possible.

$H_s = 0$, $H_d = 0$), as in Refs. [74, 75, 206], so that the unitary dynamics in the problem is governed solely by the coupling between the system and detectors determined by Hamiltonian $H_{s,d}$. For concreteness of analysis, we also constrain the detector to be a qubit initialized in a trivial state $|0\rangle$, and the system to be represented by N spins. Although a general spin S can be considered, we focus on the cases $S = 1/2$ and $S = 1$. We assume certain knowledge about the initial state of the system, which is described by the initial density matrix $\rho_s^{(\mathrm{in})}$. For the sake of simplicity, we further address the target state which is a pure state $|\psi_0\rangle$.

Although the ensuing protocol can be further generalized (see Section 7.6), we now formally fix its structure as given below:

Definition 32. *A measurement-driven state steering is a protocol that is performed to prepare a state $|\psi_0\rangle$, starting from the state $\rho_s^{(in)}$. It runs by repeating iterative cycles of the following form (see Fig. 7.1):*

1. *Prepare the detector qubit in the state $|0\rangle$.*
2. *Based on the available information, select the system-detector coupling Hamiltonian $H_{s,d}$ to be used in the next step.*
3. *Perform a system-detector evolution governed by a Hamiltonian $H_{s,d}$ for a short time interval δt: $U_{s,d} = e^{-iH_{s,d}\delta t}$.*
4. *Once the system-detector evolution is over, projectively measure the detector qubit in the Z-basis. Store the readout r for further processing.*
5. *Decide whether the protocol is to be continued or terminated. In the former case, return to step 1.*

Now, in the vast space of protocols that have such structure, we would like to emphasize the distinction between two classes of protocols: passive and active.

In a passive protocol, the stored readouts $\{r(t)\}$ from step 4 may influence the decision for protocol termination or continuation at step 5, but not the choice for the interaction Hamiltonian $H_{s,d}$ made at step 2 in the next protocol cycles. Hamiltonians $H_{s,d}$ can still be chosen differently for different iterations: e.g. for a large system, the detector qubit can be coupled to different subsystems thereof. However, $H_{s,d}$ used at each cycle in the passive protocol has to be pre-determined from the outset. If a passive protocol also has a pre-determined duration (and thus doesn't use readouts $\{r(t)\}$ at the termination step), we land in a subclass of passive protocols where the readouts don't have any influence on the protocol. We would refer to such protocols as "blind steering". For blind steering, the readouts of the detector at any given step can be averaged, i.e., following the measurement, the detector's density matrix is traced out. In this chapter, however, we will focus on the non-blind version of passive steering, where readouts are indeed employed for an informed protocol termination.

In contrast to passive protocols, in an active protocol one uses the readouts $\{r(t)\}$ to make an informed decision for the interaction Hamiltonians $H_{s,d}$ as well as for termination/continuation of the protocol. Active decision-making has to follow a certain policy, which becomes the crucial part of the protocol. For a good active policy, its adoption should result in

a significant speedup of the protocol compared to its passive counterpart. Alternatively, one can also fix the protocol runtime and aim to improve the precision of the state preparation. We focus on the former target: minimizing protocol runtime for a fixed target precision. The major challenge in this chapter is to construct such advantageous active decision-making policies. By comparing active steering with the (non-blind) passive steering as defined above, we investigate the advantage offered specifically by the directed evolution, i.e. active decision-making for $H_{s,d}$.

Before we move on to the issue of active policy constructions, let us discuss the criteria for termination of a running protocol. In general, one cannot guarantee "perfect steering," i.e., obtaining the desired target state with the fidelity of 1 in a finite number of protocol cycles. Instead, one may consider preparing the target state with *infidelity* R:

$$R\left(\rho_s^{(\mathrm{fin})}, |\psi_0\rangle\right) \equiv 1 - \langle\psi_0|\rho_s^{(\mathrm{fin})}|\psi_0\rangle, \tag{7.2.1}$$

where the state $\rho_s^{(\mathrm{fin})}$ is the final state of the system once the protocol is terminated. It is worth emphasizing that the system evolution during the protocol is probabilistic and depends on the stochastic readouts $\{r(t)\}$. It follows that different runs of the same protocol may yield different $\rho_s^{(\mathrm{fin})}$ and, thus, the infidelity R. Therefore, to characterize the protocol as a whole, we introduce the following accuracy measure:

Definition 33. *We refer to a measurement-driven state-preparation protocol as ϵ-precise, if the infidelity between the final state and the target state is bounded by ϵ for any run of the protocol:*

$$R\left(\rho_s^{(fin)}, |\psi_0\rangle\right) < \epsilon. \tag{7.2.2}$$

Given the knowledge of the readout sequence, we may simulate the quantum system state (the quantum trajectory) on a computer in parallel to the measurement run. Thus one can infer the running system state exactly (referred to as filtering in the literature [69]), and test inequality (7.2.2). This sets a trivial criterion for protocol termination, which we will apply by default to all passive and active protocols considered in this chapter. Namely, a protocol can be terminated right after the cycle when the target state infidelity becomes smaller than ϵ, thus making it an ϵ-precise protocol. Apart from controlling the precision, we are interested in the number of cycles N_c, after which the protocol has been terminated. As N_c may also differ greatly, depending on a specific run, we will characterize the protocol by $\langle N_c\rangle_{\mathrm{run}}$, where the averaging is performed over many runs.

Note that here averaging is taken over stochastic readout sequences. In reality, steering errors as well as external noise may be further contributing factors to the stochasticity. For the given target state and target precision ϵ (cf. Definition 33), our goal is to find an ϵ-precise protocol such that $\langle N_c \rangle_{\text{run}}$ is as small as possible. We will be considering this minimization as the key goal of our constructions.

7.2.2 Passive steering: Single qubit

As a simple example of a measurement-driven protocol, we consider single-qubit steering (for a more general consideration, the reader is referred to Sec. 7.2.3). For simplicity, we will assume the target state to be $|0\rangle$, and the starting state to be a perfectly mixed state: $\rho_{\text{start}} = \text{diag}(1/2, 1/2)$. A single coupling suffices to guarantee the preparation of the target state (in fact, from an arbitrary starting state) with an arbitrary precision [202]:

$$H_{s,d} = \gamma \sigma_s^- \sigma_d^+ + \text{H.c.} \tag{7.2.3}$$

Here, σ_s and σ_d are Pauli matrices acting in the system and detector spaces, respectively. By construction, a protocol that operates with only a single coupling Hamiltonian $H_{s,d}$, i.e., without a readout-based option of choosing different couplings, is considered passive. Nevertheless, even for passive protocols, one can introduce a policy based on the measurement outcomes, which would accelerate quantum-state steering.

Let us first address a protocol that runs for $N_c^{(\text{pass})}$ cycles using the coupling (7.2.3), regardless of the measurement outcomes. Under the definition given in Sec. 7.2, this would be an example of *blind* steering. In this case, the probability of obtaining a readout $r = 0$ decreases exponentially with the total number of cycles N_c. Tracing out detector outcomes (since we are blind to measurement outcomes), this results in a density matrix:

$$\rho(N_c) = \begin{pmatrix} 1 - e^{-N_c \gamma^2 \delta t^2}/2 & 0 \\ 0 & e^{-N_c \gamma^2 \delta t^2}/2 \end{pmatrix}. \tag{7.2.4}$$

Given the threshold infidelity ϵ, we need to run the protocol for $N_c^{(\text{pass})}(\epsilon)$ cycles:

$$N_c^{(\text{pass})}(\epsilon) = \frac{1}{\gamma^2 \delta t^2} \log\left(\frac{1}{2\epsilon}\right) \tag{7.2.5}$$

This characterizes the efficiency of the completely blind passive protocol [202] for the single-qubit setup.

Next, we consider passive protocols where the sequence of readouts is recorded. One then needs to interpret the measurement outcomes, which for this setup is straightforward. We note that when the readout is $r = 1$ (click event), the target state is instantly prepared [cf. Eq. (7.2.9)]. Therefore, one can terminate the protocol directly after the detector clicks for the first time: in this case, all further steps are simply redundant and do not result in any evolution of the system. This will constitute a termination-policy improvement of the passive blind protocol for this single-qubit case. If $r = 0$, i.e. no click is measured (such a null-measurement [208] event still gives the system a nudge towards the target state by measurement back-action), the protocol simply continues until a certain maximal number of cycles, $N_c^{(\max)}$. The target infidelity ϵ would be directly related to $N_c^{(\max)}$ in a way equivalent to the blind protocol runtime (7.2.5). The average runtime of the non-blind passive protocol, $\tilde{N}_c^{(\mathrm{pass})} \equiv \langle N_c \rangle_{\mathrm{run}}$, is then given by

$$\tilde{N}_c^{(\mathrm{pass})} = \frac{1}{2\gamma^2 \delta t^2} \left(1 - e^{-\gamma^2 \delta t^2 N_c^{\max}}\right) + \frac{N_c^{\max}}{2}. \tag{7.2.6}$$

This runtime is strictly smaller than the runtime for the passive blind protocol, Eq. (7.2.5), and yields a twofold speedup in the $\epsilon \to 0$ limit. It is worth emphasizing, however, that the termination policy can realistically be applied only to the case of few-body quantum states. For such a policy to be useful, a single detector click should signify that the system is in the target state. This can only be realized when the detector is coupled to all elements of the system. For a many-body system (many qubits), a natural assumption of locality rules out such a coupling: a click of the detector coupled to a subsystem of the system does not guarantee that the whole system is steered to the desired state. Nevertheless, the above simple example shows that detector readouts can be used for accelerating the state preparation. In what follows, we will focus on active feedback strategies. There, instead of protocol termination, the local-measurement outcomes are employed for choosing the most efficient sequence of further measurement cycles.

7.2.3 Selection criteria for system-detector couplings

Families of system-detector couplings

Both for the active and passive protocols, a key feature is the choice of coupling Hamiltonians $H_{s,d}$. Given the target state $|\psi_0\rangle$, it is natural to constrain this choice to a certain family $\{H_{s,d}(\mathbf{p})\}$, for a set of (discrete or

continuous) parameters $\mathbf{p}$. Deciding on the choice of $H_{s,d}$ in each protocol cycle translates into selecting the value of $\mathbf{p}$. Before discussing the policies for doing so, we address a different question: How to effectively preselect this family $\{H_{s,d}(\mathbf{p})\}$? To answer this question, it is useful to consider the following general decomposition of $H_{s,d}$:

$$H_{s,d} = V_s \sigma_d^+ + V_s^\dagger \sigma_d^- + \tilde{V}_s \sigma_d^z, \tag{7.2.7}$$

where V_s and $\tilde{V}_s$ are arbitrary system operators and matrices $\sigma_d^\pm = \frac{1}{2}(\sigma_d^x \pm i\sigma_d^y)$ act on the detector. In Eq. (7.2.7), we discard any terms of the form $\sim \mathbb{I}_d$, as those represent the internal system evolution. Furthermore, for our purposes, it is also sufficient to consider a special case of the decomposition, where $\tilde{V}_s = 0$.

With Eq. (7.2.7) in mind, let us consider the transformation of the system state ρ_s under a single cycle of the steering protocol. First, let us consider the system state transformation that is performed when the measurement outcomes are averaged over (blind measurement). In the weak measurement limit, $\delta t \to 0$, this is represented by the map:

$$\begin{aligned} &\rho_s \to \Lambda_{V_s}(\rho_s) \\ &\quad \equiv \left(1 - \frac{\delta t^2}{2} V_s^\dagger V_s\right) \rho_s \left(1 - \frac{\delta t^2}{2} V_s^\dagger V_s\right) + V_s \rho_s V_s^\dagger \delta t^2. \end{aligned} \tag{7.2.8}$$

We note that the terms of order $\mathcal{O}(\delta t^2)$ in this expression represent the standard Lindbladian jump operator. Based on the map (7.2.8) by tracing out the detector readouts after each step, one derives a Lindblad equation describing the system evolution for the blind steering [202].

Let us now turn back to our protocol, where the different measurement outcomes are discriminated. During step 4 of the protocol cycle (cf. Definition 32), there is a probability

$$p^{(\mathrm{cl})}(\rho_s, V_s) = \delta t^2 \operatorname{tr}(V_s \rho_s V_s^\dagger)$$

that a qubit flip is measured in the detector (click probability). The resulting state in the limit of small δt is then:

$$\rho_s \to \Lambda_{V_s}^{(\mathrm{cl})}(\rho_s) \equiv \frac{V_s \rho_s V_s^\dagger}{\operatorname{tr}(V_s \rho_s V_s^\dagger)}. \tag{7.2.9}$$

A "no-click" scenario occurs with probability

$$p^{(\mathrm{ncl})}(\rho_s, V_s) = 1 - p^{(\mathrm{cl})}(\rho_s, V_s),$$

and results in a state:

$$\rho_s \to \Lambda_{V_s}^{(\mathrm{ncl})}(\rho_s) \equiv \frac{\left(1-\frac{\delta t^2}{2}V_s^\dagger V_s\right)\rho_s\left(1-\frac{\delta t^2}{2}V_s^\dagger V_s\right)}{1-\delta t^2 \mathrm{tr}(V_s^\dagger V_s \rho_s)}. \tag{7.2.10}$$

Note that for the weak-measurement limit considered here ($||V_s\delta t|| \ll 1$), the click probability is parametrically smaller than that for the no-click event: a qubit flip can be recorded in the detector only rarely.

Extensivity and locality

We are now in a position to expound our considerations for the family $\{H_{s,d}(\mathbf{p})\}$ in terms of the operators $\{V_s(\mathbf{p})\}$. For a meaningful comparison between active and passive protocols, we first require that there exists a passive protocol that employs Hamiltonians $\{H_{s,d}(\mathbf{p})\}$ to reach the target state $|\psi_0\rangle$. For concreteness, we assume that the passive protocol involving all the family members is a cyclic one: each of the couplings $\{H_{s,d}(\mathbf{p})\}$ is employed one after another in a predefined manner, and the cycle is repeated once all the couplings are employed. The size of the family is restricted by the number of available system-detector couplings, which is assumed to scale up with increasing systems' size not faster than extensively.

It is then natural to demand that none of $\{H_{s,d}(\mathbf{p})\}$ can move the system state away from the target state. Given Eqs. (7.2.9) and (7.2.10), this yields a dark-state condition $V_s(\mathbf{p})|\psi_0\rangle = 0$ for every $\mathbf{p}$. This is equivalent to every operator $V_s(\mathbf{p})$ taking the following form:

$$V_s = \sum_{\alpha=1}^{D-1} v_\alpha |\psi_0\rangle\langle\psi_\alpha| + \sum_{\alpha,\beta=1}^{D-1} w_{\alpha\beta}|\psi_\beta\rangle\langle\psi_\alpha|, \tag{7.2.11}$$

where D is the Hilbert-space dimensionality of the system, and $\{|\psi_\alpha\rangle\}$ is any (many-body) basis for the system that includes $|\psi_0\rangle$ as a basis state. This general form of the system-detector coupling is, however, not realistic for many-body systems, as $D = 2^N$ grows exponentially with the number of qubits for an N-qubit system. Thus, having in mind steering of many-body states, we should further restrict the family of available steering operators.

The second condition for $\{V_s(\mathbf{p})\}$ is that these couplings can realistically be engineered in an experimental realization of the system. In this chapter, we focus on the most basic aspect of this condition: locality. One may consider two types of locality: geometric and operator (k-locality [16]).

Geometric locality of the operator V_s implies that such interaction only requires coupling the system spins that are in geometrical proximity during the experiment. A k-local operator V_s implies that only k system spins are coupled at a time. It is natural to impose the locality constraint not on the full operator V_s, but its individual terms. For example, if V_s involves all system spins, but its individual terms only couple 2 spins at a time, we will consider V_s a 2-local coupling (in line with [16]). Note that a k-local operator V_s implies an interaction Hamiltonian $H_{s,d}$ that is $(k+1)$-local.

Sufficient conditions for the coupling operators. Room for active decision-making

It is worth stressing at this point that $V_s(\mathbf{p})$ following the form given by Eq. (7.2.11) for all $\mathbf{p}$ is necessary but not sufficient for $|\psi_0\rangle$ to be the only dark state of the passive protocol. For some choices of such a family $\{V_s(\mathbf{p})\}$, a spurious final state $|\psi_0'\rangle \neq |\psi_0\rangle$ might be reached. However, this would imply a dark-state condition $V_s(\mathbf{p})|\psi_0'\rangle = 0$ (for every $\mathbf{p}$), and this should not hold for a generic (say, random-matrix-type) operator V_s, which satisfies $V_s|\psi_0\rangle = 0$. For generic coefficients v_α and $\omega_{\alpha\beta}$ in (7.2.11), one does not expect an existence of a spurious final state (for that, an extra constraint is needed, such as vanishing of certain v_α, $\omega_{\alpha\beta}$, or a specific relation between the coefficients).

One concludes that a family consisting of a single Eq. (7.2.11)-type coupling V_s is sufficient to prepare $|\psi_0\rangle$ in a passive protocol without generating a dark space. Notably, reducing the family to a single member would leave no room for active decision-making in a protocol defined by this family (an active protocol necessitates at least two operators to choose from). On the other hand, such an ultimate V_s would not generically satisfy the crucial locality conditions and, thus, would be unrealistic to implement. Natural counterexample couplings V_s' that have multiple dark states arise in the important case when V_s' acts only on a part of the system.

To construct such a counterexample, one may start from an arbitrary operator V_s that satisfies the dark-state condition $V_s|\psi_0\rangle = 0$ for a single state $|\psi_0\rangle$ in a given system. Now, consider a larger system embedding the original one and construct a different target state which is a tensor product of $|\psi_0\rangle$ and a certain auxiliary state: $|\Psi_0\rangle \equiv |\psi_0\rangle \otimes |\tilde{\psi}_0\rangle$. In this case, one may take $V_s \to V_s'$, where $V_s' = V_s \otimes \mathbb{I}_{\tilde{s}}$ still satisfies condition $V_s'|\Psi_0\rangle = 0$ relative to this new target state in the extended Hilbert space. Yet for a general starting state of the total system, the operator V_s' is obviously not sufficient to prepare the extended target state $|\Psi_0\rangle$ – also

implying the existence of spurious dark states (in fact, all states of the form $|\psi_0\rangle \otimes |\tilde{\psi}_0\rangle$ turn out to be dark, for arbitrary $|\tilde{\psi}_0\rangle$).

We see that in this "self-evident" fashion of selecting a steering operator, the condition of $|\Psi_0\rangle$ being a dark state for a single operator V_s' was not sufficient for V_s' to be able to guarantee preparation $|\Psi_0\rangle$. As discussed above, an operator V_s capable of steering a unique dark state is typically highly nonlocal, in contrast to the limited capacity of a localized operator $V_s \otimes \mathbb{I}_{\tilde{s}}$. We conclude that a family of multiple operators $\{V_s(\mathbf{p})\}$ is needed to realistically prepare a target state, once that state is sufficiently complicated. This, in turn, opens the door for gaining advantage through active decision-making.

7.3 Types of system-detector couplings

The preselected family of coupling operators $\{V_s(\mathbf{p})\}$ determines both the performance of the ensuing passive protocols and the possibilities for active policy construction. In the present section, we identify the crucial role of the commutation properties of $\{V_s(\mathbf{p})\}$. We first consider N-qubit steering protocols which employ coupling operators $\{V_s(\mathbf{p})\}$ that are mutually commuting. As a shorthand, we denote this as non-frustrated steering. We show that a realistic passive protocol of this type can be designed for product states and certain graph states. Commuting couplings also allow for a simple feedback strategy, which results in a significant speedup of the respective passive protocol. Next, we move on to passive steering protocols that are frustrated. Such frustration of local couplings naturally arises for many-body target states related to local parent Hamiltonians. We propose an explicit method of constructing a family of non-commuting operators $\{V_s(\mathbf{p})\}$ that allows to prepare such a many-body target state in a passive protocol. This forms the basis for Secs. 7.4 and 7.5, where we move on to the active versions of frustrated steering protocols.

7.3.1 Mutually commuting couplings

Here we focus on N-qubit steering protocols implemented with mutually commuting couplings $\{V_s(\mathbf{p})\}$. As will be demonstrated, a passive protocol of this type can be constructed for an arbitrary target state, yielding an asymptotically precise passive preparation. However, we find that this construction would, in general, require non-local system-detector couplings $H_{s,d}$, deeming the implementation of the protocol for many-body states impractical. We then identify an exception to this rule – a subclass of

graph states that can be obtained using local commuting couplings. For this, we discuss the constraints coming from both geometric locality, as well as k-locality. Finally, we extend the discussion from such passive protocols to their active counterparts. To achieve this, we propose a simple feedback strategy that speeds up non-frustrated steering in a substantial way.

As a trivial example of non-frustrated steering, consider an N-qubit product state as a target state, e.g., $|00..0\rangle$. The starting state will be assumed to be the perfectly mixed state. To prepare it with a steering protocol, one can use a set of couplings parameterized by the qubit number $i = 1, ..N$:

$$V_s^{(\mathrm{prod})}(i) = \gamma \sigma_i^-. \tag{7.3.1}$$

Passively alternating between the steering cycles employing $V_s^{(\mathrm{prod})}(i)$ with different i guarantees preparation of the target state with any given accuracy. This directly follows from the analysis of Secs. 7.2.2 and 7.2.3. For an active version of the protocol, partial protocol termination can be applied: if a click is registered when measuring any qubit i, the coupling $V_s^{(\mathrm{prod})}(i)$ is dropped out from the sequence of couplings that will be applied in further cycles. In other words, the steering with this "fired" coupling is terminated at this point, whereas other couplings remain active – hence the term "partial termination". Since this implies a readout-based decision on the set of steering couplings that are used at a given step, we classify this as an active steering protocol. In the $\epsilon \to 0$ limit, this strategy results in the following relation between active and passive runtimes:

$$N_c^{(\mathrm{act})}(\epsilon) = \frac{N_c^{(\mathrm{pass})}(\epsilon)}{2} + \frac{N}{2\gamma^2 \delta t^2}, \tag{7.3.2}$$

which leads to up to a substantial 2-fold speedup for the active version, similarly to Eq. (7.2.6).

Non-frustrated steering towards any target state $|\psi_0\rangle$ can in principle be designed if we allow for an arbitrary coupling set. Indeed, given a many-body unitary transformation to $|\psi_0\rangle$ from a product state $|00..0\rangle$, i.e., $|\psi_0\rangle = U_\psi |00..0\rangle$, one may formally construct a family of couplings:

$$V_s^{(U_\psi)}(i) = \gamma U_\psi \sigma_i^- U_\psi^\dagger. \tag{7.3.3}$$

Clearly, any protocol for $|\psi_0\rangle$ preparation using couplings of the form of Eq. (7.3.3) would be a unitary equivalent of the same protocol which uses couplings of Eq. (7.3.1) to prepare $|00..0\rangle$. Therefore, a passive protocol iterating over $V_s^{(U_\psi)}(i)$ for different i would successfully prepare the target

state $|\psi_0\rangle$. We also conclude that a partial-termination strategy can be applied to this coupling set with the same effect as for the product state target. Note, however, that in most cases employing $V_s^{(U_\psi)}(i)$ would not be practically feasible. Indeed, since U_ψ is a general many-body operation, the couplings $V_s^{(U_\psi)}(i)$ would involve arbitrarily non-local terms. For most N-spin states $|\psi_0\rangle$ with large N, one thus expects that the resulting $V_s^{(U_\psi)}(i)$ would break any requirement of geometric or k-locality.

This locality-violation rule can be circumvented for U_ψ which is given by a shallow circuit and thus $|\psi_0\rangle$ which is weakly entangled. As a resourceful example of such $|\psi_0\rangle$, consider a graph state defined on a generic graph G [209]:

$$|\psi_G\rangle = \left(\prod_{\substack{(j,k)\in \\ \text{edges}(G)}} U^{(\text{gr})}_{(j,k)} \right) \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right)^{\otimes N}, \tag{7.3.4}$$

$$U^{(\text{gr})}_{(a,b)} = \exp\left(\mathrm{i}\pi |00\rangle\langle 00|_{a,b}\right), \tag{7.3.5}$$

in which case

$$U_\psi = \left(\prod_{\substack{(j,k)\in \\ \text{edges}(G)}} U^{(\text{gr})}_{(j,k)} \right) \left(\prod_{j\in \text{qubits}} \exp\left(\mathrm{i}\frac{\pi}{4}\sigma_j^y\right) \right).$$

Since two-qubit rotations $U^{(gr)}_{(j,k)}$ all mutually commute, the coupling $V_s^{(U_\psi)}(i)$ acts only on spin i and on the spins j whose vertices share an edge with i in the graph G. Therefore, this coupling is $(k+1)$-local if there are k edges coming out of vertex i. Moreover, $V_s^{(U_\psi)}(i)$ is also geometrically local, if the graph G only connects the qubits which are in geometric proximity. We conclude that for the graphs satisfying the above conditions, a realistic preparation of graph states with local non-frustrated steering is possible. Such a protocol can be sped up in the same way it was possible for the product states – using active feedback via the partial-termination strategy.

For the perfectly mixed starting state, the partial-termination policy gives an optimal speed-up of a protocol driven by non-frustrated couplings $V_s^{(U_\psi)}(i)$. Indeed, the protocols in question are then equivalent to an independent set of N 1-qubit steering protocols (under the unitary transformation U_ψ). This picture, however, breaks down for a more general

starting state. Let us first consider the trivial target $U_\psi = \mathbb{1}$, $|\psi_0\rangle = |00..0\rangle$, while the starting state is itself entangled (e.g. $\frac{1}{\sqrt{2}}(|00..0\rangle + |11..1\rangle)$). In this case, the click received from a single coupling $V_s(i)$ may imply that multiple couplings can be dropped from the applied sequence, and not just $V_s(i)$ itself. This would be more optimal than the partial termination strategy outlined above. The same picture extends to the more interesting case when the target state $|\psi_0\rangle$ is entangled itself, e.g. a graph state, while the starting state is a product state. Indeed, under the unitary mapping U_ψ which takes an entangled state $|\psi_0\rangle$ to $|00..0\rangle$, the product starting state in turn becomes entangled. Hence, the previous reasoning applies and partial-termination would generally not be an optimal active policy in this situation. Instead, one may accelerate it further by applying one of the frustrated-coupling strategies outlined in the following sections.

7.3.2 Frustrated system-detector couplings

From now on, let us focus on accelerating steering protocols which employ couplings constrained by at least one of the two notions of locality. Under this premise, for target states other than the product states and states prepared by a shallow circuit, we would generally need to go beyond the non-frustrated protocols outlined above. The first question to tackle is how to design the local couplings V_s that are suitable for a passive protocol. In principle, this can be addressed on a case-by-case basis, tailoring some coupling set with a specific target state in mind. (This approach will be demonstrated for the W-state preparation in Sec. 7.5.4.) However, this is not always a straightforward task. Therefore, it is interesting to know whether one can devise a general scheme to this end. For this, we propose an approach based on a parent-Hamiltonian construction.

The parent Hamiltonian H_ψ of $|\psi_0\rangle$ has $|\psi_0\rangle$ as a non-degenerate ground state, and is constructed from $|\psi_0\rangle$ in the form:

$$H_\psi = \sum_j H_\psi^{(j)}. \tag{7.3.6}$$

Here, all the terms $H_\psi^{(j)}$, while defined as acting on the entire Hilbert space, are local in the real space given that the state $|\psi_0\rangle$ hosts a limited amount of entanglement [210] (implying an area-law dependence of the entanglement entropy accommodated in $|\psi_0\rangle$). Note that, by the construction of the parent Hamiltonian, terms $H_\psi^{(j)}$ have $|\psi_0\rangle$ as their common ground state, although they generally do not commute with each other: this is possible because of their respective ground state degeneracy. These degenerate

ground spaces of $H_\psi^{(j)}$ will be the central element of our construction of the coupling family. For the term $H_\psi^{(j)}$, which nontrivially acts on some m qubits, let us denote its m-qubit ground states as $|\phi_a^{(j)}\rangle$ and the excited states $|\theta_a^{(j)}\rangle$. Given these, we can construct the coupling operators of the following form:

$$V_s^j(\mathbf{w},\mathbf{v},\mathbf{u}) = \sum_{ab} w_{ab}|\phi_a^{(j)}\rangle\langle\theta_b^{(j)}| + \sum_{ab} v_{ab}|\theta_a^{(j)}\rangle\langle\theta_b^{(j)}| + \sum_{ab} u_{ab}|\phi_a^{(j)}\rangle\langle\phi_b^{(j)}|. \tag{7.3.7}$$

A particular example of this construction will be addressed in detail in the context of the AKLT model in Sec. 7.4 (see also Refs. [195, 202]).

For a generic (fixed) value of the parameters involved, running a passive protocol with the coupling given by Eq. (7.3.7) allows one to steer the system into the ground state of $H_\psi^{(j)}$. Alternating the coupling operators by selecting terms with different j at different measurement steps allows for steering the system into the joint ground space of all couplings $H_\psi^{(j)}$. This space is given by the target state $|\psi_0\rangle$ only, as it is the non-degenerate ground state of H_ψ. Thereby, as long as the parent Hamiltonian H_ψ is local, we have managed to construct an appropriate coupling set for a passive protocol (also see [195] for a related statement proven in more detail).

Now, let us consider an active-protocol construction. First, we note that unlike in the "non-frustrated" protocol construction, the operators V_s^j for different choices of parameter sets are, in general, not mutually commuting. This also applies to the couplings with different values of j, as $H_\psi^{(j)}$ generally do not commute. Therefore, the measurement outcome of steering by V_s^j (where, because of the locality of $H_\psi^{(j)}$, j corresponds to a certain region in real space) impacts the outcomes of steering at other locations. As a result, the partial-termination strategy cannot be applied to this coupling set, as it assumes that the respective cycles of the protocol can be considered separately. Instead, the feedback strategy for the frustrated steering should continuously coordinate the application of different couplings in the protocol. In a many-body context, this becomes a complicated navigation-type problem (cf. Ref. [205]). We devote the following two sections to the study of such possible coordination policies.

7.4 Quantum compass: Cost-function policies

One way to enable the Hilbert-space navigation is to introduce a cost function $C(\rho_s)$, which is to be minimized in the protocol. The basic example would be the infidelity $C(\rho_s) = R(\rho_s, |\psi_0\rangle)$ of the system state ρ_s to the target state $|\psi_0\rangle$, defined in Eq. (7.2.1). Achieving the global minimum $R(\rho_s, |\psi_0\rangle) = 0$ of this cost function would be equivalent to preparing the target state. In general, to calculate $R(\rho_s, |\psi_0\rangle)$, one needs to know the state of the system ρ_s. This is, in principle, feasible, as we control the system evolution given all measurements outcomes and therefore can numerically simulate it in parallel to the experiment. However, the requirement of such a simulation being done in parallel to the experiment puts a restriction on the size of the system that one can work with. For now, we will accept this limitation; finding ways to mitigate it is among the worthwhile potential extensions of the work presented in this chapter.

With a given cost function $C(\rho_s)$ at hand, we can use it to form the active decision for the coupling operator $V_s(\mathbf{p})$. The ultimate strategy is to pick $V_s(\mathbf{p})$ which brings the system to the global minimum $|\psi_0\rangle$ in the fastest expected time. For $C(\rho_s) = R(\rho_s, |\psi_0\rangle)$ this is equivalent to the ultimate strategy defined by dynamic programming [206], requiring unrealistic computation power. Therefore, the notion of the global cost function does not give any additional advantage in constructing such a strategy. Instead, one can use its cheaper version – the "greedy strategy." Specifically, one can use $V_s(\mathbf{p})$ that yields the fastest expected reduction of the cost function in a single step of the evolution:

$$V_s^{(\mathrm{greed})}(\mathbf{p}) = \mathrm{argmin}_{V_s(\mathbf{p})} R[\Lambda_{V_s(\mathbf{p})}(\rho_s)], \tag{7.4.1}$$

where $\Lambda_{V_s(\mathbf{p})}(\rho_s)$ is defined in Eq. (7.2.8) If there are multiple minima, we will assume that argmin returns a random representative among those. With only a small amount of computations needed to decide for the optimal next coupling $V_s^{(\mathrm{greed})}(\mathbf{p})$, this greedy procedure allows us to avoid the complex long-term analysis of the protocol.

As one can see from a direct implementation, the greedy minimization of the cost function can accelerate the state preparation by a large factor. To demonstrate this, we consider the example of the ground state in an AKLT spin chain as the target state. This is an entangled state of N spin-1 particles governed by the Hamiltonian H_{AKLT}:

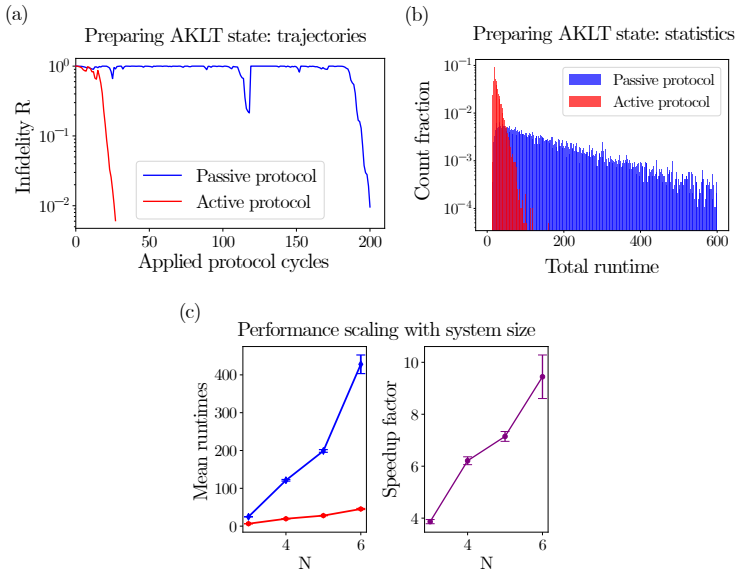

Figure 7.2: (a) Infidelity as a function of the protocol cycle for active and passive protocol runs towards a 5-spin AKLT state. These example runs are characterized by the duration similar to the mean protocol durations of respective protocols ($\approx 199 \pm 4$ for passive and $\approx 28 \pm 0.5$ for active protocol) (b) Histograms of protocol durations t for the preparation of the five-spin AKLT, with accuracy given by infidelity $R < \epsilon = 0.01$. An exponential decaying profile, characteristic of a Poissonian process, can be clearly observed (note the log scale). Note that all recorded runs for an active protocol lasted far less than the mean duration of a passive protocol (200 cycles). Each histogram was compiled from 10^4 simulated runs; the figure is truncated at 600 cycles for better presentation. (c) Scaling of the active protocol's advantage with system size N. A speedup factor tends to increase significantly as the system scales, with factor 9.5 being the estimated speedup at 6 spins. The error bars represent 95% confidence intervals due to sampling error in numerical simulation. 10^4 samples were collected to simulate both protocols $N = 3, 4, 5$, and 10^3 at $N = 6$. Similarly to the above, the infidelity threshold is $\epsilon = 0.01$.

$$H_{\mathrm{AKLT}} = \sum_i H_{i,i+1} = \sum_i \left[\vec{S}_i \cdot \vec{S}_{i+1} + \frac{1}{3} \left(\vec{S}_i \cdot \vec{S}_{i+1} \right)^2 \right], \tag{7.4.2}$$

Here, we assume periodic boundary conditions, implying a single ground

state [207]. The Hamiltonian (7.4.2) is a parent Hamiltonian (as defined in Sec. 7.3.2), where each term $H_{i,i+1}$ has four degenerate ground states $|\phi_a^i\rangle$ with eigenvalue $-2/3$, and 5 excited states $|\theta_b^i\rangle$ with energy $4/3$.

Hereafter, we consider the all-down product state as our starting state. We first design a passive steering protocol for AKLT state preparation, following the parent Hamiltonian construction from Sec. 7.3.2. For simplicity, we restrict ourselves to a less general version of Eq. (7.3.7), and use the following family of coupling operators (cf. Refs. [195, 202]):

$$V(\mathbf{c}, i) = |\phi_4^i\rangle\langle\theta_5^i| + \sum_{\alpha,\beta=1,..4} c_{\alpha\beta}|\phi_\alpha^i\rangle\langle\theta_\beta^i|, \tag{7.4.3}$$

with $c_{\alpha\beta}$ constrained to be an orthogonal matrix, to make sure that the interaction is of constant strength and thus no bias is introduced in the construction. In a passive steering protocol, we will alternate between different values of i, while drawing instances of orthogonal matrices $\mathbf{c}$ at random. For an active feedback strategy to be used on top of this, we propose a greedy policy relative to $C(\rho_s) = R(\rho_s, |\psi_0\rangle)$ to select $\mathbf{c}$. In both passive and active protocol, we assume each coupling to be applied multiple times until one either receives a click, or no-click for an asymptotically long time. Such a repeated application of a single coupling is then counted as a single protocol cycle. We take this approach for a practical purpose because simulating such protocols is more accessible numerically.

Thus simulated, the relative performance of the passive and the active policies (Fig. 7.2) shows a strong advantage of the active policy. In particular, the speedup factor is steadily increasing with system size (Fig. 7.2c), reaching the value of 9.5 for $N = 6$.

7.4.1 Discussion: orthogonality catastrophe and alternative cost functions

The approach defined above harbors a potential challenge. For the greedy procedure to be effective, it should always yield a nonzero bias in favor of a specific $V_s^{(\mathrm{greed})}(\mathbf{p})$ (or a small subset thereof). In other words, the landscape of the cost function $C(\rho_s)$ should not be flat — and some cost functions may yield better landscapes than others. In particular, applying the infidelity measure $R(\rho_s, |\psi_0\rangle)$ is, in general, fundamentally flawed. Indeed, a $(2^N - 1)$-dimensional subspace of states in the N-body Hilbert space is orthogonal to the target state. Let us consider the case when the starting state belongs to that subspace. This situation would in general not change after a single steering cycle with a local coupling $V_s(\mathbf{p})$. For our

purposes, it implies that the infidelity measure R is equal to 1 for a large manifold of states, and there might be no direction of increase that would allow us to choose an appropriate coupling. The most direct example of this can be observed when applying the greedy policy to non-frustrated steering (see Sec. 7.3.1). For simplicity, let us again take the product state of N qubits $|00..0\rangle$ as the target state, the state $|11..1\rangle$ as the starting state, and the couplings $V(i) = \sigma_i^-$ for steering. Only after such steering protocol results in N successful click events, $R(\rho_s, |\psi_0\rangle)$ gains a nonzero value. Thus before $N-1$ clicks, the greedy policy for $R^{(\mathrm{inf})}$ will not be capable of providing a meaningful decision for the next coupling. Strongly enhanced by the system size, this phenomenon is reminiscent of Anderson's orthogonality catastrophe [211].

As a remedy to this deficiency, a "subsystem infidelity" measure can be introduced:

$$R_{\mathcal{S}}(\rho_s, |\psi_0\rangle) = \sum_{\sigma\in\mathcal{S}} \left[1 - \mathrm{tr}\left(\sqrt{\sqrt{\rho_{0,\sigma}}\rho_{s,\sigma}\sqrt{\rho_{0,\sigma}}}\right)^2\right], \tag{7.4.4}$$

where $\rho_{0,\sigma}$ ($\rho_{s,\sigma}$) is the reduced density matrix of the target state (current state) with respect to subsystem σ. $\mathcal{S}$ is the family of subsystems from which σ are drawn; the choice of $\mathcal{S}$ depends per target state. In the case of the $|11..1\rangle \to |00..0\rangle$ protocol described above, the appropriate $\mathcal{S}$ would be the set of individual spins. Unlike R, such quantity $R_{\mathcal{S}}$ changes every time when a click occurs in this protocol. As a result, the greedy policy with respect to the local $R_{\mathcal{S}}$ would yield the partial-termination protocol of Sec. 7.3.1, significantly boosting the preparation of such a product state.

By continuity with the case of the product state target, such preference for $R_{\mathcal{S}}$ should extend to the weakly-entangled target states, and maybe to some highly-entangled targets. However, we did not see a manifestation of this in the case of our AKLT simulation, where using $R_{\mathcal{S}}$ as a cost function did not yield any improvement compared to R. As a likely explanation for this, the orthogonality catastrophe should become manifest only at large system sizes, where the classical simulation of the protocol is also hindered. However, we expect that some practical target states may still develop a noticeable performance difference between $R_{\mathcal{S}}$ and R, similarly to the case of the product state target. A further study of this question constitutes a promising direction for future work.

7.5 Hilbert-space orienteering map: Quantum State Machine

In this section, we present an orienteering tool that is an alternative to cost-function minimization: mapping out the steering transformations with a Quantum State Machine (QSM) construction. We then illustrate navigation in many-body Hilbert space, employing this machinery to the preparation of the highly entangled W-state of three qubits.

7.5.1 QSM generalities

Every transformation of the system's state, $\Lambda^{(\mathrm{cl})}_{V_s(\mathbf{p})}$ and $\Lambda^{(\mathrm{nc})}_{V_s(\mathbf{p})}$, associated to steering with a specific coupling $V_s(\mathbf{p})$ in a given readout scenario [click or no-click, respectively, see Eqs. (7.2.9) and (7.2.10)], can be represented with a directed graph with complex weights. For this, we notice that every such steering transformation conserves the purity of the state. Therefore, it is convenient to encode transformations $\Lambda^{(\mathrm{cl,\ ncl})}_{V_s}$ in their action on Hilbert space basis states $|\phi_\alpha\rangle$:

$$\Lambda^{(\mathrm{cl,ncl})}_{V_s}(|\phi_\alpha\rangle) = \frac{1}{\sqrt{p^{(\mathrm{cl,ncl})}}} \sum_\beta L^{(\mathrm{cl,ncl})}_{\alpha\beta} |\phi_\beta\rangle \tag{7.5.1}$$

$$L^{(\mathrm{cl})}_{\alpha\beta} = \langle\phi_\beta|\delta t V_s|\phi_\alpha\rangle, \tag{7.5.2}$$

$$L^{(\mathrm{ncl})}_{\alpha\beta} = \langle\phi_\beta|1 - \delta t^2 V_s^\dagger V_s/2|\phi_\alpha\rangle, \tag{7.5.3}$$

where $p^{(\mathrm{cl})}$ ($p^{(\mathrm{ncl})}$) is the probability of a click (non-click) readout upon this steering action. Note that in Eq. (7.5.1), we extended the action of Λ_{V_s} to pure states by a slight abuse of notation compared to Eq. (7.2.9).

The graph representation for steering action $\Lambda^{(\mathrm{cl,ncl})}_{V_s(\mathbf{p})}$, or a steering graph, is then directly obtained from the amplitudes $L^{(\mathrm{cl,ncl})}_{\alpha\beta}$. The vertices in such a graph correspond to the Hilbert space basis states, and the edges describe the steering transformations. The edges are directed and weighted with complex amplitudes, the edge $\alpha \to \beta$ being weighted with amplitude $L^{(\mathrm{cl,\ ncl})}_{\alpha\beta}$ (edges weighted with zero amplitudes are excluded from the graph). Implying this definition, we will use the notation $L^{(\mathrm{cl,\ ncl})}$ for the steering graphs themselves. For basic examples of steering graphs, please refer to Fig. 7.3.

Since the weights $L^{(\mathrm{cl})}_{\alpha\beta}$ are proportional to the matrix elements of coupling operator V_s while $L^{(\mathrm{ncl})}_{\alpha\beta}$ can be expressed via V_s as well, the graph

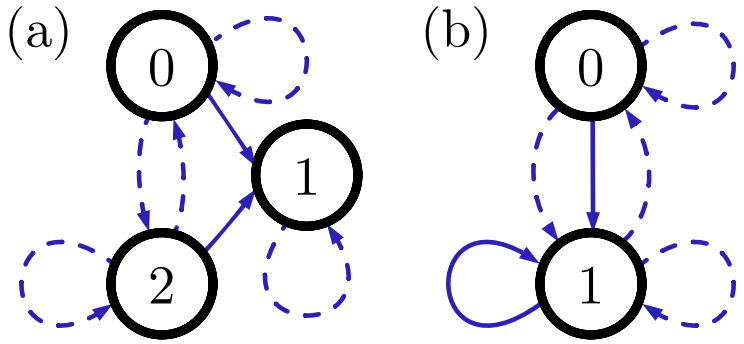

Figure 7.3: Examples of steering graphs (see definition in Sec. 7.5.1): (a) Steering graphs on a 3-level system, corresponding to the coupling $V_s = \gamma(|1\rangle\langle 0| + |1\rangle\langle 2|)$. Graph $L^{(\mathrm{cl})}$ for click action is denoted with solid arrows and the graph $L^{(\mathrm{ncl})}$ for no-click action by dashed arrows (a particular representation of graph coloring). Note that due to the identity operator in Eq. (7.5.3), every vertex is decorated with a self-loop from the $L^{(\mathrm{ncl})}$ graph. To see how the rest of $L^{(\mathrm{ncl})}$ can be deduced from $L^{(\mathrm{cl})}$ (cf. discussion in Sec. 7.5.1), consider the example of $e_{02}^{(ncl)}$ (dashed arrow from state 0 to 2). According to the graphical approach from Sec. 7.5.1, one is to follow edge $e_{01}^{(cl)}$ (solid arrow from 0 to 1) forward and then $e_{21}^{(cl)}$ (solid arrow from 2 to 1) backward - and thus manages to travel from state 0 to 2, in correspondence to $e_{02}^{(ncl)}$. (b) Steering graphs on a 2-level system, as defined by the coupling $V_s = \gamma(|1\rangle\langle 1| + |1\rangle\langle 0|)$. Following the same rule as above, inter-vertex edges of $L^{(\mathrm{ncl})}$ can be deduced from $L^{(\mathrm{cl})}$. For example, by following the edge $e_{11}^{(\mathrm{cl})}$ forward and then the edge $e_{01}^{(\mathrm{cl})}$ backward, one performs a transition from state 1 to state 0, thus reproducing the edge $e_{10}^{(\mathrm{ncl})}$ from $L^{(\mathrm{ncl})}$.

$L^{(\mathrm{ncl})}$ for the no-click action can be inferred entirely from the graph $L^{(\mathrm{cl})}$ for the click action. In particular, due to the term $\propto V_s^\dagger V_s$, graph $L^{(\mathrm{ncl})}$ contains an edge $e_{ij}^{(\mathrm{ncl})}$ from vertex v_i to v_j, if a graph $L^{(\mathrm{cl})}$ contains edges $e_{ik}^{(\mathrm{cl})}$ and $e_{jk}^{(\mathrm{cl})}$ (see Fig. 7.3). Heuristically, to yield a $L^{(\mathrm{ncl})}$-edge, one has to first follow a $L^{(\mathrm{cl})}$-edge forward, and then another $L^{(\mathrm{cl})}$-edge backward. Furthermore, due to the additional identity operator term in Eq. (7.5.3), any graph for the no-click steering action will also include self-loops on each vertex.

The steering graphs introduced above can now be used to create a Quantum State Machine. For n_V couplings $V_s(\mathbf{p})$ in the steering kit, there exist $2n_V$ graphs corresponding to steering maps $\Lambda_{V_s(\mathbf{p})}^{(\mathrm{cl},\,\mathrm{ncl})}$, because of the two possible measurement outcomes for each of the couplings. The QSM

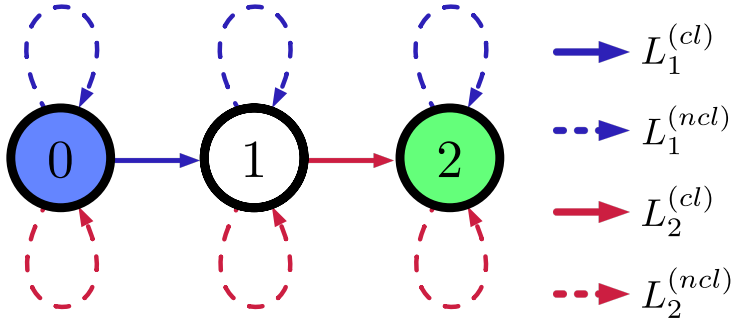

Figure 7.4: A basic example of the QSM multigraph, describing the steering kit for a three-state system. The steering options are represented by the coupling operators $V_1 = \gamma_1|1\rangle\langle 0|$ and $V_2 = \gamma_2|2\rangle\langle 1|$. The starting state is 0, marked in blue, and the target state is 2, marked in green. The optimal coordination policy of the two steering operations is straightforward: one needs to first repeatedly apply the V_1-steering until a click is obtained, and then the V_2-steering until a click is obtained. Compared to the passive steering which iterates between V_1 and V_2 regardless of measurement outcomes, this directly yields a 2-fold speedup in the average performance.

for the steering protocol is then obtained as a collection of these graphs. It can be represented as a colored multigraph, where each steering graph is represented as a single-color subgraph (Fig. 7.4). Consequently, in a QSM multigraph there may be multiple edges going from any vertex α into any other vertex β (making it a multigraph rather than a simple graph), but at most one such edge for each color.

Let us now consider our original task of finding the accelerated navigation protocol. To make use of the QSM construction in this context, we will restrict our consideration to bases $\{|\phi_\beta\rangle\}$ where one of the basis states is the target state $|\psi_0\rangle$ itself. In such a case, state $|\psi_0\rangle$ corresponds to a marked vertex in the graph, and the goal of the steering protocol becomes to drive the system state to that vertex. The goal of optimizing this protocol may then look similar to a known problem of finding the shortest path to the marked vertex on a weighted graph. This problem is standard in graph theory and can be solved as such. Can such a solution be used to design the navigation protocol?

As we will see in Sec. 7.5.2, this analogy is not complete, since the quantum evolution on the graph goes beyond the simple path-on-the-graph

picture. This aspect creates an obstacle to directly applying the graph exploration algorithms to facilitate our protocol speed-up. Fortunately, in some cases, this difficulty can be properly accounted for, as we will see in Sec. 7.5.3. In those cases, the "semi-classical heuristics" of graph exploration may indeed be applied. Finally, in Sec. 7.5.4, we will apply this approach to the preparation of the W-state, with a factor 12.5 improvement compared to the passive protocol using the active navigation protocol.

7.5.2 Quantum subgraphs in a QSM

Let us now compare our QSM navigation task to the standard problem of graph exploration. Our goal is to identify the differences between the two, which prevent us from applying the graph exploration techniques directly to QSM navigation. First of all, the state of the system in graph exploration is at all times represented by a single vertex. The system in a QSM, on the other hand, is generally represented by a superposition over multiple vertices. Furthermore, in graph exploration, the state is modified by following one of the edges. A steering action in a QSM, in contrast, corresponds to a whole collection of edges – i.e., a single steering graph in the QSM multigraph.

Some steering graphs may induce quantum effects, such as superposition and interference. For instance, the steering action whose graph contains two outgoing edges from a given vertex (e.g., vertex 0 for graph $L_1^{(\mathrm{cl})}$ in Fig. 7.5a), can create a nontrivial quantum superposition. If a state is given by a superposition of multiple vertex states, it may further undergo quantum interference. In particular, this can be facilitated by a steering action whose graph contains a vertex with two incoming edges (e.g., vertex 4 for graph $L_2^{(\mathrm{cl})}$ in Fig. 7.5a). In general, a notion of "superposition subgraphs" and "interference subgraphs" of a steering graph can be defined:

1. Superposition subgraph is a subgraph of a steering graph span by multiple (more than one) edges outgoing from a single vertex.

2. Interference subgraph is a subgraph of a steering graph span by multiple edges incoming to a single vertex.

Collectively, we will refer to such interference and superposition subgraphs of a single steering graph as its quantum subgraphs. If the quantum subgraphs are absent in the QSM, we will refer to it as a classical QSM. In other words, in a classical QSM, each vertex has at most one outgoing and at most one incoming edge of any given color.

If a QSM is classical, optimization of the navigation protocol can essentially be reduced to classical graph exploration. For a simple example of a classical QSM and the way to optimize the respective state preparation, consider the 3-level steering actions described in Fig. 7.4. Note that optimization of the classical QSM also applies to the case when the starting state is a superposition of multiple vertex states. If the steering operations contain no quantum subgraphs, the quantum superposition is equivalent to a probabilistic mixture for the sake of the protocol optimization, and the optimal navigation pattern can be extracted accordingly.

As the form of the steering graph depends on the choice of basis, it is conceivable that the number of quantum subgraphs in such a graph in some cases can be reduced by changing the basis (compare Fig. 7.5a and b). However, using a change of basis to remove all the quantum subgraphs in an arbitrary QSM is generally impossible (see Fig. 7.5).

7.5.3 Coarse-grained QSM. Semiclassical heuristic for navigation

We now focus on the steering protocols whose QSM cannot be made classical via a basis transformation. In such a case, it may still be possible to optimize it via a classical graph exploration heuristic. For that, we propose to coarse-grain the QSM by grouping subsets of its vertices into single block-vertices. The coarse-grained QSM would consist of graphs drawn between such block-vertices. The block-vertex containing the target vertex can be considered as the target block-vertex.

An inter-block edge between two block-vertices is drawn, if the original QSM has at least one edge connecting the vertices inside the respective block-vertices. For the coarse-graining to be useful for our purposes, it should be done in such a way that all of the resulting QSM graphs have a classical structure. Namely, the coarse-grained graph should not have quantum subgraphs, e.g. realizing superposition or interference between the block-vertices (in analogy to Sec. 7.5.2). To satisfy this requirement, the following rule for vertex grouping can be employed (cf. Fig. 7.6): *if two edges of the same color are simultaneously coming in or out of a given vertex, the two vertices at the other ends of these edges should be grouped within one effective block-vertex.* This rule manifestly yields basis-dependent groupings, since the very presence of quantum subgraphs in a QSM is basis-dependent. Thus, a smart choice of the basis may allow for an efficient and simpler coarse-grained graph. Designing a general explicit algorithm for finding the optimum basis for an arbitrary QSM is a highly non-trivial task. Heuristically speaking, a convenient choice of the

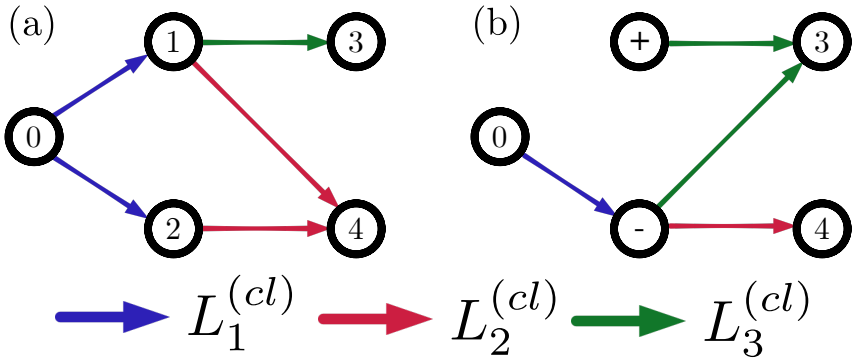

Figure 7.5: Possible configurations of quantum subgraphs in a QSM, exemplified by the 5-vertex subgraph of an example QSM. (a) The click-action graphs for the three coupling operators $V_{1,2,3}$ that form the steering kit. The operators have the form $V_1 = \gamma_1(|1\rangle - |2\rangle)\langle 0|$, $V_2 = \gamma_2|4\rangle(\langle 1| - \langle 2|)$, $V_3 = \gamma_3|3\rangle\langle 1|$. The graphs for the no-click actions are not shown, as their form can be deduced from the graphs for click actions. In the present basis, the V_1-click is manifest as a superposition, the V_2-click – as an interference, and the V_3-click corresponds to a semiclassical evolution. (b) Quantum State Machine for the steering kit from the previous panel, depicted in a different basis. The basis transformation is $|\pm\rangle = (|1\rangle \pm |2\rangle)/\sqrt{2}$. In this case, the basis transformation removes the quantum elements in the $L_{1,2}^{(cl)}$ graphs, however, it turns $L_3^{(cl)}$ into an interference element. Note that there is no basis transformation that would turn such a QSM into a classical one. This statement follows from the uniqueness of the Jordan canonical form for operators V_2 and V_3.

basis should be the one that results in the minimum number of quantum subgraphs in a QSM before coarse-graining.

For the coarse-grained graph to be effectively classical, we desire to ignore details of the system evolution inside the subspace of a given block-vertex. Specifically, we aim to view every block-vertex as an effective single state of the system and assume that every edge allows transporting the system between such block-vertex states with no obstacles. If this was directly possible, and since the coarse-grained QSM by definition contains no quantum subgraphs, optimization of its exploration would have become a classical task. However, such an approximation scheme needs more careful justification. Every block fundamentally corresponds to a Hilbert subspace, and an inter-block edge is given by a $N_1 \times N_2$ matrix of coefficients (where N_1 and N_2 are the internal dimensionalities of the

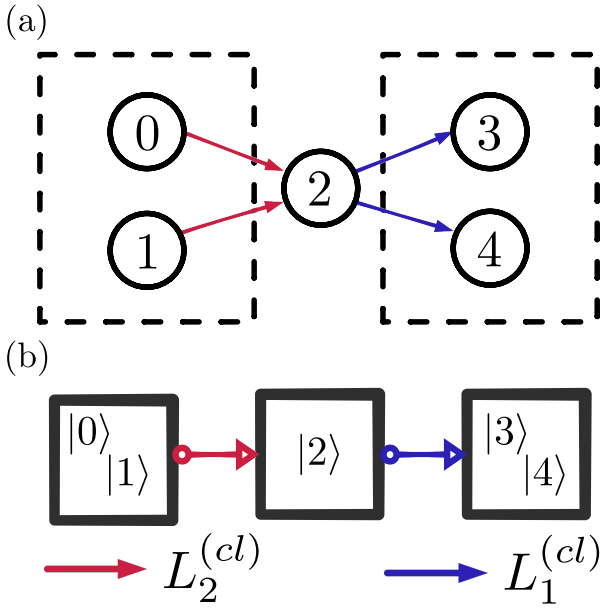

Figure 7.6: Semiclassical coarse-graining applied to a QSM. (a) A 5-state part of a QSM with two quantum subgraphs: interference subgraph realized by $L_2^{(cl)}$ and a superposition subgraph realized by $L_1^{(cl)}$. Since pairs of states $\{|0\rangle,\ |1\rangle\}$ and $\{|3\rangle,\ |4\rangle\}$ fall under conditions described in Sec. 7.5.3, these are to be grouped together in a coarse-grained QSM. (b) Simplified depiction of a coarse-grained QSM, obtained from (a).

linked blocks). Characterizing these effectively with single amplitudes may lead to erroneous navigation policies. In particular, one state internal to a block-vertex might be untouched by an inter-block edge, i.e., it only yields zero matrix elements in a matrix characterizing the edge. If the edge is outgoing, a system initialized in the said state would not be able to escape the block-vertex using that edge alone (see Fig. 7.7). This is in direct conflict with characterizing blocks and inter-block edges with single amplitudes. For an incoming edge, a similar problem may arise: some states inside a block-vertex might not get populated when that edge is activated. This may become detrimental for the navigation protocol based on a coarse-grained QSM, especially if the unavailable state in question is the final target of the protocol.

Such difficulties may be overcome, if some of the couplings given in

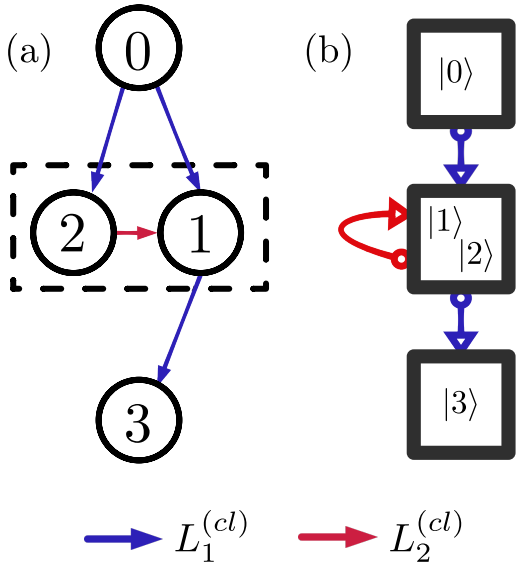

Figure 7.7: Illustration of ancillary couplings in the context of QSM coarse-graining (a) A 4-state part of a QSM that is subject to coarse-graining, featuring non-trivial actions by couplings denoted as V_1 and V_2. States $|1\rangle$ and $|2\rangle$ are to be grouped together since they are both targets in a superposition subgraph span by edges e_{02} and e_{01} of a steering graph $L_1^{(\mathrm{cl})}$. (b) The coarse-grained version of the QSM from (a). The block $\{|1\rangle, |2\rangle\}$ is connected to state $|3\rangle$ through an outgoing edge of $L_1^{(\mathrm{cl})}$. However, from microscopic point of view exemplified in (a), no population can be transferred from state $|2\rangle$ to $|3\rangle$ unless the click action $\Lambda_2^{(\mathrm{cl})}$ is realized first. Therefore, including and applying V_2 as an ancillary coupling is required for a valid semiclassical coarse-graining of this QSM.

a QSM allow for an internal mixing of the subspace (represented by a self-loop on the block-vertex in the respective $L^{(\mathrm{cl})}$-graph). Applying such a coupling in the protocol would allow one to make the block-vertex accessible to all the edges that are connected to it (see Fig. 7.7), via a sufficient number of clicks. In the scenarios described above, where additional couplings are needed to turn a block-vertex into an effective single vertex, we will refer to such couplings as ancillary couplings. Note that given a steering kit, there is no guarantee that the ancillary couplings needed for exploration of every block-vertex, are available. For simplicity, in this chapter, we restrict our further consideration to the coarse-grained

QSMs, where the ancillary couplings happen to be present wherever needed. Every block-vertex can then be made accessible to the outgoing edges, and the target state is ensured to be reachable once the target block is reached. In this case, we consider the coarse-grained QSM as effectively semiclassical.

To design an active steering policy within the coarse-grained approach, we note that the navigation protocol has the following structure. The system state can be transported between block-vertices, and eventually steered to the target block-vertex. After that, either the target state is reached already (one can obtain this information from the simulated copy of the system), or it can be reached after applying ancillary couplings on the target block-vertex. The cost of the protocol can now be broken into two parts. The first is the cost of exploring the coarse-grained graph using the inter-vertex edges. The second is the dwell time inside the block-vertices, which is spent applying the ancillary couplings. If we could find the route through the graph that minimizes the combination of these two components, it would solve our optimization problem exactly. However, because of the presence of the degrees of freedom that are internal to the block-vertices, the coarse-grained geometrical information does not allow for such a precise solution. In other words, both the inter-vertex travel time and the block-vertex dwell time depend on the microscopic details of the evolution.

Instead of studying such quantum-mechanical microscopics, we propose a semiclassical approximation to this calculation. Specifically, we assign every inter-block edge a characteristic traversal time, and every block-vertex a characteristic dwell time. For this, we use the matrices for click transitions between blocks i and j (the case of ancillary couplings given by $i = j$). Let us loosely denote these as $L^{(\mathrm{cl})}_{i,\alpha;j,\beta}$, implying that only matrix elements with states from blocks i and j are included. In that case, the effective transition amplitude between blocks i and j can be defined as operator norm $L^{(\mathrm{cl})}_{i,j} = \|L^{(\mathrm{cl})}_{i,\alpha;j,\beta}\|$, and characteristic traversal (dwell if $i = j$) time $\tau_{i,j} = (L^{(\mathrm{cl})}_{i,j})^{-2}$. Note that this reduces to the average traversal time for the case of a genuinely classical graph, with an amplitude $\gamma\delta t$ connecting two states implying duration of $\tau = \frac{1}{\gamma^2\delta t^2}$ for traversal (cf. Sec. 7.2.2).

With characteristic times $\tau_{i,j}$ assigned, the time-cost of following a specific path through this graph can be estimated as a combined characteristic time of all the edges and vertices crossed along the way. The desired path will be the one that optimizes this expected time. On the one hand, this

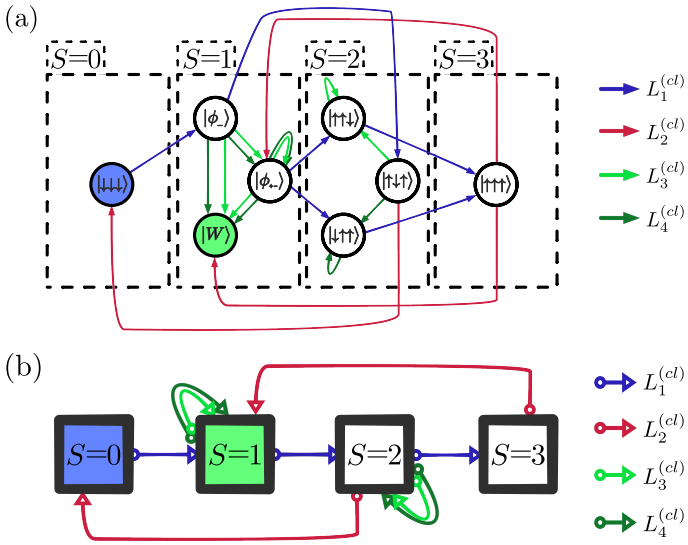

Figure 7.8: Measurement-driven navigation towards the 3-qubit W-state: QSM representation. (a) Steering with couplings Eqs. (7.5.5)-(7.5.8). The vertices in the single-excitation subspace are given by states $|W\rangle$, $|\phi_-\rangle \equiv \frac{1}{\sqrt{2}}(|100\rangle - |001\rangle)$, and $|\phi_{+-}\rangle \equiv \frac{1}{\sqrt{6}}(|100\rangle - 2|010\rangle + |001\rangle)$. (b) The coarse-grained version of the above QSM. The vertices are labeled by the excitation number. From perspective of Sec. 7.5.3, couplings 2 and 3 play the ancillary role. Indeed, those couplings mix the internal structure of the block-vertices, allowing one to eventually steer the state to the target $|W\rangle$.

may result in a different navigation protocol compared to what is optimal from the complete quantum-mechanical analysis. On the other hand, such a first-principles analysis is prohibitively hard, and we expect that our semiclassically derived protocol will still be considerably quicker than its completely passive version. One example of such an improved protocol is given below.

7.5.4 W-state preparation

To illustrate the principles of the QSM framework, we consider the coarse-graining approach to the navigation of a 3-qubit state from a trivial $|000\rangle$

state to a so-called W-state [212] that has the following form:

$$W = \frac{1}{\sqrt{3}}(|100\rangle + |010\rangle + |001\rangle). \tag{7.5.4}$$

For the steering kit, we choose the following family of couplings (assuming labels A, B, and C for the qubits):

$$V_1 = \sigma_A^+ - \sigma_C^+, \tag{7.5.5}$$
$$V_2 = \sigma_A^- \sigma_C^-, \tag{7.5.6}$$
$$V_3 = \sigma_A^- \sigma_B^+ - P_A^0 P_B^1, \tag{7.5.7}$$
$$V_4 = \sigma_B^+ \sigma_C^- - P_B^1 P_C^0. \tag{7.5.8}$$

Here, $\sigma^{\pm} = \frac{1}{2}(\sigma^x \pm i\sigma^y)$ and $P^a = |a\rangle\langle a|$, $a = 0, 1$. A passive version of the protocol would amount to blindly alternating between the steering actions with different V_i, which does yield the target state, given that the steering is applied a sufficient number of times (see Fig. 7.8.b).

To design a feedback policy, we now consider a QSM representation of the steering kit. It is shown in Fig. 7.8a. Note that this QSM has multiple quantum subgraphs. Therefore, to employ a feedback policy, it should be subjected to the subspace-clustering coarse-graining technique, as outlined in Sec. 7.5.3. It proves useful to cluster the Hilbert space by the total excitation number, which results in a semiclassical QSM, as desired (Fig. 7.8b). Given the starting state of the evolution, it is then straightforward to design the policy that leads to the target state:

1. Repeat V_1-steering until a click is obtained;
2. Repeat V_3-steering until the target state is reached (with fidelity error below ϵ).

This protocol moves the state of the system from the zero excitation state to the single-excitation subspace (part 1) and then takes the system to the W-state in that subspace (part 2). Note that this employs only two couplings out of four, a simplification that is only possible with an active steering protocol. In a passive protocol, as all couplings are employed cyclically, multiple V_1-clicks may accidentally occur before the target state is reached (Fig. 7.9a), and, therefore, other couplings are needed to reduce the excitation number back to 1. In Fig. 7.9b, the performance histogram is given for a large number of numerical trials for the active and passive protocols. Both are run with $\delta t = 0.1$ for the target fidelity error $\epsilon = 0.01$. We are primarily interested in the average runtimes of the active and the passive protocol. These are $N_{\text{act}} \simeq 365 \pm 3$ and $N_{\text{pas}} \simeq 4600 \pm 100$ cycles of the protocol, respectively, yielding a speedup factor of around 12.5.

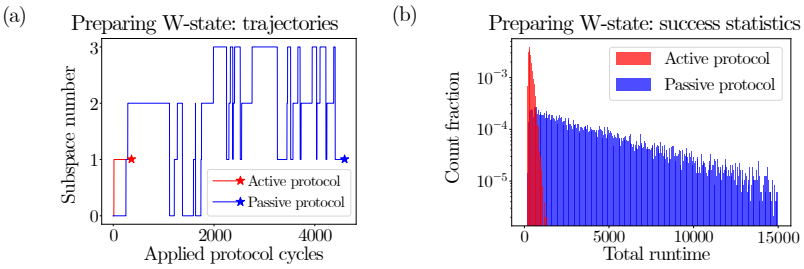

Figure 7.9: Performance of the active navigation protocol based on the QSM representation of steering towards the W-state, as described in the text. (a) Typical trajectories of the passive and active protocols, in terms of the excitation number sector that is occupied by the system state. Displayed are trajectories that yield the runtimes approximately equal to average runtimes of 365 (active) and 4600 (passive). The first V_1-click in the displayed run of the active protocol occurs as early as the 14th cycle, which is not legible from the plot. (b) The histogram over the protocol runtimes, required for a passive and an active protocol to achieve r=0.01 infidelity to the target state. Similarly to the Fig. 7.2 for AKLT target state, note the clear advantage for each recorded run of the active protocol compared to the mean duration $\simeq 4.6 \cdot 10^3$ of the passive protocol. Each histogram was obtained from 10^4 numerical simulations, and truncated at 15000 cycles for better presentation.

7.6 Discussion and conclusions

In this work, we have put forward the concept of measurement-driven active-decision steering of quantum states. We have developed steering protocols in which the measurement readouts are used to adjust the measurement protocol on-the-go, yielding significant acceleration of state preparation, with improved fidelity, compared with passive steering. The possibility of exploiting the readouts explored here is the great advantage of measurement-based steering over drive-and-dissipation (largely equivalent to "blind" steering) state preparation. While our approach has sweeping applicability, here we have chosen to focus on active measurement-driven steering as applied to the most challenging case of many-body quantum systems with entangled target states.

To satisfy physical (locality) constraints on system-detector couplings, we have proposed a scheme, based on parent Hamiltonian construction, for identifying feasible couplings. Employing such couplings, we have developed and analyzed Hilbert-space-orientation techniques for measurement-

driven steering. A central ingredient here has been to develop feedback policies based on detector readouts. One such Hilbert-space path-finding technique is based on a cost function, evaluating the running fidelity to the target state. We have shown a substantial (up to 9.5-fold) speedup of steering, employing this approach for preparation of the ground state of the AKLT model. A second protocol comprises mapping out the available measurement actions onto a Quantum State Machine (QSM), using a coarse-grained version of the corresponding graphs in Hilbert space. We have given an example of an entangled state preparation which shows acceleration by a factor of 12.5 compared to passive steering.

While we have limited ourselves here to a few examples, our schemes are of general applicability. They open the door to the design of efficient and high-quality state engineering, adiabatic state manipulation, and, possibly, quantum information processing. Moreover, steering protocols are subject to errors, both "static" (choice of steering parameters) and "dynamics" (noise). Active decision-making steering may be designed to reduce the effect of such errors. One may envision a host of directions to generalize and develop these ideas. For example, the greedy minimization of our cost function may be further improved by finding other metrics of local "steepest decent." Further, one may systematically investigate less local (less greedy) optimization of the cost function, e.g., looking n steps ahead. Another potential advantage of our protocols relies on the following observation: in the context of passive steering, one imposes constraints concerning locality (e.g., how many spins can be coupled to a local detector), and certain types of coupling terms. Given such constraints, not all target states are reachable. The introduction of active steering may overcome this handicap of target-state accessibility.

One may combine the dynamics incorporated here with the inherent unitary evolution of the system at hand (due to a system-only Hamiltonian). Consider the context of passive (blind) measurement-induced steering, which, in the continuum time limit, leads to Lindbladian dynamics. Then, the addition of Hamiltonian dynamics enriches the variability of steering, allowing, for example, to obtain mixed states by design [213]. It is intriguing to investigate how the addition of Hamiltonian dynamics extends or improves active steering, thus marrying the frameworks of closed-loop quantum control for Hamiltonian-based state preparation and active-decision measurement-based steering.

Further extensions of our approach include applications of QSM protocols to larger and more complex systems, going beyond a three-qubit setup. Optimizing such protocols may involve automatization of the creation and analysis of QSMs, e.g., for finding an optimal basis automatically, in

similarity with quantum annealing, but now at the level of measurement operators. Finally, one may envision using machine learning to find more optimized navigation protocols (see [214, 215] for related work in the context of Hamiltonian feedback and open-loop control). Given the delayed-reward setting at hand, a reinforcement learning strategy such as Q-learning [216] or SARSA [217] might be the most appropriate choice.