

Fraud Detection for Random Walks

Varsha Dani ✉🏠

Rochester Institute of Technology, Rochester, NY, USA

Thomas P. Hayes ✉

University at Buffalo, Buffalo, NY, USA

Seth Pettie ✉🏠

University of Michigan, Ann Arbor, MI, USA

Jared Saia ✉🏠

University of New Mexico, Albuquerque, NM, USA

Abstract

Traditional fraud detection is often based on finding statistical anomalies in data sets and transaction histories. A sophisticated fraudster, aware of the exact kinds of tests being deployed, might be difficult or impossible to catch. We are interested in paradigms for fraud detection that are *provably robust* against any adversary, no matter how sophisticated. In other words, the detection strategy should rely on signals in the data that are inherent in the goals the adversary is trying to achieve.

Specifically, we consider a fraud detection game centered on a random walk on a graph. We assume this random walk is implemented by having a player at each vertex, who can be honest or not. In particular, when the random walk reaches a vertex owned by an honest player, it proceeds to a uniformly random neighbor at the next timestep. However, when the random walk reaches a dishonest player, it instead proceeds to an arbitrary neighbor chosen by an omniscient Adversary.

The game is played between the Adversary and a Referee who sees the trajectory of the random walk. At any point during the random walk, if the Referee determines that a *specific* vertex is controlled by a dishonest player, the Referee accuses that player, and therefore wins the game. The Referee is allowed to make the occasional incorrect accusation, but must follow a policy that makes such mistakes with small probability of error. The goal of the adversary is to make the cover time large, ideally infinite, i.e., the walk should *never* reach at least one vertex. We consider the following basic question: how much can the omniscient Adversary delay the cover time without getting caught? Our main result is a tight upper bound on this delay factor.

We also discuss possible applications of our results to settings such as Rotor Walks, Leader Election, and Sybil Defense.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains; Mathematics of computing → Probability and statistics; Security and privacy → Intrusion detection systems

Keywords and phrases Fraud detection, random processes, Markov chains

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.36

Funding *Seth Pettie*: Supported by NSF Grants CCF-1815316 and CCF-2221980.

Jared Saia: Supported by NSF Grants CNS-2210299 and NSF NRI-2024520.

1 Introduction

Many modern fraud detection efforts look for *statistical* features of data that do not fit a known probabilistic model, or are intrinsically implausible or internally inconsistent. The Newcomb–Benford (“first digit”) Law [29, 30, 26, 28] is a well known filter for detecting fabricated data in financial records, which can be applied to detecting fraud in other numerical data, e.g., manipulated images [14, 44]. Recently uncovered frauds in social science research [35, 36, 37] can also be seen as *distribution testing* against known or unknown distributions.



© Varsha Dani, Thomas P. Hayes, Seth Pettie, and Jared Saia;
licensed under Creative Commons License CC-BY 4.0

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).

Editor: Venkatesan Guruswami; Article No. 36; pp. 36:1–36:22

Leibniz International Proceedings in Informatics



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

One weakness of this variety of fraud detection is that it preys on relatively unsophisticated fraudsters, who could easily evade detection if they were just aware of the statistical tests in advance. This critique could also be leveled against most fraud detection efforts in machine learning and information retrieval, which treat it as a pattern-matching problem [9, 13, 40, 38, 42, 31, 39].

In this paper we advance a perspective on fraud detection that differs sharply from all the work cited above. First, rather than begin with an *application domain* or a single empirical *instance* of fraud, we want to build a more general theory of fraud detection. In the most fundamental examples cited above, fraud manifests as corruption of a random process. Thus, we focus our study on abstract random processes that can be perturbed by an adversary. Furthermore, we adopt the norms of theoretical computer science, cryptography, and game theory in our adversarial model. In particular, a fraud detection mechanism should be evaluated in a *worst case* fashion, ideally against a computationally unbounded and omniscient adversary. Following Kerckhoffs' principle [23], its success should *not* depend on assuming the adversary is ignorant of the statistical tests it will be subject to.

1.1 Fraud Detection for Random Walks

Let $G = (V, E)$ be a connected, undirected graph. A random walk $(v_i)_{i \geq 0}$ is generated by placing a token at some v_0 and, in each step, letting v_i pass the token to a uniformly random neighbor $v_{i+1} \in N(v_i)$. The cover time for this walk is the time until all the vertices have been visited by the token.

Now suppose an adversary *corrupts* a set $B \subseteq V$ of up to b vertices, who may pass the token as they like. The adversary wishes to delay the cover time as much as possible, without being detected.

It is well known [2] that for any graph, the cover time is $O(mn \log n)$ with high probability. So if, after this many steps, there are vertices that have not been reached, the *existence* of corruption will be evident. However, we require a stronger form of fraud detection: a *specific* vertex must be accused. We formalize this process as the following game.

► **Definition 1** (The Random Walk Game). *Let T be a fixed time horizon and $b \leq n$ a fixed number. The game is played between two players, the Referee and the Adversary. The Adversary picks a starting vertex $v_0 \in V$ and a subset $B \subseteq V$ of (corrupt) vertices with $|B| \leq b$. A walk $(v_0, v_1, \dots, v_T)$ is constructed iteratively, with each move from an honest vertex being random, and each move from a corrupt vertex being chosen by the Adversary. If $\{v_0, v_1, \dots, v_T\} = V$, the Referee immediately wins (the vertex set has been covered). Otherwise, the Referee must specify one "accused" vertex; the Referee wins if and only if this vertex is in B .*


We are interested in the *threshold time*, $T(G, b)$, which is the minimum time T such that, with best play, the Referee wins the T -step Random Walk Game with probability at least $1 - 1/n^5$.

We note that there is nothing particularly special about the exponent 5 in the allowed error probability above, and could instead make the error probability $1/n^C$. However, for our lower bounds, we do require that C be large enough to avoid pathological examples where the Referee could accuse a random vertex of being in B and be correct just by chance.

When $b = 0$, the threshold time $T(G, 0)$ is essentially the expected cover time. More precisely, if τ is the maximum, over all starting locations, of the expected cover time of G , then

$$\tau/2 \leq T(G, 0) \leq (10 \log n) \tau$$

with the actual value depending on the specific graph.



■ **Figure 1** Comparison of the bounds on $R(G, b)$ from our main results. All the log terms have been dropped, and n has been set to 100. The blue curve is b , the general upper bound on R from Theorem 3. We note that, for every b , there is a graph for which this upper bound is tight (up to log factors). The red curve is $1 + b^2/n$, which is $\Theta(R)$ in the special case of the path, as stated in Theorem 6. Note that $1 + b^2/n$ is also the right value of $R(b)$ in the special case of the clique, if the referee is restricted to purely local strategies that make accusations only a function of the particular player's choices. The green curve gives the correct value of $R(b)$ for the clique, when the referee is allowed to make accusations based on the entire transcript. This result is given in Theorem 13.

We now introduce our main object of study.

► **Definition 2.** We define the price of corruption as the ratio

$$R(G, b) = \frac{T(G, b)}{T(G, 0)}$$

Informally, this is the factor by which an adversary with up to b corrupt vertices can increase the cover time, before the referee will be able to reliably accuse a bad player.

Our goal in this paper is to understand how much $T(G, b)$, and therefore $R(G, b)$, can depend on b . Our main result is that this dependence is at most nearly linear

► **Theorem 3** (Price of Corruption is at most nearly linear). *Let G be any graph on n vertices, and let $0 \leq b \leq n$. Then,*

$$R(G, b) = O(b \log n).$$

Moreover, there exists a family of graphs $G = G(n, b)$ for which

$$R(G, b) = \Omega(b / \log n).$$

The lower bound in Theorem 3 does not apply to all graphs. For instance, we will see that the behavior of $R(G, b)$ is more nuanced in the cases when G is a path or a clique; we examine these special cases in Sections 2 and 4

This suggests a related question: for which graphs is the Price of Corruption, $R(G, b)$, smallest? Knowing this might be helpful in applications where we have some choice about the graph on which the random walk takes place. Small-degree expander graphs seem like particularly good candidates for bounds of this type.

1.2 Related Work

Biased random walks are a mainstay of introductory courses in random processes. Azar, Broder, Karlin, Linial, and Phillips [8] studied the adversarial biasing of random walks to maximize the time spent among some target set. In their model the token moves randomly a $(1 - \epsilon)$ -fraction of the time, and is controlled by the adversary an ϵ -fraction of the time. Azar et al. [8] did not consider the problem of *detecting* such interventions or evading detection.

Our problem is inspired by the Byzantine Agreement protocols of King and Saia [24] and Huang, Pettie, and Zhu [21, 22], which achieved polynomial latency with $f = \Theta(n)$ and optimal $f < n/3$ resiliency (Byzantine corruptions), respectively. These protocols attempt to flip a fair coin via a natural distributed coin-flipping protocol. However, the adversary may interfere with the protocol by choosing coin-flip outcomes strategically, and by inducing subtle disagreements among the non-corrupt players. If such an adversary continually foils attempts to flip a fair coin, an individual Byzantine player can be identified and *blacklisted*, removing its influence over the coin flipping protocol.¹

The notion of fraud detection seems to be “in the air” these days. This year Alon, Gunby, He, Shmaya, and Solan [3] also proposed a fraud detection-type game for random walks. In their model a walk on $\mathbb{Z}$ begins near the origin and is run in perpetuity but never reaches the origin, or does not reach it infinitely often. The movement of the walk is controlled by two players, Alice and Bob, who alternate (purportedly) flipping fair coins and announcing outcomes in $\{-1, 1\}$ – but exactly one of them is a fraud. The question is how to detect which of Alice or Bob is not behaving correctly. Their fraud detection mechanism is not an “algorithm” *per se*, as it requires evaluating functions of infinitely long walks. Although our setup and the setup of [3] have some syntactic similarities, the mathematical structure of the two problems are different and lead, in some ways, to opposite conclusions.²

1.2.1 Random Walks and Dynamic Networks

Several recent results make use of random walks to solve classic problems in distributed computing over dynamically changing networks in the presence of Byzantine nodes. Problems addressed include Byzantine agreement [5]; information dissemination [34]; and leader election [6]. See also [7] for a survey of results.

The type of random walk problem considered in these results is more general than ours in that the network topology may change from step-to-step. The problem is more specific than ours in that the network is assumed to always be a regular expander; and the number of Byzantine nodes is always $O(\sqrt{n}/\log^k n)$ for some constant k .

Central to these results is a technical lemma showing that if good nodes generate random walk tokens at a certain rate, then there is a large set of nodes that have access to many well-mixed random-walk tokens. The random-walk algorithms are simple: no attempt is made to detect or identify Byzantine behavior, and the algorithms are fully distributed and scalable in terms of latency and message cost.

¹ This application illustrates why it is important to distinguish between global detection – *something* has gone wrong – and specific detection, namely, a *specific* player is corrupt w.h.p.

² Specifically, to make the cover time infinite in our model, the corrupt vertices must have some measurable bias, and the question is how long it takes to detect that bias. In the infinite Alice & Bob game [3], any biases are trivially detected (in the limit); the detector must also pay attention to negative correlations between Alice and Bob’s moves.

1.3 Organization

In Section 1.4 we review Bernstein's and Freedman's concentration inequalities. In Section 2 we analyze the random walk game on the simplest topology, an n -path P_n , and obtain nearly sharp bounds on $T(P_n, b)$. In Section 3 we generalize the detection method to work on an arbitrary graph G , and bound the price of corruption by $R(G, b) = O(b \log n)$. In Section 4 we design a fraud detection method specific to the n -clique K_n , and give nearly tight upper and lower bounds on $T(K_n, b)$. In Section 5 we discuss some possible applications of our results. We conclude with some open problems in Section 6.

1.4 Concentration Inequalities

The Referee's task is to observe the random walk, and identify vertices that are not behaving as they should. In order to do this, we need a fairly accurate idea of what the local behavior of such a random walk *should* look like. To get a handle on this, we will make use of the following concentration inequalities.

The following version of Bernstein's inequality (see [16]) will be useful in analyzing the random walk games on the path (Section 2) and the clique (Section 4).

► **Theorem 4** (Bernstein's Inequality). *Let $X_1, \dots, X_n$ be independent random variables with $|X_i - E(X_i)| \leq b$ for each $i \in [n]$, and each with variance σ_i^2 . Let $X = \sum_i X_i$, and $\sigma^2 = \sum_i \sigma_i^2$ be the variance of X . Then for all $t > 0$,*

$$\Pr(X \leq E(X) - t) \leq \exp\left(-\frac{t^2}{2\sigma^2 + (2/3)bt}\right)$$

When dealing with general graphs (Section 3) we will instead need the following extension of Freedman's inequality for martingales.

► **Theorem 5** ([10, Lem. 2]). *Suppose $X_1, \dots, X_T$ is a martingale difference sequence with $|X_t| \leq \rho$. Let $\mathbf{Var}_t X_t = \mathbf{Var}(X_t \mid X_1, \dots, X_{t-1})$. Let $V = \sum_{t=1}^T \mathbf{Var}_t X_t$ be the sum of conditional variances and $\bar{\sigma} = \sqrt{V}$. Then for any $\delta < 1/e$ and $T \geq 4$,*

$$\mathbb{P}\left(\left|\sum_{t=1}^T X_t\right| \leq 2\sqrt{\ln(1/\delta)} \max\{2\bar{\sigma}, \rho\sqrt{\ln(1/\delta)}\}\right) \geq 1 - \delta \log T.$$

2 The Path

Consider the path graph $G = (V, E)$ with vertices numbered 1 through n . Without loss of generality we can assume the token is initially at vertex 1 and never reaches vertex n . How long must a corrupted random walk be until we may accuse a corrupt vertex?

Theorem 6 gives nearly sharp bounds for this class of graphs and illustrates two qualitative features of this fraud detection model. First, although *one* corrupt vertex can make the cover time infinite it cannot do so without detection, and in fact any coalition of $b = O(\sqrt{n})$ corrupt vertices is powerless to increase the cover time by more than a constant factor, without detection. Second, there is a significant gap between the moment we detect likely corruption ($\Theta(n^2 \log n)$ time) and the moment we can confidently level an accusation at one vertex ($\tilde{\Theta}(n^3)$ time when $b = \Omega(n)$).

► **Theorem 6.** *Let G be the path of length n . Suppose the Random Walk Game on G is played for T timesteps and the adversary is allowed to corrupt up to b vertices. Then*

1. *If $T = \Omega((n^2 + nb^2) \log n)$, then there is a strategy that enables the Referee to win with probability at least $1 - \frac{1}{n^5}$. In other words,*

$$R(G, b) = O\left(1 + \frac{b^2}{n}\right)$$

2. *If $T < n^2 + nb^2$, there is an adversarial strategy such that one vertex is never visited, and no detection mechanism can identify any corrupt vertex with high probability. In other words,*

$$R(G, b) = \Omega\left(\left(1 + \frac{b^2}{n}\right) / \log(n)\right)$$

The remainder of this section constitutes a proof of Theorem 6.

Part 1 of Theorem 6. Suppose we pass the token for T time steps. For each vertex j , let X_j denote the number of times that vertex j passes the token, and Y_j the number of times j passes the token to the left. We will accuse vertex j if the number of left passes, Y_j substantially exceeds the number of right passes, $X_j - Y_j$, or more specifically, if

$$\Delta_j \stackrel{\text{def}}{=} 2Y_j - X_j \geq \sqrt{CX_j \log n}.$$

A standard application of Chernoff's bound ensures that this criterion almost certainly does not falsely accuse any good vertex. In other words after T timesteps, for each good player j ,

$$\Delta_j < \sqrt{CX_j \log n} \tag{1}$$

holds with high probability $1 - n^{-\Omega(C)}$. We may assume that (1) also holds for all bad players as well, since otherwise the algorithm will make a correct accusation.

Let $v^* = \arg \max_v X_v$ be the *mode* vertex. Since the token is passed around for T timesteps, by the pigeonhole principle, $X_{v^*} \geq T/n$. Each time v^* receives the token, it passes it to a neighbor and the token makes a round-trip excursion back to v^* . We may assume that at least $\frac{1}{3}X_{v^*}$ of these round-trip excursions are to the right of v^* , for otherwise Δ_{v^*} would already be large enough to justify accusing v^* . Define G and B to be the sets of good and bad vertices among $\{v^* + 1, v^* + 2, \dots, n - 1\}$. Since vertex n is never reached, every round-trip excursion from v^* to the right entails $G \cup B$ passing the token one more time to the left than the right.


Let Δ_G and Δ_B be the sum of this left-excess associated with G and B , respectively. Then we know $\Delta_G + \Delta_B \geq \frac{1}{3}X_{v^*}$. Let X_G and X_B denote the total number of times the token is passed by vertices in G and B , respectively.

Applying Chernoff's bound to the good vertices as a group, since $T \leq nX_{v^*}$, we have, with high probability $1 - n^{-\Omega(C)}$,

$$\Delta_G \leq \sqrt{CX_G \log n} \leq \sqrt{CT \log n} \leq \sqrt{CnX_{v^*} \log n}. \tag{2}$$

To estimate $\Delta_B = \sum_{i \in B} \Delta_i$, note that each bad vertex satisfies Eqn. (1) to avoid detection.

$$\Delta_B \leq \sum_{i \in B} \sqrt{CX_i \log n} \leq |B| \sqrt{CX_{v^*} \log n} \leq b \sqrt{CX_{v^*} \log n}. \tag{3}$$



■ **Figure 2** A biased random walk on the line graph.

where the second inequality follows from the choice of v^* as the mode. Combining Eqns.(2,3), we have

$$\frac{1}{3}X_{v^*} \leq \Delta_G + \Delta_B \leq \sqrt{CX_{v^*} \log n}(\sqrt{n} + b).$$

Squaring and rearranging terms,

$$X_{v^*} \leq 9C \log n (\sqrt{n} + b)^2.$$

Finally, since $(\sqrt{n} + b)^2 \leq 2(n + b^2)$, we have

$$T \leq nX_{v^*} \leq 18Cn(n + b^2) \log n,$$

which completes the proof of the upper bound. ◀

Part 2 of Theorem 6. For this part it is more convenient to number the vertices in reverse order: vertex 0 is the rightmost vertex and the token begins at vertex $n - 1$ and never reaches 0. See Figure 2.

In adversarial strategy $\mathcal{S}$, the adversary corrupts vertices in $[b] = \{1, \dots, b\}$ and gives vertex j a left-bias of $1/j$. Specifically, vertex j passes left with probability p_j , where

$$p_j = \begin{cases} \frac{1}{2} \left(1 + \frac{1}{j}\right) & \text{if } j \in [b] \\ 1/2 & \text{if } j > b. \end{cases}$$

This process corresponds to a reversible Markov chain on the states $\{n - 1, \dots, 1\}$ where, for $j \in [b]$, the edge between $j + 1$ and j has weight $\binom{j+1}{2}$ and all the edges to the left of b have weight $\binom{b+1}{2}$. Note that the cover time is infinite as vertex 0 is unreachable.

It follows that, for $j \in [b]$, vertex j has stationary probability proportional to j^2 , while the vertices in $\{n - 2, \dots, b + 1\}$ to the left of b all have probability proportional to $b(b + 1)$. The leftmost vertex $n - 1$ has stationary probability proportional to $b(b + 1)/2$, and vertex 0 is unreachable. See Figure 2. Summing these terms, we obtain the normalization factor N to be

$$N = b \frac{(b+1)}{2} + b(b+1)(n-b-2) + \sum_{j=1}^b j^2 = b(b+1) \left(n - b - \frac{3}{2} + \frac{2b+1}{6} \right) = \Theta(nb^2).$$

Thus, for $j \in [b]$, the stationary probability of vertex j is $j^2/N = \Theta(nj^2/b^2)$.

Define $\mathcal{S}_{-j}$ to be identical to strategy $\mathcal{S}$ except that vertex j is not corrupt, i.e., it passes left and right with probability $1/2$. We want to argue that if a corrupted random walk is too short, the false positive rate of *any* detection strategy will be intolerably large. Lemma 7 lower bounds this error.

► **Lemma 7.** *Assume the adversary picks a strategy from $\{\mathcal{S}, \mathcal{S}_{-1}, \dots, \mathcal{S}_{-b}\}$ uniformly at random. Abusing notation, let $\mathcal{S}_{-i}$ also refer to the event that strategy $\mathcal{S}_{-i}$ is chosen. Let W be the resulting corrupted random walk. Define $q = \min_{i \in [b]} \min(\rho_i, 1 - \rho_i)$, where $\rho_i = \Pr(\mathcal{S}_{-i} \mid (\mathcal{S}_{-i} \cup \mathcal{S}), W)$. Then,*

$$\forall i \in [b], \Pr(\mathcal{S}_{-i} \mid W) \geq q^2/b.$$

Proof. Note that for all $i \in [b]$:

$$\Pr(\mathcal{S}_{-i} \mid W) = \Pr(\mathcal{S}_{-i} \mid (\mathcal{S}_{-i} \cup \mathcal{S}), W) \cdot (\Pr(\mathcal{S}_{-i} \mid W) + \Pr(\mathcal{S} \mid W)).$$

Letting $\rho_i = \Pr(\mathcal{S}_{-i} \mid (\mathcal{S}_{-i} \cup \mathcal{S}), W)$, and solving for $\Pr(\mathcal{S}_{-i} \mid W)$ in the above, we get

$$\Pr(\mathcal{S}_{-i} \mid W) = \frac{\rho_i}{1 - \rho_i} \cdot \Pr(\mathcal{S} \mid W).$$

Note that $q \leq \frac{\rho_i}{1 - \rho_i} \leq 1/q$. Letting $i^* = \arg \min_{i \in [b]} \Pr(\mathcal{S}_{-i} \mid W)$, we have:

$$\Pr(\mathcal{S}_{-i^*} \mid W) \geq q \Pr(\mathcal{S} \mid W),$$

and for all $i \in [b] \setminus \{i^*\}$,

$$\Pr(\mathcal{S}_{-i} \mid W) \leq (1/q) \Pr(\mathcal{S} \mid W).$$

Hence,

$$\begin{aligned} 1 &= \Pr(\mathcal{S}_{-i^*} \mid W) + \Pr(\mathcal{S} \mid W) + \sum_{i \in [b], i \neq i^*} \Pr(\mathcal{S}_{-i} \mid W) \\ &\leq \Pr(\mathcal{S}_{-i^*} \mid W) + (1/q) \Pr(\mathcal{S}_{-i^*} \mid W) + ((b-1)/q^2) \Pr(\mathcal{S}_{-i^*} \mid W) \\ &\leq \frac{b}{q^2} \Pr(\mathcal{S}_{-i} \mid W), \end{aligned}$$

where the last inequality follows since $1 + 1/q \leq 1/q^2$. $\blacktriangleleft$

Lemma 7 says that we can assume the detector accuses the vertex $j \in [b]$ that minimizes $\Pr(\mathcal{S}_{-j} \mid (\mathcal{S}_{-j} \cup \mathcal{S}), W)$. In order to make the false positive rate small, we need the likelihood ratio

$$\begin{aligned} \frac{\Pr(W \mid \mathcal{S}_{-j})}{\Pr(W \mid \mathcal{S})} &= \frac{(1/2)^{X_j}}{p_j^{Y_j} (1 - p_j)^{X_j - Y_j}} = \left(1 - \frac{1}{j}\right)^{Y_j} \left(1 + \frac{1}{j-1}\right)^{X_j - Y_j} \\ &< \exp\left(-\frac{Y_j}{j} + \frac{X_j - Y_j}{j-1}\right) = \exp\left(-\frac{\Delta_j}{j-1} + \frac{Y_j}{j(j-1)}\right) \end{aligned}$$

to be $n^{-\Omega(C)}$. By Chernoff bounds, the likelihood ratio never gets this small until $Y_j \geq Cj^2 \log n$, so we may use this as a proxy prerequisite for accusing vertex j .

Once j is visited for the first time, the expected return time is $\Theta(nb^2/j^2)$, so *in expectation*, the criterion $Y_j \geq Cj^2 \log n$ is satisfied after another $\Theta(nb^2 \log n)$ steps. However, these return times have large variances so it is not clear that this random variable is sufficiently concentrated around its mean.³

We may assume without loss of generality that $b \in [\Omega(\sqrt{n}), n/2]$. Let W_j be the length of a random walk that begins and ends at vertex j , conditioned on moving left initially and let $E_j = \mathbb{E}(W_j)$ and $V_j = \mathbb{E}(W_j^2)$.⁴ Such a walk moves to $j+1$, makes zero or more roundtrips from $j+1$, and then returns to j . The number of roundtrips from $j+1$ is distributed geometrically, so by linearity of expectation,

$$E_j = 2 + \left(\frac{1}{1 - p_{j+1}} - 1\right) E_{j+1} = \begin{cases} 2 + \frac{j+2}{j} E_{j+1} & \text{if } j+1 \leq b, \\ 2 + E_{j+1} & \text{if } j+1 > b. \end{cases}$$

³ (Lemma 8 implies that in any graph, the visitation rate of a vertex is, with high probability, at most twice its stationary probability after a sufficiently long (corrupted) random walk, which on the line would be $\Theta(bn^2)$ steps. Since we are looking for a tight bound of $\Theta(n^2 + nb^2)$ we require a more careful analysis.)

⁴ The condition that $b \leq n/2$ implies that starting at a corrupt vertex, a roundtrip to the left is longer in expectation than a roundtrip in general.

Thus $E_b = 2(n - b - 1)$ and writing E_j , $j < b$, in terms of E_b we have a telescoping product, $E_j = \Theta(nb^2/j^2)$.

To bound the second moment V_j , suppose that a leftward roundtrip from j makes k (leftward) roundtrips from $j + 1$ before returning to j , i.e., it has length $2 + W_{j+1}^{(1)} + \dots + W_{j+1}^{(k)}$, where the $W_{j+1}^{(\cdot)}$ are independent copies of W_{j+1} . This would contribute $kV_{j+1} + (k^2 - k)E_{j+1}^2 + 2kE_{j+1} + 4$ to V_j . Thus, we can express V_j recursively as

$$\begin{aligned} V_j &= (1 - p_{j+1}) \sum_{k \geq 0} p_{j+1}^k (kV_{j+1} + (k^2 - k)E_{j+1}^2 + 2kE_{j+1} + 4) \\ &= (1 - p_{j+1}) \frac{p_{j+1}}{(1 - p_{j+1})^2} \cdot V_{j+1} + \Theta(E_{j+1}^2) \\ &= \begin{cases} \frac{j+2}{j} \cdot V_{j+1} + \Theta(E_{j+1}^2) & \text{if } j + 1 \leq b, \\ V_{j+1} + \Theta(E_{j+1}^2) & \text{if } j + 1 > b. \end{cases} \end{aligned}$$

Then $V_b = \Theta(n^3)$, and expressing V_j , $j < b$, in terms of V_b we have another telescoping product with $V_j = \Theta(n^3 b^2/j^2)$.

We claim that it is not possible to reliably accuse any vertex j in less than $M = n^2 + nb^2$ steps. In particular, once j is first visited, the length of the next $Y_j = K = Cj^2 \log n$ leftward roundtrips is not less than M . Let $W_j^{(i)}$ be the length of the i th leftward roundtrip. In expectation $W_j^{(1)} + \dots + W_j^{(K)}$ is $K \cdot E_j = \Theta(Cnb^2 \log n) = \mu$.

We may assume each $|W_j^{(i)}| \leq M$, for otherwise there's nothing to prove. Thus, by Bernstein's inequality,

$$\begin{aligned} \Pr \left(\sum_{i=1}^K W_j^{(i)} < M \right) &< \exp \left(- \frac{(\mu - M)^2}{2 \sum_{i=1}^K \mathbb{E}((W_j^{(i)})^2) + (2/3)M(\mu - M)} \right) \\ &= \exp \left(- \frac{(1 - o(1))(KE_j)^2}{2KV_j + (2/3)(1 + o(1))KE_j \cdot M} \right) \\ &= \exp \left(- \frac{\Theta(Cnb^2 \log n)^2}{\Theta(Cn^3 b^2 \log n) + \Theta(Cnb^2 \log n \cdot (n^2 + nb^2))} \right) \end{aligned}$$

and since $b = \Omega(\sqrt{n})$, $n^2 b^4 = \Omega(n^3 b^2)$,

$$= \exp(-\Omega(C \log n)) = n^{-\Omega(C)}.$$

◀

Theorem 6 gives a nearly tight characterization for the fraud detection time on paths. Qualitatively speaking, Theorem 6 shows that tracking *individual* deviations suffices to achieve near-optimal fraud detection, i.e., a vertex v is judged *solely* on the distribution of token passes to $N(v)$. Section 3 extends this type of analysis to general graphs, and obtains strong bounds for all graphs.

However, tracking *individual deviations* alone is, on some graph topologies, insufficient for optimal fraud detection. The *clique* is one such topology, which we analyze in detail in Section 4.

3 Fraud Detection on General Graphs

In this section we consider the random walk game played on an *arbitrary* connected graph G on n vertices. The Adversary can corrupt any set $\mathcal{B} \subset V$ consisting of up to b of the vertices. As in the case of the path, the Referee will watch the individual vertices and track their

apparent deviation from uniformly random behaviour. We will prove the upper bound in Theorem 3, by showing that there is a Referee strategy that guarantees that if the Adversary tries to delay the cover time by more than an $O(b \log n)$ factor, the Referee has a $1 - 1/n^5$ chance to win. To prove the lower bound we will demonstrate a family of graphs and an Adversarial strategy for which a $O(b/\log n)$ factor is achieved.

Although we eventually want to bound how much the adversary can delay the cover time, it will be convenient to analyze the Random Walk Game in terms of *hitting times*. In the next subsection we discuss some concepts and terminology relating to random walks and hitting times.

3.1 Notation

G is an undirected connected graph on n vertices. For a random walk on G , given vertices v and y , the *hitting time* from v to y is the (random) first time at which the walk, having started at v , arrives at y . The expected hitting times between all the pairs of vertices in G will be of particular interest in designing our Referee strategy.

The following quantities are solely a function of the structure graph G , not strategic considerations of the random walk game.

- π is the stationary distribution, i.e., $\pi(v) = \deg(v)/2m$.
- $H(v, y)$ is the expected hitting time to y starting from v . Let $H_{\max}(y) = \max_v H(v, y)$ and $H_{\max} = \max_y H_{\max}(y)$ be the maximum hitting times when only y is fixed, and when neither is fixed.
- For $w \in N(v)$, define $h_y(v, w)$ to be

$$h_y(v, w) = H(w, y) - H(v, y) + 1 - \frac{\mathbb{1}(v=y)}{\pi(y)}.$$

Here $\mathbb{1}(\mathcal{E})$ is the indicator variable for event $\mathcal{E}$. For $v \neq w$, this definition ensures that $h_y(v, w) - 1$ equals the change in the expected hitting time to y that results from moving across the edge $\{v, w\}$. The definition ensures that, when w is a randomly chosen neighbor of v , the quantity $h_y(v, w) - 1$ has an expected value of -1 , corresponding to the elapsing of the first time step in a random walk from v to y . The extra term, $\frac{\mathbb{1}(v=y)}{\pi(y)}$, which, when $v = y$, equals the expected excursion time from y , ensures that the expected value of $h_y(v, w)$ is zero for all $v \in V$.

- Let $\rho_y(v) = \max_{w \in N(v)} |h_y(v, w)|$, $\rho_y = \max_v \rho_y(v)$ and $\rho = \max_y \rho_y$.
- Let $\sigma_y^2(v) = \frac{1}{\deg(v)} \sum_{w \in N(v)} h_y(v, w)^2$ be the conditional variance of $H(w, y)$, conditioned on v , where as before we assume that w is a randomly chosen neighbor of v .
- Let $\mathcal{V}_\pi^y = \mathbb{E}_\pi \sigma_y^2(v) = \sum_v \sigma_y^2(v) \pi(v)$ be the average conditional variance when v is chosen from the stationary distribution. Let $\mathcal{V}_\pi = \max_y \mathcal{V}_\pi^y$ be the maximum of this average over all target vertices y .

Now consider the Random Walk Game. Let $\mathcal{G}$ be the set of good vertices and $\mathcal{B}$ be the set of bad vertices, with $|\mathcal{B}| \leq b$. Let T be the number of time steps for which the random walk game will be played, and for $t \in [0, T]$ let v_t denote the vertex holding the token at time t . Then for each t , v_{t+1} is a neighbor of v_t , and it is a uniformly random neighbor if v_t is a good vertex (i.e. $v_t \in \mathcal{G}$).

For each $v \in V$, let S_v denote the set of times when the token is at v , and let $S_{\mathcal{G}}$ denote the times when the token is with a good vertex in $\mathcal{G}$ and $S_{\mathcal{B}}$ denote the times when the token is with a bad vertex in $\mathcal{B}$. Also, let T_v , $T_{\mathcal{G}}$ and $T_{\mathcal{B}}$ denote the sizes of the corresponding sets of times. That is

$$\begin{aligned}
S_v &= \{t \mid v_t = v\}, & T_v &= |S_v|, \\
S_{\mathcal{G}} &= \{t \mid v_t \text{ is a good vertex}\}, & T_{\mathcal{G}} &= |S_{\mathcal{G}}|, \\
S_{\mathcal{B}} &= \{t \mid v_t \text{ is a bad vertex}\}, & \text{and } T_{\mathcal{B}} &= |S_{\mathcal{B}}|.
\end{aligned}$$

For a target vertex $y \in V$, we want to track the evolution of the values $H(v_t, y)$. Let

$$\Delta_t = H(v_{t+1}, y) - H(v_t, y)$$

be the change in expected hitting time at step t . Observe that $\mathbb{E}(\Delta_t \mid v_t \neq y, v_t \in \mathcal{G}) = -1$ and $\mathbb{E}(\Delta_t \mid v_t = y, v_t \in \mathcal{G}) = 1/\pi(y) - 1$. This motivates the definition of the sequence D_t^y :

$$D_t^y = \Delta_t + 1 - \frac{\mathbb{1}(v_t = y)}{\pi(y)} = h_y(v_t, v_{t+1})$$

It follows that if $v \in \mathcal{G}$ is any fixed good vertex and $y \in V$ any target, that $\mathbb{E}(D_t^y \mid v_t = v) = 0$ and moreover,

- The subsequence $(D_t^y : v_t = v)$ is a martingale difference sequence with step sizes bounded by $\rho_y(v)$,
- The subsequence $(D_t^y : v_t \in \mathcal{G})$ is a martingale difference sequence with step sizes bounded by ρ_y .

The above sequences are martingale difference sequences because, at timesteps when the token is controlled by good players, the next player is chosen fairly, and cannot be predicted in advance by the Adversary. The specific martingale difference sequence depends on the Adversary's strategy for the bad players' moves.

3.2 Referee Strategy

The referee's strategy will be based on Theorem 5 ([10, Lemma 2]), which is a version of Freedman's inequality for martingales.

Since $(D_t^y : v_t = v)$ is a martingale difference sequence with step sizes bounded by $\rho_y(v)$ whenever v is a good vertex, applying Freedman's inequality with $\delta = 1/n^C$, we know that for each good vertex v and target y ,

$$\Pr \left(\left| \sum_{t \in S_v} D_t^y \right| \geq \max \left\{ 4\sqrt{C\sigma_y^2(v)T_v \ln n}, 2C\rho_y(v) \ln n \right\} \right) \leq \frac{\log T_v}{n^C}. \quad (4)$$

With this in mind, we will accuse vertex v if $|\sum_{t \in S_v} D_t^y|$ is suspiciously large. Specifically, we will accuse v if

$$\exists y \in V. \quad \left| \sum_{t \in S_v} D_t^y \right| \geq \max \left\{ 4\sqrt{C\sigma_y^2(v)T_v \ln n}, 2C\rho_y(v) \ln n \right\}$$

By a union bound over all v, y , the probability any good vertex is mistakenly accused is at most $n^{-C+2} \log T$.

3.3 Analysis

Suppose the token passing game is played for T time steps and no player is accused by the referee of Section 3.2. Let v^* be the "stationary mode," *i.e.*, the vertex that is visited most frequently relative to its stationary probability. In particular, for all v ,

$$\frac{T_v}{\pi(v)} \leq \frac{T_{v^*}}{\pi(v^*)}.$$

36:12 Fraud Detection for Random Walks

We will denote by α the ratio between the number of times v^* is visited and the number of times it expects to be visited at stationarity. That is

$$\alpha = \frac{T_{v^*}}{T\pi(v^*)}.$$

Since v^* has been chosen to maximize the right hand side, and *some* vertex must be visited at least as often as expected, it follows that $\alpha \geq 1$. Also, we have for all v ,

$$T_v \leq \alpha T\pi(v) \quad (5)$$

Note that both v^* and α depend on the actual run of the game, so that they depend on T , the good players' randomness and the adversarial strategy. Nevertheless, we can show that when T is sufficiently large, the adversary has only a limited ability to skew who gets the token. Recall that b is the number of bad players.

► **Lemma 8.** *If $T \geq \max\{6H_{max}, 144C\mathcal{V}_\pi(1+b)\ln n, 12C\rho(1+b)\ln n\}$ then $\alpha \leq 2$. That is, for every vertex y ,*

$$T_y \leq 2T\pi(y).$$

Proof. Since the bad vertices want to avoid getting accused, based on the referee's strategy, we may assume that:

$$\forall v, y \in V. \left| \sum_{t \in S_v} D_t^y \right| \leq \max \left\{ 4\sqrt{C\sigma_y^2(v)T_v \ln n}, 2C\rho_y(v) \ln n \right\}. \quad (6)$$

Consider the sum $\sum_{t=0}^{T-1} D_t^y$. As $\sum_{t=0}^{T-1} \Delta_t$ telescopes to $H(v_T, y) - H(v_0, y)$ we have

$$\sum_{t=0}^{T-1} D_t^y = \sum_{t=0}^{T-1} \left(\Delta_t + 1 - \frac{\mathbb{1}(v_t = y)}{\pi(y)} \right) = H(v_T, y) - H(v_0, y) + T - \frac{T_y}{\pi(y)}$$

so that

$$T - \frac{T_y}{\pi(y)} \leq H(v_0, y) - H(v_T, y) + \sum_{t=0}^{T-1} D_t^y. \quad (7)$$

On the other hand, we can write

$$\left| \sum_{t=0}^{T-1} D_t^y \right| \leq \left| \sum_{t \in S_G} D_t^y \right| + \left| \sum_{v \in \mathcal{B}} \sum_{t \in S_v} D_t^y \right|.$$

Of the two sums on the right, we can deal with the first one by directly applying Freedman's inequality, since the subsequence $(D_t^y : v_t \in \mathcal{G})$ is actually a martingale difference sequence with step sizes bounded by $\rho_y = \max_v \rho_y(v)$. Thus, by Theorem 5, with error probability $(\log T)/n^C$, we have:

$$\left| \sum_{t \in S_G} D_t^y \right| \leq \max \left\{ 4\sqrt{C\mathcal{V}_\mathcal{G}^y \ln n}, 2C\rho_y \ln n \right\} \quad (8)$$

$$\leq 4\sqrt{C\mathcal{V}_\mathcal{G}^y \ln n} + 2C\rho_y \ln n, \quad (9)$$

where $\mathcal{V}_{\mathcal{G}}^y$ is the sum of the conditional variances of the steps of the martingale. It is bounded by

$$\begin{aligned}\mathcal{V}_{\mathcal{G}}^y &= \sum_{t \in S_{\mathcal{G}}} \mathbf{Var}_t D_t^y = \sum_{v \in \mathcal{G}} \sum_{t \in S_v} \mathbf{Var}(D_t^y | v_t = v) \\ &= \sum_{v \in \mathcal{G}} \sum_{t \in S_v} \sigma_y^2(v) \\ &= \sum_{v \in \mathcal{G}} \sigma_y^2(v) T_v \leq \alpha T \sum_{v \in \mathcal{G}} \sigma_y^2(v) \pi(v),\end{aligned}$$

where the last line follows from equation (5). Plugging this back into (9), we get

$$\left| \sum_{t \in S_{\mathcal{G}}} D_t^y \right| \leq 4 \sqrt{2\alpha T \ln n \sum_{v \in \mathcal{G}} \sigma_y^2(v) \pi(v)} + 4\rho_y \ln n. \quad (10)$$

To bound the corresponding term of the bad players we use apply Eqns. (5) and (6).

$$\begin{aligned}\left| \sum_{v \in \mathcal{B}} \sum_{t \in S_v} D_t^y \right| &\leq \sum_{v \in \mathcal{B}} \max \left\{ 4\sqrt{C\sigma_y^2(v)T_v \ln n}, 2C\rho_y(v) \ln n \right\} \\ &\leq \sum_{v \in \mathcal{B}} \left(4\sqrt{C\sigma_y^2(v)T_v \ln n} + 2C\rho_y(v) \ln n \right) \\ &\leq \sum_{v \in \mathcal{B}} \left(4\sqrt{C\sigma_y^2(v)T\alpha\pi(v) \ln n} + 2C\rho_y(v) \ln n \right) \\ &\leq 4\sqrt{C\alpha T \ln n} \left(\sum_{v \in \mathcal{B}} \sqrt{\sigma_y^2(v)\pi(v)} \right) + 2Cb\rho_y \ln n \\ &\leq 4\sqrt{C\alpha T \ln n} \sqrt{b \sum_{v \in \mathcal{B}} \sigma_y^2(v)\pi(v)} + 2Cb\rho_y \ln n\end{aligned} \quad (11)$$

where (11) follows from the Cauchy-Schwarz inequality. By Cauchy-Schwarz again,

$$\sqrt{\sum_{v \in \mathcal{G}} \sigma_y^2(v)\pi(v)} + \sqrt{b \sum_{v \in \mathcal{B}} \sigma_y^2(v)\pi(v)} \leq \sqrt{(1+b) \sum_v \sigma_y^2(v)\pi(v)} = \sqrt{\mathcal{V}_{\pi}^y(1+b)}. \quad (12)$$

Combining (10), (11), and (12), we obtain

$$\begin{aligned}\left| \sum_{t=0}^{T-1} D_t^y \right| &\leq 4 \sqrt{C\alpha T \ln n \sum_{v \in \mathcal{G}} \sigma_y^2(v)\pi(v)} + 2C\rho_y \ln n + 4 \sqrt{C\alpha T b \ln n \sum_{v \in \mathcal{B}} \sigma_y^2(v)\pi(v)} + 2Cb\rho_y \ln n \\ &\leq 4 \sqrt{C\alpha T \mathcal{V}_{\pi}^y(1+b) \ln n} + 2C\rho_y(1+b) \ln n.\end{aligned} \quad (13)$$

Now, plugging Eqn. (13) back into (7), and noting that $|H(v_0, y) - H(v_T, y)| \leq H_{\max}(y)$, we have that for every target y ,

$$\begin{aligned}\left| T - \frac{T_y}{\pi(y)} \right| &\leq |H(v_1, y) - H(v_{T+1}, y)| + \left| \sum_{t=0}^{T-1} D_t^y \right| \\ &\leq H_{\max}(y) + 4 \sqrt{C\alpha T \mathcal{V}_{\pi}^y(1+b) \ln n} + 2C\rho_y(1+b) \ln n.\end{aligned} \quad (14)$$

36:14 Fraud Detection for Random Walks

Recall that for the stationary mode vertex v^* , $T_{v^*} = \alpha T \pi(v^*)$, where $\alpha > 1$. Dividing by T and fixing $y = v^*$, we have

$$\begin{aligned} \alpha - 1 &\leq \frac{H_{\max}(v^*)}{T} + 4\sqrt{\frac{C\alpha\mathcal{V}_{\pi}^{v^*}(1+b)\ln n}{T}} + \frac{2C\rho_{v^*}(1+b)\ln n}{T} \\ &\leq \frac{H_{\max}}{T} + 4\sqrt{\frac{C\alpha\mathcal{V}_{\pi}(1+b)\ln n}{T}} + \frac{2C\rho(1+b)\ln n}{T} \\ &\leq \frac{1}{6} + \frac{\sqrt{\alpha}}{3} + \frac{1}{6} \end{aligned} \tag{15}$$

where (15) follows because $T \geq \max\{6H_{\max}, 144C\mathcal{V}_{\pi}(1+b)\ln n, 12C\rho(1+b)\ln n\}$. Thus α satisfies the quadratic inequality

$$\alpha - \frac{\sqrt{\alpha}}{3} - \frac{4}{3} \leq 0,$$

which implies $\sqrt{\alpha} \leq 4/3$ and hence $\alpha \leq 16/9 < 2$. Substituting this back into Eqn. (5) proves the lemma. $\blacktriangleleft$

The proof of Lemma 8 shows that for sufficiently large T we can drive α arbitrarily close to 1. Moreover the proof actually shows something even stronger. Let $\hat{\pi}$ denote the empirical distribution of how often each vertex is visited. By definition, for all y ,

$$\hat{\pi}(y) = \frac{T_y}{T}.$$

Using the fact that $\alpha < 2$ in Eqn. (14), we see that for all y ,

$$\left| T - \frac{T_y}{\pi(y)} \right| \leq H_{\max}(y) + 4\sqrt{2CT\mathcal{V}_{\pi}^y(1+b)\ln n} + 2C\rho_y(1+b)\ln n$$

Dividing by T ,

$$\left| 1 - \frac{\hat{\pi}(y)}{\pi(y)} \right| \leq \frac{H_{\max}(y)}{T} + 4\sqrt{\frac{2C\mathcal{V}_{\pi}^y(1+b)\ln n}{T}} + \frac{2C\rho_y(1+b)\ln n}{T}.$$

We restate this as a Corollary of Lemma 8.

► **Corollary 9.** *If $T \geq \max\{6H_{\max}, 144C\mathcal{V}_{\pi}(1+b)\ln n, 12C\rho(1+b)\ln n\}$ then for every vertex y ,*

$$\left| 1 - \frac{\hat{\pi}(y)}{\pi(y)} \right| \leq \frac{H_{\max}(y)}{T} + 4\sqrt{\frac{2C\mathcal{V}_{\pi}^y(1+b)\ln n}{T}} + \frac{2C\rho_y(1+b)\ln n}{T}.$$

Corollary 9 actually implies that it is impossible for the adversary to prolong the game for this many time steps without detection. If some vertex $x \in V$ has never passed the token, then $\hat{\pi}(x) = 0$ and $1 - \hat{\pi}(x)/\pi(x) = 1$. By Corollary 9, if no accusations yet have been leveled, then

$$\begin{aligned} 1 &\leq \frac{H_{\max}(x)}{T} + 4\sqrt{\frac{2C\mathcal{V}_{\pi}^x(1+b)\ln n}{T}} + \frac{2C\rho_x(1+b)\ln n}{T} \\ &\leq \frac{1}{6} + \frac{\sqrt{2}}{3} + \frac{1}{6} \\ &< 1. \end{aligned}$$

which is a contradiction. Thus we have shown that

► **Theorem 10.** *If $T = \Omega(H_{\max} + b(\mathcal{V}_\pi + \rho) \log n)$ then the Referee of Section 3.2 wins with probability at least $1 - 1/n^5$. In other words,*

$$T(G, b) = O(H_{\max} + b(\mathcal{V}_\pi + \rho) \log n)$$

We can simplify the expression of Theorem 10 as follows. Since for all y , $\rho_y \leq H_{\max}(y)$ it follows that $\rho \leq H_{\max}$. Furthermore, for any y the stationary conditional variance can be bounded by $\mathcal{V}_\pi^y \leq 2\mathbb{E}_\pi H(\cdot, y) \leq 2H_{\max}(y)$. (For completeness, these last inequalities are proved in Appendix A.) Thus, $\mathcal{V}_\pi + \rho \leq 3H_{\max}$. Also, for any graph, $H_{\max} = O(mn)$. This establishes the following Corollary.

► **Corollary 11.** *For any graph G and any number b of bad players,*

$$T(G, b) = O(bH_{\max} \log n) = O(bmn \log n).$$

The following corollary also follows directly from the above bounds and Corollary 9.

► **Corollary 12.** *For any graph G , any number b of bad players, and a walk that lasts for T steps with no accusations, if*

$$T = \Omega(mnb \log n)$$

then, for every vertex y ,

$$\left| 1 - \frac{\hat{\pi}(y)}{\pi(y)} \right| = O\left(\sqrt{\frac{mnb \log n}{T}}\right).$$

To relate these results back to the *price of corruption*, we note that the maximum expected cover time is clearly at least $H_{\max}$, and therefore $T(G, 0) \geq H_{\max}$. Moreover, repeatedly applying Markov's inequality shows that regardless of the starting vertex, after $6H_{\max} \log n$ steps, the probability that a particular vertex is unreachable, is at most $1/n^6$. Taking a union bound over all the vertices, after $6H_{\max} \log n$ steps, the probability that there is an unreachable vertex is at most $1/n^5$ and therefore $T(G, 0) = O(H_{\max} \log n)$, and $R(G, b) = O(b \log n)$.

This bound on $R(G, b)$ is close to tight, as witnessed by the class of *Ball & Chain* graphs. Let $BC_{n,b}$ consist of an $(n - b)$ -clique attached to a b -path; we assume $b \leq n/2$. Starting from a vertex in the “ball,” the cover time is $H_{\max} = \Theta(n^2b)$, thus, the zero-corruption game threshold is $T(BC_{n,b}, 0) = \Theta(n^2b \log n)$. We now need to lower bound $T(BC_{n,b}, b) = \Omega(n^2b^2)$. The corrupt vertices will lie only on the chain, and bias the walk slightly towards the ball, as in the proof of Theorem 6. Let u be the common vertex of the ball and chain. By Theorem 6, the walk restricted to the chain takes $\Omega(b^3)$ time steps, with high probability. Vertex u sees the token at least as often as in a truly random walk, which would be at least $\Omega(b^2)$ times. Each time u takes the token from the chain, it returns it to the chain after $\Theta(n^2)$ steps, in expectation, walking around the ball. Hence $T(BC_{n,b}, b) = \Omega(n^2b^2)$ and $R(BC_{n,b}, b) = \Omega(b/\log n)$.

Putting this all together, we have established Theorem 3.

4 The Clique

In this section, we analyze the Random Walk Game on the clique K_n . Here, every starting vertex is equivalent, and the hitting time to any vertex is a geometric random variable with mean $n - 1$. Thus $H_{\max} = n - 1$. Moreover, the cover time for the clique is essentially the

coupon collector problem, and therefore the maximum expected cover time is $O(n \log n)$, and the cover time is at most $\beta n \log n$ with probability $1 - 1/n^{\beta-1}$. The results of Section 3 tell us that $T(K_n, b) = O(bn \log n)$. When $b = \Omega(n)$, that is an upper bound of $O(n^2 \log n)$. In this section, we show that in fact, we can use the structure of the clique to get a much better bound. To get a sense of why fraud detection is faster for the clique, consider the game from the Adversary's perspective, and suppose the adversary wants to select a vertex that will not be reached. In a graph where there are low degree vertices, the adversary can surround such a vertex with corrupted vertices, all of whom always pass the token to one of their other neighbors. But in the clique, unless the Adversary takes over $n - 1$ vertices, every vertex has some good neighbors, who will pass it the token every n th time they get it, on average. This makes the Adversary's task much more difficult.

Theorem 13 gives near-tight bounds on the fraud detection time for cliques. The Referee's strategy differs from the strategy for the path or for a general graph, in that we take the entire trajectory of the walk into account when judging how a vertex v passes the token.

► **Theorem 13.** *Consider the Random Walk Game played on an n -clique, in which the adversary controls b vertices.*

1. *There is a Referee strategy that enables the Referee to win with high probability after $T(K_n, b) = O(\frac{n^2 \log n \log(n/(n-b))}{n-b})$ steps.*
2. *Moreover, there is an adversarial strategy for b corrupted players such that any accusation within $O(\frac{n^2 \log n}{n-b})$ steps cannot be correct with probability $1 - n^{-5}$, so that $T(G, b) = \Omega(\frac{n^2 \log n}{n-b})$.*

The remainder of this section constitutes a proof of Theorem 13.

Let C be a sufficiently large constant and G, B be the sets of good and bad players. If the G -players collectively pass the token $Cn \ln n$ times then the game will end naturally with high probability $1 - n^{-C+1}$, regardless of what other actions are taken by B .

Suppose the path taken by the token in T steps is $P = (v_1, v_2, \dots, v_T)$. When the token is at v_i , define L_i ("low" vertices) to be the set of vertices that have passed the token less than $2C|G|^{-1}n \ln n$ times. In the beginning $|L_1| = n$ and once $|L_i| \leq |G|/2$ at least $|G|/2$ vertices are not in L and the game has already ended, with high probability.

We partition time into stages, where stage $j \in [0, \log(2n/|G|)]$ covers the time that $|L_i| \in (n/2^{j+1}, n/2^j]$. Fix some stage j and let X_v be the number of times v passes the token in stage j and Y_v be the number of times v passes it to an L -vertex. Note that if v is good, Y_v is the sum of X_v indicator variables each with mean at least $2^{-(j+1)}$ and variance less than 2^{-j} . By Bernstein's inequality, $\Pr(Y_v < 2^{-(j+1)}X_v - t) < \exp(-\frac{t^2}{2 \cdot 2^{-j}X_v + (2/3)t})$. We will accuse v whenever $Y_v \leq 2^{-(j+1)}X_v - \sqrt{C2^{-(j+1)}X_v \ln n}$. Thus, with probability $n^{-\Omega(C)}$ no good vertex is accused. Suppose that stage j lasts for $T_j = 4C|G|^{-1}n^2 \ln n$ steps without any vertex being accused. Then:

$$\begin{aligned}
 \sum_{v \in V} Y_v &\geq \sum_{v \in V} \left(2^{-(j+1)}X_v - \sqrt{C2^{-(j+1)}X_v \ln n} \right) \\
 &\geq 2^{-(j+1)}T_j - \sqrt{2^{-(j+1)}T_j \cdot Cn \ln n} && \text{(Cauchy-Schwarz)} \\
 &\geq 2^{-(j+2)}T_j && \text{(Since: } Cn \ln n = T_j|G|/(4n) \leq T_j2^{-(j+1)}/4) \\
 &= (n/2^{j+1}) \cdot (2C|G|^{-1}n \ln n).
 \end{aligned}$$

However, if this were true then the number of L -vertices would have already shrunk to less than $n/2^{j+1}$, ending stage j . Thus, stage j cannot last for $4Cn^2|G|^{-1} \ln n$ steps without accusing a vertex of corruption. In total the number of steps before an accusation is $O(\frac{n^2 \log n \log(n/|G|)}{|G|}) = O(\frac{n^2 \log n \log(n/(n-b))}{n-b})$.

Turning to the lower bound, suppose we are aiming to make a correct accusation with probability $1 - n^{-C}$. Suppose the adversary picks a set $B \subseteq V$ uniformly at random with $|B| = b$. Under strategy $\mathcal{S}$ it corrupts B and under strategy $\mathcal{S}_{-j}$, $j \in B$, it corrupts $B - \{j\}$. In either case, whenever a corrupt vertex v has the token it passes it to a neighbor in B uniformly at random. The adversary chooses its strategy uniformly at random from $\{\mathcal{S}\} \cup \{\mathcal{S}_{-j}\}_{j \in B}$. Let $\mathcal{E}$ be the event that, after a walk of length $T = n^2 \ln n / (n - b)$, every vertex in B has only passed the token to others in B . Since, by Chernoff bounds, each vertex in B sees the token less than $3n \ln n / (n - b)$ times with probability $1 - o(1)$, we have

$$\Pr(\mathcal{E}) \geq (1 - o(1))(1 - (n - b)/n)^{3n \ln n / (n - b)} = \Omega(n^{-3}).$$

Moreover, $\Pr(\mathcal{S}_{-j} \mid (\mathcal{S}_{-j} \cup \mathcal{S}), \mathcal{E}) = q = 1/2$ since once we condition on $\mathcal{E}$, $\mathcal{S}_{-j}$ and $\mathcal{S}$ behave identically. By Lemma 7, the probability of error is at least $q^2/b = 1/(4b)$ after conditioning on $\mathcal{E}$, hence at least $\Omega(n^{-3}b^{-1})$ with no conditioning. For $C > 4$ this bound does not meet the desired n^{-C} error bound.

This concludes the proof of Theorem 13. It says that a coalition of $(1 - \epsilon)n$ bad vertices can delay the hitting time by $\Omega(\epsilon^{-1}n \ln n)$ and $O(\epsilon^{-1} \log \epsilon^{-1}n \ln n)$, i.e., no asymptotic delay at all when ϵ is constant. This is quite different than the line graph, in which a tiny minority of $\omega(\sqrt{n})$ bad vertices can asymptotically delay the hitting time.

5 Applications

5.1 Rotor Walks and Derandomization

Our results show that if all nodes pass the token in a way that is *locally* balanced across their neighbors, then the resulting *global* random walk has good cover time. The local balance condition can be ensured even without making any random choices. For example, in the *rotor walk* algorithm [32, 17, 20], every node passes the token to each of its neighbors in a round-robin fashion whenever it receives the token. A rotor walk ensures that every node satisfies the referee of Section 3.2.

Thus, Corollaries 11 and 12 directly apply to rotor walks when we set $b = n$. They give results analogous to Theorems 2 and 3 of [20]. In particular, Corollary 11 bounds the cover time of rotor walks, and Corollary 12 bounds the occupation frequencies. Our results are weaker than Theorems 2 and 3 of [20] in that they only apply to walks on unweighted, undirected graphs. But, they are stronger in that they apply to a broader class of derandomization techniques: for example, any routing works that ensures token passing is locally balanced across neighbors as specified by the referee of Section 3.2.

5.2 Leader Election

Leader election is a fundamental problem in distributed computing [12, 11, 27, 33, 25, 45]. Consider a simple communication model common to blockchains: there is a public key infrastructure (PKI) over the players, and communication occurs synchronously via a broadcast primitive that enables any player to send to all other players in the network (See [18, 15, 19]). Further, assume there is a publicly-known connected and regular graph G that has m edges and n nodes.

Corollary 12 enables us to perform repeated leader elections such that after $T = O(mn^2 \log n)$ elections, the fraction of good players elected approximates the fraction of good players, or else at least one bad player is caught.

The algorithm to achieve this is simple. First, the player with the token is the leader for that turn. This leader chooses one of its neighbors in G to pass the token to, and broadcasts a cryptographically signed message giving their choice. The PKI prevents equivocation, and synchronous communication forces some choice to be made, or else the current leader is known to be bad. Since all players learn all choices of the other players, each player can individually enforce the referee strategy of Section 3.2.

5.3 Sybil Defense

Consider a graph G with n nodes and $m = \Theta(n)$ edges, where (1) the bad and good nodes are separated by a cut with only α crossing edges; and (2) the subgraph induced by the good nodes is an expander. We want for almost all good nodes v , that node v learns a set of players S_v such that (1) S_v contains almost all of the good nodes; and (2) S_v contains “few” bad nodes. A simple algorithm is for each node to start a random walk at each of its edges, and for each of these walks to continue for $\Theta(\sqrt{n \ln n})$ steps. Then, for each node v , S_v is the set of all nodes w such that there is some node in the intersection of the nodes visited by walks starting at v and the nodes visited by walks starting at w .

This problem and algorithm is inspired by random-walk based Sybil defenses prevalent in the academic literature [43, 41, 4, 1], particularly the work of Yu, Kaminsky, Gibbons and Flaxman [43]. The graph represents a social network where the good nodes and the Sybil nodes are typically separated by a “small” number of crossing edges.

We can extend our referee and Corollary 12 to handle the algorithm described here that creates many random walks. Each edge has probability $1/m$ in the stationary distribution, and the initial steps in the algorithm above are also distributed uniformly over the edges. Thus, as the number of steps increases, each edge is visited $\Theta(\sqrt{n \ln n})$ times. This is true no matter what choices are made by the Sybil nodes, provided none of them are caught by the referee.

Thus, the total number of steps on the α crossing edges should be $\Theta(\alpha\sqrt{n \ln n})$. Call a random walk *bad* if it crosses one of the crossing edges and *good* otherwise. Then, there are at most $\Theta(\alpha\sqrt{n \ln n})$ bad walks. In particular, assuming $\alpha = o(\sqrt{n/\log(n)})$, the vast majority of the random walks starting on good nodes visit only good nodes.

Since the graph induced by the good nodes is an expander, with high probability, each pair of good random walks starting at two good nodes will intersect. Let $\mathcal{G}$ be the set of good nodes and $\mathcal{B}$ be the set of bad nodes. Then, by the above, there is a set $\mathcal{G}' \subseteq \mathcal{G}$ such that $|\mathcal{G}'| = \Omega(n - \alpha\sqrt{n \ln n})$ and for all $v \in \mathcal{G}'$, $\mathcal{G}' \subseteq S_v$ and $|S_v \cap \mathcal{B}| = O(\alpha\sqrt{n \ln n})$.

Thus, if $\alpha = o(\sqrt{n/\ln n})$, and we say that node u trusts node v if $v \in S_u$, we can say the following. There is a set, $\mathcal{G}'$ of all but a $o(1)$ fraction of the good nodes such that: all nodes in $\mathcal{G}'$ mutually trust each other; and every node in $\mathcal{G}'$ has a $o(1)$ fraction of Sybil nodes among the nodes it trusts.

6 Conclusion and Open Problems

It is well known that real-world fraud can sometimes be discovered by looking for statistical anomalies in data sets or transaction records. However, these statistical tests [29, 30, 26, 28, 35, 36, 37] work best on *unsophisticated* fraudsters, and may not work against adversaries who operate with full knowledge of the specific statistical tests.

In this paper we advocated for an approach to fraud detection that is *abstract*, *robust* against sophisticated adversaries, and *rigorous* in its quantitative guarantees.⁵ We illustrated how rigorous fraud detection against powerful adversaries can work in a simple abstract setting, namely *random walks on undirected graphs* in which *vertices* can be corrupted by the adversary; cf. [3, 22]. There are several directions for future research.

- One of our findings is that there can be a large delay between the time to detect the existence of fraud, w.h.p., and the time to make an accurate *accusation*, w.h.p. One could explore less strict notions of “accurate” accusation. In some contexts it may be fine to accuse a set $S \subseteq V$, such that 90% of S is corrupt, w.h.p.
- Given a specific graph G , we may be interested in the gap between its cover time and the fraud detection time against an adversary controlling b vertices. Up to log-factors we understand this gap on the path and clique, and know the extremal bound for arbitrary graphs, which is attained by the Ball and Chain graph. However, it is an open problem to efficiently *compute* this gap-factor for a specific G , or to bound it in terms of natural parameters of G , e.g., diameter.
- There are several algorithmic problems from the adversary’s perspective. Given a graph G and budget b , which b vertices should be corrupted to maximize the time of detection? To lower bound the detection time, we considered adversaries that corrupt vertices by simply changing the transition probabilities for their incident edges; such adversarial strategies are *Markovian*. Is there a specific graph for which all Markovian strategies are *suboptimal*? If so, it would be interesting to see what a superior non-Markovian strategy would look like.
- A natural direction is to consider random walks on directed, strongly connected graphs.

In general, the fraud detection paradigm can be introduced into the analysis of essentially any random process where it is conceivable that some or all of the randomness is being controlled by an adversary to achieve an unlikely outcome.

References

- 1 Muhammad Al-Qurishi, Mabrook Al-Rakhani, Atif Alamri, Majed Alrubaian, Sk Md Mizanur Rahman, and M Shamim Hossain. Sybil defense techniques in online social networks: a survey. *IEEE Access*, 5:1200–1219, 2017.
- 2 Romas Aleliunas, Richard M. Karp, Richard J. Lipton, László Lovász, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 218–223, 1979. doi:10.1109/SFCS.1979.34.
- 3 Noga Alon, Benjamin Gunby, Xiaoyu He, Eran Shmaya, and Eilon Solan. Identifying the deviator. *CoRR*, abs/2203.03744, 2022. doi:10.48550/arXiv.2203.03744.
- 4 Lorenzo Alvisi, Allen Clement, Alessandro Epasto, Silvio Lattanzi, and Alessandro Panconesi. Communities, random walks, and social sybil defense. *Internet Mathematics*, 10(3-4):360–420, 2014.
- 5 John Augustine, Gopal Pandurangan, and Peter Robinson. Fast byzantine agreement in dynamic networks. In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*, pages 74–83, 2013.

⁵ In the context of some purportedly random process, *fraud* is defined as effecting a particular outcome that is statistically unlikely by corrupting elements of the random process.

- 6 John Augustine, Gopal Pandurangan, and Peter Robinson. Fast byzantine leader election in dynamic networks. In *International Symposium on Distributed Computing*, pages 276–291. Springer, 2015.
- 7 John Augustine, Gopal Pandurangan, and Peter Robinson. Distributed algorithmic foundations of dynamic networks. *ACM SIGACT News*, 47(1):69–98, 2016.
- 8 Yossi Azar, Andrei Z. Broder, Anna R. Karlin, Nathan Linial, and Steven J. Phillips. Biased random walks. *Combinatorica*, 16(1):1–18, 1996. doi:10.1007/BF01300124.
- 9 Nikesh Bajaj, Tracy Goodluck Constance, Marvin Rajwadi, Julie A. Wall, Mansour Moniri, Cornelius Glackin, Nigel Cannings, Chris Woodruff, and James Laird. Fraud detection in telephone conversations for financial services using linguistic features. *CoRR*, abs/1912.04748, 2019. arXiv:1912.04748.
- 10 Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Proceedings of the 21st Annual Conference on Learning Theory-COLT 2008*, pages 335–342. Omnipress, 2008.
- 11 Michael Ben-Or, M Linial, and Michael Saks. *Collective coin flipping and other models of imperfect randomness*. IBM Thomas J. Watson Research Division, 1989.
- 12 Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 408–416. IEEE, 1985.
- 13 Romain Bertrand, Petra Gomez-Krämer, Oriol Ramos Terrades, Patrick Franco, and Jean-Marc Ogier. A system based on intrinsic features for fraudulent document detection. In *Proceedings 12th International Conference on Document Analysis and Recognition (ICDAR)*, pages 106–110, 2013. doi:10.1109/ICDAR.2013.29.
- 14 Nicolò Bonettini, Paolo Bestagini, Simone Milani, and Stefano Tubaro. On the use of Benford’s law to detect GAN-generated images. In *Proceedings 25th International Conference on Pattern Recognition (ICPR)*, pages 5495–5502, 2020. doi:10.1109/ICPR48806.2021.9412944.
- 15 Jing Chen and Silvio Micali. Algorand: A secure and efficient distributed ledger. *Theoretical Computer Science*, 777:155–183, 2019.
- 16 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. URL: <http://www.cambridge.org/gb/knowledge/isbn/item2327542/>.
- 17 Ioana Dumitriu, Prasad Tetali, and Peter Winkler. On playing golf with two balls. *SIAM Journal on Discrete Mathematics*, 16(4):604–615, 2003.
- 18 Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 281–310. Springer, 2015.
- 19 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th symposium on operating systems principles*, pages 51–68, 2017.
- 20 Alexander E Holroyd and James Propp. Rotor walks and markov chains. *Algorithmic probability and combinatorics*, 520:105–126, 2010.
- 21 Shang-En Huang, Seth Pettie, and Leqi Zhu. Byzantine agreement in polynomial time with near-optimal resilience. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*, pages 502–514, 2022. doi:10.1145/3519935.3520015.
- 22 Shang-En Huang, Seth Pettie, and Leqi Zhu. Byzantine agreement with optimal resilience via statistical fraud detection. *CoRR*, abs/2206.15335, 2022. doi:10.48550/arXiv.2206.15335.
- 23 A. Kerckhoffs. La cryptographie militaire. *Journal des Sciences Militaires*, pages 161–191, 1883.
- 24 Valerie King and Jared Saia. Byzantine agreement in expected polynomial time. *J. ACM*, 63(2):13:1–13:21, 2016. doi:10.1145/2837019.

- 25 Valerie King, Jared Saia, Vishal Sanwalani, and Erik Vee. Scalable leader election. In *SODA*, volume 6, pages 990–999, 2006.
- 26 Alex Ely Kossovsky. *Benford's Law: Theory, the General Law of Relative Quantities, and Forensic Fraud Detection Applications*. World Scientific, 2014. doi:10.1142/9089.
- 27 Nathan Linial. *Games computers play: Game-theoretic aspects of computing*. Citeseer, 1992.
- 28 Steven J. Miller. *Benford's Law: Theory and Applications*. Princeton University Press, Princeton, N.J., 2015.
- 29 Mark J. Nigrini. *Digital analysis using Benford's Law*. Global Audit Publications, 2000.
- 30 Mark J. Nigrini. *Benford's Law: Applications for Forensic Accounting, Auditing, and Fraud Detection*. Wiley, Hoboken, N.J., 2012.
- 31 Shashank Pandit, Duen Horng Chau, Samuel Wang, and Christos Faloutsos. Netprobe: a fast and scalable system for fraud detection in online auction networks. In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, pages 201–210, 2007. doi:10.1145/1242572.1242600.
- 32 Vyatcheslav B Priezzhev, Deepak Dhar, Abhishek Dhar, and Supriya Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Physical Review Letters*, 77(25):5079, 1996.
- 33 Alexander Russell and David Zuckerman. Perfect information leader election in $\log^* n + o(1)$ rounds. *Journal of Computer and System Sciences*, 63(4):612–626, 2001.
- 34 Atish Das Sarma, Anisur Rahaman Molla, and Gopal Pandurangan. Distributed computation in dynamic networks via random walks. *Theoretical Computer Science*, 581:45–66, 2015.
- 35 Uri Simonsohn. Just post it: The lesson from two cases of fabricated data detected by statistics alone. *Psychological Science*, 24(10):1875–1888, 2013.
- 36 Uri Simonsohn, Joseph P. Simmons, and Leif D. Nelson. Better p -curves: Making p -curve analysis more robust to errors, fraud, and ambitious p -hacking, a reply to Ulrich and Miller (2015). *Journal of Experimental Psychology*, 144(6):1146–1152, 2015.
- 37 Uri Simonsohn, Joseph P. Simmons, and Leif D. Nelson. Datacolada 98: Evidence of fraud in an influential field experiment about dishonesty, 2021. URL: <http://datacolada.org/98>.
- 38 Niek Tax, Kees Jan de Vries, Mathijs de Jong, Nikoleta Dosoula, Bram van den Akker, Jon Smith, Olivier Thuong, and Lucas Bernardi. Machine learning for fraud detection in e-commerce: A research agenda. *CoRR*, abs/2107.01979, 2021. arXiv:2107.01979.
- 39 Tian Tian, Jun Zhu, Fen Xia, Xin Zhuang, and Tong Zhang. Crowd fraud detection in internet advertising. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*, pages 1100–1110. ACM, 2015. doi:10.1145/2736277.2741136.
- 40 Chen Wang, Yingdong Dou, Min Chen, Jia Chen, Zhiwei Liu, and Philip S. Yu. Deep fraud detection on non-attributed graph. *CoRR*, abs/2110.01171, 2021. arXiv:2110.01171.
- 41 Wei Wei, Fengyuan Xu, Chiu C Tan, and Qun Li. Sybildefender: A defense mechanism for sybil attacks in large social networks. *IEEE transactions on parallel and distributed systems*, 24(12):2492–2502, 2013.
- 42 Chang Xu and Jie Zhang. Collusive opinion fraud detection in online reviews: A probabilistic modeling approach. *ACM Trans. Web*, 11(4):25:1–25:28, 2017. doi:10.1145/3098859.
- 43 Haifeng Yu, Michael Kaminsky, Phillip B Gibbons, and Abraham Flaxman. Sybilguard: defending against sybil attacks via social networks. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 267–278, 2006.
- 44 João G. Zago, Fabio L. Baldissera, Eric A. Antonelo, and Rodrigo T. Saad. Benford's law: what does it say on adversarial images? *CoRR*, abs/2102.04615, 2021. arXiv:2102.04615.
- 45 Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. Rapidchain: Scaling blockchain via full sharding. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 931–948, 2018.

A

 Stationary Conditional Variances

Fix any target y and let $H(v)$ be short for $H(v, y)$. The stationary conditional variance $\mathcal{V}_\pi^y$ is

$$\mathcal{V}_\pi^y = \sum_{v \in V} \pi(v) \left(\frac{1}{\deg(v)} \sum_{w \in N(v)} (H(w) - H(v))^2 - \left(\frac{1}{\deg(v)} \sum_{w \in N(v)} H(w) - H(v) \right)^2 \right).$$

This is a centered second moment, and is therefore always less than the corresponding uncentered second moment, which is better known as the Dirichlet form.

$$\mathcal{E}(H, H) = \sum_{v \in V} \sum_{w \in N(v)} \frac{\pi(v)}{\deg(v)} (H(v) - H(w))^2.$$

Since what we are about to say applies to arbitrary reversible Markov chains, we will switch notations accordingly. Let P be the transition matrix for any reversible Markov chain on state space V , with stationary distribution π . Reversible means that every pair of states v, w satisfies the *detailed balance* condition,

$$\pi(v)P(v, w) = \pi(w)P(w, v).$$

In this setting, the Dirichlet form $\mathcal{E}$ can be defined by either of the expressions below. Here, $f, g : V \rightarrow \mathbb{R}$.

$$\mathcal{E}(f, g) = \sum_{v, w \in V} \pi(v)P(v, w)(f(v) - f(w))^2 = 2 \cdot \sum_{v, w \in V} \pi(v)P(v, w)f(v)(f(v) - f(w)).$$

Specializing to the case where $f = g = H$, where recall that H is the hitting time to a fixed target state $y \in V$, we find that

$$\begin{aligned} \mathcal{E}(H, H) &= 2 \sum_{v, w \in V} \pi(v)P(v, w)H(v)(H(v) - H(w)) \\ &= 2 \sum_{v \in V} \pi(v)H(v) \left(1 - \frac{\mathbb{1}(v = y)}{\pi(y)} \right) \\ &= 2 \sum_{v \in V} \pi(v)H(v) && \text{since } H(y) = 0 \\ &= 2 \mathbb{E}_\pi H. \end{aligned}$$

Hence, for any y , $\mathcal{V}_\pi^y \leq 2 \mathbb{E}_\pi H$, and so $\mathcal{V}_\pi \leq 2H_{\max}$.