

# (No) Quantum Space-Time Tradeoff for USTCON

Simon Apers 

CNRS, IRIF, Paris, France

Stacey Jeffery 

CWI & QuSoft, Amsterdam, The Netherlands

Galina Pass 

Korteweg-de Vries Institute for Mathematics & QuSoft, University of Amsterdam, The Netherlands  
Faculty of Computer Science, Ruhr University Bochum, Germany

Michael Walter  

Faculty of Computer Science, Ruhr University Bochum, Germany

---

## Abstract

Undirected  $st$ -connectivity is important both for its applications in network problems, and for its theoretical connections with logspace complexity. Classically, a long line of work led to a time-space tradeoff of  $T = \tilde{O}(n^2/S)$  for any  $S$  such that  $S = \Omega(\log(n))$  and  $S = O(n^2/m)$ . Surprisingly, we show that quantumly there is no nontrivial time-space tradeoff: there is a quantum algorithm that achieves both optimal time  $\tilde{O}(n)$  and space  $O(\log(n))$  simultaneously. This improves on previous results, which required either  $O(\log(n))$  space and  $\tilde{O}(n^{1.5})$  time, or  $\tilde{O}(n)$  space and time. To complement this, we show that there is a nontrivial time-space tradeoff when given a lower bound on the spectral gap of a corresponding random walk.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Quantum computation theory; Theory of computation  $\rightarrow$  Design and analysis of algorithms

**Keywords and phrases** Undirected  $st$ -connectivity, quantum walks, time-space tradeoff

**Digital Object Identifier** 10.4230/LIPIcs.ESA.2023.10

**Related Version** *Full Version*: <https://arxiv.org/abs/2212.00094>

**Funding** *Stacey Jeffery*: Supported by ERC STG grant 101040624-ASC-Q, NWO Klein project number OCENW.Klein.061, and ARO contract no W911NF2010327. SJ is a CIFAR Fellow in the Quantum Information Science Program.

*Galina Pass*: Supported by the National Agenda for Quantum Technologies (NAQT), as part of the Quantum Delta NL programme.

*Michael Walter*: Supported by the European Research Council (ERC) through ERC Starting Grant 101040907-SYMOPTIC, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972, the Federal Ministry of Education and Research (BMBF) through project Quantum Methods and Benchmarks for Resource Allocation (QuBRA), and NWO grant OCENW.KLEIN.267.

**Acknowledgements** Part of this work was initiated at the 2022 QOPT (QuantERA ERA-NET Cofund 2022-25) workshop hosted at Université Libre de Bruxelles.

## 1 Introduction

For an undirected graph  $G = (X, E)$  on  $n = |X|$  vertices and  $m = |E|$  edges, with  $s, t \in X$ ,  $st$ -connectivity or USTCON is the problem of deciding whether  $s$  and  $t$  are in the same component. This problem has applications in many other graph and network problems, and is of theoretical importance for its connection with space complexity (see e.g. [23]). In particular, USTCON is complete for the class *symmetric logspace*, SL, which was shown to be equal to *logspace*, L, by exhibiting a classical deterministic logspace algorithm for USTCON [22]. In this paper, we consider quantum algorithms for this problem.



© Simon Apers, Stacey Jeffery, Galina Pass, and Michael Walter;  
licensed under Creative Commons License CC-BY 4.0

31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 10;  
pp. 10:1–10:17



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

There are different versions of the problem USTCON depending on how  $G$  is accessed. If  $G$  is given as an adjacency matrix, we denote the problem  $\text{USTCON}_{\text{mat}}$ . If  $G$  is given as an array of arrays, one for each vertex, enumerating the neighbours, we denote the problem  $\text{USTCON}_{\text{arr}}$ .<sup>1</sup> If one only cares about space complexity, these problems are equivalent, but the same is not true of time complexity: adjacency queries can simulate an array query, and vice versa, in logspace, but there is a non-negligible time overhead.

A classical deterministic algorithm based on breadth-first search or depth-first search can solve  $\text{USTCON}_{\text{arr}}$  in  $\tilde{O}(m)$  time, using  $\tilde{O}(n)$  space. Using a random walk, the space complexity can be improved to  $O(\log(n))$ , at the expense of  $\tilde{O}(nm)$  time complexity [2]. A series of works [11, 7, 14, 5, 16] culminated in a space-time tradeoff for  $\text{USTCON}_{\text{arr}}$  of  $T = \tilde{O}(n^2/S)$  queries for any space bound  $S = \Omega(\log(n))$  and  $S = O(n^2/m)$ , due to Kosowski [19]. While there is no matching time-space lower bound, it is unlikely that this tradeoff can be significantly improved (see [19, Section 5.1 of arXiv v2] for a discussion). Kosowski's algorithm is based on using Metropolis-Hastings random walks to find connections between  $S$  sampled vertices and  $s, t$  until it becomes possible to conclude that  $s$  and  $t$  are connected. For comparison, in the adjacency matrix model, the randomized query complexity of  $\text{USTCON}_{\text{mat}}$  is  $\tilde{\Theta}(n^2)$  and there is no space-time tradeoff.

A quantum algorithm of Dürr, Heiligman, Høyer and Mhalla [13] for CONNECTIVITY can be adapted to solve  $\text{USTCON}_{\text{mat}}$  in  $\tilde{O}(n^{1.5})$  time and  $\text{USTCON}_{\text{arr}}$  in  $\tilde{O}(n)$  time, both of which are optimal up to polylog factors. Both of these algorithms use  $\tilde{O}(n)$  space, of which all but  $O(\log(n))$  can be classical space (assuming quantum RAM access). A subsequent quantum algorithm for  $\text{USTCON}_{\text{mat}}$  due to Belovs and Reichardt uses  $\tilde{O}(n^{1.5})$  time, but only  $O(\log(n))$  space [9], which is optimal in terms of both space and time. It is also possible to solve  $\text{USTCON}_{\text{arr}}$  in  $O(\log(n))$  space and  $\tilde{O}(\sqrt{nm})$  time, using a quantum walk (see for example [8]). This quantum walk algorithm requires a quantum version of array access to the input graph, which we refer to as  $\text{USTCON}_{\text{qw}}$  in the next section.

## 1.1 Summary of results

We describe new quantum walk algorithms for  $\text{USTCON}_{\text{arr}}$ . These algorithms consider a quantum walk version of the adjacency array model, in which the input graph is accessed by a quantum analogue of classical random walk steps. Recall that in the adjacency array model, we assume that for any vertex  $u$ , we can query, for any  $i \in [d_u]$ , the  $i$ -th neighbour of  $u$ ,  $v_i(u)$ . Then a random walk step can be performed from state  $u$  by sampling a uniform  $i \in [d_u]$ , and then computing  $v_i(u)$ , which becomes the current state. In the *quantum walk access model*, we assume that for any vertex  $u$ , we can prepare a uniform superposition over the neighbours of  $u$ . While these models are not identical, they are very similar, and in Section 3, we formally define the models, and show that quantum walk access can be simulated in the array model with polylogarithmic overhead under reasonable additional assumptions.

Letting  $\text{USTCON}_{\text{qw}}$  denote the  $st$ -connectivity problem in the quantum walk access model, we present a one-sided error quantum algorithm that solves  $\text{USTCON}_{\text{qw}}$  in time  $\tilde{O}(n)$  and space  $O(\log(n))$ . Perhaps surprisingly, this means that  $\text{USTCON}_{\text{qw}}$  admits *no* nontrivial tradeoff between space and time in the quantum setting – a single algorithm can solve this problem optimally in terms of both time and space (see Theorem 15 for the formal result).

---

<sup>1</sup> There are variations on the details of this model. For now, we allow  $\text{USTCON}_{\text{arr}}$  to stand in for multiple variations of the array access model, but precise details of the variations can be found in Section 3.

■ **Table 1** A summary of classical (randomized) and quantum time and space complexities for  $\text{USTCON}$  in the adjacency matrix and adjacency array models. The classical results for  $\text{USTCON}_{\text{mat}}$  follow from (1) the  $\log(n)$ -space result for  $\text{USTCON}_{\text{arr}}$  with an  $n/d$  overhead for finding neighbours of the current vertex in a  $d$ -regular graph; and (2) BFS.

USTCON <sub>mat</sub>		
	Time	TS-tradeoffs
Classical	$\tilde{\Theta}(n^2)$	$S = O(\log(n)), T = \tilde{O}(n^3/d)$ $S = \tilde{O}(n), T = \tilde{O}(n^2)$
Quantum	$\tilde{\Theta}(n^{1.5})$	$S = O(\log(n)), T = \tilde{O}(n^{1.5})$ [9]
USTCON <sub>arr</sub>		
	Time	TS-tradeoffs
Classical	$\tilde{\Theta}(m)$	$T = \tilde{O}(\max\{n^2/S, m\})$
Quantum	$\tilde{\Theta}(n)$	$S = O(\log(n)), T = \tilde{O}(n^{1.5})$ $S = T = \tilde{O}(n)$ [13]
		<b>This work:</b> $S = O(\log(n)), T = \tilde{O}(n)$

► **Theorem 1 (Informal).** *There is a  $O(\log(n))$ -space quantum algorithm that decides  $\text{USTCON}_{\text{qw}}$  with one-sided error in  $\tilde{O}(n)$  time.*

In this paper, when we say *time*, we are counting: (1) quantum gates (unitaries that act on at most a constant number of qubits); (2) quantum walk queries to  $G$ ; and (3) (quantum) random access (QCRAM) operations (QCRAM is used in our second algorithm only, see below). Inspired by [19], our algorithm is based on a quantum walk search for  $t$  starting from  $s$  using a random walk that can be interpreted as a Metropolis-Hastings random walk.

Because of the close relationship between  $\text{USTCON}$  and classical logspace, we can consider what this means for logspace problems in general. It does not mean that more space does not reduce the quantum time complexity of *any* problem, but it is interesting to consider: in what settings do we get a non-trivial time-space tradeoff? We consider one such setting: when we are given a promise on the spectral gap or mixing time of the random walk on  $G$  (see Section 2.2). In that case, we prove the following theorem (see Theorem 17 for the formal result).

► **Theorem 2 (Informal).** *Suppose whenever  $s$  and  $t$  are connected, the random walk spectral gap is at least  $\delta > 0$ . For any  $S \in \Omega(\log(n))$ , there is a quantum algorithm that decides  $\text{USTCON}_{\text{qw}}$  with bounded error in  $O(S)$  space and  $T = \tilde{O}\left(\frac{S}{\delta} + \sqrt{\frac{n}{\delta S}}\right)$  time.*

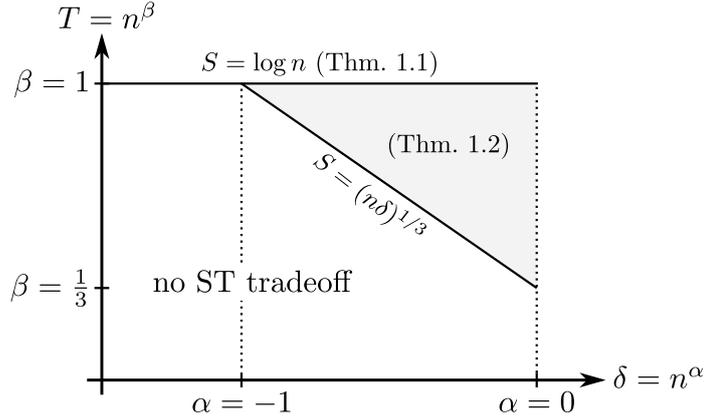
The time bound decreases monotonically for  $S \in \Omega(\log n)$  until  $S \in O((n\delta)^{1/3})$ , at which point it reaches time complexity  $T = \tilde{O}(n^{1/3}/\delta^{2/3})$ . We leave it as an open problem to prove a matching lower bound (at least for some values of  $\delta$ ), which would prove that in certain regimes, it is not possible to achieve optimal time and space simultaneously.

Our algorithm takes inspiration from [3]. In fact, with some imagination, one can derive a similar (but incomparable) time-space tradeoff for  $\text{USTCON}_{\text{qw}}$  from that work: for  $1 \leq S \leq m$ , the algorithm in [3] can be adapted to use space  $\tilde{O}(S)$  and time  $T \in \tilde{O}(S + \sqrt{m/(\delta S)})$ , with  $\delta$  the random walk spectral gap.

## 10:4 (No) Quantum Space-Time Tradeoff for USTCON

In the space bound  $S$  of both algorithms in Theorems 1 and 2, only  $O(\log(n))$  memory needs to be actual quantum workspace (i.e., qubits). The remaining  $O(S)$  memory can be classical RAM in the first algorithm and QCRAM in the second algorithm, that is, classical RAM that is queryable at a quantum superposition of addresses. We discuss the latter in Section 2.4.

We summarize our results in Figure 1. For  $S = \log(n)$ , the algorithm of Theorem 2 has a worse time complexity than the algorithm of Theorem 1, whenever  $\delta < \frac{1}{n}$ . We leave it as an open problem to give a single algorithm that is optimal for all  $\delta$ .



■ **Figure 1** Quantum space-time tradeoffs for USTCON, with axes representing the time complexity and spectral gap promise (up to polylog-factors). The grey area represents the regime in which a non-trivial tradeoff is achieved. Theorem 1 (upper line) corresponds to the regime with space  $S = O(\log n)$  and time  $T = \tilde{O}(n)$ . Theorem 2 (grey area) corresponds to the regime with a promise on  $\delta$ , and interpolates between  $S = O(\log n)$  and  $T = \tilde{O}(n)$ , and  $S = O((n\delta)^{1/3})$  and  $T = \tilde{O}(n^{1/3}/\delta^{2/3})$ .

### Organization

The remainder of this paper is organized as follows. We describe preliminaries in Section 2 and Section 3. In Section 4, we prove Theorem 1 by exhibiting a quantum algorithm for  $\text{USTCON}_{\text{qw}}$  that is optimal in both time and space. For completeness, we also include a proof of a corresponding lower bound in Section 4.1. In Section 5, we prove Theorem 2 exhibiting a quantum time-space tradeoff when given a promise on the spectral gap.

## 2 Preliminaries

We first give some general notation. For a positive integer  $k$ , we let  $[k] = \{1, \dots, k\}$ . Throughout this work,  $n$  denotes the number of vertices and  $m$  the number of edges of the input graph. For any function  $f$ , we let  $\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$ .

### 2.1 Probability theory

A (*probability*) *distribution* on a finite set  $X$  is a non-negative function  $\sigma: X \rightarrow \mathbb{R}_{\geq 0}$  such that  $\sum_{v \in X} \sigma(v) = 1$ . Its *support* is defined as  $\text{supp}(\sigma) := \{v \in X : \sigma(v) > 0\}$ . We will implicitly identify such  $\sigma$  with *row* vectors, as is customary in the random walk literature. To any distribution  $\sigma$ , we also associate a quantum state  $|\sigma\rangle := \sum_{v \in X} \sqrt{\sigma(v)} |v\rangle$ . Measuring  $|\sigma\rangle$  in the standard basis returns a sample from  $\sigma$ .

For any distribution  $\sigma$  on  $X$ , and any subset  $M \subseteq X$ , we will let  $\sigma(M) = \sum_{u \in M} \sigma(u)$ . We let  $\sigma_M$  denote the *normalized restriction* of  $\sigma$  to  $M$ , defined by  $\sigma_M(u) = \sigma(u)/\sigma(M)$  for all  $u \in M$  and  $\sigma_M(u) = 0$  elsewhere.

Finally, the *total variation distance* between two distributions  $\sigma$  and  $\tau$  on  $X$  is defined as

$$\|\sigma - \tau\|_{\text{TV}} := \frac{1}{2} \sum_{u \in X} |\sigma(u) - \tau(u)| = \max_{A \subseteq X} |\sigma(A) - \tau(A)|.$$

## 2.2 Random walks

Fix an undirected graph  $G = (X, E)$  with  $n = |X|$  vertices and  $m = |E|$  edges. We take  $E \subseteq \binom{X}{2}$ , that is, edges  $e \in E$  are subsets  $e = \{u, v\} = \{v, u\}$  of pairs of vertices. We will let

$$N(u) := \{v \in X : \{u, v\} \in E\}$$

denote the neighbourhood of  $u \in X$ , and  $d_u = |N(u)|$  the degree of  $u$ . For convenience we assume that all vertices have positive degree.

Fix edge weights given by a symmetric matrix  $W \in \mathbb{R}_{\geq 0}^{X \times X}$  such that  $W_{u,v} = W_{v,u}$  for all  $u, v \in X$ , and  $W_{u,v} > 0$  if and only if  $\{u, v\} \in E$ . Then  $\bar{G} = (X, E, W)$  defines a *weighted graph*. When no  $W$  is given, the graph is *unweighted* and we let  $W_{u,v} = 1$  for all  $\{u, v\} \in E$ . For  $u \in X$ , define  $w_u = \sum_{v \in X} W_{u,v}$ . The corresponding (*weighted*) *random walk* is the reversible Markov chain on  $X$  with *transition matrix*  $P \in \mathbb{R}_{\geq 0}^{X \times X}$  given by

$$P_{u,v} = \begin{cases} \frac{W_{u,v}}{w_u} & \text{if } \{u, v\} \in E \\ 0 & \text{otherwise} \end{cases} \quad \forall u, v \in X. \quad (1)$$

This means that the probability of moving from the vertex  $u$  along an edge to a neighbouring vertex  $v$  is proportional to the edge's weight. In the unweighted case, this is called the *simple random walk*; in each step it simply moves to a neighbouring vertex chosen uniformly at random.

Let  $\pi \in \mathbb{R}_{> 0}^X$  be the distribution defined by

$$\pi(u) = \frac{w_u}{\mathcal{W}(G)} \quad \forall u \in X,$$

where  $\mathcal{W}(G) = \sum_{u \in X} w_u = \sum_{u,v \in X} W_{u,v}$ . In the unweighted case,  $\pi$  is proportional to the degree. The distribution  $\pi$  is a *stationary* distribution of the random walk, i.e.,  $\pi P = \pi$  (it is a left eigenvector of  $P$  with eigenvalue 1).

In fact, when the graph  $G$  is connected,  $\pi$  is also the *unique* stationary distribution of  $P$ . If in addition the graph is not bipartite, then all other eigenvalues have absolute value strictly less than one. That is, if  $1 = \lambda_1 \geq \dots \geq \lambda_n \geq -1$  are the eigenvalues of  $P$  then the (*absolute spectral gap*)  $\gamma_* = \gamma_*(G) := \min\{1 - |\lambda_j| : j = 2, \dots, n\} = \min\{1 - \lambda_2, 1 + \lambda_n\}$  is strictly positive. Importantly, the inverse of the spectral gap bounds the random walk's *mixing time*, that is, the time required for convergence to the stationary distribution:

► **Theorem 3** ([21, Thm. 12.4]). *Assume  $G$  is connected and not bipartite. Let  $\varepsilon > 0$  and*

$$t \geq \frac{1}{\gamma_*} \log \left( \frac{1}{\varepsilon \pi_{\min}} \right),$$

where  $\pi_{\min} = \min_{u \in X} \pi(x)$ . Then  $\|\sigma P^t - \pi\|_{\text{TV}} \leq \varepsilon$  for any distribution  $\sigma$  on  $X$ .

Conversely, it is known that  $t \geq (\frac{1}{\gamma_*} - 1) \log(\frac{1}{2\varepsilon})$  is necessary to ensure mixing from an arbitrary initial distribution [21, Thm. 12.5]. In the unweighted case, we have  $\pi_{\min} \geq \frac{d_{\min}}{n d_{\max}} \geq \frac{1}{n^2}$ , so the former shows that Theorem 3 is tight up to  $\log(n)$  factors in that case.

Finally, for any  $s, t \in X$  we let  $\mathcal{H}_{s,t}$  denote the *hitting time* from  $s$  to  $t$ , which is the expected number of steps needed to reach  $t$  in a random walk starting from  $s$ . We let  $\mathcal{C}_{s,t} = \mathcal{H}_{s,t} + \mathcal{H}_{t,s}$  denote the *commute time* between  $s$  and  $t$  – the expected number of steps needed to reach  $t$  and then return to  $s$  in a random walk starting from  $s$ . These quantities are finite if and only if  $s$  and  $t$  are in the same component of  $G$ . More generally, the commute time  $\mathcal{C}_{s,M}$  from  $s$  to a subset  $M \subseteq X$  is the expected number of steps needed to reach any vertex in  $M$  and then return to  $s$  in a random walk starting from  $s$ .

### 2.3 Quantum walk search algorithms

Quantum walk search refers to the use of quantum walks to find certain “marked” elements on a graph. We will use quantum walk search to search for a vertex connected to  $t$  in the connected component of  $S$ . Specifically, we will use the following special case of [4, Thm. 13].<sup>2</sup>

► **Theorem 4.** *Let  $P$  be a random walk on a weighted graph with vertex set  $X$ ,  $M \subseteq X$  a subset of “marked” vertices, and  $s \in X$ . Let  $\mathcal{C}$  be the (quantum) time complexity to check for a given  $u \in X$  whether  $u \in M$ , let  $\mathcal{U}$  be the time complexity of implementing the weighted quantum walk oracle*

$$|u\rangle |0\rangle \mapsto \sum_{v \in N(u)} \sqrt{P_{u,v}} |u\rangle |v\rangle.$$

*in space  $O(\log(n))$ . Let  $\mathcal{C}$  be a known upper bound on the commute time  $\mathcal{C}_{s,M}$  in the case where  $s$  and  $M$  are connected (and in particular  $M \neq \emptyset$ ). Then there is a quantum algorithm that, if  $M \neq \emptyset$  and  $s$  is connected to  $M$ , finds an element of  $M$  with probability at least  $2/3$ . If  $M = \emptyset$  or  $s$  is not connected to  $M$ , then the algorithm outputs a vertex not in  $M$ . The algorithm has time complexity  $O(\sqrt{\log(\mathcal{C})} \log(n) + \sqrt{\mathcal{C} \log(\mathcal{C}) \log(\log(\mathcal{C}))}(\mathcal{C} + \mathcal{U}))$  and space complexity  $O(\log(n))$ .*

### 2.4 Quantum RAM

Our algorithm will exploit the given space by saving sets of vertices which will be either connected to  $s$  or to  $t$ . For our quantum algorithm to access this space, we assume access to a so-called quantum-classical random access memory or *QCRAM*. This refers to a memory that only stores classical information, but can be queried at a superposition of addresses. More specifically, an  $R$ -bit QCRAM stores a string of bits  $q \in \{0, 1\}^R$  so that the following operations are supported in time  $\text{polylog}(R)$ :

1. *UPDATE*( $i, x$ ): store  $x \in \{0, 1\}$  in the  $i$ -th bit (i.e., set  $q_i = x$ ).
2. *QUERY*: for any superposition  $\sum_i \alpha_i |i\rangle |s_i\rangle$ , it maps

$$\sum_i \alpha_i |i\rangle |s_i\rangle \mapsto \sum_i \alpha_i |i\rangle |s_i \oplus q_i\rangle.$$

As was first described by Kerenidis and Prakash [18], using such a QCRAM we can set up a data structure to generate quantum superpositions over elements in the QCRAM. We will use the following formulation based on [3].

<sup>2</sup> To see that this follows from [4, Thm. 13], note that when  $|\sigma\rangle = |s\rangle$ , the cost to set up  $|\sigma\rangle$  is  $\log(n)$  and the value  $\mathcal{C}_{\sigma,M}$  from [4] is exactly the commute time from  $s$  to  $M$  [4, Thm. 4].

► **Lemma 5.** Fix integer parameters  $\ell$  and  $k$ . Using an  $O(k\ell \log(\ell))$ -bit QCRAM, there is a data structure,  $D$ , that stores up to  $\ell$  elements  $x \in \{0, 1\}^k$  with associated integer weights,  $c_x$ , of bounded absolute value for some  $\text{poly}(\ell)$  bound, and supports the following operations in time  $O(k \cdot \text{polylog}(k\ell))$  per operation:

1. insertion or deletion of a pair  $(x, c_x)$ ,
2. quantum queries of the form “Is  $x \in D$ ?”,
3. preparation of the quantum state  $\frac{1}{\sqrt{\sum_{x \in D} c_x}} \sum_{x \in D} \sqrt{c_x} |x\rangle$ .

### 3 USTCON and the Quantum Walk model

In this section we define the undirected  $st$ -connectivity problem (USTCON). The input to this problem is an undirected graph  $G = (X, E)$ . Classically, there are various ways this input may be given, which may change the complexity of the problem. For example, in the *adjacency array model* (defined below), it is possible to randomly sample a neighbour of any vertex  $u$  in  $O(1)$  queries to  $G$  (assuming access to the vertex degrees), facilitating a random walk on  $G$ , whereas if  $G$  is given as an *adjacency matrix*, a random walk step is not so simple.

We will work in a quantum walk analogue of the adjacency array model. We assume that  $G$  can be accessed via the *quantum walk oracle* that for every  $u \in X$  outputs a uniform superposition over its neighbours:<sup>3</sup>

$$\mathcal{O}_W : |u\rangle |0\rangle \mapsto \frac{1}{\sqrt{d_u}} \sum_{v \in N(u)} |u\rangle |v\rangle. \quad (2)$$

Formally, we describe  $\text{USTCON}_{\text{qw}}$  in terms of the input and output.

► **Problem 6** ( $\text{USTCON}_{\text{qw}}$ ). Given access to an undirected graph  $G = (X, E)$  via the quantum walk oracle  $\mathcal{O}_W$ , and two vertices  $s, t \in X$ , decide whether  $s$  and  $t$  are in the same connected component of  $G$ .

To compare our work with classical results on  $\text{USTCON}_{\text{arr}}$ , we describe an implementation of the quantum walk oracle defined above based on adjacency array access to a graph. Let  $u \in X$  and  $i \in [d_u]$ . We assume that for each vertex  $u$  there is a fixed numbering of its neighbours from 1 to  $d_u$ . In the *adjacency array model*, two types of queries are allowed:

- Degree query  $\mathcal{O}_D : |0\rangle |u\rangle \mapsto |d_u\rangle |u\rangle$
- Neighbour query  $\mathcal{O}_N : |u\rangle |i\rangle |0\rangle \mapsto |u\rangle |i\rangle |v_i(u)\rangle$

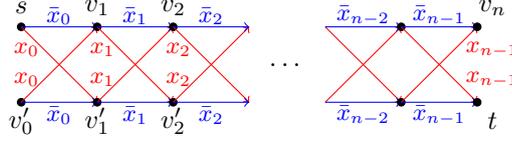
In the *sorted adjacency array model* we additionally assume that for every vertex  $u \in X$  its neighbours are sorted: for any  $i, j \in [d_u]$ , if  $i < j$  then  $v_i(u) < v_j(u)$ . In particular, this allows us to check with  $O(\log(n))$  queries whether a given pair of vertices  $u, v$  are adjacent. We define the USTCON-problem in this model as follows.

► **Problem 7** ( $\text{USTCON}_{\text{s-arr}}$ ). Given access to an undirected graph  $G = (X, E)$  via the sorted adjacency array model and two vertices  $s, t \in X$ , decide whether  $s$  and  $t$  are in the same connected component of  $G$ .

The question that we consider is how many sorted adjacency array queries to the graph it takes to implement the quantum walk oracle  $\mathcal{O}_W$ .

<sup>3</sup> Note that this is exactly the quantum walk oracle defined in Theorem 4, specialized to unweighted graphs.

## 10:8 (No) Quantum Space-Time Tradeoff for USTCON



■ **Figure 2** The parity graph. We include an edge labelled by “ $x_i$ ” (in red) if and only if  $x_i = 1$ , and an edge labelled “ $\bar{x}_i$ ” (in blue) if and only if  $x_i = 0$ , meaning that for each vertex we include exactly one of the two incoming edges, and exactly one of the two outgoing edges. The resulting graph has  $s$  and  $t$  connected if and only if  $\text{PARITY}(x) = 1$ .

► **Lemma 8.** *The quantum walk oracle  $\mathcal{O}_W$  for an unweighted graph  $G$  (Equation (2)) with maximum degree  $d_{\max}$  can be implemented with  $O(\log(d_{\max}))$  queries in the sorted adjacency array model, and  $\tilde{O}(1)$  other elementary operations and space.*

The proof of this lemma is deferred to the appendix.

It follows that any quantum algorithm solving  $\text{USTCON}_{\text{qw}}$  in  $T$  time and  $S = \Omega(\log(n))$  space can solve  $\text{USTCON}_{\text{s-arr}}$  in  $\tilde{O}(T)$  time and  $O(S)$  space.

### 4 Time- and space-optimal quantum algorithm

In Section 4.2 and Section 4.3, we give an algorithm for  $\text{USTCON}_{\text{qw}}$  that is optimal in both time and space. For completeness, we first give a time lower bound in Section 4.1.

#### 4.1 Lower bound

The proof of the following lower bound follows the lines of the proof of an analogous lower bound for the strong connectivity problem described in [13]. The proof is via a reduction from  $\text{PARITY}$ .

► **Problem 9** ( $\text{PARITY}$ ). *Given oracle access to a string  $x \in \{0, 1\}^n$  via  $\mathcal{O}_x : |i\rangle |b\rangle \mapsto |i\rangle |b \oplus x_i\rangle$ , return  $\bigoplus_{i=0}^{n-1} x_i$ .*

► **Lemma 10** ([6, 15]). *The bounded error quantum query complexity of  $\text{PARITY}$  is  $\Omega(n)$ .*

We use Lemma 10 and a reduction from parity, using the parity graph illustrated in Figure 2, to show the following result. The detailed proof can be found in the appendix.

► **Theorem 11.** *The bounded error quantum query complexity of  $\text{USTCON}_{\text{s-arr}}$  and  $\text{USTCON}_{\text{qw}}$  is  $\Omega(n)$ .*

#### 4.2 Metropolis-Hastings walk

In this section, we consider an unweighted simple graph  $G$ . The algorithm that we propose involves a quantum walk on a modified weighted version of  $G$  that we call  $G' = (X', E', W)$ . We start by describing the construction of  $G'$  that was introduced in [19, arXiv v2].

► **Definition 12** (Metropolis-Hastings walk). *For any graph  $G = (X, E)$ , the corresponding Metropolis-Hastings walk is the random walk on the weighted graph  $G' = (X', E', W)$  defined as follows. For every  $u \in X$ , we include a corresponding vertex  $x_u$  in  $X'$ . In addition, for every edge  $\{u, v\} \in E$ , we add a new vertex  $x_{u,v}$  that splits the edge into two new edges. Formally:*

$$X' = \{x_u : u \in X\} \cup \{x_{u,v} : \{u,v\} \in E, u < v\}$$

$$E' = \{\{x_u, x_{u,v}\} : \{u,v\} \in E\}.$$

For every edge  $\{x_u, x_{u,v}\} \in E'$ , we define the edge weight  $W_{x_u, x_{u,v}} = \frac{1}{d_u}$ . These weights define transition probabilities for the random walk on  $G'$ .

The above has been called the *cautious walk* in [19, arXiv v2], while *Metropolis-Hastings-type* walks are walks in which neighbours are sampled and accepted with some probability. Our terminology is motivated by the following observation. If we start with a vertex  $u \in X$  and take two steps of the walk of Definition 12, then we arrive at another vertex  $v \in X$ , which is either the same or a neighbour of  $u$  in  $G$ . The walk on  $G$  defined this way has the following alternative description: sample a uniformly random neighbour and accept it with probability  $\frac{1/d_v}{1/d_u + 1/d_v} = \frac{1}{1+d_v/d_u}$ . This is precisely a random walk that falls into the Metropolis-Hastings framework, justifying our terminology. The precise choice of acceptance probabilities is sometimes called the *Glauber choice* in the literature (e.g., [20]). We note that a later version of [19] uses another choice of Metropolis-Hastings walk, but of our purposes we find it convenient to stick to the walk as defined above.

While the hitting time of a random walk between two vertices in  $G$  may be as high as  $O(n^3)$ , in  $G'$  it is at most  $O(n^2)$  [19, Lemma 2 of arXiv v2]:

► **Lemma 13** ([19]). *Let  $G = (X, E)$  be any unweighted graph, and  $G'$  the corresponding (weighted) Metropolis-Hastings graph as in Definition 12. For any  $u, v \in X$  connected by a path,  $\mathcal{H}_{u,v}(G') \leq 18n^2$ .*

In order to apply Theorem 4 to  $G'$ , we need to upper bound  $\mathbf{U}$ , the cost of implementing the weighted quantum walk oracle. For  $u \in X'$  the oracle is defined as

$$U : |x\rangle |0\rangle \mapsto \sum_{y \in N(x)} \sqrt{P'_{xy}} |x\rangle |y\rangle, \quad \forall x \in X'$$

where  $P'_{xy}$  is the probability of walking from  $x$  to  $y$  defined by the edge weights.

► **Lemma 14.** *The weighted quantum walk oracle  $U$  for the Metropolis-Hastings walk  $G'$  can be implemented with  $\tilde{O}(1)$  degree queries  $\mathcal{O}_D$ ,  $\tilde{O}(1)$  applications of the quantum walk operator  $\mathcal{O}_W$  on the graph  $G$ , and  $\tilde{O}(1)$  additional gates.*

*Therefore,  $U$  can be implemented with  $\tilde{O}(1)$  queries to  $G$  in the sorted adjacency array model, and  $\tilde{O}(1)$  additional gates.*

**Proof.** Note that the first statement implies the second one due to Lemma 8. Therefore, we will only prove the first statement. Consider the following encoding of the vertices of  $G'$ .

$$X' = \{(u, 0) : u \in X\} \cup \{(u, v) : \{u, v\} \in E, u < v\} \subseteq X \times (X \cup \{0\}),$$

where 0 is a null symbol not contained in  $X$ . The first set of this union corresponds to original vertices of  $G$  and the second one corresponds to the added ones.

To implement  $U$  on  $|x\rangle |0\rangle$ , we first compute a bit in an ancilla register  $A$  that is  $|0\rangle_A$  if  $x = (u, 0)$  for some  $u$ , and  $|1\rangle_A$  otherwise. We will condition on this value.

First, conditioned on  $|0\rangle_A$ , our implementation proceeds as follows, for  $u \in X$ :

$$|x_u\rangle |0\rangle = |u, 0\rangle |0\rangle \xrightarrow{J_{\rightarrow}^{\mathcal{O}_W}} \frac{1}{\sqrt{d_u}} \sum_{v \in N(u)} |u\rangle |v\rangle |u, v\rangle$$

$$\xrightarrow{J'} \frac{1}{\sqrt{d_u}} \sum_{v \in N(u)} |u, 0\rangle |u, v\rangle = \frac{1}{\sqrt{d_u}} \sum_{v \in N(u)} |x_u\rangle |x_{u,v}\rangle = U |x_u\rangle |0\rangle$$

## 10:10 (No) Quantum Space-Time Tradeoff for USTCON

where  $J$  is a unitary that acts as  $|u, v\rangle |0\rangle \mapsto |u, v\rangle |u, v\rangle$ , and  $J'$  as  $|u, v\rangle |u, v\rangle \mapsto |u, 0\rangle |u, v\rangle$ , each of which can be implemented with  $O(\log(n))$  controlled-NOT gates (since every vertex is described by  $O(\log(n))$  bits).

Next, conditioned on  $|1\rangle_{A'}$ , our implementation proceeds as follows, for  $\{u, v\} \in E$  with  $u < v$ , and  $|0\rangle_{A'}$  a fresh ancilla:

$$\begin{aligned}
|0\rangle_{A'} |x_{u,v}\rangle |0\rangle &= |0\rangle_{A'} |u, v\rangle |0\rangle \\
&\xrightarrow{1} \left( \sqrt{\frac{1/d_u}{1/d_u + 1/d_v}} |0\rangle + \sqrt{\frac{1/d_v}{1/d_u + 1/d_v}} |1\rangle \right)_{A'} |u, v\rangle |0\rangle \\
&\xrightarrow{2} \sqrt{\frac{1/d_u}{1/d_u + 1/d_v}} |0\rangle_{A'} |u, v\rangle |u\rangle + \sqrt{\frac{1/d_v}{1/d_u + 1/d_v}} |1\rangle_{A'} |u, v\rangle |v\rangle \\
&\xrightarrow{3} |0\rangle_{A'} \left( \sqrt{\frac{1/d_u}{1/d_u + 1/d_v}} |u, v\rangle |u\rangle + \sqrt{\frac{1/d_v}{1/d_u + 1/d_v}} |u, v\rangle |v\rangle \right) \\
&= |0\rangle_{A'} \left( \sqrt{\frac{W_{x_{u,v}, x_u}}{w(x_{u,v})}} |x_{u,v}\rangle |x_u\rangle + \sqrt{\frac{W_{x_{u,v}, x_v}}{w(x_{u,v})}} |x_{u,v}\rangle |x_v\rangle \right) \\
&= |0\rangle_{A'} U |x_{u,v}\rangle |0\rangle,
\end{aligned}$$

where we use the following mappings:

- 1: Query degrees for  $u$  and  $v$  into a new ancilla register, perform the rotation controlled on the degrees (cf. [17]), and then uncompute the degrees ( $O(1)$  degree queries to  $G$ ).
- 2: Controlled on the first register, select one of the two vertices to copy into the last register ( $O(\log(n))$  Toffoli gates).
- 3: Flip the bit in  $A'$  if the second vertex of  $|u, v\rangle$  is the same as the one written in the third register ( $O(\log(n))$  elementary gates).

To complete the proof, note that we can uncompute the bit in ancilla  $A$ , because the register containing  $|x\rangle$  has not been changed.  $\blacktriangleleft$

### 4.3 The algorithm

We can solve  $\text{USTCON}_{\text{qw}}(G)$  using Algorithm 1. This leads to our main theorem of this section.

■ **Algorithm 1** Quantum algorithm for  $\text{USTCON}_{\text{qw}}$  with optimal time and space.

---

Apply the algorithm from Theorem 4 to the Metropolis-Hastings walk  $P'$  with  $M = \{t\}$ , using Lemma 14 to implement the quantum walk oracle for  $G'$ . If the algorithm returns  $t$ , output “connected”, and otherwise output “disconnected”.

---

► **Theorem 15.** *There exists a  $O(\log(n))$ -space quantum algorithm that decides  $\text{USTCON}_{\text{qw}}$  and  $\text{USTCON}_{\text{s-arr}}$  with bounded one-sided error in  $\tilde{O}(n)$  gates and queries.*

**Proof.** Let  $X_s \subseteq X$  denote the connected component of  $s$ . If  $t \in X_s$ , the algorithm will output  $t$  with probability at least  $2/3$ , in which case our algorithm will output the correct answer, “connected”. If  $t \notin X_s$ , then the algorithm will output an element of  $X_s$  with probability 1, in which case, our algorithm will output the correct answer “disconnected”. This establishes correctness of Algorithm 1 with one-sided error.

To analyze the complexity, note that by Lemma 14 we have  $U = \tilde{O}(1)$ . For any  $u \in X_s$ , we can check if  $u \in M$  by checking if  $u = t$  in complexity  $C = O(\log(n)) = \tilde{O}(1)$ . To complete the analysis, we need only upper bound the commute time between  $s$  and  $t$  when  $t \in X_s$ . Since  $C_{s,t} = \mathcal{H}_{s,t} + \mathcal{H}_{t,s}$ , by Lemma 13, we have  $C_{s,t} \leq 36n^2 =: \mathcal{C}$ . Thus, referring to Theorem 4, the complexity of our algorithm is:

$$\tilde{O}\left(\sqrt{\mathcal{C} \log(\mathcal{C}) \log(\log(\mathcal{C}))}(\mathcal{C} + U)\right) = \tilde{O}(n). \quad \blacktriangleleft$$

## 5 Time-space tradeoff for bounded spectral gap

In this section we revisit the problem of undirected  $st$ -connectivity in the setting where one is given a lower bound on the spectral gap of the random walk. As discussed in Section 2.2, such a bound is tightly related to the mixing time of the walk. We will give a quantum algorithm that exhibits a nontrivial time-space tradeoff in this setting.

Our discussion will be general and apply to random walks on weighted graphs as defined in Equation (1). This is useful since the spectral gaps and mixing times of random walks on  $G$  with different edge weights are in general *not* comparable. E.g., on the lollipop graph (an  $n$ -vertex clique connected to an  $O(n)$ -vertex path) the mixing time of the unweighted random walk is  $\Theta(n^3)$  [10], while it is  $O(n^2)$  for the Metropolis-Hastings walk.<sup>4</sup> On the other hand, on an  $n$ -vertex star graph the unweighted random walk has mixing time  $O(1)$  while the Metropolis-Hastings walk has mixing time  $\Theta(n)$ . Thus, while the specific edge weights do not affect whether  $s$  and  $t$  are connected, they do impact the algorithm. Throughout this section, we assume some fixed edge weights are given, and we do not try to optimize for “good” edge weights. More specifically, we assume access to a *weighted quantum walk oracle* that for every vertex outputs a superposition of its neighbours, with squared amplitudes proportional to the edge weights:

$$\mathcal{O}_W: |u\rangle |0\rangle \mapsto \sum_{v \in N(u)} \sqrt{\frac{W_{u,v}}{w_u}} |u\rangle |v\rangle = \sum_{v \in N(u)} \sqrt{P_{u,v}} |u\rangle |v\rangle \quad \forall u \in X$$

Moreover, we assume access to the weighted vertex degrees  $w_u$  and that these degrees are of bounded absolute value for some  $\text{poly}(n)$  bound. This will allow us to generate the state  $|\pi_{X'}\rangle$  for any subset  $X' \subseteq X$  stored in QCRAM.

► **Problem 16** ( $\text{USTCON}_{\text{qw},\delta}$ ). *Given access to an undirected weighted graph via the quantum walk oracle  $\mathcal{O}_W$ , two vertices  $s, t \in X$ , and the promise that either  $s$  and  $t$  are disconnected or the spectral gap of the transition matrix of the walk is at least some  $\delta > 0$ , decide which is the case.*

Our main result of this section is the following:

► **Theorem 17.** *Fix  $\delta \geq 0$ . Let  $\mathcal{G}_n$  be a family of undirected weighted graphs  $G = (X, E, W)$  with  $n = |X|$ , such that  $\gamma_*(G) \geq \delta$  whenever  $s$  and  $t$  are connected. Then for any  $S = \Omega(\log(n))$ , there is a quantum algorithm that decides  $\text{USTCON}_{\text{qw},\delta}$  on  $\mathcal{G}_n$  with bounded error in  $O(S)$  space – of which  $O(\log(n/\delta))$  is quantum memory, and the remainder is QCRAM – and  $T = \tilde{O}\left(\frac{S}{\delta} \log\left(\frac{1}{\pi_{\min}}\right) + \sqrt{\frac{n}{\delta S}}\right)$  queries to  $\mathcal{O}_W$ , elementary gates, and QCRAM queries.*

<sup>4</sup> This follows from the  $O(n^2)$  upper bound on the maximum hitting time of the Metropolis-Hastings walk (Lemma 13), and the fact that the maximum hitting time upper bounds the mixing time [21, Lemma 10.2].

## 10:12 (No) Quantum Space-Time Tradeoff for USTCON

Note that we can assume  $\delta \geq 1/n$ . If  $\delta < 1/n$  then  $T \approx S/\delta$ . There is no time-space tradeoff, and it is always faster to run the Metropolis-Hastings algorithm (Algorithm 1).

The algorithm is stated below as Algorithm 2. It consists of three stages. We fix some parameter  $p$ , which denotes the number of “pebbles”, or vertices the algorithm will keep track of (so  $S = O(p \log(n))$ ). First, we run  $O(p)$  classical random walks starting from  $s$ , each of length  $\ell = O\left(\frac{1}{\delta} \log\left(\frac{n}{p\pi_{\min}}\right)\right)$ . This allows us to sample a set  $L$  of  $O(p)$  points from  $X_s$ , the connected component of  $s$  (the big-O notation suppresses a universal constant that is given in the proof). Since  $\ell$  is at least the mixing time of  $G$  (see Theorem 3), assuming  $s$  and  $t$  are connected, each point is sampled (approximately) from  $\pi$ . We do the same from  $t$  to get a random subset  $M \subseteq X_t$  connected to  $t$ .

Next, we use  $L$  and  $M$  to prepare (up to some error) the states  $|\pi_{X_s}\rangle$  and  $|\pi_{X_t}\rangle$ , using *inverse quantum walk search*, which we describe in more detail in Section 5.2. If  $s$  and  $t$  are in the same connected component, then  $|\pi_{X_s}\rangle = |\pi_{X_t}\rangle$ , and otherwise, the states are orthogonal. The final step is to distinguish these two cases using a SWAP test [1, Claim 1]. This roughly follows an earlier approach in [3], the main difference being that we sample the sets  $L$  and  $M$  using a random walk (which allows us to exploit the gap promise), while in [3] the sets are constructed using a breadth-first search.

■ **Algorithm 2** Quantum algorithm for  $\text{USTCON}_{\text{qw}}$  with a tradeoff.

---

**Seed set:** Run  $O(p)$  classical random walks from  $s$  and  $O(p)$  classical random walks from  $t$ , each for  $O\left(\frac{1}{\delta} \log\left(\frac{n}{p\pi_{\min}}\right)\right)$  steps. Let  $L$  and  $M$  denote the respective sets of endpoints, without duplicates. If  $L \cap M \neq \emptyset$ , return “connected”.<sup>5</sup>

**State preparation:** Run inverse quantum walk search from  $|\pi_L\rangle$  and  $|\pi_M\rangle$  for time  $\tilde{O}\left(\sqrt{\frac{n}{\delta p}}\right)$  to prepare  $|\pi_{X_s}\rangle$  and  $|\pi_{X_t}\rangle$ , respectively, to precision  $1/8$ .

**SWAP test:** Do a SWAP test on the resulting states. If the test returns “0”, return “connected”, otherwise return “disconnected”.

---

If we specialize Theorem 17 to the unweighted graph case, we get the following corollary.

► **Corollary 18.** *For any  $S > 0$ , there is a quantum algorithm that solves  $\text{USTCON}_{S\text{-arr}}$  with a promise  $\delta > 0$  on the random walk spectral gap using space  $S$  and time  $T \in \tilde{O}(S/\delta + \sqrt{n}/(\delta S))$ .*

An analogous result holds for the Metropolis-Hastings walk described in Section 4.2 given a promise on its spectral gap, since we showed that the corresponding quantum walk oracle can also be efficiently implemented.

In the remainder of this section we will analyze each stage of Algorithm 2.

### 5.1 Analysis of step 1: Seed set

Recall that the first stage results in random sets  $L = \{x_1, \dots, x_{cp}\}$  and  $M = \{y_1, \dots, y_{cp}\}$ , where  $x_1, \dots, y_{cp}$  are the endpoints of independent random walks starting at  $s$  or  $t$ , respectively, and  $c > 0$  is some universal constant that we will choose later. Since we run those random walks for  $O\left(\frac{1}{\delta} \log\left(\frac{n}{p\pi_{\min}}\right)\right)$  steps, by Theorem 3 it follows that the  $x_j$  are independent

---

<sup>5</sup> We use the following strategy to check whether  $L \cap M \neq \emptyset$ : Sort  $L$  and  $M$  and check for every element of  $L$  if it is present in  $M$ . This takes  $\tilde{O}(p)$  time and space.

samples drawn from a distribution  $\tilde{\pi}$  such that

$$\|\tilde{\pi} - \pi\|_{\text{TV}} \leq \frac{p}{8n},$$

where  $\pi$  is the stationary distribution on  $X_s$ . If  $s$  and  $t$  are connected then  $X_s = X_t$  and the samples  $y_j$  are similarly drawn from a distribution that is  $p/(8n)$ -close to  $\pi$ . Here we prove that this implies lower bounds on the stationary measure of the sets  $L$  and  $M$ .

► **Proposition 19.** *There exists a universal constant  $c > 0$  such that the following holds. Let  $p \in [n]$  and assume  $L \subseteq X$  is a random set obtained by sampling  $cp$  independent elements from a distribution  $\tilde{\pi}$  such that  $\|\tilde{\pi} - \pi\|_{\text{TV}} \leq \frac{p}{8n}$  (and removing duplicates). Then,  $\Pr(\pi(L) \geq \frac{p}{8n}) \geq \frac{9}{10}$ .*

The proof of the proposition uses the following lemma, which formalizes the intuition that adding a random element to a low-probability subset should increase the probability.

► **Lemma 20.** *Let  $X$  be a set of cardinality  $n$ ,  $A \subseteq X$  an arbitrary fixed subset, and let  $b$  be drawn at random from an arbitrary distribution  $\sigma$  on  $X$ . Then:*

$$\Pr\left(\sigma(A \cup \{b\}) > \sigma(A) + \frac{1 - \sigma(A)}{2n}\right) \geq \frac{1 - \sigma(A)}{2}.$$

**Proof.** Say  $b$  is *bad* if  $b \in A$  or  $\sigma(b) \leq \frac{1 - \sigma(A)}{2n}$ , and *good* otherwise. Then

$$\Pr\left(\sigma(A \cup \{b\}) > \sigma(A) + \frac{1 - \sigma(A)}{2n}\right) = \Pr(b \text{ is good}) = 1 - \Pr(b \text{ is bad}).$$

We can compute:

$$\begin{aligned} \Pr(b \text{ is bad}) &= \sigma\left(A \cup \left\{x \in X : \sigma(x) \leq \frac{1 - \sigma(A)}{2n}\right\}\right) \\ &\leq \sigma(A) + n \cdot \frac{1 - \sigma(A)}{2n} = \frac{1 + \sigma(A)}{2}, \end{aligned}$$

from which the claim follows. ◀

**Proof of Proposition 19.** Let  $x_1, x_2, \dots$  denote samples drawn independently at random from  $\tilde{\pi}$ . For any integer  $T \geq 1$ , define  $L_T := \{x_1, \dots, x_T\}$  as the set consisting of the first  $T$  samples (with duplicates removed), as well as  $L_0 := \emptyset$ . We say that the  $T$ -th sample is a *success* if

$$\tilde{\pi}(L_{T-1}) \geq \frac{1}{2} \quad \text{or} \quad \tilde{\pi}(L_T) \geq \tilde{\pi}(L_{T-1}) + \frac{1}{4n},$$

and a *failure* otherwise. Let  $T_j$  denote the index of the  $j$ -th success, with  $T_0 := 0$ . Then, clearly,

$$\tilde{\pi}(L_{T_p}) \geq \min\left\{\frac{1}{2}, \frac{p}{4n}\right\} = \frac{p}{4n} \quad \text{and hence} \quad \pi(L_{T_p}) \geq \tilde{\pi}(L_{T_p}) - \frac{p}{8n} \geq \frac{p}{8n}.$$

On the other hand, note that by Lemma 20, the  $T$ -th sample is a success with probability at least  $\frac{1}{4}$ , even if we condition on all prior samples. In particular, the probability that there are  $k$  failures in a row is at most  $(\frac{3}{4})^k$ . Therefore,

$$\mathbf{E}[T_j - T_{j-1}] = \sum_{k=1}^{\infty} k \Pr(T_j = T_{j-1} + k) \leq \sum_{k=1}^{\infty} k \left(\frac{3}{4}\right)^k = 12$$

## 10:14 (No) Quantum Space-Time Tradeoff for USTCON

and hence, by Markov's inequality,

$$\Pr(T_p > cp) \leq \frac{1}{cp} \mathbf{E}[T_p] = \frac{1}{cp} \sum_{j=1}^p \mathbf{E}[T_j - T_{j-1}] \leq \frac{12p}{cp} = \frac{1}{10}$$

provided we choose  $c := 120$ . Since the random set  $L$  in the statement of the lemma is defined by taking  $cp$  many samples, we obtain that

$$\Pr\left(\pi(L) \geq \frac{p}{8n}\right) \geq \Pr(T_p \leq cp) \geq 1 - \frac{1}{10} = \frac{9}{10}. \quad \blacktriangleleft$$

### 5.2 Analysis of step 2: State preparation

Now we turn to the analysis of the quantum walk search routine in step 2 of the algorithm. We rely on the following proposition from [3], which formalizes the idea of “inverse quantum walk search”.<sup>6</sup>

► **Proposition 21** ([3, Proposition 1]). *Consider a subset  $A \subseteq X$  of a (connected) graph  $G$ , and let  $\delta$  be a lower bound on the spectral gap of a random walk  $P$  on  $G$  with stationary distribution  $\pi$ . From  $|\pi_A\rangle$ , we can generate a state  $|\tilde{\pi}\rangle = |\pi\rangle + |\Gamma\rangle$  with  $\|\Gamma\|_2 \leq \epsilon$  using an expected number of calls to the weighted quantum walk oracle*

$$O\left(\frac{1}{\sqrt{\pi(A)\delta}} \log\left(\frac{1}{\pi(A)\epsilon}\right)\right),$$

and  $O(1/\sqrt{\pi(A)\delta})$  reflections around  $|\pi_A\rangle$ . The algorithm uses space logarithmic in  $n$ ,  $1/\delta$ ,  $1/\pi(A)$  and  $1/\epsilon$ .

The proposition implies the following.

▷ **Claim 22.** Step 2 of Algorithm 2 prepares  $1/8$ -approximations of  $|\pi_{X_s}\rangle$  and  $|\pi_{X_t}\rangle$  with probability at least  $9/10$  in time complexity  $O\left(\sqrt{\frac{n}{p\delta}} \log\left(\frac{n}{p}\right)\right)$ .

### 5.3 Analysis of step 3: SWAP test

In the last step of our algorithm, we wish to decide whether  $|\pi_{X_s}\rangle = |\pi_{X_t}\rangle$  or whether they are orthogonal. For this we use the SWAP test.

▷ **Claim 23.** Step 3 of Algorithm 2 decides whether  $|\pi_{X_s}\rangle = |\pi_{X_t}\rangle$  or whether they are orthogonal with constant probability in time  $\tilde{O}(1)$ .

*Proof.* Using a single copy of two states  $|\psi\rangle$  and  $|\psi'\rangle$ , and  $O(\log(n))$  additional gates, the SWAP test returns “0” with probability  $(1 + |\langle\psi|\psi'\rangle|^2)/2$  and “1” with probability  $(1 - |\langle\psi|\psi'\rangle|^2)/2$  [1, Claim 1].

By Claim 22, in step 2 we prepared states  $|\tilde{\pi}_{X_s}\rangle = |\pi_{X_s}\rangle + |\Gamma_L\rangle$  and  $|\tilde{\pi}_{X_t}\rangle = |\pi_{X_t}\rangle + |\Gamma_M\rangle$  such that  $\|\Gamma_L\|_2, \|\Gamma_M\|_2 \leq 1/8$  with probability  $9/10$ . By a triangle inequality, this implies that

$$\left| |\langle\tilde{\pi}_{X_s}|\tilde{\pi}_{X_t}\rangle| - |\langle\pi_{X_s}|\pi_{X_t}\rangle| \right| < 1/3,$$

and so  $|\langle\tilde{\pi}_{X_s}|\tilde{\pi}_{X_t}\rangle| > 2/3$  if  $s$  and  $t$  are connected, but  $|\langle\tilde{\pi}_{X_s}|\tilde{\pi}_{X_t}\rangle| < 1/3$  otherwise. The SWAP test distinguishes these cases with constant probability.  $\triangleleft$

<sup>6</sup> [3, Proposition 1] only proves the proposition for simple random walks, however it trivially extends to random walks on weighted graphs.

## 5.4 Proof of Theorem 17

In this section, we prove Theorem 17, which we restate here for convenience.

► **Theorem 17.** Fix  $\delta \geq 0$ . Let  $\mathcal{G}_n$  be a family of undirected weighted graphs  $G = (X, E, W)$  with  $n = |X|$ , such that  $\gamma_*(G) \geq \delta$  whenever  $s$  and  $t$  are connected. Then for any  $S = \Omega(\log(n))$ , there is a quantum algorithm that decides  $USTCON_{\text{qw},\delta}$  on  $\mathcal{G}_n$  with bounded error in  $O(S)$  space – of which  $O(\log(n/\delta))$  is quantum memory, and the remainder is QCRAM – and  $T = \tilde{O}\left(\frac{S}{\delta} \log\left(\frac{1}{\pi_{\min}}\right) + \sqrt{\frac{n}{\delta S}}\right)$  queries to  $\mathcal{O}_W$ , elementary gates, and QCRAM queries.

**Proof.** We first analyze the space complexity of Algorithm 2. Step 1 is purely classical, and uses  $O(p \log(n))$  space to store the  $O(p)$  vertices in  $L$  and  $M$ , with each random walk using  $O(\log(n))$  space. We can implement step 2 using the algorithm referred to in Proposition 21 using  $O(\log(n))$  qubits of space, but this requires that the  $O(p \log(n))$  classical space used to store  $L$  and  $M$  in step 1 is QCRAM. Finally, step 3 just uses  $O(\log(n))$  quantum space. Thus, the claimed space complexity follows if we set  $S = p \log(n)$ .

Next, we analyze the time complexity. Every random walk of step 1 adds  $O\left(\frac{1}{\delta} \log\left(\frac{n}{p\pi_{\min}}\right)\right)$  to the time complexity. The total number of walks is  $O(p)$ . Checking whether  $L \cap M = \emptyset$  is  $O(p)$  as this is the total number of samples. Hence, the overall complexity of the first step is  $\tilde{O}\left(\frac{p}{\delta} \log\left(\frac{1}{\pi_{\min}}\right)\right)$ . By Claim 22, the complexity of step 2 is  $O\left(\sqrt{\frac{n}{p\delta}} \log\left(\frac{n}{p}\right)\right)$ . Finally, the SWAP test in step 3 uses only  $O(\log(n))$  gates, since the states being compared are  $O(\log(n))$ -qubit states. Hence, the total time complexity of Algorithm 2 is

$$T = \tilde{O}\left(\frac{p}{\delta} \log\left(\frac{1}{\pi_{\min}}\right) + \sqrt{\frac{n}{\delta p}}\right) = \tilde{O}\left(\frac{S}{\delta} \log\left(\frac{1}{\pi_{\min}}\right) + \sqrt{\frac{n}{\delta S}}\right),$$

since  $S = \tilde{O}(p)$ .

Finally, for the correctness of the algorithm, by Claim 23, Algorithm 2 distinguishes between the case where  $|\pi_{X_s}\rangle$  and  $|\pi_{X_t}\rangle$  are equal and the case where they are orthogonal with bounded error. If  $X_s = X_t$  (i.e.  $s$  and  $t$  are connected) then the states are equal, and if  $X_s \cap X_t = \emptyset$  (i.e.  $s$  and  $t$  are not connected) then they are orthogonal. ◀

---

### References

- 1 Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 20–29, 2003.
- 2 Romas Aleliunas, Richard M. Karp, Richard J. Lipton, Laszlo Lovasz, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 218–223, 1979. doi:10.1109/SFCS.1979.34.
- 3 Simon Apers. Quantum walk sampling by growing seed sets. In *Proceedings of the 27th Annual European Symposium on Algorithms (ESA)*, volume 144, pages 9:1–9:12. Springer, 2019.
- 4 Simon Apers, András Gilyén, and Stacey Jeffery. A unified framework of quantum walk search. In *Proceedings of the 38th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 6:1–6:13, 2020. arXiv:1912.04233 doi:10.4230/LIPIcs.STACS.2021.6.
- 5 Greg Barnes and Uriel Feige. Short random walks on graphs. In *Proceedings of the 1993 ACM Symposium on the Theory of Computing (STOC)*, pages 728–737, 1993.
- 6 Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS’98. arXiv:quant-ph/9802049 doi:10.1145/502090.502097.

- 7 Paul Beame, Allan Borodin, Prabhakar Raghavan, Walter L Ruzzo, and Martin Tompa. Time-space tradeoffs for undirected graph traversal. In *Proceedings of the 1990 IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 429–438. IEEE, 1990.
- 8 Aleksandrs Belovs. Quantum walks and electric networks. [arXiv:1302.3143](https://arxiv.org/abs/1302.3143), 2013.
- 9 Aleksandrs Belovs and Ben W. Reichardt. Span programs and quantum algorithms for  $st$ -connectivity and claw detection. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, pages 193–204, 2012. doi:10.1007/978-3-642-33090-2\_18.
- 10 Graham Brightwell and Peter Winkler. Maximum hitting time for random walks on graphs. *Random Structures & Algorithms*, 1(3):263–276, 1990.
- 11 Andrei Z Broder, Anna R Karlin, Prabhakar Raghavan, and Eli Upfal. Trading space for time in undirected  $st$ -connectivity. In *Proceedings of the 1989 ACM Symposium on the Theory of Computing (STOC)*, pages 543–549, 1989.
- 12 Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998.
- 13 Christoph Dürr, Mark Heiligman, Peter Høyer, and Mehdi Mhalla. Quantum query complexity of some graph problems. *SIAM Journal on Computing*, 35(6):1310–1328, 2006. Earlier version in ICALP’04. [arXiv:quant-ph/0401091](https://arxiv.org/abs/quant-ph/0401091) doi:10.1137/050644719.
- 14 Jeff Edmonds. Time-space trade-offs for undirected  $st$ -connectivity on a JAG. In *Proceedings of the 1993 ACM Symposium on the Theory of Computing (STOC)*, pages 718–727, 1993.
- 15 Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Limit on the speed of quantum computation in determining parity. *Physical Review Letters*, 81(24):5442, 1998. [arXiv:quant-ph/9802045](https://arxiv.org/abs/quant-ph/9802045) doi:10.1103/PhysRevLett.81.5442.
- 16 Uriel Feige. A randomized time-space tradeoff of  $\tilde{O}(m\hat{R})$  for USTCON. In *Proceedings of the 1993 IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 238–246. IEEE, 1993.
- 17 Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112*, 2002.
- 18 Iordanis Kerenidis and Anupam Prakash. Quantum recommendation systems. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 49:1–49:21. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- 19 Adrian Kosowski. Faster walks in graphs: A  $\tilde{O}(n^2)$  time-space trade-off for undirected  $s$ - $t$  connectivity. In *Proceedings of the 2013 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1873–1883, 2013. [arXiv:1204.1136v2](https://arxiv.org/abs/1204.1136v2) doi:10.1137/1.9781611973105.133.
- 20 Jessica Lemieux, Bettina Heim, David Poulin, Krysta Svore, and Matthias Troyer. Efficient quantum walk circuits for metropolis-hastings algorithm. *Quantum*, 4:287, 2020.
- 21 David A Levin, Yuval Peres, and Elizabeth L Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2017. second edition.
- 22 Omer Reingold. Undirected connectivity in log-space. *Journal of the ACM*, 55(4), 2008. doi:10.1145/1391289.1391291.
- 23 Avi Wigderson. The complexity of graph connectivity. In *International Symposium on Mathematical Foundations of Computer Science*, pages 112–132. Springer, 1992.

## A Appendix

**Proof of Lemma 8.** Assume first that there is an additional type of query allowed, namely:

$$\text{Index query: } \mathcal{O}_I : |u\rangle |v\rangle |0\rangle \mapsto |u\rangle |v\rangle |i\rangle \quad (3)$$

for  $u, v \in X$  and  $i \in [d_u]$  such that  $v_i(u) = v$ . Then  $\mathcal{O}_W$  can be implemented using  $\mathcal{O}_I$ ,  $\mathcal{O}_D$ , and  $\mathcal{O}_N$  as follows. Let  $F_d$  denote the Fourier transform over  $\mathbb{Z}_d$ , and let  $F = \sum_{d=1}^n |d\rangle \langle d| \otimes F_d$ , which can be implemented (to any inverse polynomial precision) in  $O(\log(n))$  gates [12]. Then for any  $u \in X$ , we implement:

$$\begin{aligned}
|0\rangle |u\rangle |0\rangle |0\rangle &\xrightarrow{\mathcal{O}_D} |d_u\rangle |u\rangle |0\rangle |0\rangle \xrightarrow{F} \frac{1}{\sqrt{d_u}} |d_u\rangle \sum_{i=1}^{d_u} |u\rangle |i\rangle |0\rangle \\
&\xrightarrow{\mathcal{O}_N} \frac{1}{\sqrt{d_u}} |d_u\rangle \sum_{i=1}^{d_u} |u\rangle |i\rangle |v_i(u)\rangle \xrightarrow{\mathcal{O}_I^\dagger \mathcal{O}_D^\dagger} \frac{1}{\sqrt{d_u}} \sum_{i=1}^{d_u} |u\rangle |v_i(u)\rangle = \mathcal{O}_W |u\rangle |0\rangle.
\end{aligned}$$

To complete the proof, note that the index query operator  $\mathcal{O}_I$  only requires  $O(\log(d_u))$  sorted adjacency array queries, since the neighbours are sorted and this makes it possible to perform binary search for  $i$  such that  $v_i(u) = v$ . ◀

**Proof of Theorem 11.** We will reduce the PARITY problem to  $\text{USTCON}_{\text{s-arr}}$ . Since PARITY requires  $\Omega(n)$  queries by Lemma 10, and the quantum walk oracle  $\mathcal{O}_W$  can be implemented using  $\tilde{O}(1)$  sorted array queries by Lemma 8, this reduction will prove the statement of the theorem.

Let  $x \in \{0, 1\}^n$  be an input of PARITY. The corresponding output would be  $\bigoplus_{i=0}^{n-1} x_i$ . Given this, we need to build a  $\text{USTCON}_{\text{s-arr}}$  input that can be queried with a constant number of queries to  $x$ . Consider an undirected graph  $G = (X, E)$  defined as follows (see also Figure 2):

$$\begin{aligned}
X &= \{s = v_0, v'_0, v_1, v'_1, \dots, v_{n-1}, v'_{n-1}, v_n, t = v'_n\} \\
E &= \{\{v_i, v'_{i+1}\}, \{v'_i, v_{i+1}\} : i \in \{0, \dots, n-1\}, x_i = 1\} \\
&\quad \cup \{\{v_i, v_{i+1}\}, \{v'_i, v'_{i+1}\} : i \in \{0, \dots, n-1\}, x_i = 0\}.
\end{aligned}$$

In this setting,  $\bigoplus_{i=0}^{n-1} x_i = 1$  if and only if  $s$  and  $t$  are connected in the graph  $G$ .

Next, we describe how to implement queries  $\mathcal{O}_D$  and  $\mathcal{O}_N$  to  $G$  as required by the  $\text{USTCON}_{\text{s-arr}}$  problem, using queries to  $\mathcal{O}_x$ . Consider the following encoding of vertices of  $G$ . For  $i \in \{0, \dots, n\}$ , we let  $v_i = (i, 0)$ , and  $v'_i = (i, 1)$ . That is, for a vertex  $(i, b)$ ,  $i \in \{0, \dots, n\}$  encodes the ‘‘column’’ and  $b \in \{0, 1\}$  encodes the ‘‘row’’. Assume that the vertices are ordered lexicographically, i.e.

$$(i, b_i) < (j, b_j) \iff i < j \text{ or } i = j, b_i < b_j.$$

Queries to  $G$  are described according to this ordering.

1. Degree queries,  $\mathcal{O}_D$ , are trivial in this case as  $d_{v_0} = d_{v'_0} = d_{v_n} = d_{v'_n} = 1$ , and all other degrees are 2.
2. Since every vertex has degree at most 2, we explicitly describe neighbour queries,  $\mathcal{O}_N$  for indices 1 and 2 such that the ordering assumption holds:

- $\mathcal{O}_N : |i\rangle |b\rangle |1\rangle |0\rangle |0\rangle \mapsto |i\rangle |b\rangle |1\rangle |i-1\rangle |b\rangle \xrightarrow{\mathcal{O}_x} |i\rangle |b\rangle |1\rangle |i-1\rangle |b \oplus x_{i-1}\rangle, \forall 0 < i \leq n,$
- $\mathcal{O}_N : |i\rangle |b\rangle |2\rangle |0\rangle |0\rangle \mapsto |i\rangle |b\rangle |2\rangle |i+1\rangle |b\rangle \xrightarrow{\mathcal{O}_x} |i\rangle |b\rangle |2\rangle |i+1\rangle |b \oplus x_i\rangle, \forall 0 \leq i < n.$

It can be seen from the formulas that queries to  $G$  can be implemented with a constant number of queries to the parity input  $x$ . This implies the  $\Omega(n)$  lower bound in  $\text{USTCON}_{\text{s-arr}}$ .

For  $\text{USTCON}_{\text{qw}}$ , note that the graph has bounded degree, and so by Lemma 8 we can simulate a query in this model using  $O(1)$  queries in the sorted adjacency array model. This implies a similar  $\Omega(n)$  lower bound for this model. ◀