

Reachability Analysis of Quantum Markov Decision Processes

Shenggang Ying and Mingsheng Ying

Tsinghua University, China and University of Technology, Sydney, Australia
yingshenggang@gmail.com, Mingsheng.Ying@uts.edu.au, yingmsh@tsinghua.edu.cn

Abstract

We introduce the notion of quantum Markov decision process (qMDP) as a semantic model of nondeterministic and concurrent quantum programs. It is shown by examples that qMDPs can be used in analysis of quantum algorithms and protocols. We study various reachability problems of qMDPs both for the finite-horizon and for the infinite-horizon. The (un)decidability and complexity of these problems are settled, or their relationships with certain long-standing open problems are clarified. We also develop an algorithm for finding optimal scheduler that attains the supremum reachability probability.

Categories and Subject Descriptors F.1.1 [Computation by Abstract Devices]: Models of Computation; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages - Program Analysis

General Terms Algorithms, Theory, Verification

Keywords Quantum programming, semantic model, Markov decision process, reachability

1. Introduction

As a generalisation of Markov chains, Markov decision processes (MDPs) stemmed from operations research in 1950's. Now they have been successfully applied in various areas such as economics and finance, manufacturing, control theory, robotics, artificial intelligence and machine learning. Also, effective analysis and resolution techniques for MDPs like linear programming have been developed in the last six decades. Since Vardi [39] proposed to adopt MDPs as a model of concurrent probabilistic programs, MDPs have been widely used in analysis and verification of randomised algorithms and probabilistic programs (see, for instance, [24]) as well as model checking of probabilistic computing systems [2].

In this paper we introduce the notion of quantum Markov decision process (qMDP) as a model of nondeterministic and concurrent quantum programs. Research on quantum programming has been intensively conducted in the last 18 years since Knill [21] introduced the Quantum Random Access Machine model for quantum computing and proposed a set of conventions for writing quantum pseudocode. The research includes design of quantum programming languages, e.g. QCL [26], qGCL [32], QPL [33] and

Quipper [14], semantic models of quantum programs [10], and verification of quantum programs [40] (we refer the reader to [11] for basic ideas of quantum programming and an excellent survey on the early works in this area). In particular, quantum Markov chains were defined in [41, 42] for modelling sequential quantum programs. This paper extends quantum Markov chains considered in [41, 42] to qMDPs so that we can model nondeterministic and concurrent quantum programs [43, 46].

A classical MDP consists of a set S of states and a set Act of actions. Each action $\alpha \in Act$ is modelled by a probabilistic transition function $P_\alpha : S \times S \rightarrow [0, 1]$ with $P_\alpha(s, s')$ being the probability that the system moves from state s to s' after action α . A MDP allows not only probabilistic choice between the system states as a result of performing an action but also a nondeterministic choice between actions: there may be more than one action enabled on entering a state s . Thus, the notion of scheduler was introduced to resolve the nondeterministic choice between the enabled actions. A scheduler selects the next action according to the previous and current states of the system. A qMDP is defined as quantum generalisation of MDP with the set S of states replaced by a Hilbert space \mathcal{H} which always serves as the state space of a quantum system in physics. Now each action $\alpha \in Act$ is described by a super-operator \mathcal{E}_α in \mathcal{H} . Super-operators were recognised by physicists as the most general mathematical formalism of physically realisable operations in quantum mechanics [25]. They were also adopted as denotational semantics of quantum programs by Selinger [33] and D'Hont and Panangaden [10] in their pioneering works on quantum programming.

A major conceptual difference between classical MDPs and qMDPs comes from the notion of scheduler. The information used by a scheduler in a MDP to select the next action is the state of the system. In the quantum case, however, we choose to introduce a series of measurements at the middle of the evolution of a qMDP and to define a scheduler as a function that selects the next action according to the outcomes of these measurements.

This paper focuses on the aspect of qMDPs more related to program analysis and verification, namely reachability analysis. As in the case of classical MDPs, we consider the reachability probability of a subspace B of the state Hilbert space of a qMDP with a fixed scheduler and the supremum reachability probability of B over all schedulers. Although the definition of reachability probabilities in qMDPs looks similar to that of classical MDPs, their behaviours are very different; for example, a MDP has an optimal scheduler that can achieve the supremum reachability probability for all initial states. But it is not the case in a qMDP even for a given initial state. It is also interesting to observe the difference between the behaviour of qMDPs and that of quantum Markov chains. It was proved in [42] that a quantum Markov chain eventually reaches a subspace B for any initial state if the ortho-complement B^\perp of B in the state Hilbert space \mathcal{H} contains no bottom strongly connected components (BSCCs). The corresponding notion of BSCC in a qMDP is invariant subspace. However, it is possible that in

a qMDP B^\perp contains no invariant subspaces but for some schedulers, B is reached by a probability smaller than 1.

As indicated in Subsection 2.7, some problems in the analysis of quantum algorithms can be properly formulated as the reachability problem of qMDPs. We believe that it will be inevitable to develop effective techniques for reachability analysis of qMDPs with applications in quantum program analysis and verification as quantum algorithm and program design become more and more sophisticated.

The aspects of qMDPs more related to decision making and machine learning are left for future research. In the last few years, it has been found that probabilistic programming is very useful in machine learning for describing probabilistic distributions and Bayesian inference (see, for instance, [13]). On the other hand, it was realised recently that a major application area of quantum computing might be machine learning and big data analytics. We expect that qMDPs will serve as a bridge between the researches on quantum programming and quantum machine learning.

Contribution of the paper: This paper studies (un)decidability and complexity of reachability analysis for qMDPs. In the case of finite-horizon, it is proved that both quantitative reachability and qualitative reachability of qMDPs are undecidable. In the case of infinite-horizon, we show that it is EXPTIME-hard to decide whether the supremum reachability probability of a qMDP is 1, and if it is smaller than 1, then the supremum reachability probability is uncomputable. It is further proved that a qMDP has an optimal scheduler for reaching an invariant subspace of its state Hilbert space if and only if the ortho-complement of the target subspace contains no invariant subspaces. This result enables us to develop an algorithm for finding an optimal scheduler. We also consider the problem whether a qMDP always reach an invariant subspace with probability 1, no matter what the scheduler is. A connection between this problem and a long-standing open problem - the joint spectral radius problem [8, 16, 37] - is observed.

Related work: Before this paper, a very interesting paper by Barry, Barry and Aaronson [3] was recently posted at <http://arxiv.org/abs/1406.2858> where the notion of quantum partially observable Markov decision process was introduced. It was proved in [3] that reachability of a goal state is undecidable in the quantum case but decidable in the classical case. The undecidability in the quantum case is similar to our Theorem 3.2, but they are not the same since we consider reachability of invariant subspaces rather than a single state. Other results in [3] and ours are unrelated.

Organisation of the paper: The rest of this paper is organised as follows. Section 2 gives formal definitions of qMDPs and their reachability probabilities and invariant subspaces. It also presents several examples to illustrate how can quantum algorithms and protocols be modelled as qMDPs and to show some essential differences between qMDPs and classical MDPs as well as quantum Markov chains. All main results obtained in the paper are stated in Section 3. Sections 4 and 5 are devoted to prove the results for finite-horizon and infinite-horizon, respectively. A brief conclusion is drawn in Section 6.

2. Definitions and Examples

2.1 Basics of Quantum Theory

For convenience of the reader, we very briefly recall some basic notions in quantum theory with the main aim being fixing notations; see [25] for details. In this paper we always assume that the state Hilbert space is d -dimensional, i.e. $\mathcal{H} = \mathbb{C}^d$ where \mathbb{C} is the field of complex numbers. We use the Dirac notation and assume that $\{|i\rangle\}_{i=1}^d$ is an orthonormal basis of \mathcal{H} . Then we have $\mathcal{H} = \text{span}\{|i\rangle\}$, a pure state in \mathcal{H} can be written as $|\psi\rangle = \sum \alpha_i |i\rangle$ with $\sum |\alpha_i|^2 = 1$, and a mixed state is represented by a den-

sity matrix in \mathcal{H} , i.e. a semi-definite positive $d \times d$ matrix with trace 1. Write $\mathcal{D}(\mathcal{H})$ for the set of all density matrices in \mathcal{H} . The identity matrix is denoted I . If a density matrix can be written as $\rho = \sum p_i |\psi_i\rangle\langle\psi_i|$, where $\langle\psi_i|$ stands for the transpose conjugate of $|\psi_i\rangle$, then its support is $\text{supp}(\rho) = \text{span}\{|\psi_i\rangle : p_i > 0\}$.

The evolution of a closed quantum system is described by a $d \times d$ unitary matrix: $|\phi\rangle \mapsto U|\phi\rangle$. A super-operator $\mathcal{E} : \mathcal{D}(\mathcal{H}) \rightarrow \mathcal{D}(\mathcal{H})$ depicts the dynamics of a system which is realised with noise or interacts with its environment, and it can always be represented by $\mathcal{E}(\rho) = \sum E_i \rho E_i^\dagger$ where all E_i are $d \times d$ matrices with $\sum E_i^\dagger E_i = I$ and E_i^\dagger denotes the conjugate transpose of E_i . The $d^2 \times d^2$ matrix $M = \sum (E_i \otimes E_i^*)$ is called the matrix representation of \mathcal{E} .

A quantum measurement in \mathcal{H} is described by a set of $d \times d$ matrices $M = \{M_{m_1}, \dots, M_{m_k}\}$ with $\sum M_{m_i}^\dagger M_{m_i} = I$, where m_i 's denote the possible outcomes. If we perform measurement M on a quantum system which is currently in state ρ , then the probability that we get outcome m_i is $p_i = \text{tr}(M_{m_i}^\dagger M_{m_i} \rho)$ and the system's after-measurement state is $\rho_i = M_{m_i} \rho M_{m_i}^\dagger / p_i$ whenever the outcome is m_i . A measurement $P = \{P_{m_1}, \dots, P_{m_k}\}$ is projective if $P_{m_i} P_{m_j} = \delta_{ij} P_{m_i}$.

2.2 Quantum Markov Decision Processes

In this subsection, we formally define our notions of qMDPs and their schedulers.

Definition 2.1. A qMDP is a 4-tuple $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M}, Q \rangle$, where:

- \mathcal{H} is a d -dimensional Hilbert space, called the state space. The dimension of \mathcal{H} is also called the dimension of \mathcal{M} , i.e. $\dim \mathcal{M} = \dim \mathcal{H} = d$.
- Act is a finite set of action names. For each $\alpha \in \text{Act}$, there is a corresponding super-operator \mathcal{E}_α that is used to describe the evolution of the system caused by action α .
- \mathbf{M} is a finite set of quantum measurements. We write Ω for the set of all possible observations; that is,

$$\Omega = \{O_{M,m} : M \in \mathbf{M} \text{ and } m \text{ is a possible outcome of } M\}.$$

Intuitively, $O_{M,m}$ indicates that we perform the measurement M on the system and obtain the outcome m .

- $Q : \text{Act} \cup \mathbf{M} \rightarrow 2^{\text{Act} \cup \mathbf{M}}$ is a mapping. For each $\alpha \in \text{Act}$ (or $M \in \mathbf{M}$), $Q(\alpha)$ (resp. $Q(M)$) stands for the set of the available actions or measurements after α (resp. M) is performed. For the trivial case that $Q(\alpha) = \text{Act} \cup \mathbf{M}$ for all α , Q will be omitted, and the qMDP \mathcal{M} will be simply written as a triple $\langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$.

Definition 2.2. A scheduler for a qMDP \mathcal{M} is a function

$$\mathfrak{S} : (\text{Act} \cup \Omega)^* \rightarrow \text{Act} \cup \mathbf{M}.$$

For any sequence $\sigma = \alpha_1 \dots \alpha_n \in (\text{Act} \cup \Omega)^*$, $\mathfrak{S}(\sigma)$ indicates the next action or measurement after actions or observations $\alpha_1 \dots \alpha_n$ happen.

As pointed out in the introduction, a scheduler in a qMDP selects the next action based on the outcomes of performed measurements. Actually, in the above definition the performed actions are also recorded as a part of the information for such a selection. This design decision is motivated by several examples in Subsection 2.7. We now describe the evolution of a qMDP \mathcal{M} with an initial state $\rho \in \mathcal{D}(\mathcal{H})$ and a scheduler \mathfrak{S} . For simplicity, we write $W = (\text{Act} \cup \Omega)^*$. For each word $w \in W$, the state $\rho_w^\mathfrak{S}$ of the qMDP \mathcal{M} and probability $p_w^\mathfrak{S}$ that this state is reached in \mathcal{M} after sequence w of actions or observations are defined by induction on the length of w :

- $\rho_\epsilon^\mathfrak{S} = \rho$ and $p_\epsilon^\mathfrak{S} = 1$, where ϵ is the empty word.
- If $\mathfrak{S}(w) = \alpha \in Act$, then $\rho_{w\alpha}^\mathfrak{S} = \mathcal{E}_\alpha(\rho_w^\mathfrak{S})$ and $p_{w\alpha}^\mathfrak{S} = p_w^\mathfrak{S}$. (Note that all the super-operators \mathcal{E}_α ($\alpha \in Act$) are assumed to be trace-preserving.)
- If $\mathfrak{S}(w) = \alpha = O_{M,m} \in \Omega$, then

$$\rho_{w\alpha}^\mathfrak{S} = M_m \rho_w^\mathfrak{S} M_m^\dagger / \text{tr}(M_m \rho_w^\mathfrak{S} M_m^\dagger)$$

$$\text{and } p_{w\alpha}^\mathfrak{S} = p_w^\mathfrak{S} \cdot \text{tr}(M_m \rho_w^\mathfrak{S} M_m^\dagger).$$

Furthermore, for each $n \geq 0$, we can define the global state of the qMDP \mathcal{M} at step n according to scheduler \mathfrak{S} by

$$\rho(n, \mathfrak{S}) = \sum_{w \in W \text{ s.t. } |w|=n} p_w^\mathfrak{S} \rho_w^\mathfrak{S}.$$

For a subspace B of \mathcal{H} , the probability that B is reached at step n with initial state ρ and scheduler \mathfrak{S} is defined by

$$\Pr(\rho(n, \mathfrak{S}) \models B) = \text{tr}(P_B \rho(n, \mathfrak{S})) \quad (1)$$

where P_B is the projection onto B .

2.3 Invariant Subspaces

A key notion used in reachability analysis of quantum Markov chains [42] is BSCC. A counterpart of BSCC in qMDPs is the notion of (common) invariant subspace. Let B be a subspace of Hilbert space \mathcal{H} . We say that B is invariant under a super-operator \mathcal{E} if $\text{supp}(\mathcal{E}(\rho)) \subseteq B$ for all density matrices ρ with $\text{supp}(\rho) \subseteq B$. Moreover, B is invariant under a measurement $M = \{M_1, \dots, M_k\}$ if $\text{supp}(M_i \rho M_i^\dagger) \subseteq B$ for all $1 \leq i \leq k$ and all ρ with $\text{supp}(\rho) \subseteq B$.

Definition 2.3. Let $\mathcal{M} = \langle \mathcal{H}, Act, \mathbf{M}, Q \rangle$ be a qMDP and B a subspace of \mathcal{H} . If B is invariant under super-operator \mathcal{E}_α for all $\alpha \in Act$, and it is invariant under all measurement $M \in \mathbf{M}$, then B is called an invariant subspace of \mathcal{M} .

The probability that an invariant subspace is reached is a non-decreasing function of the number of steps.

Theorem 2.1. Let \mathcal{M} be a qMDP with initial state ρ and B an invariant subspace of \mathcal{M} . Then for any scheduler \mathfrak{S} and $n \geq 0$, we have:

$$\Pr(\rho(n+1, \mathfrak{S}) \models B) \geq \Pr(\rho(n, \mathfrak{S}) \models B).$$

Proof. Induction on n by using Theorem 1 in [42]. \square

2.4 Reachability Probability

The reachability probability of finite-horizon was defined in equation (1). Now we define the reachability probability of infinite-horizon.

Definition 2.4. Let \mathcal{M} be a qMDP with state Hilbert space \mathcal{H} , ρ an initial state, \mathfrak{S} a scheduler for \mathcal{M} and B a subspace of \mathcal{H} . Then reachability probability of B in \mathcal{M} starting in ρ with scheduler \mathfrak{S} is defined by

$$\Pr^\mathfrak{S}(\rho \models \diamond B) = \lim_{n \rightarrow \infty} \Pr(\rho(n, \mathfrak{S}) \models B). \quad (2)$$

It is worth noting that, in general, the limit in the above equation does not necessarily exist. However, we have:

Lemma 2.1. If B is an invariant subspace of \mathcal{M} , then for any initial state ρ and any scheduler \mathfrak{S} , the reachability probability $\Pr^\mathfrak{S}(\rho \models \diamond B)$ always exists.

Proof. Since $\Pr(\rho(n, \mathfrak{S}) \models B)$ is bounded by 1, the conclusion follows immediately from Theorem 2.1. \square

Definition 2.5. Let \mathcal{M} be a qMDP with state Hilbert space \mathcal{H} , ρ an initial state and B a subspace of \mathcal{H} . Then supremum reachability probability of B in \mathcal{M} starting in ρ is defined by

$$\Pr^{\text{sup}}(\rho \models \diamond B) = \sup_{\mathfrak{S}} \Pr^\mathfrak{S}(\rho \models \diamond B). \quad (3)$$

If scheduler \mathfrak{S}_0 satisfies that $\Pr^{\mathfrak{S}_0}(\rho \models \diamond B) = \Pr^{\text{sup}}(\rho \models \diamond B)$, then \mathfrak{S}_0 is called the optimal scheduler for the initial state ρ .

2.5 A Difference between Classical and Quantum Markov Decision Processes

It is well-known [2, Lemma 10.102] that there exists a memoryless scheduler \mathfrak{S}_0 that is optimal for all initial states. In the quantum case, however, it is possible that no optimal scheduler exists even for a fixed initial state.

Example 2.1. Consider a quantum Markov decision process $\mathcal{M} = \langle \mathcal{H}, Act, \mathbf{M} \rangle$, where $\mathcal{H} = \text{span}\{|1\rangle, |2\rangle, |3\rangle, |4\rangle\}$, $\mathbf{M} = \emptyset$, $Act = \{\alpha, \beta\}$ and

$$\mathcal{E}_\alpha(\rho) = (|2\rangle\langle 1|\rho|1\rangle\langle 2| + |1\rangle\langle 1|\rho|1\rangle\langle 1|)/2 + |2\rangle\langle 2|\rho|2\rangle\langle 2| + |3\rangle\langle 3|\rho|3\rangle\langle 3| + |4\rangle\langle 4|\rho|4\rangle\langle 4|,$$

$$\mathcal{E}_\beta(\rho) = |4\rangle\langle 1|\rho|1\rangle\langle 4| + |3\rangle\langle 2|\rho|2\rangle\langle 3| + |3\rangle\langle 3|\rho|3\rangle\langle 3| + |4\rangle\langle 4|\rho|4\rangle\langle 4|.$$

Let $\rho_0 = |1\rangle\langle 1|$ and $B = \text{span}\{|3\rangle\}$. Then

$$\Pr^{\mathcal{P}}(\rho_0 \models \diamond B) < \sup_{\mathfrak{S}} \Pr^\mathfrak{S}(\rho_0 \models \diamond B) = 1$$

for all schedulers \mathcal{P} . Indeed, if $\mathcal{P} = \alpha^\omega$, then $\Pr^\mathfrak{S}(\rho_0 \models \diamond B) = 0$. Let $\mathcal{P} \neq \alpha^\omega$ be a scheduler and let k be the first index such that $a_k = \beta$ where $\mathcal{P} = a_1 a_2 \dots$. Then $\Pr^{\mathcal{P}}(\rho_0 \models \diamond B) = 1 - 0.5^{k-1} < 1$.

One reason for nonexistence of the optimal scheduler is that the current state of a quantum system usually cannot be known exactly from the outside, and thus we often have no enough information to choose the next action in a scheduler for a qMDP. In the above example, whence we know the exact state of the system, we can choose an appropriate action to reach the target state: if the state is $|1\rangle$, we take α , and if the state is $|2\rangle$, we take β . However, consider the case where the first action is α . The state of the system will become $\rho_1 = (|1\rangle\langle 1| + |2\rangle\langle 2|)/2$. Then we do not know it is in $|1\rangle$ or $|2\rangle$ exactly, and we cannot decide which action should be taken.

However, the above is not the only reason for nonexistence of the optimal scheduler. As shown in the following example, it is still possible that a qMDP has no the optimal scheduler when we know exactly its state.

Example 2.2. Let $\mathcal{M} = \langle \mathcal{H}, Act, \mathbf{M} \rangle$ be a qMDP, $\rho_0 = |1\rangle\langle 1|$ an initial state and $B = \text{span}\{|4\rangle\}$, where

- $\mathcal{H} = \text{span}\{|1\rangle, |2\rangle, |3\rangle, |4\rangle\}$;
- $Act = \{a, b\}$ and $\mathbf{M} = \emptyset$;
- $\mathcal{E}_a = A_1 \cdot A_1^\dagger + A_2 \cdot A_2^\dagger + A_3 \cdot A_3^\dagger$, where

$$A_1 = \begin{pmatrix} \cos \theta & \sin \theta & & \\ -\sin \theta & \cos \theta & & \\ & & 0 & \\ & & & 0 \end{pmatrix},$$

$$\theta = 0.6, A_2 = |3\rangle\langle 3| \text{ and } A_3 = |4\rangle\langle 4|;$$

- $\mathcal{E}_b = \sum_{i=1}^4 C_i \cdot C_i^\dagger$, where $C_1 = |3\rangle\langle 1|$, $C_2 = |4\rangle\langle 2|$, $C_3 = |3\rangle\langle 3|$, $C_4 = |4\rangle\langle 4|$.

Since $\theta = 0.6$, the set $\{A_1^n |1\rangle : n \in \mathbb{N}\}$ is dense on the circle $\{a|1\rangle + b|2\rangle : a, b \in \mathbb{R}, a^2 + b^2 = 1\}$. For any $\epsilon > 0$, there exists n , such that $\mathcal{E}_a^n(|1\rangle\langle 1|) = |\psi_n\rangle\langle \psi_n|$ with $|\langle 2|\psi_n\rangle| > \sqrt{1-\epsilon}$.

Thus $\Pr^{\mathfrak{S}}(\rho_0 \models B) > 1 - \epsilon$ for $\mathfrak{S} = (a^n b)^\omega$. This leads to $\Pr^{\text{sup}}(\rho_0 \models B) = 1$. But since $A_1^n|1\rangle \neq |2\rangle$ for any m , there is no optimal scheduler.

In the above example, we have complete information about the state of the system after \mathcal{E}_a : it is always a superposition $a|1\rangle + b|2\rangle$ of $|1\rangle, |2\rangle$. But this does not help to derive an optimal scheduler because only $|2\rangle$ can reach the supremum 1.

2.6 A Difference between Quantum Markov Chains and Decision Processes

It was shown in [42] that a quantum Markov chain will eventually reach a subspace B for any initial state if there is no BSCC contained in the ortho-complement B^\perp . The following question asks whether a similar conclusion holds for qMDPs.

Problem 2.1. Let \mathcal{M} be a qMDP with state space \mathcal{H} , \mathfrak{S} a given scheduler for \mathcal{M} and B a subspace of \mathcal{H} . Suppose that \mathcal{M} has no invariant subspace contained in B^\perp . Will \mathcal{M} reach B eventually, i.e. $\Pr^{\mathfrak{S}}(\rho \models B) = 1$ for all initial states ρ ?

This question is negatively answered by the following example.

Example 2.3. Let $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ with $\mathcal{H} = \text{span}\{|1\rangle, |2\rangle, |3\rangle\}$, $\text{Act} = \{a, b\}$ and $\mathbf{M} = \emptyset$. The super-operators corresponding to a and b are defined as follows:

$$\mathcal{E}_a(\rho) = |3\rangle\langle 1|\rho|1\rangle\langle 3| + |1\rangle\langle 2|\rho|2\rangle\langle 1| + |3\rangle\langle 3|\rho|3\rangle\langle 3|,$$

$$\mathcal{E}_b(\rho) = |2\rangle\langle 1|\rho|1\rangle\langle 2| + |3\rangle\langle 2|\rho|2\rangle\langle 3| + |3\rangle\langle 3|\rho|3\rangle\langle 3|$$

for any density operator ρ . Let $B = \text{span}\{|3\rangle\}$. It is easy to see that \mathcal{E}_a and \mathcal{E}_b have no common invariant subspace in B^\perp . We consider initial state $\rho_0 = (|1\rangle\langle 1| + |2\rangle\langle 2|)/2$ and two schedulers $\mathfrak{S}_1 = (ab)^\omega$ and $\mathfrak{S}_2 = (ab)^k aa(ab)^\omega$ for some k . Then we have $\Pr^{\mathfrak{S}_1}(\rho_0 \models B) = 1/2$, but $\Pr^{\mathfrak{S}_2}(\rho_0 \models B) = 1$.

2.7 Quantum Algorithms and Protocols as qMDPs

In this subsection, we show how can the existing quantum algorithms and communication protocols be seen as examples of qMDP by analysing their structures. The early quantum algorithms and protocols can be roughly classified into three classes:

1. The first class applies a sequence of unitary operators followed by a measurement. If the outcome of measurement is desirable, the algorithm terminates. Otherwise, the algorithm is re-initialized and executed again; see Figure 1(a). Examples include the famous quantum order-finding and factoring algorithms [25], the Grover search algorithm [15], several quantum-walk-based algorithms [9, 20, 34] and the algorithm for solving the expectation value of some operators of systems of linear equations [18].
2. The second class repeatedly applies an action-measurement loop until success; see Figure 1(b). One example is the routing algorithm based on a many-measurement quantum walk in [20].
3. The structure of the third class looks like a decision tree; see Figure 2. Examples are quantum teleportation [25], one-way quantum computer [29]. These examples always terminate.

Recently, several algorithms have been developed with the structures different from Figures 1 and 2. For example, a modified quantum factoring algorithm was experimentally realised in [23], where in order to reduce the number of necessary entangled qubits, the ancilla (control) qubits are recycled. The structure of this algorithm is shown in Figure 3. Another example is the quantum Metropolis sampling [36]. This algorithm can be used to prepare the ground or thermal state of a quantum system. The structure of this algorithm for reaching the ground state is shown in Figure




Figure 1. \mathcal{E}_U represents the one or several sequential unitary operators. ρ_0 is the initial state. \mathcal{E}_{init} represents re-initializing, i.e., restarting the algorithm. \mathcal{E}_{skip} means maintaining the result for further application. M represents measurements with observation o_T standing for success and o_F for failure.




Figure 2. Structure of quantum decision trees.

4. It consists of decisions dependent on the history of actions and measurement outcomes as well as repeated loops until success.

As indicated by Figures 1-4, all of the algorithms and protocols mentioned above can be seen as qMDPs. Here we only elaborate the qMDP model of quantum Metropolis sampling.




Figure 3. Structure of modified quantum factoring algorithm.




Figure 4. Structure of quantum Metropolis sampling in [36].

Example 2.4. The qMDP $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ for the quantum Metropolis algorithm [36] is defined as follows:

- The state Hilbert space is the tensor product of five spaces, $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_{E1} \otimes \mathcal{H}_{E2} \otimes \mathcal{H}_a \otimes \mathcal{H}_c$, where
 1. \mathcal{H}_S is the Hilbert space of the original system, whose ground state is the target.
 2. \mathcal{H}_{E1} and \mathcal{H}_{E2} are ancilla spaces, used to represent the energies of the states in \mathcal{H}_S , where \mathcal{H}_{E1} represents the energy before updating in each round and \mathcal{H}_{E2} represents the new energy after updating.
 3. \mathcal{H}_a is 2-dimensional with its basis states represent the success or failure of eigenstate updating.
 4. \mathcal{H}_c is used to implement the probabilistic choice of actions C .
- Act consists of actions in the form of \mathcal{E}_* in Figure 4, where \mathcal{E}_{RC} stands for probabilistic choice of unitary operators C in [36].
- \mathbf{M} consists of measurements in the form of M_* in Figure 4. Ω is the set of observations.

The task of the algorithm is actually to find a scheduler that reaches the ground state in this qMDP. One such scheduler is illustrated in Figure 4.

Various generalisations and variants of quantum Metropolis sampling have been proposed, e.g. quantum rejection sampling [27], quantum-quantum Metropolis sampling [44] and complementing quantum Metropolis algorithm [30]. An experiment for preparing thermal states was realised [45] by employing some ideas from quantum Metropolis sampling. The correctness of quantum Metropolis algorithm and its variants can actually be seen as a reachability problem for qMDPs. This motivates us to systematically develop techniques for reachability analysis of qMDPs.

2.8 A Concurrent Quantum Program

As one more example of qMDP, we consider a simple concurrent quantum program consisting of n processes. Every process is a quantum loop. We assume a yes/no measurement $M = \{P_0, P_1\}$ in the state Hilbert space \mathcal{H} , which is projective; that is, both P_0 and P_1 are projections. For each $1 \leq i \leq n$, the i th process behaves as follows: it performs the measurement M , if the outcome is 0, then it executes a unitary transformation U_i and enter the loop again; if the outcome is 1 then it terminates. Note that the loop guard (termination condition) of the n processes are the same, but their loop bodies, namely unitary transformations U_i , are different.

This concurrent quantum program can be modelled as a qMDP \mathcal{M} with $\text{Act} = \{1, 2, \dots, n\}$. For each $i \in \text{Act}$, the action super-operator \mathcal{E}_i is defined by

$$\mathcal{E}_i(\rho) = P_1 \rho P_1 + U_i P_0 \rho P_0 U_i^\dagger$$

for all density matrices ρ . If P_1 is the projection onto the subspace B of \mathcal{H} , then the overall termination probability of the concurrent program with initial state ρ is the supremum reachability $\Pr^{\text{sup}}(\rho \models \diamond B)$. The following proposition provides us with a method for computing this termination probability. We write $\bar{\mathcal{E}}$ for the average super-operator of \mathcal{E}_i ($1 \leq i \leq n$); that is,

$$\bar{\mathcal{E}} = \frac{1}{n} \sum_{i=1}^n \mathcal{E}_i.$$

We further define

$$\mathcal{E}_\infty = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \bar{\mathcal{E}}^i.$$

(It was shown in [41] that \mathcal{E}_∞ can be computed by Jordan decomposition of the matrix representation of $\bar{\mathcal{E}}$.)

Proposition 2.1. 1. The overall termination probability

$$\Pr^{\text{sup}}(\rho \models \diamond B) = 1 - \text{tr}(\rho P_C),$$

where $C = \text{supp}(\mathcal{E}_\infty(I_{\mathcal{H}}))$ and P_C is the projection onto C .

2. There is a string $s \in \text{Act}^*$ such that the scheduler $\mathfrak{S} = s^\omega$ can attain the overall termination probability; that is,

$$\Pr^{\mathfrak{S}}(\rho \models \diamond B) = \Pr^{\text{sup}}(\rho \models \diamond B).$$

Proof. Let Y_a be an invariant subspace included in B^\perp of \mathcal{E}_a . Since $Y_a \perp B$, we have $Y_a \supseteq \mathcal{E}_a(Y_a) = U_a Y_a$. As $\dim Y_a = \dim(U_a Y_a)$, we have $Y_a = U_a Y_a$. Since unitary operators preserves the orthogonality, we have $\text{tr}(\rho P_{Y_a}) = \text{tr}(\mathcal{E}_a(\rho) P_{Y_a})$. If we write $C = \text{supp}(\mathcal{E}_\infty(I_{\mathcal{H}}))$, then C is invariant by definition and we have $\Pr^{\mathfrak{S}'}(\rho \models \diamond C) = \text{tr}(\rho P_C)$ for any scheduler \mathfrak{S}' . By Theorem 3.6 below, there exists $\mathfrak{S} = s^\omega$ such that $\Pr^{\mathfrak{S}}(\rho \models \diamond B \cup C) = 1$. So,

$$\Pr^{\mathfrak{S}}(\rho \models \diamond B) = \Pr^{\text{sup}}(\rho \models \diamond B) = 1 - \text{tr}(\rho P_C). \quad \square$$

3. Statement of Main Results

The aim of this paper is to study decidability and complexity of reachability analysis for qMDPs. For readability, we summarise the main results in this section but postpone their proofs to the sequent sections.

3.1 Results for the Finite-Horizon

We first examine the case of finite-horizon and consider the following:

Problem 3.1. Given a qMDP \mathcal{M} , an initial state ρ , a subspace B of \mathcal{H} and $0 \leq p \leq 1$, are there a scheduler \mathfrak{S} and a non-negative integer n such that

$$\Pr(\rho(n, \mathfrak{S}) \models B) \Delta p$$

where $\Delta \in \{>, \geq, <, \leq\}$?

Theorem 3.1. Problem 3.1 is undecidable for any Δ .

Now let us consider a qualitative variant of Problem 3.1.

Problem 3.2. Given a qMDP \mathcal{M} with the state Hilbert space \mathcal{H} and a subspace B of \mathcal{H} .

1. Are there a scheduler \mathfrak{S} and an integer n such that $\text{supp}(\rho(n, \mathfrak{S})) \subseteq B$ for all initial states ρ ?
2. Are there a scheduler \mathfrak{S} and an integer n such that $\text{supp}(\rho_0(n, \mathfrak{S})) \subseteq B$ for a given initial state ρ_0 ?

The counterpart of Problem 3.2.2 for classical MDPs can be stated as follows: given a MDP \mathcal{M} with a finite set S of states, an initial state s_0 and $B \subseteq S$, decide whether there exists a scheduler \mathfrak{S} and an integer N such that for any possible sequence of states $s_0 s_1 s_2 \dots$ under \mathfrak{S} , there exists $j < N$ such that $s_j \in B$. The polynomial-time decidability of this problem immediately follows the fact that an optimal scheduler for maximum reachability problem of a MDP can be found in polynomial time [2]. The only thing we need to do is to check whether there exists a cycle in all states reachable from s_0 in $S \setminus B$, if $\Pr(s_0 \models \diamond B) = 1$. The same result is true for the counterpart of Problem 3.2.1 for classical MDPs. This idea also applies to partially observable MDPs with a technique for reducing them to MDPs [1].

However, undecidability of Problem 3.2 was proved in [4] if subspace B is allowed to be not invariant. We prove undecidability of the problem for invariant subspace B and thus significantly improve the main result of [4].

Theorem 3.2. Both Problems 3.2.1 and 3.2.2 with $|\text{Act}| + |\mathbf{M}| \geq 2$ and B invariant are undecidable.

3.2 Results for the Infinite-Horizon

Let us turn to the case of infinite-horizon. If the target subspace B is allowed to be not an invariant subspace, then the limit in equation (2) does not necessarily exist, and we consider the corresponding upper limit:

Theorem 3.3. *Given a qMDP \mathcal{M} , an initial state ρ_0 and a subspace B (not necessarily invariant of \mathcal{M}), then it is undecidable to determine whether*

$$\sup_{\mathfrak{S}} \limsup_{n \rightarrow \infty} \Pr(\rho(n, \mathfrak{S}) \models B) = 1.$$

In the remainder of this section, we only consider invariant subspace B of \mathcal{M} , since the supremum reachability probability is not well-defined for those subspaces that are not invariant (see Definition 2.4 and Lemma 2.1). As for classical MDPs, a major reachability problem for qMDPs is the following:

Problem 3.3. *Given a qMDP \mathcal{M} , an initial state ρ_0 and an invariant subspace B .*

1. *Decide whether $\Pr^{\text{sup}}(\rho_0 \models \diamond B) = 1$.*
2. *Furthermore what is the exact value of $\Pr^{\text{sup}}(\rho_0 \models \diamond B)$?*

Theorem 3.4. *Given a qMDP \mathcal{M} , an initial state ρ_0 and an invariant subspace B .*

1. *It is EXPTIME-hard to decide whether $\Pr^{\text{sup}}(\rho_0 \models \diamond B) = 1$ even for \mathcal{M} whose actions are all unitary.*
2. *The value of $\Pr^{\text{sup}}(\rho_0 \models \diamond B)$ is uncomputable, if $\Pr^{\text{sup}}(\rho_0 \models \diamond B) < 1$.*

For a special class of super-operators and measurements operators, Theorem 3.4.1 can be significantly improved:

Theorem 3.5. *Let \mathcal{M}, ρ_0, B be as in Theorem 3.4. We assume:*

1. *for each $\alpha \in \text{Act}$, $\mathcal{E}_\alpha(\rho) = \sum_{i=1}^{n_\alpha} A_{\alpha i} \rho A_{\alpha i}^\dagger$ with all $A_{\alpha i}$ being of the form $a \cdot |\varphi\rangle\langle\psi|$;*
2. *for $M_\beta \in \mathbf{M}$, $M_\beta = \{M_{\beta 1}, \dots, M_{\beta k_\beta}\}$ with all $M_{\beta j}$ being also of form $a \cdot |\varphi\rangle\langle\psi|$.*

Let $N = \max(\{n_\alpha : \alpha \in \text{Act}\} \cup \{k_\beta : M_\beta \in \mathbf{M}\})$. Then whether $\Pr^{\text{sup}}(\rho_0 \models \diamond B) = 1$ can be decided in time $O(\text{poly}(|\text{Act}| + |\mathbf{M}|)2^N)$.

A variant of Problem 3.3 is the following:

Problem 3.4. *Given a qMDP \mathcal{M} , and an invariant subspace B , is there a scheduler \mathfrak{S} , such that $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ for all initial states ρ ?*

The difference between this problem and Problem 3.3 is that the initial state is arbitrary in the former but it is fixed in the latter. It is worth noting that the counterparts of these two problems for classical MDPs are similar because they have only a finite number of states which can be checked one by one. However, the quantum versions are very different due to the fact that the state Hilbert space of a qMDP is a continuum. It is also worth carefully comparing this problem with Problem 2.1: scheduler \mathfrak{S} is given in the latter, whereas we want to find a special scheduler \mathfrak{S} in the former.

Theorem 3.6. *For a given qMDP $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ and an invariant subspace B of \mathcal{M} , the following two statements are equivalent:*

1. *There exists a scheduler \mathfrak{S} such that $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ for all initial states ρ ;*
2. *There is no invariant subspace C of \mathcal{M} included in B^\perp .*

Furthermore, if there is no invariant subspace C of \mathcal{M} included in B^\perp , then there exists an optimal finite-memory scheduler $\mathfrak{S} = s^\omega$ with $s \in \text{Act}^$.*

Based on the above theorem, we develop Algorithm 1 for checking existence of the optimal scheduler, of which the correctness and complexity are presented in the next theorem.

input : A quantum Markov decision process \mathcal{M} , the Hilbert space \mathcal{H} , a subspace $B \subseteq \mathcal{H}$

output: A string $s \in (\text{Act} \cup \mathbf{M})^*$
* $s = \epsilon$ means no such scheduler.*

```

begin
   $s = \epsilon$ ;
   $\text{Act}' \leftarrow \text{Act}$ ;
   $E \leftarrow \mathcal{E}_i$ , for all  $i \in \text{Act}$ ;
   $t \leftarrow |\text{Act}'|$ ;
  for any  $M \in \mathbf{M}$  do
     $t \leftarrow t + 1$ ;
     $\mathcal{E}_t \leftarrow \sum M_i \cdot M_i^\dagger$ , for all measurement operators
       $M_i$  of  $M$ ;
     $E \leftarrow E \cup \{\mathcal{E}_t\}$ ;
     $\text{Act}' \leftarrow \text{Act}' \cup \{t\}$ ;
  end
   $\mathcal{F} = \frac{1}{t} \sum_{\mathcal{E}_i \in E} \mathcal{E}_i$ ;
  if  $\mathcal{F}(I_B) \not\subseteq B$  then
    return  $s$ ;
  end
   $\mathcal{G} \leftarrow \mathcal{F}|_{B^\perp}$ ;
   $N \leftarrow$  null space of  $\mathcal{G}(x) - x = 0$ ;
  if  $N = \{0\}$  then
    return  $s$ ;
  end
   $d \leftarrow \dim(\mathcal{H})$ ;
   $S \leftarrow \cup_{i=1}^{d-1} \text{Act}'^i$ ;
   $T \leftarrow B^\perp$ ;
   $b \leftarrow \dim(T)$ ;
  while  $b > 0$  do
    for  $v \in S$  do
       $w \leftarrow s \cdot v$ ;
       $Y \leftarrow (\mathcal{E}_w^*(T^\perp))^\perp$ ;
      if  $\dim(Y) < b$  then
         $b \leftarrow \dim(Y)$ ;
         $s \leftarrow w$ ;
        break;
      end
    end
  end
  return  $s$ ;
end

```

Algorithm 1: Find an optimal scheduler

Theorem 3.7. 1. *Algorithm 1 returns ϵ if there is no such a scheduler, otherwise it returns $s \neq \epsilon$ such that $\mathfrak{S} = s^\omega$ is an optimal scheduler.*

2. *The time complexity of Algorithm 1 is $O(d^9 t^d)$, where $d = \dim \mathcal{H}$ and $t = |\text{Act}| + |\mathbf{M}|$.*

We now consider another variant of Problem 3.3, where not only the initial state ρ but also the scheduler \mathfrak{S} can be arbitrary.

Problem 3.5. *Given a qMDP \mathcal{M} , and an invariant subspace B , is the reachability probability always 1, i.e. $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ for all initial states ρ and all schedulers \mathfrak{S} ?*

For this problem, we only have an answer in a special case.

Theorem 3.8. *Let \mathcal{M} be a qMDP with $\mathbf{M} = \emptyset$ and B an invariant subspace of \mathcal{H} . Then $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ holds for all schedulers \mathfrak{S} and all initial states ρ if and only if it holds for all initial states and all schedulers of the form $\mathfrak{S} = s^\omega$ with $|s| \leq L_d$, where L_d is inductively defined as follows:*

- $L_0 = 1$ and $K_0 = k$, where $k = |\text{Act}|$.
- $L_{i+1} = (K_i + 1)L_i$ and $K_{i+1} = k^{L_{i+1}}$ for any $i \geq 0$.

We can develop an algorithm to check whether $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ holds for all initial states and all schedulers \mathfrak{S} . By Theorem 3.8, we only need to examine all schedulers of the form $\mathfrak{S} = s^\omega$ with $|s| \leq L_d$. There are totally K_d such schedulers, and for each one, it costs at most $O(d^6 L_d)$ arithmetic operations to check the conclusion. Thus, the complexity of the algorithm is $O(d^6 L_d K_d)$. For the special class of qMDPs considered in Theorem 3.5, we can significantly reduce this complexity.

Theorem 3.9. *Let \mathcal{M} , B and N be as in Theorem 3.5. Then whether $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$ holds for all schedulers \mathfrak{S} and all initial states ρ can be decided in time $O(\text{poly}(|\text{Act}| + |\mathbf{M}|)2^N)$.*

To conclude this section, we point out a link from Problems 3.4 and 3.5 to a long-standing problem in matrix analysis and control theory, namely the joint spectral radius problem [8, 16, 37]. For a given set of square matrices $\Sigma = \{A_i : i = 1, \dots, t\}$, the discrete linear inclusion of Σ is defined to be the set

$$DLI(\Sigma) = \{x_n = A_{s_n} \cdots A_{s_1} x_0 : x_0 \in \mathbb{C}^d, n \geq 0 \text{ and } A_{s_j} \in \Sigma (1 \leq j \leq n)\}.$$

The set Σ is said to be absolutely asymptotically stable (AAS) if $\lim_{n \rightarrow \infty} A_{s_n} \cdots A_{s_1} = 0$ for any infinite sequences $A_{s_1} A_{s_2} \cdots$ in Σ . The joint spectral radius and lower spectral radius of Σ are defined as

$$\bar{\rho}(\Sigma) = \limsup_{k \rightarrow \infty} \bar{\rho}_k(\Sigma), \quad \underline{\rho}(\Sigma) = \liminf_{k \rightarrow \infty} \underline{\rho}_k(\Sigma)$$

respectively, where for every $k \geq 1$,

$$\bar{\rho}_k(\Sigma) = \sup\{\|A\|^{1/k} : A \in \Sigma^k\},$$

$$\underline{\rho}_k(\Sigma) = \inf\{\|A\|^{1/k} : A \in \Sigma^k\}.$$

It is known [8, 16] that Σ is AAS if and only if the joint spectral radius $\bar{\rho}(\Sigma) < 1$. It was shown in [37] that unless $P = NP$, there are no polynomial-time approximate algorithms for computing $\bar{\rho}(\Sigma)$. The problem “ $\underline{\rho}(\Sigma) < 1$ ” and “ $\bar{\rho}(\Sigma) \leq 1$ ” were proved to be undecidable in [5, 37]. However, the problem whether “ $\bar{\rho}(\Sigma) < 1$ ” is decidable is still open although the notion of joint spectral radius was introduced more than fifty years ago [31].

Theorem 3.10. *Let \mathcal{M} be a qMDP with $\mathbf{M} = \emptyset$ and B an invariant subspace of \mathcal{M} . For each $\alpha \in \text{Act}$, let M_α be the matrix representation of $P_T \mathcal{E}_\alpha(\cdot) P_T$, where $T = B^\perp$. We write $\Sigma_{\mathcal{M}} = \{M_\alpha : \alpha \in \text{Act}\}$. Then:*

1. $\underline{\rho}(\Sigma_{\mathcal{M}}) < 1$ if and only if there exists a scheduler \mathfrak{S} such that for any initial state ρ , it holds that $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$.
2. $\bar{\rho}(\Sigma_{\mathcal{M}}) < 1$ if and only if for any scheduler \mathfrak{S} and any initial state ρ , it holds that $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$.

4. Finite-Horizon Problems

In this section, we prove the theorems for finite-horizon stated in Subsection 3.1.

4.1 Proof of Theorem 3.1

We prove this theorem by an easy reduction from the emptiness problem of cut-point languages for probabilistic finite au-

tomata (PFA) to Problem 3.1. For a given MO-1gQFA $\mathcal{M}' = \langle \mathcal{H}, \Sigma, \rho_0, \{\mathcal{E}_\sigma\}_{\sigma \in \Sigma}, P_{acc} \rangle$ [19], we can construct a qMDP $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ such that $\text{Act} = \Sigma$, and $\mathbf{M} = \Omega = \emptyset$. Let $B = \text{supp}(P_{acc})$. Then these exist \mathfrak{S} and n such that $\Pr(\rho(n, \mathfrak{S}) \models B) \Delta p$ if and only if there exists a word $\sigma_1 \cdots \sigma_n$ such that $\text{tr}(P_{acc} \mathcal{E}_{\sigma_n} \circ \cdots \circ \mathcal{E}_{\sigma_1}(\rho_0)) \Delta p$. Since MO-1gQFA can simulate any PFA [19] and the emptiness problem for PFA is undecidable [4], Problem 3.1 is undecidable too.

4.2 Proof of Theorem 3.2

Our proof technique is a reduction from the matrix mortality problem to Problem 3.2. The matrix mortality problem can be simply stated as follows:

- Given a finite set of matrices $G = \{M_i \in \mathbb{Z}^{n \times n} : i \in \{1, 2, \dots, k\}\}$, is there any sequence j_1, \dots, j_m such that $M_{j_m} M_{j_{m-1}} \cdots M_{j_1} = 0$?

It is known [17, Theorem 3.2] that the matrix mortality problem is undecidable for $k = 2$.

We now prove Theorem 3.2. For a set G of matrices as above, we construct a qMDP $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ from it as follows:

- The state space is $\mathcal{H} = \text{span}\{|1\rangle, \dots, |2n\rangle\}$.
- Let $\text{Act} = \{1, 2, \dots, n\}$. For each $i \in \text{Act}$, we construct a super-operator \mathcal{E}_i from M_i :

$$\mathcal{E}_i(\rho) = A_i \rho A_i^\dagger + B_i \rho B_i^\dagger + C_i \rho C_i^\dagger,$$

where

$$A_i = \begin{pmatrix} a_i & 0 \\ 0 & 0 \end{pmatrix}, B_i = \begin{pmatrix} 0 & 0 \\ 0 & I_{n \times n} \end{pmatrix}, C_i = \begin{pmatrix} 0 & 0 \\ c_i & 0 \end{pmatrix},$$

and

$$a_i = \frac{1}{r_i} M_i, \quad c_i = \sqrt{I - M_i^\dagger M_i / r_i^2}.$$

In the defining equation of a_i, r_i is a positive integer such that $I - M_i^\dagger M_i / r_i^2 \geq 0$.

- $\mathbf{M} = \emptyset$.

Now, it is easy to show that for any state

$$\sigma = \begin{pmatrix} \sigma_a & * \\ * & \sigma_b \end{pmatrix},$$

we have

$$\mathcal{E}_i(\sigma) = \begin{pmatrix} a_i \sigma_a a_i^\dagger & 0 \\ 0 & \sigma_b + c_i \sigma_a c_i^\dagger \end{pmatrix}.$$

Therefore, for any initial state

$$\rho_0 = \begin{pmatrix} \rho_a & * \\ * & \rho_b \end{pmatrix},$$

it holds that

$$\rho(m, \mathfrak{S}) = \begin{pmatrix} A \rho_a A^\dagger & 0 \\ 0 & * \end{pmatrix},$$

where $A = a_{j_m} \cdots a_{j_1}$. Now let $B = \text{span}\{|n+1\rangle, \dots, |2n\rangle\}$. Then

$\forall \rho_0, \exists n, \mathfrak{S}$ s.t. $\rho(n, \mathfrak{S}) \subseteq B$

$$\Leftrightarrow \exists j_m, \dots, j_1 \text{ s.t. } a_{j_m} \cdots a_{j_1} = 0$$

$$\Leftrightarrow \exists j_m, \dots, j_1 \text{ s.t. } M_{j_m} \cdots M_{j_1} / (r_{j_m} \cdots r_{j_1})^2 = 0$$

$$\Leftrightarrow \exists j_m, \dots, j_1 \text{ s.t. } M_{j_m} \cdots M_{j_1} = 0.$$

Since the matrix mortality problem is undecidable for $k = 2$, Problem 3.2.1 with $\mathbf{M} = \emptyset$ and B invariant is undecidable for dimension $|\text{Act}| + |\mathbf{M}| \geq 2$.

Note in the above reduction, $A\rho_a A^\dagger$ will always be rational, if ρ_0 is rational. Since for any $\sigma \geq 0$, $\text{supp}(\sigma) \subseteq B$ holds if and only if $\sigma_a = 0$, we only compute the upper left corner and leave c_i as a symbol in the lower right corner when computing $\rho(m, \mathfrak{S})$. (There are at most $O(m)$ c_i 's.) Thus this reduction does not employ any operation on irrational numbers.

The reduction still works if we fix the initial state ρ_0 to be $I/2n$, which is one special case of Problem 3.2.2. Therefore Problem 3.2.2 is undecidable too.

5. Infinite-Horizon Problems

In this section, we prove the theorems for infinite-horizon stated in Subsection 3.2.

5.1 Proof of Theorem 3.3

This theorem can be proved by reduction from the value 1 problem of probabilistic automata on finite words in [12]. The value 1 problem asks whether $\sup_{w \in \Sigma^*} \Pr(q_0 \xrightarrow{w} F) = 1$ for a probabilistic finite automaton, where q_0 is the initial state, F is the set of accept states and w is a finite word over the input symbols Σ . We can reduce this automaton to a qMDP with $\mathbf{M} = \emptyset$, $Act = \Sigma$, $\rho_0 = |q_0\rangle\langle q_0|$ and $B = \text{span}\{|q\rangle : q \in F\}$. The reduction technique is the similar as in the proof of Theorem 3.1. Thus we have

$$\sup_{s \in Act^*} \Pr^{s^\omega}(\rho(|s|, s^\omega) \models B) = 1 \quad (4)$$

is undecidable. Since $\mathbf{M} = \emptyset$, all schedulers are of form $\mathfrak{S} = s^\omega$ or $\mathfrak{S} \in Act^\omega$. Therefore equation (4) is equivalent to

$$\sup_{\mathfrak{S}} \limsup_{n \rightarrow \infty} \Pr^{\mathfrak{S}}(\rho(n, \mathfrak{S}) \models B) = 1.$$

This completes the proof.

5.2 Proof of Theorem 3.4

We prove part 1 of the theorem by a reduction from an EXPTIME-complete game in [35] to the problem of deciding whether $\Pr^{\text{sup}}(\rho_0 \models \diamond B) = 1$. Some ideas are similar to those used in [6, 28].

- The game is a two-player game on a propositional formula $F(X, Y)$ in the conjunctive normal form (CNF). Player 1 (resp. 2) changes at most one variable in X (resp. Y) at each move, alternately. Once F becomes true, Player 1 wins.

It is known [35] that the following problem is EXPTIME-complete: given an input string w encoding a position of this game, decide whether Play 1 has a strategy to win definitely, where a position is a tuple $(\tau, F(X, Y), \alpha)$, where $\tau \in \{1, 2\}$ denotes the current player, F is a formula, and α is an assignment.

Now we start to construct the reduction. Let $X = \{x_1, \dots, x_n\}$, $Y = \{y_1, \dots, y_n\}$, $\alpha \in \{0, 1\}^{n+m}$ and $F = \bigwedge_{i=1}^c C_i$, where $C_i = \bigvee_{j=1}^k z_{i,j}$, and $z_{i,j}$ is one of $x_t, \neg x_t, y_t, \neg y_t$ for some t . We define a qMDP as follows:

State space. The state space $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_C \otimes \mathcal{H}_F \otimes \mathcal{H}_R$, where $\mathcal{H}_S = \mathcal{H}_2^{\otimes(n+m)}$, $\mathcal{H}_C = \mathcal{H}_{k+1}^{\otimes c}$, $\mathcal{H}_F = \mathcal{H}_{c+1}$, $\mathcal{H}_R = \mathcal{H}_{m+2}$, where $\mathcal{H}_i = \text{span}\{|0\rangle, \dots, |i\rangle\}$. The intuition behind the definition of these spaces is:

- \mathcal{H}_S encodes the assignment α ;
- \mathcal{H}_C is the work space for clauses;
- \mathcal{H}_F is the work space for the formula;
- \mathcal{H}_R encodes the randomness of Player 2's choice.

Initial state. The initial state is $|\psi_0\rangle = |\alpha(x_1)\rangle \cdots |\alpha(y_m)\rangle |0_C\rangle |0_F\rangle |0_R\rangle$. We will see that the state of the system can always be represented in such a separable form during the computation of this qMDP.

S_1 : Unitary operators for modelling actions by Player 1. Since Player 1 can change at most 1 valuable, there are $n + 1$ choices/actions:

- Do nothing: this can be described by the identity operator I ;
- Change the i -th valuable x_i : this can be realised by the NOT gate $X = |0\rangle\langle 1| + |1\rangle\langle 0|$ operator on i -th space of \mathcal{H}_S , i.e., $U_{1,i} = U_{x_i} \otimes I_C \otimes I_F \otimes I_R$, where

$$U_{x_i} = I_{S,1} \otimes \cdots \otimes I_{S,i-1} \otimes X \otimes I_{S,i+1} \otimes \cdots \otimes I_{S,n+m}.$$

All these operators can be represented in this form using space $O(n(n+m+c+k))$.

S_2 : Randomness of Player 2's choice. First we split the state $|0\rangle\langle 0|$ in \mathcal{H}_R into $\frac{1}{\sqrt{m+1}} \sum_{i=1}^{m+1} |i\rangle$ by a unitary

$$U_R = I_S \otimes I_C \otimes I_F \otimes \left(\frac{1}{\sqrt{m+1}} \sum_{i=1}^{m+1} |i\rangle\langle 0| + \cdots \right).$$

Then we apply

$$U_2 = \sum_{i=1}^m U_{y,i} \otimes I_C \otimes I_F \otimes |i\rangle\langle i| + I_S \otimes I_C \otimes I_F \otimes |m+1\rangle\langle m+1| + I_S \otimes I_C \otimes I_F \otimes |0\rangle\langle 0|.$$

At last, we apply a measurement $M_R = \{M_{R,i} : i = 0, \dots, m+1\}$, where $M_{R,i} = I_S \otimes I_C \otimes I_F \otimes |0\rangle\langle i|$. These step can be encoded in space $O(m^2(n+m+c+k))$.

S_C : Checking the formula. This can be done by the following steps:

1. First, we check each clause. A clause $C_i = \bigvee_{j=1}^k z_{i,j}$ is checked via each of its literals. For instance, if $z_{i,j}$ is x_t , we apply

$$U_{z_{i,j}} = I_{S,1} \otimes \cdots \otimes I_{S,t-1} \otimes |1\rangle\langle 1| \otimes I_{S,t+1} \otimes \cdots \otimes I_{S,n+m} \otimes U_{\text{shift},i} \otimes I_F \otimes I_R$$

$$+ I_{S,1} \otimes \cdots \otimes I_{S,t-1} \otimes |0\rangle\langle 0| \otimes I_{S,t+1} \otimes \cdots \otimes I_{S,n+m} \otimes I_C \otimes I_F \otimes I_R,$$
 where $U_{\text{shift},i} = I_{C,1} \otimes \cdots \otimes I_{C,i-1} \otimes (\sum_{v=0}^{k-1} |v+1\rangle\langle v| + |0\rangle\langle k|) \otimes I_{C,t+1} \otimes \cdots \otimes I_{C,c}$ is the shift operator on subspace \mathcal{H}_{C_i} . The case of $z_{i,j}$ being $\neg x_t, y_t, \neg y_t$ is similar. This step means that $z_{i,j}$ is true, and we shift one level in \mathcal{H}_{C_i} .
2. Second, we compute the value of the whole formula. This is similar the first step. If the state is \mathcal{H}_{C_i} is shifted at least once; that is, it is not $|0\rangle\langle 0|$, then we shift \mathcal{H}_F once.
3. Third, we take a projective measurement $P_F = \{P_1, P_0\}$ on \mathcal{H}_F , where $P_1 = I_S \otimes I_C \otimes |c\rangle\langle c| \otimes I_R$ represents the fact that all c clauses are true, i.e. F is true, and

$$P_0 = I_S \otimes I_C \otimes \sum_{i=0}^{c-1} |i\rangle\langle i| \otimes I_R,$$

indicates that F is false. If the outcome is 1, we terminate.

4. Forth, we undo the first two steps if the result is false. Let U denote the unitary operator of the first two steps. If the projective measurement gives result 0, the state remains unchanged because of the separable form of the initial state. Thus we can apply U^\dagger to undo U .

The above four steps can be represented in space $O(k^2c(n+m+c+k))$.

Schedulers. If the input $\tau = 1$, i.e. Player 1 first moves, then we execute sequence $(S_1 S_C S_2 S_C)^\omega$ of steps; otherwise $(S_2 S_C S_1 S_C)^\omega$. This is realised by the mapping Q in Definition 2.1. The decision is made in step S_1 (Player 1's turn).

Target and reachability probability. The target is to reach the outcome 1; that is, P_1 appears in S_C . Because of the separable form of the initial state, the state of the system is of the form $|\psi\rangle = |\alpha'(x_1)\rangle \cdots |\alpha'(y_m)\rangle |0_C\rangle |0_F\rangle |0_R\rangle$ after each step. Thus any step can be computed in polynomial time of n, m, c, k . Therefore, this is a polynomial time reduction. Furthermore, it is easy to see that Player 1 has a “forced win” strategy if and only if there is a scheduler (for decisions in step S_1) with reachability probability is 1.

Remark: The target space B may not be invariant. But we can easily modify the space \mathcal{H}_F so that B becomes invariant. What we need to do is:

- extend \mathcal{H}_F to $k + 2$ level space;
- change P_1 to $|k + 1\rangle\langle k|$, and add $P_2 = |k + 1\rangle\langle k + 1|$;
- make all unitary operators to be a controlled operator by \mathcal{H}_F .

After the modification, the system state remains unchanged in each decision branching unless it reaches the target.

We now turn to prove part 2 of the theorem; that is, $\text{Pr}^{\text{sup}}(\rho_0 \models \diamond B)$ is uncomputable. This can be done simply by a reduction from probabilistic automata on infinite words. In [7], it was shown that the following quantitative value problem is undecidable: for any $\epsilon > 0$, does there exist a word w such that the reachability probability in acceptance absorbing automata is greater than $\lambda - \epsilon$, for a given rational number $0 < \lambda < 1$. We reduce this problem to the supremum reachability problem for qMDPs. The reduction technique is similar to the proofs of Theorems 3.1 and 3.3. Since the automata are acceptance absorbing, B is invariant. Thus, it is undecidable whether there exists $\mathfrak{S} \in \text{Act}^\omega$, such that $\text{Pr}^{\mathfrak{S}}(\rho_0 \models \diamond B) > \lambda - \epsilon$. Since this is equivalent to decide $\sup_{\mathfrak{S}} \text{Pr}^{\mathfrak{S}}(\rho_0 \models \diamond B) = \lambda$, we complete the proof.

5.3 Proof of Theorem 3.5

By the assumption, \mathcal{E}_α can be written as $\mathcal{E}_\alpha(\rho) = \sum_{i=1}^{n_\alpha} a_{\alpha,i} A_{\alpha,i} \rho A_{\alpha,i}^\dagger$, where $A_{\alpha,i} = a_{\alpha,i} |\varphi_{\alpha,i}\rangle\langle \psi_{\alpha,i}|$. Then for any state ρ , we have $\mathcal{E}_\alpha(\rho) = \sum_i c_{\alpha,i} |\varphi_{\alpha,i}\rangle\langle \varphi_{\alpha,i}|$ for some $c_{\alpha,i} \geq 0$. Define $Y(\alpha, \rho) \triangleq \text{supp}(\mathcal{E}_\alpha(\rho)) = \text{span}\{|\varphi_{\alpha,i}\rangle : c_{\alpha,i} > 0\}$. It is easy to see that there are at most 2^{n_α} different $Y(\alpha, \rho)$'s ranging over all ρ for an given α . Then the total number of $Y(\alpha, \rho)$'s with all actions α is at most $|\text{Act}|2^N$. Similarly, we define $Z(\beta, j, \rho) \triangleq \text{supp}(M_{\beta j} \rho M_{\beta j}^\dagger)$. If probability $\text{tr}(M_{\beta j} \rho M_{\beta j}^\dagger) > 0$, then $Z(\beta, j, \rho) = \text{span}\{|\varphi_{\beta j}\rangle\}$. Otherwise it equals $\{0\}$. The total number of $Z(\beta, i, \rho)$'s is at most $|\mathbf{M}|N$. Thus there are at most $(|\text{Act}| + |\mathbf{M}|)2^N$ possible different supports of resulting states. Let Y to be the set of all these supports. Now we reduce this problem to the supremum-1 reachability problem of a classical Markov decision process $\mathcal{M}' = \langle S, \text{Act}', T, s_0 \rangle$:

- each state corresponds to a possible support, i.e. $S = \{s_y : y \in Y\} \cup \{s_0\}$;
- $s_0 = s_{\{\text{supp}(\rho_0)\}}$;
- $\text{Act}' = \text{Act} \cup \{\beta : M_\beta \in \mathbf{M}\}$;
- for each $\alpha \in \text{Act}$, the transition function T maps s_x to s_y with probability 1, where $\mathcal{E}_\alpha(x) = y$;
- for each $M_\beta \in \mathbf{M}$, T maps s_x to s_y with probability $1/l(\beta, x)$, where $y \in \{z : z = Z(\beta, i, x)\}$ and $l(\beta, y)$ is the number of elements in this set;
- the target states $B' = \{s_y : y \subseteq B \wedge y \in Y\}$.

For this classical Markov decision process, it is known [2] that there is an optimal memoryless scheduler \mathfrak{S}_0 such that

$$P_{\max} \triangleq \text{Pr}^{\mathfrak{S}_0}(s_0 \models \diamond B') = \text{Pr}^{\text{sup}}(s_0 \models \diamond B').$$

If $P_{\max} = 1$, then \mathfrak{S}_0 can be converted to a scheduler of \mathcal{M} , whose decisions are based on supports of states. We immediately have $\text{Pr}^{\text{sup}}(\rho_0 \models \diamond B) = 1$. Conversely, if $\text{Pr}^{\text{sup}}(\rho_0 \models \diamond B) = 1$, then for any $\epsilon > 0$ there exists a history-dependent scheduler \mathfrak{S}_ϵ convertible to that of \mathcal{M}' such that $\text{Pr}^{\mathfrak{S}_\epsilon}(s_0 \models \diamond B') > 1 - \epsilon$. Thus $P_{\max} = 1$. This completes the reduction. The proof is finished by the fact from [2] that the maximum reachability of a classical MDP can be solved in polynomial time of the size of \mathcal{M}' .

5.4 Proofs of Theorems 3.6 and 3.7

We first present several technical lemmas. For a super-operator \mathcal{E} , we define:

$$X_\mathcal{E} = \text{span}\left(\bigcup\{\text{supp}(\rho) : \text{tr}(P_B \mathcal{E}(\rho)) = 0\}\right). \quad (5)$$

Since B is invariant, $X_\mathcal{E}$ is obviously a subspace of B^\perp .

Lemma 5.1. *For any density operator ρ , $\text{tr}(P_B \mathcal{E}(\rho)) = 0$ if and only if $\text{supp}(\rho) \subseteq X_\mathcal{E}$.*

Proof. The “only if” part is by definition. We now prove the “if” part. If $\text{supp}(\rho) \subseteq X_\mathcal{E}$, then there exist $\sigma_1, \dots, \sigma_k$ with $\text{supp}(\sigma_i) \subseteq X_\mathcal{E}$ and $\text{supp}(\rho) \subseteq \bigvee \text{supp}(\sigma_i)$, i.e. $\rho \leq \gamma \sum \sigma_i$ for some $\gamma > 0$. Thus

$$\begin{aligned} \gamma \sum \sigma_i - \rho \geq 0 &\Rightarrow \mathcal{E}(\gamma \sum \sigma_i - \rho) \geq 0 \\ &\Rightarrow P_B \mathcal{E}(\gamma \sum \sigma_i - \rho) P_B \geq 0 \\ &\Rightarrow P_B \mathcal{E}(\gamma \sum \sigma_i) P_B \geq P_B \mathcal{E}(\rho) P_B. \end{aligned}$$

By definition, we have $P_B \mathcal{E}(\gamma \sum \sigma_i) P_B = 0$ and $P_B \mathcal{E}(\rho) P_B \geq 0$. Therefore $P_B \mathcal{E}(\rho) P_B = 0$. This implies $\text{supp}(\rho) \subseteq X$. \square

We now consider a special qMDP $\mathcal{M} = \langle \mathcal{H}, \text{Act}, \mathbf{M} \rangle$ without measurements: $|\mathbf{M}| = \emptyset$. We write $\mathcal{E}_s = \mathcal{E}_{s_k} \circ \dots \circ \mathcal{E}_{s_2} \circ \mathcal{E}_{s_1}$ for a finite sequence $s = s_1 s_2 \dots s_k \in \text{Act}^*$. Since $(\mathcal{H}, \mathcal{E}_s)$ can be seen a quantum Markov chain, we know from [42] that $\text{Pr}^{\mathfrak{S}}(\rho \models \diamond B) = 1$ if and only if there is no invariant subspace in B^\perp , where $\mathfrak{S} = s^\omega$ is a periodic scheduler. For any $s \in \text{Act}^*$, we simply write X_s for $X_{\mathcal{E}_s}$ defined by equation (5) from super-operator \mathcal{E}_s .

Lemma 5.2. *Let $\mathfrak{S} = s^\omega$. If $X_s = \{0\}$, then $\text{Pr}^{\mathfrak{S}}(\rho \models \diamond B) = 1$ for any ρ .*

Proof. We prove it by contradiction. Suppose $X_s = \{0\}$ and $\text{Pr}^{\mathfrak{S}}(\rho \models \diamond B) < 1$ for some ρ . Since $(\mathcal{H}, \mathcal{E}_s)$ is a quantum Markov chain, the scheduler \mathfrak{S} is a actually repeated application of \mathcal{E}_s , we have from Theorems 4 and 6 in [42] that $\text{Pr}^{\mathfrak{S}}(\rho \models \diamond B) < 1$ if and only if there exists a (non-empty) BSCC C of $T = B^\perp$ under \mathcal{E}_s . Corresponding to this BSCC, there exists a minimal fixed point state ρ with $\mathcal{E}_s(\rho) = \rho$ and $\text{supp}(\rho) = C \subseteq T = B^\perp$. By definition, we get $\{0\} \subsetneq \text{supp}(\rho) \subseteq X_s = \{0\}$. A contradiction! \square

Lemma 5.3. *For any $s, v \in \text{Act}^*$ and $w = sv$, we have $X_w \subseteq X_s$. In particular, if $\dim X_w = \dim X_s$, then we have $\mathcal{E}_s(X_s) \subseteq X_v$.*

Proof. For any ρ with $\text{supp}(\rho) \subseteq X_w$, we have $0 = P_B \mathcal{E}_w(\rho) P_B = P_B \mathcal{E}_v(\mathcal{E}_s(\rho)) P_B$. Thus $\text{supp}(\mathcal{E}_s(\rho)) \subseteq X_v \subseteq B^\perp$. This implies $P_B \mathcal{E}_s(\rho) P_B = 0$, and $\text{supp}(\rho) \subseteq X_s$. Therefore, it holds that $X_w \subseteq X_s$. We now turn to prove the second part. If $\dim X_w = \dim X_s$, then for any ρ with $\text{supp}(\rho) \subseteq X_s$, we have $\text{supp}(\rho) \subseteq X_w$. This means $\text{supp}(\mathcal{E}_s(\rho)) \subseteq X_v$ as B is invariant. \square

Now we are ready to prove Theorems 3.6 and 3.7.

Proof of Theorem 3.6. The proof of (1) \Rightarrow (2) is easy. Suppose that there is an invariant subspace C of \mathcal{M} included in B^\perp , then $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 0$ for any ρ in C and for any scheduler \mathfrak{S} .

We now prove (2) \Rightarrow (1). For the special case of $\mathbf{M} = \emptyset$, assume that there is no invariant subspace C of \mathcal{M} included in B^\perp . Let $D = \{\dim X_u : u \in \text{Act}^*\}$ and let $d_{\min} = \min D$. Then there exists $s \in \text{Act}^*$ such that $\dim X_s = d_{\min}$. We assert that $d_{\min} = 0$. Indeed, if $d_{\min} > 0$, then for each word $v \in \text{Act}^*$, we put $w = sv$. By Lemma 5.3, we have $X_w \subseteq X_s$. Then it follows from the definition of d_{\min} that $X_w = X_s$. As a consequence, $Y \triangleq \mathcal{E}_s(X_s) \subseteq X_v$. This implies $\mathcal{E}_v(Y) \subseteq B^\perp$. For a super-operator \mathcal{E} , we write $\mathcal{R}_{\mathcal{E}}(Y)$ for the transitive closure of Y under \mathcal{E} , i.e.

$$\mathcal{R}_{\mathcal{E}}(Y) \triangleq \bigvee_{i=0}^{d-1} \mathcal{E}^i(Y).$$

Let $\mathcal{F} = \frac{1}{t} \sum_{\alpha \in \text{Act}} \mathcal{E}_\alpha$, where $t = |\text{Act}|$. We have:

$$\mathcal{R}_{\mathcal{F}}(Y) = \bigvee_{i=0}^{d-1} \bigvee_{x \in \text{Act}^i} \mathcal{E}_x(Y) \subseteq B^\perp.$$

It is clear that $\mathcal{R}_{\mathcal{F}}(Y)$ is invariant under \mathcal{F} , and thus invariant under any \mathcal{E}_i . So, $\mathcal{R}_{\mathcal{F}}(Y)$ is an invariant subspace of \mathcal{M} included in B^\perp under \mathcal{M} . This contradicts to the assumption. So, we have $d_{\min} = 0$, and it follows from Lemma 5.2 that $\mathfrak{S} = s^\omega$ is a optimal scheduler.

For the general case of $\mathbf{M} \neq \emptyset$, we define a super-operator $\mathcal{E}_{M_\beta} = \sum M_m \cdot M_m^\dagger$ for each $M_\beta \in \mathbf{M}$. Furthermore, we can construct a new qMDP $\mathcal{M}' = \langle \mathcal{H}, \text{Act}', \mathbf{M}' \rangle$ with $\text{Act}' = \text{Act} \cup \{\beta : M_\beta \in \mathbf{M}\}$ and $\mathbf{M}' = \emptyset$. Then we complete the proof by applying the above argument to \mathcal{M}' . \square

It is worth noting that the optimal scheduler given in the proof of the above theorem depends on which measurement is chosen in each step but not its outcome.

Proof of Theorem 3.7. The design idea of Algorithm 1 is to see whether there exists an invariant subspace of B^\perp under super-operator

$$\mathcal{F} = \frac{1}{K} \left(\sum_{\alpha \in \text{Act}} \mathcal{E}_\alpha + \sum_{M \in \mathbf{M}} \sum_{M_i \in M} M_i \cdot M_i^\dagger \right),$$

where $K = |\text{Act}| + |\mathbf{M}|$. A crucial part of the algorithm is to compute X_s for each $s \in \text{Act}^*$. By definition, we have $\mathcal{E}_s(V) \subseteq B^\perp$ whenever $V = \text{supp}(\rho) \subseteq X_s$. Therefore,

$$\begin{aligned} X_s &= \text{span} \left(\bigcup \{ \text{supp}(\rho) : P_B \mathcal{E}_s(\rho) P_B = 0 \} \right) \\ &= \bigvee \{ V : \mathcal{E}_s(V) \subseteq B^\perp \} = \mathcal{E}_s^{-1}(B^\perp) = (\mathcal{E}_s^*(B))^\perp, \end{aligned}$$

where \mathcal{E}^* stands for the dual of super-operator \mathcal{E} , i.e. $\mathcal{E}^* = \sum A_i^\dagger \cdot A_i$ when $\mathcal{E} = \sum A_i \cdot A_i^\dagger$.

1. The correctness of the algorithm is essentially based on the proof of Theorem 3.6. Here we give a detailed argument. The algorithm returns $s = \epsilon$ at the first two “return” statements where B is not invariant or there is an invariant subspace of \mathcal{M} included in B^\perp . Otherwise b is initialized as $b > 0$, and the algorithm enters the “while” loop. During the loop, b must decrease at least 1. If not, we have found some s such that $b_s > 0$, and for any $v \in \text{Act}'^*$, it holds that $b_{s \cdot v} = b_s$. By Lemma 5.3, we have $X_s = X_{s \cdot v}$ and $\mathcal{E}_s(X_s) \subseteq X_v \subseteq B^\perp$ for all v . Therefore, $\mathcal{E}_s(X_s)$ is an invariant subspace of \mathcal{M} included in B^\perp , which is a contradiction. So, b will be 0 finally and $\mathfrak{S} = s^\omega$ is then an optimal scheduler.

2. We note that the algorithm will run the “while” loop at most d times and each time it will run the “for” loop within the body

of the “while” loop at most t^d times. So the length of s will be at most d^2 , as it increases at most d in each running of the “while” loop. In the “for” loop, the complexity mainly comes from computing \mathcal{E}_w . It costs at most $O(d^8)$ because the length of w (i.e. the number of matrix multiplications) is at most $O(d^2)$ and each matrix multiplication costs $O(d^6)$. So the complexity of the algorithm is $O(d \cdot t^d \cdot d^8) = O(d^9 t^d)$. \square

5.5 Proofs of Theorems 3.8 and 3.9

We first introduce an auxiliary tool.

Definition 5.1. For any sequence $s \in \text{Act}^*$, its repetition degree $\text{rd}(s)$ is inductively defined as follows:

1. If there does not exist $t \in \text{Act}^+$ and $a, b, c \in \text{Act}^*$, such that $s = a \cdot t \cdot b \cdot t \cdot c$, then $\text{rd}(s) = 0$.
2. In general, $\text{rd}(s) = \max\{\text{rd}(t) + 1 : s = a \cdot t \cdot b \cdot t \cdot c, t \in \text{Act}^+, a, b, c \in \text{Act}^*\}$.

It is clear that $\text{rd}(s) = 0$ for any $s \in \{\epsilon\} \cup \text{Act}$. The following lemma provides a way to estimate the repetition degree $\text{rd}(s)$.

Lemma 5.4. Let \mathcal{M} be a qMDP with $\mathbf{M} = \emptyset$ and B an invariant subspace of \mathcal{H} . Assume $|\text{Act}| = k$ and $\dim \mathcal{H} = d$. Then for any sequence $s \in \text{Act}^+$ and any $x \geq 0$,

$$|s| \geq L_x \Rightarrow \text{rd}(s) \geq x.$$

Here, L_x is as the same as in Theorem 3.8.

Proof. We prove it by induction on x . For the case of $x = 0$, it is obvious. For $x = 1$, assume s is a sequence with length $|s| \geq L_1 = k + 1$. Since there is only k possible actions, there must be two different integers $p, q \in [1, k + 1]$ such that $s_p = s_q$. Then by definition, $\text{rd}(s) \geq 1$.

Now we suppose that for all $x \leq i$ we have $|t| \geq L_x \Rightarrow \text{rd}(t) \geq x$. Assume $|s| \geq L_{i+1} = (K_i + 1)L_i$. Then s can be rewritten as $s = v_1 \cdots v_{K_i+1} \cdots$, where for $u \in [1, K_i + 1]$, $v_u = s_{(u-1) \cdot L_i + 1} \cdots s_{u \cdot L_i}$ is a subsequence of length L_i . Since there are only $K_i = k_i^{L_i}$ different possible sequences of length L_i , there must be two different integers $p, q \in [1, K_i + 1]$ such that $v_p = v_q$. By induction assumption, we $\text{rd}(v_p) \geq i$. Therefore, $\text{rd}(s) \geq i + 1$. This completes the proof. \square

Now we can establish a connection between $\text{rd}(s)$ and $\dim X_s$.

Lemma 5.5. Let \mathcal{M} be a qMDP with $\mathbf{M} = \emptyset$ and B an invariant subspace of \mathcal{H} . If for any $s' \in \text{Act}^*$ with $0 < |s'| \leq L_q$ and $q = \max_{a \in \text{Act}} \dim(X_a)$, and for any initial state ρ , the scheduler $\mathfrak{S} = s'^\omega$ satisfies $\Pr^{\mathfrak{S}}(\rho \models \diamond B) = 1$, then for any sequence $s \in \text{Act}^*$ with $|s| \leq L_q$, there exists a non-empty subsequence v of $s = f \cdot v \cdot g$ such that $\dim X_v \leq \max\{q - \text{rd}(s), 0\}$.

Proof. We prove it by induction on $\text{rd}(s)$.

(1) For s with $\text{rd}(s) = 0$ and $0 < |s| \leq L_q$, we have $X_s \subseteq X_{s_1}$ by Lemma 5.3. So, $\dim X_s \leq \dim X_{s_1} \leq q$.

(2) Suppose for any $s' \in \text{Act}^+$ with $\text{rd}(s') = i$ and $|s'| \leq L_q$, there exists a non-empty subsequence v of s' , such that $\dim X_v \leq \max\{q - i, 0\}$. Now assume s is a sequence with $\text{rd}(s) = i + 1$ and $|s| \leq L_q$. If $\dim X_s = 0$, the claim is true. Otherwise, by definition, there exists a non-empty subsequence t of s such that $s = a \cdot t \cdot b \cdot t \cdot c$ and $\text{rd}(t) = i$. By the induction assumption, there exists a non-empty subsequence u of $t = f \cdot u \cdot g$ such that $\dim X_u \leq q - i$. Here $\dim X_u > 0$, since $\dim X_s > 0$. Therefore, s can be rewritten as $s = a \cdot f \cdot u \cdot g \cdot b \cdot f \cdot u \cdot g \cdot c$. Let $f' = a \cdot f$, $v = u \cdot g \cdot b \cdot f \cdot u$ and $g' = g \cdot c$. Now we prove $\dim X_v \leq q - i - 1$. Since $\dim X_u \leq q - i$ and $X_v \subseteq X_u \neq \emptyset$, we only need to prove $X_v \subsetneq X_u$. We do this

by refutation. Suppose $X_v = X_u$. Then by Lemma 5.3, we have $X_u = X_{u \cdot g \cdot b \cdot f}$ and $\mathcal{E}_{u \cdot g \cdot b \cdot f}(X_u) = \mathcal{E}_{u \cdot g \cdot b \cdot f}(X_{u \cdot g \cdot b \cdot f}) \subseteq X_u$. Thus, X_u is an invariant subspace under super-operator $\mathcal{E}_{u \cdot g \cdot b \cdot f}$. As $X_u \perp B$, by definition, we have $\Pr^{\mathfrak{S}}(\rho_0 \models B) = 0$ for $\mathfrak{S} = (u \cdot g \cdot b \cdot f)^\omega$ and $\rho_0 = I_{X_u} / \dim X_u$. Since $|s| \leq L_q$, we have $|u \cdot g \cdot b \cdot f| \leq |s| \leq L_q$. This is a contradiction! Therefore, it must be that $X_v \subsetneq X_u$, and we complete the proof. \square

Now we can prove Theorems 3.8 and 3.9.

Proof of Theorem 3.8. We only need to prove the ‘‘if’’ part because the ‘‘only if’’ is obvious. Assume that $\Pr^{\mathfrak{S}}(\rho \models \Diamond B) = 1$ holds for any initial state and any scheduler $\mathfrak{S} = s^\omega$ with $|s| \leq L_q$, where $q = \max_{a \in Act} \dim(X_a)$. By Lemma 5.4, we have $\text{rd}(s) \geq q$ for all s with $|s| = L_q$. Furthermore, by Lemma 5.5 and the assumption, we have $\dim X_s \leq \max\{q - \text{rd}(s), 0\} = 0$ for any sequence s with $|s| = L_q$. Thus $\text{tr}(P_B \mathcal{E}_s(\rho)) > 0$ for any ρ . Since $\text{tr}(P_B \mathcal{E}_s(\rho)) = \text{tr}(\mathcal{E}_s^*(P_B)\rho)$ and $\mathcal{E}_s^*(P_B) = U_s D_s U_s^\dagger$ where $D_s = \text{diag}\{\lambda_{s,1}, \dots, \lambda_{s,d}\}$, we have $\lambda_{s,i} > 0$ for any i . Then $\text{tr}(P_B \mathcal{E}_s(\rho)) \geq m_s > 0$ for any trace-1 operator ρ , where $m_s = \min \lambda_{s,i}$. Consequently, for any scheduler \mathfrak{S} , it holds that

$$\Pr^{\mathfrak{S}}(\rho \models \Diamond B) \geq 1 - \lim_{t \rightarrow \infty} (1 - m)^t = 1,$$

where $m = \min_{|s|=L_q} m_s > 0$. This completes the proof by $q \leq d = \dim \mathcal{H}$. \square

Proof of Theorem 3.9. This proof is similar to the proof of Theorem 3.5. We can construct a classical MDP with $S = \{s_x : x \in Y\}$ and check whether $\Pr^{\mathfrak{S}}(s_x \models \Diamond B) = 1$ for all s_x by noting the following two simple facts:

- for any initial state ρ and any scheduler \mathfrak{S} , the support of the resulting state after first action/measurement will be in Y ;
- for any $s_x \in S$, we can construct an initial state $\rho = P_x / \text{tr}(P_x)$.

\square

5.6 Proof of Theorem 3.10

Let \mathcal{M} be a qMDP with state Hilbert space \mathcal{H} and B an invariant subspace of \mathcal{M} . For each $\alpha \in Act$, we define a new super-operator: $\mathcal{F}_\alpha(\cdot) = P_T \mathcal{E}_\alpha(\cdot) P_T$ from \mathcal{E}_α , where $T = B^\perp$ is the orthocomplement of B in \mathcal{H} and P_T is the projection operator onto T . Furthermore, let M_α be the matrix representation of \mathcal{F}_α .

Lemma 5.6. *Let \mathcal{M} be a qMDP with $\mathbf{M} = \emptyset$ and B an invariant subspace of \mathcal{M} . Then:*

1. *The following two statements are equivalent:*
 - (a) *There exists a scheduler \mathfrak{S} such that $\Pr^{\mathfrak{S}}(\rho \models \Diamond B) = 1$ for all initial states ρ .*
 - (b) *There exists $\alpha_1 \alpha_2 \dots \in Act^\omega$ such that $\lim_{n \rightarrow \infty} M_{\alpha_n} \dots M_{\alpha_1} = 0$.*
2. *The following two statements are equivalent:*
 - (a) *For any scheduler \mathfrak{S} and any initial state ρ , it holds that $\Pr^{\mathfrak{S}}(\rho \models \Diamond B) = 1$.*
 - (b) *For any $\alpha_1 \alpha_2 \dots \in Act^\omega$, it holds that $\lim_{n \rightarrow \infty} M_{\alpha_n} \dots M_{\alpha_1} = 0$.*

Proof. 1. It is obvious that (b) \Rightarrow (a) because $\text{tr}(\rho(n, \mathfrak{S})) = 1$ and the probability in T goes to 0. We now prove (a) \Rightarrow (b). Suppose that \mathfrak{S} is a scheduler required in (a). Let $T = \text{span}\{|1\rangle, \dots, |k\rangle\}$

and $B = \text{span}\{|k+1\rangle, \dots, |d\rangle\}$. As $\mathbf{M} = \emptyset$, \mathfrak{S} is a sequence of actions, i.e. $\mathfrak{S} = s_1 s_2 \dots$ with $s_i \in Act$ for all i . Since

$$\begin{aligned} \mathcal{E}_\alpha(\rho) &= \sum E_{\alpha,i} \rho E_{\alpha,i}^\dagger \\ &= \sum \begin{pmatrix} a_{\alpha,i} & 0 \\ c_{\alpha,i} & b_{\alpha,i} \end{pmatrix} \begin{pmatrix} \rho_T & * \\ * & \rho_B \end{pmatrix} \begin{pmatrix} a_{\alpha,i}^\dagger & c_{\alpha,i}^\dagger \\ 0 & b_{\alpha,i}^\dagger \end{pmatrix} \\ &= \sum \begin{pmatrix} a_{\alpha,i} \rho_T a_{\alpha,i}^\dagger & * \\ * & * \end{pmatrix} \end{aligned}$$

and

$$\begin{aligned} \mathcal{F}_\alpha(\rho) &= P_T \left(\sum \begin{pmatrix} a_{\alpha,i} & 0 \\ c_{\alpha,i} & b_{\alpha,i} \end{pmatrix} \rho \begin{pmatrix} a_{\alpha,i}^\dagger & c_{\alpha,i}^\dagger \\ 0 & b_{\alpha,i}^\dagger \end{pmatrix} \right) P_T \\ &= \sum \begin{pmatrix} a_{\alpha,i} & 0 \\ 0 & 0 \end{pmatrix} \rho \begin{pmatrix} a_{\alpha,i}^\dagger & 0 \\ 0 & 0 \end{pmatrix}, \end{aligned}$$

we have $\sigma(n, \mathfrak{S}) \equiv P_T \rho(n, \mathfrak{S}) P_T = \mathcal{F}_{s_n} \dots \mathcal{F}_{s_1}(\rho)$. Moreover, as $\Pr^{\mathfrak{S}}(\rho \models \Diamond B) = 1$, we have $\lim_{n \rightarrow \infty} \text{tr}(\sigma(n, \mathfrak{S})) = 0$. As $\sigma(n, \mathfrak{S})$ is a density operator, it follows that $\lim_{n \rightarrow \infty} \sigma(n, \mathfrak{S}) = 0$.

Let $\mathcal{G}_n(\cdot) \triangleq \mathcal{F}_{s_n} \dots \mathcal{F}_{s_1}(\cdot)$. Since $\mathcal{G}_n(\cdot)$ is completely positive, we have $\mathcal{G}_n(\rho) \leq \mathcal{G}_n(I)$ as $I \geq \rho$ for any density operator ρ . If we use the matrix norm

$$\|A\| = \sup_{\|x\|_2=1} \|Ax\|_2 = \sqrt{\lambda_{\max}(A^\dagger A)},$$

then it holds that $\|\rho\| = \lambda_{\max}(\rho) \leq \|\sigma\|$ when $\rho \leq \sigma$. As a consequence, we obtain

$$\|\mathcal{G}_n(I/d)\| < \frac{\epsilon}{4d} \Rightarrow \|\mathcal{G}_n(\rho)\| \leq \|\mathcal{G}_n(I)\| < \frac{\epsilon}{4}.$$

For any matrix R , we have $R = a_+ - a_- + i(b_+ - b_-)$, where $a_+, a_-, b_+, b_- \geq 0$ and $a_+ a_- = b_+ b_- = 0$. Furthermore,

$$\|a_+\| \leq \|a_+ - a_-\| = \left\| \frac{R + R^\dagger}{2} \right\| \leq \frac{\|R\| + \|R^\dagger\|}{2} = \|R\|.$$

The first inequality is because a_+ and a_- are both positive and their supports are orthogonal. Therefore, we have

$$\begin{aligned} \forall \epsilon > 0, \exists N \in \mathbb{N}, \forall n > N, \forall R \in \mathbb{M}_n(\mathbb{C}), \\ \|\mathcal{G}_n(R)\| &\leq \|\mathcal{G}_n(a_+)\| + \|\mathcal{G}_n(a_-)\| + \|\mathcal{G}_n(b_+)\| + \|\mathcal{G}_n(b_-)\| \\ &< \epsilon \|R\|. \end{aligned}$$

Thus, for the matrix represents \mathbf{A}_n of \mathcal{G}_n , it holds that $\lim_{n \rightarrow \infty} \mathbf{A}_n = 0$, and we complete the proof of part 1.

2. We actually proved that for each scheduler \mathfrak{S} and its corresponding sequence $\mathbf{A}_1, \mathbf{A}_2, \dots$,

$$\forall \rho, \Pr^{\mathfrak{S}}(\rho \models \Diamond B) = 1 \Leftrightarrow \lim_{n \rightarrow \infty} \mathbf{A}_n = 0$$

in the proof of part 1. Hence, the conclusion of part 2 follows immediately. \square

With the help of the above lemma, we are now able to prove Theorem 3.10.

Proof of Theorem 3.10. 1. If $\underline{q}(\Sigma_{\mathcal{M}}) < 1$, then by definition, there exists a sequence $\mathbf{A}_1, \mathbf{A}_2, \dots$ such that

$$\lim_{n \rightarrow \infty} \|\mathbf{A}_n\|^{1/n} \leq \underline{q}(\Sigma_{\mathcal{M}}) + \epsilon < 1.$$

This implies $\lim_{n \rightarrow \infty} \mathbf{A}_n = 0$. Conversely, if there exists $\mathbf{A}_1, \mathbf{A}_2, \dots$ such that $\lim_{n \rightarrow \infty} \mathbf{A}_n = 0$, then we can find $A \in \Sigma_{\mathcal{M}}^m$ with $\|A\|^{1/m} < 1$ for some m . Thus, $\underline{q}(\Sigma_{\mathcal{M}}) \leq \lim_{n \rightarrow \infty} \|A^n\|^{1/nm} < 1$.

2. By Theorem 3.10 in [8], we know that $DLI(\Sigma_{\mathcal{M}})$ is AAS if and only if $\underline{q}(\Sigma_{\mathcal{M}}) < 1$. Together with Lemma 5.6, it completes the proof. \square

6. Conclusions

In this paper, we introduced the notion of quantum Markov decision process (qMDP). Several examples were presented to illustrate how can qMPD serve as a formal model in the analysis of nondeterministic and concurrent quantum programs. The (un)decidability and complexity of a series of reachability problems for qMDPs were settled, but several others left unsolved (the exact complexity of Problem 3.3.1 and the general case of Problem 3.5).

Developing automatic tools for reachability analysis of qMDPs is a research line certainly worth to pursue because these tools can be used in verification and analysis of programs for future quantum computers. Another interesting topic for further studies is applications of qMDPs in developing machine learning techniques for quantum physics and control theory of quantum systems.

References

- [1] C. Baier, N. Bertrand, and M. Gröber. On decision problems for probabilistic Büchi automata. In *FOSSACS*, pages 287-301, 2008.
- [2] C. Baier and J. Katoen. *Principles of model checking*. MIT Press, Cambridge, Massachusetts, 2008.
- [3] J. Barry, D. T. Barry, and S. Aaronson. Quantum POMDPs. *arXiv:1406.2858*.
- [4] V. D. Blondel, E. Jeandel, P. Koiran, and N. Portier. Decidable and undecidable problems about quantum automata. *SIAM J. Comput.*, 34(6):1464-1473, 2005.
- [5] V. D. Blondel and J. N. Tsitsiklis. The boundedness of all products of a pair of matrices is undecidable. *Syst. Control Lett.*, 41(2):135-140, 2000.
- [6] K. Chatterjee, L. Doyen, and T. A. Henzinger. Qualitative analysis of partially observable Markov decision processes. In *MFCS*, pages 258-269, 2010.
- [7] K. Chatterjee and H. A. Henzinger. Probabilistic automata on infinite words: decidability and undecidability results. In *ATVA* pages 1-16, 2010.
- [8] D. Cheban and C. Mammana. Absolute asymptotic stability of discrete linear inclusions. *Bul. Acad. Stiine Repub. Mold. Mat.*, 1: 43-68, 2005.
- [9] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman. Exponential algorithmic speedup by a quantum walk. In *STOC*, pages 59-68, 2003.
- [10] E. D'Hondt and P. Panangaden. Quantum weakest preconditions. *Math. Struct. Comp. Sci.*, 16(03):429-451, 2006.
- [11] S. Gay. Quantum programming languages: survey and bibliography. *Math. Struct. Comp. Sci.*, 16(04):581-600, 2006.
- [12] H. Gimbert and Y. Oualhadj. Probabilistic automata on finite words: decidability and undecidable problems. In *ICALP*, pages 527-538, 2010.
- [13] A. D. Gordon, T. Graepel, N. Rolland, C. Russo, J. Borgstrom and J. Guiver. Tabular: a schema-driven probabilistic programming language. In *POPL*, pages 321-334, 2014.
- [14] A. S. Green, P. LeFanu Lumsdaine, N. J. Ross, P. Selinger and B. Valiron. Quipper: a scalable quantum programming language. In *PLDI*, pages 333-342, 2013.
- [15] L. K. Grover. A fast quantum mechanical algorithm for database search. In *STOC*, pages 212-219, 1996.
- [16] L. Gurvits. Stability of discrete linear inclusion. *Linear Algebra Appl.*, 231:43-60, 1995.
- [17] V. Halava. *Decidable and undecidable problems in martix theory*. TUCS Technical Report No 127, Turku Centre for Computer Science, 1997.
- [18] A. W. Harrow, A. Hassidim, and S. Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103:150502, 2009.
- [19] M. Hirvensalo. Various aspects of finite quantum automata. In *DLT*, pages 21-33, 2008.
- [20] J. Kempe. Discrete quantum walks hit exponentially faster. *Probab. Theory Relat. Fields*, 133(2):215-235, 2005.
- [21] E. H. Knill. *Conventions for quantum pseudocode*. Technical Report LAUR-96-2724, Los Alamos National Laboratory, 1996.
- [22] S. Lloyd, M. Mohseni and P. Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv:1307.0411*.
- [23] E. Martín-López, A. Laing, T. Lawson, R. Alvarez, X. Zhou, and J. L. O'Brien. Experimental realization of Shor's quantum factoring algorithm using qubit recycling. *Nat. Photon.*, 6:773-776, 2012.
- [24] D. Monniaux. Abstract interpretation of programs as Markov decision processes. *Sci. Comput. Program.*, 58(1-2):179-205, 2005.
- [25] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge, 2000.
- [26] B. Ömer. *Structural quantum programming*. Ph.D. Thesis, Technical University of Vienna, 2003.
- [27] M. Ozols, M. Roetteler, and J. Roland. Quantum rejection sampling. In *ITCS*, pages 290-308, 2012.
- [28] C. H. Papadimitiou, and J. N. Tsitsiklis. The complexity of Markov decision processes. *Math. Oper. Res.*, 12(3):441-450, 1987.
- [29] R. Raussendorf and H. J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86:5188, 2010.
- [30] A. Riera, C. Gogolin, and J. Eisert. Thermalization in nature and on a quantum computer. *Phys. Rev. Lett.*, 108:080402, 2012.
- [31] G. C. Rota and W. G. Strang. A note on the joint spectral radius. *Indag. Math.*, 22(4):379-381, 1960.
- [32] J. W. Sanders and P. Zuliani. Quantum programming. In *MPC*, pages 88-99, 2000.
- [33] P. Selinger. Towards a quantum programming language. *Math. Struct. Comp. Sci.*, 14(4):527-586, 2004
- [34] N. Shenvi, J. Kempe, and K. B. Whaley. Quantum random-walk search algorithm. *Phys. Rev. A*, 67: 052307, 2003.
- [35] L. J. Stockmeyer and A. K. Chandra. Provably difficult combinatorial games. *SIAM J. Comput.*, 8(2):151-174, 1979.
- [36] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete. Quantum Metropolis sampling. *Nature*, 471:87-90, 2011.
- [37] J. N. Tsitsiklis and V. D. Blondel. The Lyapunov exponent and joint spectral radius of pairs of matrices are hard – when not impossible – to compute and to approximate. *Math. Control Signals Systems*, 10(1):31-40, 1997.
- [38] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang. Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414:883-887, 2001.
- [39] M. Y. Vardi. Automatic verification of probabilistic concurrent finite state programs. In *FOCS*, pages 327-338, 1985.
- [40] M. S. Ying. Floyd-Hoare logic for quantum programs. *TOPLAS*, 33(6): article No. 19, 2011.
- [41] M. S. Ying, N. K. Yu, Y. Feng, and R. Y. Duan. Verification of quantum programs. *Sci. Comput. Program.*, 78(9):1679-1700, 2013.
- [42] S. G. Ying, Y. Feng, N. K. Yu, and M. S. Ying. Reachability probabilities of quantum Markov chains. In *CONCUR*, pages 334-348, 2013.
- [43] N. K. Yu and M. S. Ying. Reachability and termination analysis of concurrent quantum programs. In *CONCUR*, pages 69-83, 2012.
- [44] M. H. Yung and A. Aspuru-Guzik. A quantum-quantum Metropolis algorithm. *Proc. Natl. Acad. Sci. USA*, 109:754-759, 2012.
- [45] J. Zhang, M. H. Yung, R. Laflamme, A. Aspuru-Guzik, and J. Baugh. Digital quantum simulation of the statistical mechanics of a frustrated magnet. *Nat. Commun.*, 3: Article No. 880.
- [46] P. Zuliani. Nondeterministic quantum programming. In: *Proceedings of the 2nd International Workshop on Quantum Programming Languages (QPL)*, pages 179-195, 2004.