

One time-traveling bit is as good as logarithmically many

Ryan O’Donnell*

A. C. Cem Say†

July 18, 2014

Abstract

We consider computation in the presence of closed timelike curves (CTCs), as proposed by Deutsch. We focus on the case in which the CTCs carry *classical* bits (as opposed to qubits). Previously, Aaronson and Watrous showed that computation with polynomially many CTC bits is equivalent in power to PSPACE. On the other hand, Say and Yakaryılmaz showed that computation with just 1 classical CTC bit gives the power of “postselection”, thereby upgrading classical randomized computation (BPP) to the complexity class BPP_{path} and standard quantum computation (BQP) to the complexity class PP. It is natural to ask whether increasing the number of CTC bits from 1 to 2 (or 3, 4, etc.) leads to increased computational power. We show that the answer is no: randomized computation with logarithmically many CTC bits (i.e., polynomially many CTC states) is equivalent to BPP_{path} . (Similarly, quantum computation augmented with logarithmically many *classical* CTC bits is equivalent to PP.) Spoilsports with no interest in time travel may view our results as concerning the robustness of the class BPP_{path} and the computational complexity of sampling from an implicitly defined Markov chain.

1 On time travel

We begin with a discussion of time travel. Readers not interested in this concept may skip directly to Section 2, wherein we define the problem under consideration in a purely complexity-theoretic manner, with no reference to time travel.

Kurt Gödel [Göd49] was the first to point out that Einstein’s theory of general relativity is consistent with the existence of *closed timelike curves* (CTCs), raising the theoretical possibility of time travel. Any model of time travel must deal with the “Grandfather Paradox”, wherein a trip to the past causes a chain of events that leads to a future in which that very trip does not take place. Assume that a time-traveler changes the state of the universe at the earlier end t_0 of a time loop from state s to some different state s' . Then just what is the state of the universe at time t_0 : is it s or s' ? Seeing a logical inconsistency in this scenario, most thinkers of earlier generations concluded that time travel to the past must be impossible. There is, however, a way out. In an influential paper [FMN⁺90], Friedman *et al.* suggested Nature might allow CTCs as long as they do not “change the past”, an idea that has come to be known as the *Novikov self-consistency principle*. The main two rivaling models of time travel — the “Deutschian model” (which we study in this work), and the “postselected CTC model” from [LMGP⁺10] (which is mentioned in Section 4.2) — both conform to the Novikov self-consistency principle.

*Department of Computer Science, Carnegie Mellon University. Work performed while the author was at the Boğaziçi University Computer Engineering Department, supported by Marie Curie International Incoming Fellowship project number 626373.

†Boğaziçi University Computer Engineering Department.

In the model put forward by Deutsch [Deu91], the universe need not be in a single deterministic state at time t_0 . Rather, the state of the universe should be viewed as a probability distribution over several states (possibly even quantum states) like s and s' in the example above. The requirement that the past should not change is fulfilled by stipulating that Nature sets the state x of the portion of the universe affected by the CTC at time t_0 to a fixed point of the operator f describing the evolution in the CTC (meaning $x = f(x)$).

To take the traditional example, suppose a deranged scientist can access a CTC to the past century, and he sends through it a bomb that is programmed to kill his grandfather (who is only a child back then). We consider two (classical) states of the universe at the time of the bomb's arrival: state 1 is "grandfather dies" and state 2 is "grandfather lives". We assume the universe proceeds deterministically from that point on: if the grandfather is killed, then in the future no bomb is sent through the CTC; conversely, if the grandfather lives, then the deranged scientist is born and *does* send the bomb back in time. We can model this evolution by a 2-state Markov chain with the following transition matrix (that happens to be deterministic):

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

In Deutsch's model, Nature sets the state of the universe to be the stationary distribution for this chain:

$$\begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}.$$

That is, the bomb arrives to kill the grandfather with probability $\frac{1}{2}$.

1.1 Computation with CTCs

As the reader can see, Nature performs a kind of computation here, determining the stationary distribution of the Markov chain that has been arranged within the CTC by the deranged scientist. It is natural to wonder if Nature's power can be effectively harnessed by a computational device. Indeed, Deutsch [Deu91] pointed out that in general his model involves Nature solving an NP-hard problem; later, Brun [Bru03] discussed the possibility of using CTCs to solve the Factoring problem efficiently. The first clear model of computation with Deutschian CTCs was proposed by Bacon [Bac04]. Both Deutsch and Bacon consider sending *qubits* through a CTC. However, as pointed out by Aaronson [Aar05a], it is also very interesting (and simpler) to consider only classical bits passing through a CTC. Indeed, as far as we aware, there are no results showing that time-traveling quantum bits confer a computational advantage over time-traveling classical bits. Therefore, in the rest of this section we will sketch the Deutschian model of computation with classical CTC bits, and mention prior work. A formal complexity-theoretic definition of the model (with no reference to time travel) is given in Section 2.

Suppose that a computational agent \mathcal{A} has access to a CTC which is "wide" enough to support the transmission of w bits. Thus the physical object being sent through the CTC can be in one of $S = 2^w$ states. (We may also more generally consider values of S that are not powers of 2.) Let us think of \mathcal{A} as a classical polynomial-time randomized Turing machine (though it might be of another type; e.g., a BQP-machine). Say that \mathcal{A} is trying to decide if a given input $x \in \{0, 1\}^n$ belongs to language L . The algorithm \mathcal{A} can read the w bits in the CTC, perform some computation, and then send a new string of w bits through the CTC. Since the incoming and outgoing bit strings can be in one of $2^w = S$ states, and since \mathcal{A} is a randomized algorithm, the operation of \mathcal{A} on the CTC constitutes an S -state Markov chain M_x , which depends on the input $x \in \{0, 1\}^n$ to the L -decision problem.

In the Deutschian model, we assume that once the Markov chain M_x is defined, Nature sets the distribution of the bits in the CTC to *some* stationary distribution of M_x . We emphasize that it’s merely *some* stationary distribution (at least one of which always exists) — we don’t assume that M_x must have a unique stationary distribution. (Now is a good time to mention that if \mathcal{A} is allowed to send qubits along the CTC, then its operation constitutes a *quantum channel*. It is also known [Deu91, Wol12] that every quantum channel has at least one stationary mixed state, and we assume Nature selects one.) Finally, given that the incoming CTC bits are now presumed to be in a stationary distribution for the Markov chain M_x , the algorithm \mathcal{A} effectively gets one *sample* from this stationary distribution. Using this sample, the algorithm \mathcal{A} can output its decision on whether or not $x \in L$. When thinking of \mathcal{A} as a BPP-type machine, this decision should be correct with probability at least $\frac{2}{3}$.

1.2 Prior work

Bacon [Bac04] considered the case of a 1-qubit CTC, though his construction actually works equally well with a CTC supporting just 1 classical bit. However, Bacon’s model was also more generous in that he allowed 1-bit CTC computations as “subroutines” within polynomial-time algorithms; in effect, he allowed the use of $\text{poly}(n)$ many 1-(qu)bit CTCs. Bacon showed that in this model one can efficiently solve any NP problem. Subsequently, Aaronson and Watrous [Aar05a, AW09] investigated the model in which the CTC supports $\text{poly}(n)$ many bits (i.e., $S = 2^{\text{poly}(n)}$ many states). They showed that this model is extremely powerful: if \mathcal{A} ’s computational power is anywhere between AC^0 and PSPACE (including the most usual choices of BPP or BQP), the result is that the model becomes equivalent in power to PSPACE. Actually, this result was not even the main one in their paper; their main result is that if $\text{poly}(n)$ many CTC *qubits* are allowed, then the power of the model is still *only* that of PSPACE.

Regarding the difference between using a 1-bit CTC polynomially many times, and using a $\text{poly}(n)$ -bit CTC once, Aaronson [Aar05a] remarked, “It is difficult to say which model is the more reasonable!” One can argue that both models are rather impractical in that they require constructing new/wider CTCs as the input length increases.¹ Indeed, the main technical question left open at the end of Aaronson and Watrous’s work was to understand the computational power of the more realistic “narrow” CTCs; e.g., one-time-use CTCs that can only transmit a single bit, or a bounded number of bits. In this direction, Say and Yakaryilmaz [SY12] showed that augmenting standard complexity models with access to a 1-bit CTC is exactly equal in power to augmenting them with “postselection” [Aar05b] (defined in Section 4). In particular, this shows that classical randomized computation with a 1-bit CTC is equivalent to the complexity class BPP_{path} , and quantum computation with a 1-bit CTC is equivalent to the complexity class PP. We recall in further detail the class BPP_{path} in Section 4. For now, suffice it to say that it contains NP and coNP, is likely equal to $\text{P}_{\parallel}^{\text{NP}}$, and is very likely to be much smaller than PSPACE. In particular, randomized computation with access to a 1-bit CTC can efficiently solve the SAT problem; this is discussed below in Example 2.5.

To summarize, the aforementioned results show that for classical polynomial-time randomized computation, adding a 1-bit CTC gives the power of BPP_{path} and adding a $\text{poly}(n)$ -bit CTC gives the power of PSPACE. What about in between (presuming of course that $\text{BPP}_{\text{path}} \neq \text{PSPACE}$)? Sticking with the more “realistic” end of the spectrum, this is the question motivating our work:

Question. *Are 2-bit (or 3-bit, 4-bit etc.) CTCs more powerful than 1-bit CTCs?*

¹Bearing in mind the comment concerning practical considerations in the final paragraph of Bacon’s work [Bac04].

2 Formal computational complexity statements

In what follows we formally define the complexity model of computing with CTCs. Our definitions are equivalent to those in [AW09, SY12]; however we phrase them differently, in terms of Markov chains. Informally and in brief, $\text{BPP}_{\text{CTC}[w]}$ is the class of languages decidable by efficient randomized algorithms that are allowed to set up a 2^w -state Markov chain and then freely get one sample from the chain’s stationary distribution.

Definition 2.1. Let M be an S -state Markov chain. A *state-transition oracle* \mathcal{M} for M is any algorithm that takes as input a state $i \in [S]$ and outputs the state resulting from taking one random step in M starting from state i . Most typically we think of $S = 2^w$ and \mathcal{M} as being implemented by a w -bit-input/output *standard randomized circuit*; i.e., one with AND, OR, NOT, and “probability- $\frac{1}{2}$ coin-flip” gates. We might also consider *standard quantum circuits* \mathcal{M} in which Hadamard and Toffoli gates (which are universal [Shi03]) are also used.

Definition 2.2. Let $w = w(n)$ be a “width” parameter. Consider a deterministic polynomial-time Turing Machine \mathcal{A} that, on input $x \in \{0, 1\}^n$, outputs the description of two standard randomized circuits, \mathcal{M}_x and \mathcal{D}_x . The circuit \mathcal{M}_x should have w input and output bits, thereby defining a state-transition oracle for a Markov chain M_x on $S = 2^w$ states. The “decision circuit” \mathcal{D}_x should have w input bits and one output bit. We suppose computation proceeds as follows: First, an arbitrary stationary distribution π for M_x is chosen. Next, a sample $i \sim \pi$ is chosen from this distribution and is fed as input to \mathcal{D}_x . Finally, \mathcal{D}_x ’s output gate is considered to be the overall output of \mathcal{A} ’s computation. We define $\text{BPP}_{\text{CTC}[w]}$ to be the class of all languages L such that there exists an \mathcal{A} as above with the following property: for every x (and every stationary distribution π for M_x),

$$\Pr_i[\mathcal{A} \text{ outputs } 1] \geq \frac{2}{3} \quad \text{when } x \in L, \quad \Pr_i[\mathcal{A} \text{ outputs } 1] \leq \frac{2}{3} \quad \text{when } x \notin L.$$

We may also analogously define $\text{BQP}_{\text{CTC}[w]}$ in case \mathcal{M}_x and \mathcal{D}_x are allowed to be standard quantum circuits.

Remark 2.3. We warn the reader that our notation $\text{C}_{\text{CTC}[w]}$ is different from that in [AW09, SY12], in that “CTC[w]” signifies a CTC carrying w *classical* bits. We suggest notation such as $\text{BQP}_{\text{QCTC}[w]}$ for the case of CTCs carrying w qubits; however we neither define nor consider CTC-qubits in this paper (except in a concluding open problem).

Remark 2.4. There is nothing special about considering Markov chains with S states where S is a power of 2. However we stick with the above notation for simplicity, and for consistency with similar complexity class definitions such as that of $\text{P}^{\text{NP}[w]}$ (polynomial-time computation with $2^w - 1$ parallel queries to an NP oracle).

Example 2.5. Following [SY12], let us show that $\text{NP} \subseteq \text{BPP}_{\text{CTC}[1]}$. Equivalently, we illustrate how SAT can be solved by a “1-bit CTC algorithm” \mathcal{A} , which can set up a 2-state Markov chain and get a sample from its stationary distribution. On input an n -variable CNF formula ϕ , algorithm \mathcal{A} constructs a state-transition circuit \mathcal{M}_ϕ for a certain 2-state Markov chain M_ϕ . Think of state 0 of M_ϕ as meaning “no evidence that ϕ is satisfiable” and state 1 as meaning “evidence that ϕ is satisfiable”. The operation of \mathcal{M}_ϕ is as follows: On input state i , \mathcal{M}_ϕ first chooses a uniformly random string $\mathbf{y} \in \{0, 1\}^n$ and checks if it satisfies ϕ . If \mathbf{y} is satisfying, \mathcal{M}_ϕ outputs state 1. If \mathbf{y} is unsatisfying, then \mathcal{M}_ϕ outputs state 0 with probability $\epsilon := 2^{-n^2}$ and outputs its input state i with

probability $1 - \epsilon$. It is clear that \mathcal{A} can write down \mathcal{M}_ϕ 's description in deterministic polynomial time. One can now check that the resulting Markov chains M_ϕ are as follows:

if ϕ is satisfiable, M_ϕ has transition matrix $\begin{bmatrix} 1 - 2^{-n} & 2^{-n} \\ \epsilon' & 1 - \epsilon' \end{bmatrix}$ (where $\epsilon' := (1 - 2^{-n})\epsilon \approx 2^{-n^2}$);
if ϕ is unsatisfiable, M_ϕ has transition matrix $\begin{bmatrix} 1 & 0 \\ \epsilon & 1 - \epsilon \end{bmatrix}$.

Now if ϕ is unsatisfiable, state 0 is absorbing and it's clear that the (unique) stationary distribution π of M_ϕ is entirely concentrated on state 0. On the other hand, suppose ϕ is satisfiable. Then since the $0 \rightarrow 1$ transition probability of M_ϕ is much higher (relatively speaking) than the $1 \rightarrow 0$ transition probability, the long-term (i.e., stationary) distribution π of M_ϕ will be almost entirely concentrated on state 1. (More precisely, π will put only probability $\frac{\epsilon'}{2^{-n} + \epsilon'} \approx 2^{-n^2 + n}$ on state 0.) We now stipulate that for every ϕ , algorithm \mathcal{A} outputs the same 1-bit decision circuit \mathcal{D}_ϕ , which on input $i \sim \pi$ simply outputs i . From the above discussion, we see that this correctly indicates whether $\phi \in \text{SAT}$ except with negligible error probability.

The following theorems concerning $\text{BPP}_{\text{CTC}[w]}$ have previously been shown:

Theorem 2.6. (Aaronson–Watrous [AW09].) $\text{BPP}_{\text{CTC}[\text{poly}(n)]} = \text{PSPACE}$.

(Indeed $\text{P}_{\text{CTC}[\text{poly}(n)]} = \text{BQP}_{\text{QCTC}[\text{poly}(n)]} = \text{PSPACE}$.)

Theorem 2.7. (Say–Yakaryılmaz [SY12].) $\text{BPP}_{\text{CTC}[1]} = \text{BPP}_{\text{path}}$.

(Indeed, adding 1 CTC-bit generally confers the power of “postselection”, discussed in Section 4. For example, it also holds that $\text{BQP}_{\text{CTC}[1]} = \text{PostBQP} = \text{PP}$.)

As mentioned in Section 1.1 (and discussed further in Section 4), BPP_{path} is likely equal to $\text{P}_{\parallel}^{\text{NP}} = \text{P}^{\text{NP}[O(\log n)]}$ (as it's contained in $\text{BPP}_{\parallel}^{\text{NP}}$), and is very likely to be much smaller than PSPACE.

2.1 Our theorem

Paraphrasing the above two theorems, we have that Markov chains with 2 states ($w = 1$) give the power of BPP_{path} , and Markov chains with exponentially many states ($w = \text{poly}(n)$) give the power of PSPACE. What about in between? Are 3-state or 4-state ($w = 2$) Markov chains more powerful than 2-state chains? To take an analogy from another family of complexity classes, we remind the reader that it's widely believed that

$$\text{P}^{\text{NP}[1]} \subsetneq \text{P}^{\text{NP}[2]} \subsetneq \text{P}^{\text{NP}[3]} \subsetneq \dots$$

The main result of this paper is that in apparent contrast, “the hierarchy collapses” for $\text{BPP}_{\text{CTC}[w]}$; polynomially many states ($w = O(\log n)$) confer no more advantage than 2 states.

Main theorem. $\text{BPP}_{\text{CTC}[O(\log n)]} \subseteq \text{PostBPP} = \text{BPP}_{\text{path}}$; thus

$$\text{BPP}_{\text{CTC}[1]} = \text{BPP}_{\text{CTC}[2]} = \text{BPP}_{\text{CTC}[3]} = \dots = \text{BPP}_{\text{CTC}[O(\log n)]} = \text{BPP}_{\text{path}}.$$

It will be clear from our proof that more generally, $O(\log n)$ CTC bits still only confer the power of postselection, and in particular $\text{BQP}_{\text{CTC}[O(\log n)]} = \text{PostBQP} = \text{PP}$. Our main theorem may also be seen as further demonstration of the robustness and naturalness of the class BPP_{path} .

2.2 Proof techniques

Here we briefly outline the proof of our theorem, with the actual proof being given in Section 5. Let’s return to Example 2.5, which shows that $\text{SAT} \in \text{BPP}_{\text{CTC}[1]}$. One might ask, why doesn’t the proof show that $\text{SAT} \in \text{BPP}$? After all, the algorithm \mathcal{A} simply constructs a 2-state Markov chain M and then takes a sample from its stationary distribution. Why doesn’t \mathcal{A} simply exactly solve for M ’s stationary distribution? The trouble of course is that even though \mathcal{A} constructed M itself, in some sense M is still only “implicitly defined” from \mathcal{A} ’s point of view. \mathcal{A} cannot directly access the transition probabilities of M (doing so requires \mathcal{A} to solve an NP-complete problem); rather, \mathcal{A} can only “simulate” M , by use of the state-transition matrix \mathcal{M} it constructed. Naively, this still might not seem like a problem; given the ability to simulate M , couldn’t \mathcal{A} find a (near-)stationary distribution π for M simply by simulating it for a long time? The trouble here is that even though M only has 2 states, it has some transition probabilities that are “exponentially small” (in n). Furthermore, the stationary distribution of M can be extremely sensitive to the *relative* exponential smallness of these transition probabilities — Example 2.5 illustrates exactly this.

Our proof that $\text{BPP}_{\text{CTC}[w]} \subseteq \text{PostBPP} = \text{BPP}_{\text{path}}$ for $w \leq O(\log n)$ in some sense follows Say and Yakaryılmaz’s proof [SY12] in the case of $w = 1$. They essentially observed that using the power of postselection (discussed further in Section 4), a randomized algorithm can get an exact sample from the stationary distribution of a Markov chain given only a state-transition oracle for it. Their proof of this was greatly facilitated by the fact that 2-state Markov chains are easy to analyze: If the $0 \rightarrow 1$ transition probability is p and the $1 \rightarrow 0$ transition probability is q , then the stationary distribution is

$$\pi = \begin{bmatrix} \frac{q}{p+q} \\ \frac{p}{p+q} \end{bmatrix}.$$

(This presumes we don’t have $p = q = 0$, an important issue that we discuss later.) For our main theorem, we need a similar postselecting algorithm for general poly(n)-state Markov chains. The key technical tool for this will be the *Markov Chain Tree Theorem*, apparently first proved by Hill [Hil66], and called by Aldous [Ald02] “the most often rediscovered result in probability theory”; see also [WF70, Shu75, KV80, LR83, AT89, Bro89, Pok99]. We state here the version for irreducible chains:

Markov Chain Tree Theorem. *Let M be an S -state irreducible Markov chain with transition matrix $(p_{ij})_{i,j \in [S]}$. Let G_M be the underlying strongly connected digraph for M in which (i, j) is a directed edge if and only if $p_{ij} > 0$. Recall that a rooted arborescence T in G_M is a collection of edges forming a rooted spanning tree in which all edges are directed toward the root vertex. We write*

$$\|T\| = \prod_{(i,j) \in T} p_{ij}.$$

Let \mathcal{T}_i denote the set of all arborescences in G_M rooted at $i \in [S]$, and write $\mathcal{T} = \cup_i \mathcal{T}_i$. Then if π denotes the (unique) stationary distribution of M , we have the formula

$$\pi_i = \frac{\sum_{T \in \mathcal{T}_i} \|T\|}{\sum_{T \in \mathcal{T}} \|T\|}.$$

We add that this theorem plays an important role in the theory of exact sampling from unknown Markov chains [AGT92, LW95, PW98, Wil96]. That theory is concerned with a problem similar to ours; however, there are two main differences: i) That theory involves only traditional algorithms,

and therefore by necessity the running time may be exponential if the chain’s mixing time is exponential. By contrast, we are using postselecting algorithms and therefore have the chance to run in polynomial time. ii) That theory is concerned with *exact* sampling from the stationary distribution. By contrast, we actually only need approximate sampling from the stationary distribution.

Finally, we mention one challenge for our proof that at first seems like a technicality but in fact proves to be quite a nuisance: There is no promise in the definition of $\text{BPP}_{\text{CTC}[w]}$ that the Markov chains M_x be irreducible. This is precisely the “ $p = q = 0$ issue” elided in the discussion of 2-state stationary distributions above. We overcome this difficulty by proving a somewhat technical lemma that allows us to perturb general Markov chains into irreducible ones.

2.3 Outline of the remainder of the paper

In Section 3 we prove the aforementioned technical lemma on Markov chain perturbations, allowing us to work only with irreducible Markov chains. In Section 4 we recall BPP_{path} and postselection in more detail, and we also describe the “restarting” view of postselection (from [YS13]) that will be helpful in the proof of our main theorem. Finally, we give the proof of the main theorem in Section 5, and then end with an open question.

3 A technical lemma on Markov chain perturbations

Our main theorem essentially requires us to give a $\text{PostBPP} = \text{BPP}_{\text{path}}$ algorithm for sampling from *some* stationary distribution of a given Markov chain M , specified by a state-transition oracle circuit \mathcal{M} . We would like to do this using the Markov Chain Tree Theorem; however, that theorem requires M to be irreducible. In fact there is an extension of the theorem to reducible Markov chains, due to Leighton and Rivest [LR83]; however, its use requires determination of the communication classes of the graph G_M , which in turn requires determining which transition probabilities are 0. This is already NP-hard in the 2-state case — consider again Example 2.5, in which ϕ is satisfiable if and only if M_ϕ is irreducible. We resolve this difficulty by showing that we can “blindly” perturb M to a necessarily irreducible chain M' with the property that the unique stationary distribution of M' is close to *some* stationary distribution of M . (Recall that $\text{BPP}_{\text{CTC}[w]}$ algorithms are required to work given access to *any* stationary distribution for M .) Here by “blindly” we mean that we can easily construct a state-transition circuit \mathcal{M}' for M' using \mathcal{M} as a black box.

The idea behind the perturbation is a simple and natural one; it was implicitly used in [SY12] in the 2-state case, and it even occurs in the “random surfer” model of PageRank [PBMW99]. Specifically, for each state we introduce a tiny probability ϵ of jumping to a uniformly random other state. We will make ϵ small even in comparison with all nonzero transition probabilities in M ; these probabilities are at worst exponentially small, since \mathcal{M} is a polynomial-size circuit. With such a small ϵ , it is somewhat intuitive that the perturbation should yield an irreducible chain M' whose (unique) stationary distribution is close to a mixture of the stationary distributions of the communication classes of M . The below technical lemma makes this intuition precise. We should add that since there is significant prior work on Markov chain perturbations (see, e.g., [Sch68, CM01]), it’s possible that the lemma is not new. However we were unable to find anything like it in the literature, all of which seemed to make the assumption that the original chain is *already* irreducible.

Before stating and proving the lemma, we introduce one piece of notation:

Notation 3.1. Let B be a matrix (possibly a vector or even just a scalar) with rational entries. We write $\langle B \rangle$ for the number of bits needed to represent B ; see, e.g., [GLS88] for precise details.

Lemma 3.2. *Let M be an S -state Markov chain ($S \geq 2$) with rational transition matrix K . Let $0 < \epsilon \leq 2^{-\langle K \rangle^c}$, where $c \in \mathbb{Z}^+$ is a certain universal constant. Define*

$$K' = (1 - \epsilon)K + \epsilon \frac{1}{S}J,$$

where J is the $S \times S$ all-1's matrix. In other words, K' is the transition matrix for the Markov chain M' defined as follows: With probability $1 - \epsilon$, take a step according to M ; with probability ϵ , jump to a uniformly random state. Let π' be the stationary distribution for M' , which is unique because M' is irreducible. Then there exists a stationary distribution π for M satisfying

$$\|\pi - \pi'\|_1 \leq 2^{\langle K \rangle^c} \epsilon.$$

Proof. Let $L = I - K$ be the Laplacian for M . Similarly define L' , so $L' = (1 - \epsilon)L + \epsilon(I - \frac{1}{S}J)$. Then

$$0 = \pi' L' = (1 - \epsilon)\pi' L + \epsilon\pi'(I - \frac{1}{S}J) \implies \pi' L = \frac{\epsilon}{1 - \epsilon}\pi'(I - \frac{1}{S}J) \implies \|\pi' L\|_2 \leq \frac{\epsilon}{1 - \epsilon}. \quad (1)$$

The last inequality here is justified as follows (writing $u = \frac{1}{S} [1 \ 1 \ \dots \ 1]$):

$$\|\pi'(I - \frac{1}{S}J)\|_2^2 = \|u - \pi'\|_2^2 = \|u\|_2^2 - 2\langle u, \pi' \rangle + \|\pi'\|_2^2 = \frac{1}{S} - 2 \cdot \frac{1}{S} + \|\pi'\|_2^2 \leq -\frac{1}{S} + \|\pi'\|_1^2 = 1 - \frac{1}{S} \leq 1.$$

Write $\pi' = \pi_1 + \pi_2$, where π_1 is in the kernel of L and π_2 is in the range. Note that $\pi_1 K = \pi_1$; our goal will be to show that π_2 is very small, so that we may take $\pi \approx \pi_1$.

Since π_2 is in the range of L ,

$$\pi_2 = \pi_2 L L^+ = \pi' L L^+, \quad (2)$$

where L^+ is the Moore–Penrose pseudoinverse of L . We claim that $\|L^+\| \leq 2^{\langle K \rangle^b}$ for some universal $b \in \mathbb{Z}^+$; this is because there is a strongly polynomial time algorithm for computing L^+ from L (equivalently, from K), and the spectral norm of a matrix is at most its dimension times its maximum entry (in absolute value) [GLS88, Chapters 0.1, 1.3, 1.4]². Combining this with (1), (2) we get

$$\|\pi_2\|_2 = \|\pi' L L^+\|_2 \leq 2^{\langle K \rangle^b} \|\pi' L\|_2 \leq 2^{\langle K \rangle^b} \frac{\epsilon}{1 - \epsilon} \leq 2 \cdot 2^{\langle K \rangle^b} \epsilon =: \delta. \quad (3)$$

Recall that $\pi_1 = \pi' - \pi_2$, where π' is a nonnegative vector with $\|\pi'\|_1 = 1$. From (3) we deduce:

$$\|\pi_1 - \pi'\|_1 \leq \sqrt{S}\delta; \quad (\pi_1)_i \geq -\delta \quad \forall i \in [S]; \quad \|\pi_1\|_1 \geq 1 - \|\pi_2\|_1 \geq 1 - \sqrt{S}\delta. \quad (4)$$

Further, π_1 is in the kernel of L . By Perron–Frobenius theory (specifically, see [Rot75, Theorem 3.1, Corollary 3.5], [And11, Chapter 3.5.1]) it is known that there is a basis for the kernel of L consisting of stationary distributions ν_1, \dots, ν_s for M with disjoint supports ($1 \leq s \leq S$). Thus

$$\pi_1 = \alpha_1 \nu_1 + \alpha_2 \nu_2 + \dots + \alpha_s \nu_s \quad (5)$$

for some α_j 's satisfying $\alpha_1 + \dots + \alpha_s = 1$. Some of the α_j 's may be negative; however we must always have $\alpha_j \geq -S\delta$, as otherwise $\alpha_j \nu_j$ would have an entry smaller than $-\delta$, contradicting (4) (recall that the ν_j 's have disjoint supports). Let us form $\widetilde{\pi}_1 \geq 0$ by dropping from (5) all contributions $\alpha_j \nu_j$ with α_j nonpositive. As there are at most S such j , this changes the ℓ_1 -norm by at most $S^2\delta$;

²This reference only discusses computing the usual inverse of a nonsingular matrix, but it is clear from the Gaussian Elimination-based algorithm how to extend it to pseudoinverses.

we therefore conclude that $\widetilde{\pi}_1 \geq 0$ is a positive linear combination of stationary distributions for M satisfying

$$\|\widetilde{\pi}_1 - \pi'\|_1 \leq (S^2 + \sqrt{S})\delta \leq 2S^2\delta = 4S^2 \cdot 2^{\langle K \rangle^b} \epsilon \leq \frac{1}{2} \cdot 2^{\langle K \rangle^c} \epsilon, \quad (6)$$

the last inequality by taking $c > b$ large enough (and using $\langle K \rangle \geq S$). Finally, we normalize $\widetilde{\pi}_1$ to make it a stationary distribution for M , setting $\pi = \widetilde{\pi}_1 / \|\widetilde{\pi}_1\|_1$ (note that the denominator $\|\widetilde{\pi}_1\|_1$ is nonzero; in fact, it's at least $\frac{1}{2}$ by (6) and the assumption $\epsilon \leq 2^{-\langle K \rangle^c}$). Then

$$\|\pi - \widetilde{\pi}_1\|_1 = \left\| \pi - \frac{\widetilde{\pi}_1}{\|\widetilde{\pi}_1\|_1} \cdot \pi \right\|_1 = \left| 1 - \frac{\|\widetilde{\pi}_1\|_1}{\|\pi\|_1} \right| \cdot \|\pi\|_1 = \left| 1 - \|\widetilde{\pi}_1\|_1 \right| \leq \frac{1}{2} \cdot 2^{\langle K \rangle^c} \epsilon, \quad (7)$$

the inequality following from (6). Now combining (6) and (7) yields $\|\pi - \pi'\|_1 \leq 2^{\langle K \rangle^c} \epsilon$, as needed. \square

4 BPP_{path}, postselection, and restarts

In this section we describe three different viewpoints on the class BPP_{path}.

The complexity class BPP_{path} was originally defined by Han, Hemaspaandra, and Thierauf [HHT93, HHT97], in a paper also concerned with certain cryptographic problems. (It was also independently defined much later in a paper by Aspnes, Fischer, Fischer, Kao, and Kumar [AFF⁺01, AFF⁺04] on the computational complexity of the stock market.) We quote Fortnow's explanation of the original definition when he named it "Complexity Class of the Week" [For03]:

"Let us call a nondeterministic Turing machine M balanced if for every input x , all of its computational paths have the same length. [We can define the] class BPP as follows: L is in BPP if there is a balanced nondeterministic polynomial-time M such that:

- If x is in L then there are at least twice as many accepting as rejecting paths of $M(x)$.
- If x is not in L then there are at least twice as many rejecting as accepting paths of $M(x)$.

Suppose we use the same definition without the "balanced" requirement. This gives us the class BPP_{path}."

Interestingly, the analogous class "PP_{path}" — for which $x \in L$ iff $M(x)$ has more accepting than rejecting (unbalanced) paths — was defined much earlier in Simon's 1975 thesis [Sim75]. Simon showed that PP_{path} is equal to the class PP (which had recently been defined by Gill [Gil74]). By way of contrast, BPP_{path} is very unlikely to equal BPP, as it is known [HHT93] that BPP_{path} contains both MA and P_{||}^{NP}. BPP_{path} is also known [HHT93] to be contained in BPP_{||}^{NP}. Indeed, under the standard complexity assumptions used to derandomize AM, Shaltiel and Umans [SU06] showed that BPP_{path} = P_{||}^{NP}. For a related class known as SBP, which sits between MA and BPP_{path}, see [BGM06, BGM03].

Another characterization of BPP_{path} was given by Aaronson, via the notion of adding *postselection* (see [Aar05b]) to a complexity class. Suppose that you have a probabilistic algorithm that can end in three kinds of final states: accepting, rejecting, and indecisive. We assume the probability of ending in a decisive state is guaranteed to be nonzero. "Postselection" refers to the (nonrealistic) ability to *condition* the computation on ending in a decisive state. This yields probabilistic computation with just the usual two kinds of final states. For example, one says that $L \in \text{PostBPP}$ if there is a polynomial-time randomized Turing machine as described above which, for each input x , gives the correct answer about $x \in L$ with probability at least $\frac{2}{3}$, *conditioned* on not ending in an indecisive state. More generally, if C is a probabilistic or quantum complexity class, $\text{Post}C$ is the class of languages decided by C -machines with the ability to postselect on ending

in a decisive state. In [Aar04a], Aaronson proved that $\text{PostBQP} = \text{PP}$; later [Aar04b], he observed that $\text{PostBPP} = \text{BPP}_{\text{path}}$.

In the derivation of our main theorem we will prefer a third perspective on BPP_{path} and post-selection, introduced by Yakaryılmaz and Say [YS13]: that of randomized algorithms with *restarts*. For some probabilistic complexity class C , suppose again that we have C -machines that can end in one of three states: accept, reject, or indecisive. We think of the third state as the *restart* state, imagining that whenever the C -machine enters such a state, it immediately restarts its computation from the initial configuration, using no information that it may have gathered up to that point.³ As observed in [YS13], the class of languages decided by such a machine is again PostC . In particular, BPP_{path} is the class of languages that are decided by bounded-error probabilistic polynomial-time Turing machines with this ability to restart. This perspective seems most useful for algorithm design. As an illustration, we believe it is fairly “obvious” that the following restarting-algorithm decides SAT with very high probability, thereby showing $\text{NP} \subseteq \text{BPP}_{\text{path}}$:

“On input formula ϕ with n variables, randomly choose an assignment \mathbf{y} . If \mathbf{y} satisfies ϕ , accept. Otherwise, restart with probability $1 - 2^{-n^2}$ and reject with probability 2^{-n^2} .”

(8)

The reader may compare (8) with the 1-bit CTC algorithm for SAT from Example 2.5, which also shows $\text{NP} \subseteq \text{BPP}_{\text{path}}$ in light of Say and Yakaryılmaz’s Theorem 2.7, $\text{BPP}_{\text{CTC}[1]} = \text{BPP}_{\text{path}}$. For another characterization of BPP_{path} in terms of time-traveling bits, see Section 4.2.

Remark 4.1. For all three definitions of BPP_{path} described above, it is easy to see that the “ $\frac{2}{3}$ cutoff” for success could equivalently be an “ α cutoff” for any fixed constant $\frac{1}{2} < \alpha < 1$, just as is the case for the class BPP .

4.1 Remarks on random coins for BPP_{path} algorithms

In informal descriptions of randomized algorithms, it’s typical to make statements like, “Next, with probability $\frac{1}{3}$ the algorithm. . .” Such statements sweep a well-known, minor detail under the rug; namely, the traditional BPP model (based on nondeterministic branching) only has “access” to probability- $\frac{1}{2}$ coin flips. Of course, this is not an essential problem, since one can simulate a $\frac{1}{3}$ -biased coin flip to error δ in time $O(\log \frac{1}{\delta})$, and δ needn’t be smaller than $1/\text{poly}(n)$. Here we remark that in the context of *restarting* algorithms, the problem is not just inessential, it’s literally no problem at all:

Lemma 4.2. *A restarting randomized algorithm can simulate a p -biased coin flip exactly in time $O(\langle p \rangle)$, and can simulate a uniformly random draw from $[n]$ exactly in time $O(\log n)$.*

Proof. We give the simplest example, leaving the general case for the reader. Suppose we wish to draw $\mathbf{r} \sim [3]$ uniformly at random. We toss two probability- $\frac{1}{2}$ coins, forming a 2-bit integer $0 \leq \mathbf{r} \leq 3$. Then if $\mathbf{r} = 0$, we restart. \square

On the other hand, it’s also important to remember that time $\Omega(\langle p \rangle)$ is also *required* to flip a p -biased coin; for example, in algorithm (8) the step “reject with probability 2^{-n^2} else restart” takes n^2 time steps.

³Note that the new algorithm obtained in this manner will in general have unbounded runtime, even if C is a time-bounded class.

4.2 Postselected CTCs

A physically inequivalent alternative to Deutsch’s model of time travel is the paradigm of postselected CTCs [LMGP⁺10], wherein the time-traveler is forbidden by Nature from changing his own past: The grandfather simply cannot be killed by the time-traveling bomb, precisely because he flourished and went on to have grandchildren! Any outcome of the time-travel device that has the effect of changing the past is assigned a probability of zero, and the remaining outcomes are postselected.

Postselected CTCs are especially straightforward to use as a computational aid, since their very definition indicates that they may be used to realize a computer with the power of postselection: Read a single bit b from the CTC, run a probabilistic algorithm with final states classified into the accepting, rejecting, and indecisive categories as described above, flip b to $1 - b$ if and only if you reach an indecisive state, and send that bit back through the CTC at the end. This has the effect of creating a Grandfather Paradox precisely for the computational paths ending with indecision, which would then never occur according to the model. It has been shown [LMGP⁺10, BW12, SY12] that the class of languages decided by C-machines which have been augmented with a (postselected) CTC register of width 1 is PostC , and this characterization does not change if one increases the width of the CTC register. We therefore have yet another characterization of BPP_{path} , as the class of languages decided by polynomial-time probabilistic algorithms augmented with a postselected CTC of any width.

5 Proof of the main theorem

We now give the proof of our main theorem, that $\text{BPP}_{\text{CTC}[O(\log n)]} \subseteq \text{BPP}_{\text{path}}$.

Proof. Let $L \in \text{BPP}_{\text{CTC}[O(\log n)]}$; say L is defined by algorithm \mathcal{A} as in Definition 2.2. Thus there are constants $c_1, c_2, c_3 \in \mathbb{N}$ such that on inputs $x \in \{0, 1\}^n$, algorithm \mathcal{A} outputs state-transition circuits \mathcal{M}_x of size $O(n^{c_1})$ defining Markov chains M_x on $S = O(n^{c_2})$ states, as well as decision circuits \mathcal{D}_x of size $O(n^{c_3})$. Further, we have that for each x , if π is *any* stationary distribution for M_x and $\mathbf{i} \sim \pi$, then

$$\Pr[\mathcal{D}_x(\mathbf{i}) = 1_{\{x \in L\}}] \geq \frac{2}{3}. \tag{9}$$

Our goal will be to define a polynomial-time randomized restarting algorithm \mathcal{R} that has

$$\Pr[\mathcal{R}(x) = 1_{\{x \in L\}}] \geq 0.65 \tag{10}$$

for all x . As discussed in Section 4, this will show that $L \in \text{BPP}_{\text{path}}$, as required.

On input $x \in \{0, 1\}^n$, the first step of algorithm \mathcal{R} involves invoking the technical Lemma 3.2, with a certain $\epsilon = 2^{-\text{poly}(n)}$ to be described later. More precisely, \mathcal{R} first simulates \mathcal{A} to get state-transition oracle circuit \mathcal{M}_x for Markov chain M_x . It then constructs a state-transition oracle circuit \mathcal{M}'_x for the irreducible perturbed chain M'_x , using the description of \mathcal{M}_x in a black-box fashion. Let π' denote the stationary distribution for M'_x and let K denote the transition matrix for M_x . By definition, K is square matrix of dimension $S \leq O(n^{c_2})$ in which each entry is an integer multiple of $2^{-\text{size}(M_x)} = 2^{-O(n^{c_1})}$. It follows that $\langle K \rangle \leq O(n^{c_4})$ for some constant $c_4 \in \mathbb{N}$. We will now specify that $\epsilon = 2^{-bn^b}$ for a sufficiently large constant b depending on c_1, c_2 , and the c of Lemma 3.2, so that the lemma implies

$$\|\pi - \pi'\|_1 \leq .01 \tag{11}$$

for *some* stationary distribution π of M_x . It is easy to see that with this choice of ϵ , algorithm \mathcal{R} can construct \mathcal{M}'_x from \mathcal{M}_x in $\text{poly}(n)$ time. (Here we use the fact that ϵ is “only” exponentially small; cf. the last paragraph in Section 4.1.)

The remainder of the proof is devoted to showing that \mathcal{R} can obtain an *exact* sample $\mathbf{r} \sim \pi'$. Having shown this, we only need to let \mathcal{R} simulate \mathcal{A} to get \mathcal{D}_x , and then output $\mathcal{D}_x(\mathbf{r})$. Then combining (11) with (9) shows that (10) holds for all x , as required.

We now exhibit the subroutine which the restarting-algorithm \mathcal{R} will use to obtain an exact sample $\mathbf{r} \sim \pi'$ (the unique stationary distribution of irreducible chain M'_x):

- Choose a uniformly random labeled, rooted, undirected tree \mathbf{T} on vertex set $[S]$. This can be done exactly (i.e., with each \mathbf{T} occurring with probability $1/S^{S-1}$) in $\text{poly}(S) = \text{poly}(n)$ time, by choosing a uniformly random Prüfer Code [Prü18, Wik14] in $[S]^{S-2}$, converting it to a tree \mathbf{T} , and then choosing a random vertex \mathbf{r} of \mathbf{T} to be the root.⁴
- Make \mathbf{T} into a rooted arborescence $\vec{\mathbf{T}}$ by directing all edges toward the root \mathbf{r} .
- For each directed edge $(i, j) \in \vec{\mathbf{T}}$, simulate $\mathcal{M}'_x(i)$ and “check” if the output is j . If the check fails, restart.
- If all $S - 1$ checks pass, halt with output \mathbf{r} .

The fact that this subroutine restarts with probability strictly less than 1 follows from the fact that M'_x is irreducible; indeed, its underlying digraph $G = G_{M'_x}$ is the complete digraph. The probability P_r that this subroutine outputs $\mathbf{r} = r$ without encountering any restarts is precisely

$$P_r = \sum_{r\text{-rooted arborescences } T} \frac{1}{S^{S-1}} \prod_{(i,j) \in T} p_{ij},$$

where p_{ij} denotes the transition probability from i to j in the Markov chain M'_x . It follows that the probability of \mathcal{R} finally outputting r (when restarts are taken into account) is

$$\frac{P_r}{\sum_{r \in [S]} P_r} = \frac{\sum_{r\text{-rooted arborescences } T} \prod_{(i,j) \in T} p_{ij}}{\sum_{\text{arborescences } T} \prod_{(i,j) \in T} p_{ij}}.$$

By the Markov Chain Tree Theorem, this is indeed precisely the probability $\pi'(r)$ of r under the stationary distribution of M'_x . \square

We conclude this section by observing that besides the power of restarting, algorithm \mathcal{R} only really needed the power to simulate the state-transition circuits \mathcal{M}_x and the decision circuits \mathcal{D}_x . For example, if these were standard quantum circuits, it would suffice for \mathcal{R} to be a quantum algorithm. Thus we may also conclude $\text{BQP}_{\text{CTC}[\mathcal{O}(\log n)]} \subseteq \text{PostBQP} = \text{PP}$.

We also add that Say and Yakaryılmaz [SY12] studied various models of *finite automata* augmented with 1-bit CTCs; they showed that this augmentation causes both probabilistic and quantum finite automata to become as powerful as their respective postselected versions. The technique used in the proof of our main result can be simplified easily to show that no additional gain arises when these machines are augmented with larger constant-width CTCs.

⁴Here we may use restarting to get exactly random samples from $[S]$; see Lemma 4.2.

6 Conclusion

A very interesting open question left by our work is one also raised at the end of [SY12]:

What is the computational power conferred by 1 time-traveling *qubit*?

Answering this question precisely would seem to require a good understanding of stationary density matrices for 1-qubit quantum channels. As mentioned in Section 1.1, we are not aware of any work showing that time-traveling qubits confer more computational power than time-traveling bits.

Acknowledgment

The first author would like to thank the Boğaziçi University Computer Engineering Department for their hospitality.

References

- [Aar04a] Scott Aaronson. Is quantum mechanics an island in Theoryspace? Technical Report quant-ph/0401062, arXiv, 2004. 4
- [Aar04b] Scott Aaronson. *Limits on Efficient Computation in the Physical World*. PhD thesis, University of California, Berkeley, 2004. 4
- [Aar05a] Scott Aaronson. NP-complete problems and physical reality. *ACM SIGACT News*, 36(1):30–52, 2005. 1.1, 1.2
- [Aar05b] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461(2063):3473–3482, 2005. 1.2, 4
- [AFF⁺01] James Aspnes, David Fischer, Michael Fischer, Ming-Yang Kao, and Alok Kumar. Towards understanding the predictability of stock markets from the perspective of computational complexity. In *soda01*, pages 745–754, 2001. 4
- [AFF⁺04] James Aspnes, David Fischer, Michael Fischer, Ming-Yang Kao, and Alok Kumar. Towards understanding the predictability of stock markets from the perspective of computational complexity. In *New Directions in Statistical Physics*, pages 129–151. Springer Berlin Heidelberg, 2004. 4
- [AGT92] Søren Asmussen, Peter Glynn, and Hermann Thorisson. Stationarity detection in the initial transient problem. *ACM Transactions on Modeling and Computer Simulation*, 2(2):130–157, 1992. 2.2
- [Ald02] David Aldous. Reversible Markov chains and random walks on graphs, 2002. <http://www.stat.berkeley.edu/~aldous/RWG/book.pdf>. 2.2
- [And11] David Anderson. Introduction to stochastic processes with applications in the bio-sciences, 2011. <http://www.math.wisc.edu/~anderson/605F13/Notes/StochBio.pdf>. 3
- [AT89] Venkat Anantharam and Pantelis Tsoucas. A proof of the Markov chain tree theorem. *Statistics & Probability Letters*, 8(2):189–192, 1989. 2.2

- [AW09] Scott Aaronson and John Watrous. Closed timelike curves make quantum and classical computing equivalent. *Proceedings of the Royal Society A*, 465(2102):631–647, 2009. [1.2](#), [2](#), [2.3](#), [2.6](#)
- [Bac04] Dave Bacon. Quantum computational complexity in the presence of closed timelike curves. *Physical Review A*, 70(3):032309, 2004. [1.1](#), [1.2](#), [1](#)
- [BGM03] Elmar Böhler, Christian Glaßer, and Daniel Meister. Small bounded-error computations and completeness. Technical Report 69, Electronic Colloquium on Computational Complexity, 2003. [4](#)
- [BGM06] Elmar Böhler, Christian Glaßer, and Daniel Meister. Error-bounded probabilistic computations between MA and AM. *Journal of Computer and System Sciences*, 72(6):1043–1076, 2006. [4](#)
- [Bro89] Andrei Broder. Generating random spanning trees. In *focs89*, pages 442–447, 1989. [2.2](#)
- [Bru03] Todd Brun. Computers with closed timelike curves can solve hard problems efficiently. *Foundations of Physics Letters*, 16(3):245–253, 2003. [1.1](#)
- [BW12] Todd Brun and Mark Wilde. Perfect state distinguishability and computational speedups with postselected closed timelike curves. *Foundations of Physics*, 42(3):341–361, 2012. [4.2](#)
- [CM01] Grace Cho and Carl Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335(1-3):137–150, 2001. [3](#)
- [Deu91] David Deutsch. Quantum mechanics near closed timelike lines. *Physical Review D*, 44(10):3197–3217, 1991. [1](#), [1.1](#)
- [FMN⁺90] John Friedman, Michael Morris, Igor Novikov, Fernando Echeverria, Gunnar Klinkhammer, Kip Thorne, and Ulvi Yurtsever. Cauchy problem in spacetimes with closed timelike curves. *Physical Review D*, 42(6):1915–1930, 1990. [1](#)
- [For03] Lance Fortnow. Complexity class of the week: BPP_{path} , February 2003. <http://blog.computationalcomplexity.org/2003/02/complexity-class-of-week-bpppath.html>. [4](#)
- [Gil74] John Gill. Computational complexity of probabilistic Turing machines. In *stoc74*, pages 91–95, 1974. [4](#)
- [GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer–Verlag, 1988. [3.1](#), [3](#)
- [Göd49] Kurt Gödel. An example of a new type of cosmological solutions of Einstein’s field equations of gravitation. *Reviews of Modern Physics*, 21(3):447–450, 1949. [1](#)
- [HHT93] Yenjo Hem, Lane Hemaspaandra, and Thomas Thierauf. Threshold computation and cryptographic security. In *Proceedings of the 4th Annual International Symposium on Algorithms and Computation*, pages 230–239, 1993. [4](#)

- [HHT97] Yenjo Han, Lane Hemaspaandra, and Thomas Thierauf. Threshold computation and cryptographic security. *SIAM Journal on Computing*, 26(1):59–78, 1997. 4
- [Hil66] Terrell Hill. Studies in irreversible thermodynamics IV: diagrammatic representation of steady state fluxes for unimolecular systems. *Journal of Theoretical Biology*, 10(3):442–459, 1966. 2.2
- [KV80] Hans-Helmut Kohler and Eva Vollmerhaus. The frequency of cyclic processes in biological multistate systems. *Journal of Mathematical Biology*, 9(3):275–290, 1980. 2.2
- [LMGP⁺10] Seth Lloyd, Lorenzo Maccone, Raul Garcia-Patron, Vittorio Giovannetti, and Yutaka Shikano. The quantum mechanics of time travel through post-selected teleportation. Technical Report 1007.2615, arXiv, 2010. 1, 4.2
- [LR83] F. Thomson Leighton and Ronald Rivest. The Markov Chain Tree Theorem. Technical Report TM-249, Massachusetts Institute of Technology, 1983. 2.2, 3
- [LW95] László Lovász and Peter Winkler. Exact mixing in an unknown Markov chain. *Electronic Journal of Combinatorics*, 2:Research Paper 15, 1995. 2.2
- [PBMW99] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. Technical Report SIDL-WP-1999-0120, Stanford, 1999. 3
- [Pok99] Piotr Pokarowski. Directed forests with application to algorithms related to Markov chains. *Applicationes Mathematicae*, 26(4):395–414, 1999. 2.2
- [Prü18] Heinz Prüfer. Neuer beweis eines satzes über permutationen. *Archiv der Mathematik und Physik*, 27:742–744, 1918. 5
- [PW98] James Propp and David Wilson. How to get a perfectly random sample from a generic markov chain and generate a random spanning tree of a directed graph. *Journal of Algorithms*, 27(2):170–217, 1998. 2.2
- [Rot75] Uriel Rothblum. Algebraic eigenspaces of nonnegative matrices. *Linear Algebra and its Applications*, 12:281–292, 1975. 3
- [Sch68] Paul Schweitzer. Perturbation theory and finite Markov chains. *Journal of Applied Probability*, 5:401–413, 1968. 3
- [Shi03] Yaoyun Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computing. *Quantum Information & Computation*, 3(1):84–92, 2003. 2.1
- [Shu75] Bruno Shubert. A flow-graph formula for the stationary distribution of a Markov chain. *Institute of Electrical and Electronics Engineers. Transactions on Systems, Man, and Cybernetics*, SMC-5(5):565–566, 1975. 2.2
- [Sim75] Janos Simon. *On some central problems in computational complexity*. PhD thesis, Cornell University, 1975. TR 75-224. 4
- [SU06] Ronen Shaltiel and Christopher Umans. Pseudorandomness for approximate counting and sampling. *Computational Complexity*, 15(4):298–341, 2006. 4

- [SY12] A. C. Cem Say and Abuzer Yakaryılmaz. Computation with multiple CTCs of fixed length and width. *Natural Computing*, 11(4):579–594, 2012. [1.2](#), [2](#), [2.3](#), [2.5](#), [2.7](#), [2.2](#), [3](#), [4.2](#), [5](#), [6](#)
- [WF70] Alexander Wentzell and Mark Freidlin. On small random perturbations of dynamical systems. *Russian Mathematical Surveys*, 25(1):1–55, 1970. [2.2](#)
- [Wik14] Wikipedia. Prüfer sequence, July 2014. http://en.wikipedia.org/wiki/Prufer_sequence. [5](#)
- [Wil96] David Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 296–303, 1996. [2.2](#)
- [Wol12] Michael Wolf. Quantum channels & operations: guided tour, 2012. <http://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf>. [1.1](#)
- [YS13] Abuzer Yakaryılmaz and A. C. Cem Say. Proving the power of postselection. *Fundamenta Informaticae*, 123(1):107–134, 2013. [2.3](#), [4](#)