The Upper Bound on Time Complexity of Collisions Among Random Walks on a Graph

Zhiyu Liu *

March 13, 2011

Abstract

On a connected, undirected graph G, a token at some vertex v of G is said to be taking a *random walk* on G if, whenever it takes a move, it moves with equal probability to any of the neighbors of v. Suppose that there are several tokens on G, and at each tick of the clock only one token makes a move. When two tokens meet at some vertex, they collide and merge into one token.

Consider the following problem: Suppose that there are k tokens on a graph, and a demon decides which of them is to make a move at each time. The demon is trying to keep the tokens apart as long as possible. What is the expected time M before k tokens collide and merge into one?

In this paper, we show that for k = 2, in the worst case the expected time $M = (\frac{4}{27} + o(1))n^3$.

1 Introduction

An application of token random walks in distributeds computing is the self-stabilizing token management schemes, first proposed by Israeli and Jalfon [2]. A protocol for a distributed computing network is said to be self-stabilizing if, no matter what initial configuration it starts from, it eventually enters a "legal" configuration by itself and resumes normal operation.

In a token management scheme, only one processor at a time has the permission to execute its operation on shared memory. This processor is said to "possess the token". The token is an abstract object that is passed from processor to processor.

A self-stabilizing token management scheme needs to recover from two types of illegal configurations: (1) no processor possesses the token and (2) several processors possess tokens.

^{*}Dartmouth College, Hanover NH.

Israeli and Jalfon give a simple protocol which can avoid configuration (1). In Israeli and Jalfon's protocol, each processor has a shared register with an integer value in it. A processor knows that it possesses the token if the value in its register is at least as great as the value in any of its neighbors' registers. Thus, at any time, there always exists at least one processor that has the greatest value and hence possesses the token.

If several processors have the "local" maximum values in their registers, each of them will possess a token. To reduce the number of tokens to one, this token management scheme also makes sure that:(a) at each time, only one of the processors that possess tokens is allowed to execute its operation on shared memory, (b) after this processor finishes its operation, it passes the token to any of its neighboring processors with equal probability, and (c) if several tokens are held by one processor, these tokens merge to be only one token.

However, at each time, any of these "token holders" may be allowed to execute its operation. Thus, we can imagine a demon who, in the worst case, is trying to keep the tokens apart as long as possible. We must assume that at each time, the demon can decide which of those token-possessing processors to make the next move. Thus, the question becomes the following: What is the expected time M before several tokens collide into one?

Coppersmith, Tetali and Winkler [1] prove that if there are only two tokens, the expected time M has a polynomial upper bound in general case, namely,

 $M \leq \left(\frac{4}{27}\right)n^3$ plus lower-order terms

2 Preliminaries

The hitting time H(u, v) from u to v is defined as the expected number of moves for a random walk on a graph G beginning at vertex u to reach vertex v for the first time.

Similarly, we define the meeting time M(u, v) as the expected number of moves before two tokens placed initially at vertices u and v collide, given optimal play by the demon in deciding at each time which token moves.

For convenience, we define $H(\overline{u}, v)$ and $M(\overline{u}, v)$ to be the average of all H(w, v) and M(u, v), respectively, where w is a neighbor of u. Two simple but important observations are

$$H(w, v) = 1 + H(\overline{u}, v),$$

$$M(w,v) = 1 + \max{\{M(\overline{u},v), M(u,\overline{v})\}},\$$

for all distinct vertices u and v. That is to say, both the hitting time and the meeting time decline by 1 after each step.

3 Result

Now we are going to prove the upper bound on the maximum meeting time.

Lemma 3.1. For any vertices x, y, and z of a connected, undirected graph G, we have

$$H(x, y) + H(y, z) + H(z, x) = H(x, z) + H(z, y) + H(y, z).$$

Proof. Let $x, v_1, v_2, ..., v_k, x$ be the outcome of a random walk. Its probability is

$$\frac{1}{d(x)}\prod_{i=1}^k \frac{1}{d(v_i)},$$

where $d(v_i)$ is the degree of the vertex v_i . Note that this value is the same as the probability of the reverse walk $x, v_k, v_k - 1, ..., v_1, x$. The left-hand side of the equation in the lemma is the expected time for a random walk that goes from x to y, then to z and back to x, and similarly for the right. Symmetric argument can show that the number of x-to-y-to-z-to-x tours is the same as the number of x-to-z-to-y-to-x tours. It follows that the expected time of the two types of tours from x to x are the same, proving the lemma.

Lemma 3.2. On any graph G, the vertex-relation given by

$$u \leq v$$
 if and only if $H(u, v) \leq H(v, u)$

is transitive, i.e., constitutes a preorder on all vertices of G.

Proof. The proof is immediately from the equation in Lemma 3.1. \Box

Lemma 3.2 implies that in any graph there is always a vertex t such that $H(v,t) \ge H(t,v)$ for any vertex v. We call t the *hidden* vertex.

Theorem 3.3. Let t be a hidden vertex of a connected, undirected graph G. For any pair vertices of u and v in G, we have

$$M(u,v) \le H(u,v) + H(v,t) - H(t,v).$$

Proof. First, we define a *potential function* f in accordance with the right-hand side of the inequality in the theorem, namely,

$$f(u, v) = H(u, v) + H(v, t) - H(t, v) = H(v, u) + H(u, t) - H(t, u)$$

by Lemma 3.1. Thus, f(u, v) if symmetric, and

$$f(u, v) = 1 + f(\bar{u}, v) = 1 + f(u, \bar{v}).$$

That is to say, no matter which token makes the next move, f will decline by 1. Since t is a hidden vertex, we know that f is nonnegative.

Assume that the theorem is false and let s be the maximum value of M(u, v) - f(u, v). Among all pairs of u and v realizing s, choose one of minimum distance. Note that the distance cannot be 0, since $f(u, u) \ge 0 = M(u, u)$. Without loss of generality, assume the demon chooses the token at u to make a move. Then we have

$$M(u, v) = f(u, v) + s = 1 + f(\bar{u}, v) + s$$

 $\ge 1 + M(\bar{u}, v) = M(u, v),$

Note that at least one neighbor x of u is closer to v than u is, and u and v are the closest pair realizing s. Thus, we know f(x, v) + s > M(v), and hence $f(\bar{x}, v) + s > M(\bar{x}, v)$. Therefore, the inequality above is strict, namely, M(u, v) > M(u, v), a contradiction. Hence, the theorem is true.

Let us define *commute time* C(x, y) between vertices x and y as follows:

$$C(x,y) = H(x,y) + H(y,x).$$

With Chandra et al.'s [3] conclusion, Coppersmith, Tetali and Winkler [1] prove the following lemma:

Lemma 3.4. In any graph G of n vertices, where $n \ge 13$, for any three distinct vertices x, y, and z, we have

$$C(x,y) + C(y,z) + C(z,x) \le \frac{8}{27}n^3 + \frac{8}{3}n^2 + \frac{4}{9}n - \frac{592}{27}.$$

Theorem 3.5. In any connected, undirected graph G of n vertices, where $n \ge 13$, the maximum meeting time M is bounded by

$$\frac{4}{27}n^3 + \frac{4}{3}n^2 + \frac{2}{9}n - \frac{296}{27}.$$

Proof. Let t be a hidden vertex. By Theorem 3.3, we have

$$M(u,v) \le H(u,v) + H(v,t) - H(t,v) = H(v,u) + H(u,t) - H(t,u).$$

It is obvious that

$$M(u, v) \le H(u, v) + H(v, t) + H(t, v),$$

 $M(u, v) \le H(v, u) + H(u, t) + H(t, u).$

Thus,

$$\begin{split} 2M(u,v) &\leq H(u,v) + H(v,t) + H(t,v) + H(v,u) + H(u,t) + H(t,u) \\ &= C(u,v) + C(v,t) + C(t,v). \end{split}$$

If t is distinct from u and v, Lemma 3.4 gives the theorem. if t = u, then we have

$$M(u, v) \le H(u, v) + H(v, u) - H(u, v) = H(v, u) < C(u, v).$$

Now select a vertex t' different from u and v. We know that

$$C(v, t') + C(t', u) \ge C(u, v).$$

Thus,

$$2M(u, v) \le C(u, v) + C(v, t') + C(t', v).$$

Hence Lemma 3.4 suffices.

References

- D. Coppersmith, P. Tetali, P. Winkler(1993), Collisions among random walks on a graph, SIAM J Discrete Math Vol., No. 3, pp. 363–374.
- [2] A. Israeli and M. Jalfon(1990), Token management schemes and random walks yield self stabilizing mutual exclusion, Proc. 9th Annual ACM Symposium on Principles of Distributed Computing, Quebec City, Canada, pp. 119–131.
- [3] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P.Tiwari(1989), The electrical resistance of a graph captures its commute and cover times, Proc. 21st Annual ACM Symposium on Theory of Computing, Seattle, WA, pp. 574–586.