Available online at www.sciencedirect.com
SCIENCE DIRECT®
Theoretical Computer Science 332 (2005) 63–81
Theoretical Computer Science
www.elsevier.com/locate/tcs

New bounds for randomized busing☆

Steven S. Seiden[a], Peter P. Chen[a,*], R.F. Lax[b], J. Chen[a], Guoli Ding[b]

[a] *Department of Computer Science, 298 Coates Hall, LSU, Baton Rouge, LA 70803, USA*
[b] *Department of Mathematics, LSU, Baton Rouge, LA 70803, USA*

Received 29 August 2004; received in revised form 30 August 2004; accepted 23 September 2004
Communicated by A. Fiat

Abstract

We consider anonymous secure communication, where parties not only wish to conceal their communications from outside observers, but also wish to conceal the very fact that they are communicating. We consider the bus framework introduced by Beimel and Dolev (J. Cryptology 16 (2003) 25), where messages are delivered by a bus traveling on a random walk. We generalize this idea to consider more than one bus. We show that if w buses are allowed, then the expected delivery time for a message can be decreased from $\Theta(n)$ to $\Theta(n/\sqrt{w})$ in the case of a complete graph. Additionally, we introduce a class of graphs called r-partite directed collars and obtain analogous bounds on the expected delivery time for these graphs. We also propose several new features that resolve possible shortcomings in the systems proposed by Beimel and Dolev.

Keywords: Anonymous communication; Randomized busing; Random walk; Hitting time; Complete graph; r-partite directed collar

1. Introduction

Suppose we have a communication network, modeled by a graph G, composed of n vertices and m edges (or arcs, in the case of a directed graph). Messages are passed through this network, so that the various nodes can communicate with each other. A well-studied

☆ Research supported by AFOSR Grant No. F49620-01-1-0264 and NSF Grant No. 0326387.

* Corresponding author.
E-mail address: chen@bit.csc.lsu.edu (P.P. Chen).

0304-3975/$ - see front matter © 2004 Published by Elsevier B.V.
doi:10.1016/j.tcs.2004.09.039

problem is that of how to encrypt messages, so that even if an outside observer is able to intercept messages, the information being passed remains secret. A different and less well-studied problem is the following: Suppose we wish to conceal not only the contents of a message, but its point of origin and destination. We might imagine that the communications network is a military network for country *A*, over which critical orders are transmitted. We might wish to conceal which node is the command center, so that an enemy, say country *B*, does not know where to attack. Further, we may wish to conceal the fact that orders of some kind are being transmitted, as this may alert country *B* to a coming attack from *A*. This is known as the *anonymous communication* problem.

Previous results: The anonymous communication problem was first explored by Chaum, who proposed and analyzed a basic approach called a *mix* [5]. Mixes are further explored in [15–17]. Another approach to anonymous communication is to use generic secure multi-party function evaluation [3,4,7,6,12]. However, such schemes can be very inefficient [2]. To solve some of the problems with these methods, two further schemes have been proposed. The first is the *xor-tree* scheme developed by Dolev and Ostrovsky [10]. The second is the *bus* scheme introduced by Beimel and Dolev [2]. In this paper, we focus on the bus scheme.

Beimel and Dolev actually propose several different busing schemes. These schemes can be classified as either *deterministic* or *randomized*. Their main focus is on deterministic schemes, whereas our main focus shall be on randomized schemes. A drawback of the deterministic schemes of Beimel and Dolev is as follows: In all of the deterministic protocols proposed by these authors, the route a message takes through the network is fixed. If an enemy cuts a particular edge, or corrupts messages at a particular node, this could lead to the situation where the communication path between two nodes is unusable. The protocols have no possibility of exploring alternative paths. Essentially, in these protocols, it is possible to discern the general communication pattern, and thus disrupt it, even though it is not possible to know exactly who is communicating with whom. This criticism is also true of xor-trees [10]. As we shall see in the next section, there are several other shortcomings with the bus schemes proposed in [2].

Our results: The aforementioned problems with deterministic busing lead us to explore further the randomized busing protocol proposed in [2]. In this protocol, messages are delivered by a single bus traveling on a random walk in G. If, for instance, G is complete then the expected delivery time is $\Theta(n)$. We show that if G is complete and there are $w \leqslant n$ buses, then the expected delivery time for a message can be reduced to $O(n/\sqrt{w})$. We further show that this result is tight, that the expected delivery time is lower bounded by $\Omega(n/\sqrt{w})$. This is somewhat surprising, as one might hope for linear speed-up; i.e., a bound of $\Theta(n/w)$. We then define a new class of graphs called r-partite directed collars and we obtain analogous bounds on the delivery time for this class of graphs. We also propose several new features that overcome problems in the original bus system. We show that for an appropriate choice of parameters these new features do not impact the expected delivery time in the case of a complete graph.

2. Background

Before we present our results, we briefly describe the family of protocols presented in [2], which our method builds upon. To get complete details, the reader should see the

original paper. The basic idea explored in [2] is explained using the metaphor of a public transportation system. We think of the nodes of the communication system as being "bus stops" and of there being one or more "buses" that travel from stop to stop. Each bus has "seats" $s_{i,j}$, $1 \leqslant i \leqslant n$, $1 \leqslant j \leqslant n$, each of which can hold a message.

When the bus arrives at node k, seats $s_{k,j}$, $1 \leqslant j \leqslant n$, are all modified. If node k wishes to send a message to node j, then the message is encoded and placed in $s_{k,j}$. Otherwise $s_{k,j}$ is filled with random bits. We assume a public key cryptosystem is used, so that node k uses the public key of node j to encrypt the message that is placed in seat $s_{k,j}$; but, other cryptosystems, such as a symmetric key system, are possible. A basic assumption is that it is computationally intractable to tell encrypted messages from random bits. Further, node k checks each seat $s_{i,k}$, $1 \leqslant i \leqslant n$, for incoming messages. Each message $s_{i,k}$ is decrypted by node k using its private key. If the result is garbage, then it is ignored. Otherwise, node k receives the message.

Note that if β is the security parameter used to encrypt each message, then the "bus" is a message of size at least $O(\beta n)$. This is because there is a dedicated "seat" on the bus for each of the n nodes, and each seat is occupied by an encrypted message that is of size at least β.

Different schemes are distinguished by the number of buses and the patterns in which they travel. The simplest scheme is to have a single bus that follows a Hamiltonian cycle of G. A more communication intensive scheme involves having $2m$ buses traveling at each time step with a bus traversing each edge in each direction. Messages are relayed from bus to bus until they reach their destination. In order for this to work, each node must maintain a routing table that indicates where a message should go next in order to reach a particular destination (in [2] the routes are always shortest paths). An intermediate protocol involves using the preceding method on some subset of the edges in G (in fact the first scheme mentioned is just the case where the subgraph is a Hamiltonian cycle).

A basic problem with the schemes we have just described is that the path that a message follows through the network is fixed. If an enemy is able to disrupt messages along the path between two nodes (say by cutting an edge completely or replacing selected seats on a bus with random bits), then it can effectively cut communication between them. Another problem is that the schemes described so far require some sort of global control; i.e., nodes must either know how to route messages to their destination, which requires global knowledge of the network, or in the case of a Hamiltonian cycle this cycle must somehow be established, which again requires global knowledge.

To overcome the first problem, Beimel and Dolev proposed routing a bus randomly. The route the bus follows is a random walk on G. Specifically, at each time step, if the bus is at node u, then we pick a neighbor v of u uniformly and randomly, and send the bus along the edge (u, v). This overcomes the problem of edge failure, since a message will simply not travel through disabled edges. As long as G remains connected, a message will eventually reach its destination (with probability one). Randomized busing also eliminates the need for global routing tables to be stored in each node. However, it introduces a number of new problems:

1. The position of the bus is a random variable. When a node wants to send a message, it has to wait for the bus to arrive first. There is no absolute guarantee on how long this will take.

2. The time a message takes to travel from its source to its destination is also a random variable. Although it is possible to show that this travel time is reasonable with high probability, there is no guarantee that a message will ever reach its destination.
3. Some sort of global control is still required to initialize the system; i.e., the nodes have to agree where and when the bus will start traveling.
4. If the bus "crashes", meaning it reaches a node and the bus or node becomes disabled before the bus departs, either through accident or malicious behavior, then there is no way for the system to rectify or even recognize this situation.

In this paper, we present a number of modifications to the random walk busing scheme that seek to rectify these problems.

First, however, we make a comment about problem 2. Even in the case that buses travel on deterministic paths, and there is no chance of buses being crashed or corrupted, there is some very small probability of mis-communication. This is because we use random bits to fill the unused seats of the bus. There is a small probability that these random bits will decrypt to some message that seems plausible to the receiver. This problem could be overcome by having the sender check the random bits that fill each unused seat to see if they, in fact, decrypt to give a valid message, but one may not want to add this extra overhead to the scheme.

3. Our schemes

We assume we are dealing with a "listening adversary," who can monitor all communication links (either statically or dynamically). As in [2], we assume this adversary is honest-but-curious, meaning it cannot change, delete, or add any messages, or change the state of any node. (Beimel and Dolev [2] also consider the case of a Byzantine adversary in the context of a fixed routing scheme.) Also, as in [2], we assume semantic security; i.e., messages are encrypted, say by a public key cryptosystem, so that an eavesdropper cannot effectively distinguish between encryptions of any pair of messages.

We consider two schemes that extend the basic idea of randomized busing introduced by Beimel and Dolev [2]:

Multiple persistent buses. In this scheme there are a constant number w of buses on random walks in the system.

Multiple perishable buses. In this scheme, there are multiple buses on random walks. These buses are periodically created and destroyed as part of the protocol. This creation and destruction occurs in such a way that the anonymity of the scheme is preserved.

This second scheme requires no global knowledge of the network or coordination, other than a global clock. The first scheme requires some coordination to start the buses initially, but is easier to analyze.

Under the multiple perishable buses protocol, the creation of new buses is facilitated as follows: Each node is equipped with a single counter. At each time step the counter is incremented modulo k, where k is a parameter of the system. This counter is initialized to a random value in $\{0, \ldots, k-1\}$. If at some time step the counter achieves a value of 0, then a bus is created, and started on a random walk.

The destruction of buses is accomplished by using a "time to live" counter that is part of the bus. When the bus is created, this counter is initialized to a value ℓ. At each time step,

this counter is decremented. When it reaches 0, the bus is destroyed.

When a node u wishes to send a message to another node v, it waits until a bus arrives or is created at u, and places the message on the bus. (If a bus arrives at u that already is carrying a message for v, the node u should check to see that v has acknowledged receipt of this message before u overwrites this message with a new message. This will not cause significant problems in the case of a complete graph since the probability that a bus that picked up a message at u visits v before returning to u is $n/2(n-1) > 1/2$ by [13, Proposition 2.3].)

Since the bus might never arrive at v (if it dies before it arrives), the message can be put on more than one bus. In particular, we consider placing the message on the first λ buses that reach u. Care must be taken to ensure that different copies of the message have different encrypted texts. This might be accomplished by padding each copy with a string of random bits. When v receives the message, it should send an acknowledgment to u via the same method.

The first thing to note about this protocol is that it preserves anonymity, as long as the messages are indistinguishable from random bits. Further, as advertised, no global coordination is required. In fact, the protocol does not necessarily fail even if different nodes use different values of the parameters ℓ and k. We can also accommodate nodes being "reset"; i.e., disabled either through accident or malevolence and then restarted at a later point in time. In this case, we just reinitialize the counter randomly.

We do require that each node know its own "identity" and that of any other nodes it wishes to send messages to, so that it can place messages in the correct seats of the bus. Further, given that a public key cryptosystem is used to encrypt messages, the sender must have the receiver's public key. We also must have some upper bound on the number of nodes in G to determine the number of seats on the bus.

However, unlike the situation in the random walk protocol of [2], we do have some guarantee on how long we will wait before a message can be sent, since a bus is created at each node every k time steps. Further, if a bus crashes, the system does not fail.

The parameters ℓ, λ and k need to be tuned to provide a system that

- Avoids congestion. If many buses arrive at a node simultaneously, this could cause problems for the system. In this situation, one possible solution is to drop some buses randomly.
- Assures a high probability of quick message delivery. Obviously, the longer a bus lives, the more likely it is to deliver messages successfully. Similarly, with more buses in the system, we can expect buses to arrive at a sender more frequently.

Two types of graphs are considered in our randomized busing schemes. We first focus on complete graphs in which there is a direct communication link between any pair of nodes. Then we introduce a new type of graph called r-partite directed collars, which are natural generalizations of ring graphs, and we consider randomized busing for r-partite directed collars. A directed graph G is called an r-partite directed collar if the vertices V in G are partitioned into nonempty sets $V_1, V_2, \ldots, V_r$ such that the arcs in G consist of arcs from every node in V_i to every node in V_{i+1} for $i = 1, 2, \ldots, r-1$ and from every node in V_r to every node in V_1.

Remark 1. *Justifications for randomized busing schemes for complete graphs.* One may

argue why we need the randomized busing protocols for communication in the case of complete graphs since every pair of nodes is directly connected. However, our goal of anonymity could not be achieved if two nodes simply communicated directly whenever one had a message to send to the other.

Anonymity could be achieved if one used a naive scheme in which every node sent every other node a (possibly nonsensical) message at each time period, but this would result in significant congestion with $O(n)$ buses both arriving at and departing from every node at every time step. We will show that in our scheme, even when the number of buses equals the number of nodes, it is highly probable that no more than $O(\ln n / \ln \ln n)$ buses are at any single node at any time in the case of a complete graph. Thus, although our scheme involves significant communication complexity if we employ $O(n)$ buses, there will be advantages if the message complexity is high since each node will be dealing with fewer messages than in the naive scheme.

As we shall see in , the multiple perishable buses scheme will work well on a complete graph if we take $k = \ell = n^{7/2}$ and $\lambda = \sqrt{n}$.

Another possible criticism to considering the complete graph case may be that complete graphs are too specialized to be useful in real life communication networks. While we recognize that complete graphs are a special subclass of all graphs, it is not inconceivable to find in the real world some communication networks with complete graph topology. Moreover, the study of busing schemes in complete graphs forms the basis for studying randomized busing in r-partite directed collars, which are closely related to the ring topology widely used in real communication networks.

Remark 2. *Justifications for randomized busing scheme for directed collar graphs.* The r-partite directed collar graph is a natural generalization of the ring graph. The ring topology is a very commonly used topology in real communication networks. Understanding anonymous messaging in directed collar networks has potential applications for real world communication network security.

Similar to the case of complete graphs, for r-partite directed collars, the naive method, which requires every node in a block V_j to send a message to every node in the next block V_{j+1} at each time step, although achieving anonymity, would result in congestion. Our busing scheme alleviates the congestion problem for arbitrary directed collars, and avoids the problem in the case of directed collars with each block having equal size.

4. Mathematical preliminaries

We assume the reader is familiar with basic probability theory. For an introduction to such material we refer the reader to the books Feller [11] and Motwani and Raghavan [14].

We briefly review some relevant material regarding standard random walks [1,13,14]. For $v, u \in V(G)$, the *hitting time* $h(u, v)$ is the expected number of steps for a random walk starting at u to reach v. We define the *maximum hitting time* to be $h^* = \max_{u,v \in V(G)} h(u, v)$. For $v \in V(G)$, the *cover time from* v, denoted $C(v)$, is the expected number of steps in a random walk starting at v that reaches every vertex. The cover time of G is $C^* = \min_{v \in V(G)} C(v)$. The hitting time seems to be the most relevant parameter in our situation.

However, the cover time is obviously an upper bound on the maximum hitting time, and, as we shall see, in the worst case, they have the same value asymptotically. A random walk on a graph is a special type of Markov chain. If the exact topology of G is known and fixed, then the hitting time can be calculated directly, using the theory of Markov chains.

If G is complete, then it is not hard to see that h^* is $\Theta(n)$, whereas C^* is $\Theta(n \log n)$. In the case that G is a lollipop graph (a clique connected to a path; see [1,14]), then $h^* = C^* = \Theta(n^3)$. In general, for any graph G we have $h^* \leqslant C^* \leqslant 2m(n-1) < n^3$. If we know more about G, better bounds are possible. For instance it is possible to show that

$$mR \leqslant C^* \leqslant 2e^3 mR \ln n + n,$$

where R is the resistance of G [13,14].

We will make use of the following inequalities (cf. [14, Proposition B.3]):

$$\left(1 + \frac{x}{i}\right)^i \leqslant e^x \quad \text{for } i > 0 \text{ and } x \geqslant -i, \tag{1}$$

$$x + 1 \leqslant e^x \quad \text{for all } x. \tag{2}$$

5. Multiple persistent buses: complete graphs

We begin by providing some analysis of the multiple persistent bus protocol where there are w buses on random walks on a complete graph G. For technical reasons, we shall consider only $w \leqslant n$, and we assume throughout that $n \geqslant 3$. Our main result shows that the expected delivery time of a message is $\Theta(n/\sqrt{w} + n/\lambda)$ for all $1 \leqslant \lambda \leqslant w \leqslant n$. Admittedly, this analysis does not take into account the sort of destructive attacks on the system that we mentioned earlier; however, we should first verify that the system works well under normal conditions. We also consider the question of congestion.

We begin by considering the situation where w buses initially located at u walk randomly through G. Define the *w-bus hitting time* $h_w(u, v)$ to be the expected number of steps for one of the w randomly walking buses to reach v. Then

Lemma 5.1. $(n-2)/w \leqslant h_w(u, v) \leqslant (n-2)/w + 1$.

Proof. Let $\sigma^i = \{\sigma^i_j\}$, for $1 \leqslant i \leqslant w$, be the sequence of vertices visited by the ith bus. So, for $j \geqslant 0$, the vertex visited by the ith bus after j steps is σ^i_j.

First consider the question "Given a fixed vertex $v \neq u$, what is the minimum index X such that $\sigma^i_X = v$ for some i?" Note that $\mathrm{E}[X] = h_w(u, v)$. We define a new sequence ς from these w sequences as follows:

$$\varsigma_j = \sigma^{(j \bmod w)+1}_{\lfloor j/w \rfloor} \qquad \text{for } j \geqslant 0.$$

Note that $\varsigma_j = u$ for $0 \leqslant j < w$ and $\varsigma_j \neq \varsigma_{j-w}$ for all $j \geqslant w$. Now consider the question "Given a fixed vertex $v \neq u$, what is the minimum index Y such that $\varsigma_Y = v$?" Note that $X = \lfloor Y/w \rfloor \leqslant Y/w$ and therefore $h_w(u, v) \leqslant \mathrm{E}[Y/w] = \mathrm{E}[Y]/w$, since w is constant.

Further note that $h_w(u,v) = \mathrm{E}[X] \geqslant (\mathrm{E}[Y]/w) - 1$. For all $0 \leqslant i < w$, we have $\varsigma_i \neq v$. For all $i \geqslant w$, given that $\varsigma_j \neq v$ for all $0 \leqslant j < i$, the probability that $\varsigma_i = v$ is just $1/(n-1)$. Using the fact that G is complete, it follows that $Y - w + 1$ is a geometric random variable with parameter $1/(n-1)$. Hence, we have $\mathrm{E}[Y - w + 1] = n - 1$. We conclude that $(n-2)/w \leqslant h_w(u,v) \leqslant (n-2)/w + 1$. □

Similar results can be shown for cover time, but we do not need them here.

Suppose the graph G is r-regular, but not complete. Consider the following experiment: Choose a vertex u' different from v at random. What is the probability p that a bus at this chosen vertex will arrive at v on the next step? There are r neighbors of v, so the probability that we chose one of these neighbors is $r/(n-1)$. Given that we chose a neighbor of v, the probability that the bus goes to v on the next step is $1/r$ (since u' also has r neighbors). Therefore, we have $p = (r/(n-1))(1/r) = 1/(n-1)$. However, the above proof of Lemma 5.1 does *not* remain valid for regular graphs. The problem is that the random variable $Y - w + 1$ would not be a geometric random variable if the graph is not complete. Indeed, if the bus went from u' to $v' \neq v$, then the probability that the bus would go from v' to v on the next step would depend on whether v' is or is not a neighbor of v, so we do not have a sequence of independent trials. (Of course, it is clear that Lemma 5.1 could not hold for regular graphs since, in the case that $w = 1$, the hitting time for an arbitrary regular graph is not $O(n)$. For example, one may consider a "cycle" of cliques with the same number of vertices in each one, where an edge is removed from each clique to allow the connection of adjacent cliques without violating regularity, see [9, p. 306]. Such a graph is sometimes called a necklace.) However, we will show in the next section that we can obtain similar results for a class of graphs we call directed collars.

We now consider how buses arrive at a node u. In particular, we consider the question, "Starting at a particular time t, how many steps do we have to wait before i distinct buses arrive at u?" Call this value A_i.

Proposition 5.1. $\mathrm{E}[A_i] \leqslant \dfrac{i(n-1)}{w-i+1}$.

Proof. Define B_j to be the number of steps between the arrival of the $(j-1)$st distinct bus and the arrival of the jth distinct bus.

Then $A_i = \sum_{j=1}^{i} B_j$ and so $\mathrm{E}[A_i] = \sum_{j=1}^{i} \mathrm{E}[B_j]$. Consider B_j. For any fixed bus that has not yet arrived at u, the probability that it does not arrive at u on a given time step is $(n-2)/(n-1)$. The probability that none of the $w - j + 1$ buses that have yet to arrive at u will arrive on the next step is

$$q_j = \left(\frac{n-2}{n-1}\right)^{w-j+1},$$

since the buses move independently. Note that B_j could equal 0 if the $(j-1)$st and jth bus arrive simultaneously, and we have $\Pr[B_j = 0] = 1 - q_j$. For B_j to equal 1, it must be the case that none of the $w - j + 1$ buses yet to arrive at u arrives at the same time as the $(j-1)$st bus and one of them arrives at u at the next time step. Hence, $\Pr[B_j = 1] = (1 - q_j)q_j$. In general, we have $\Pr[B_j = b] = (1 - q_j)q_j^b$. (So B_j is a geometric random variable that

can take the values $0, 1, \ldots$.)

From the definition of expected value, we have

$$\mathrm{E}[B_j] = \sum_{b=0}^{\infty} b(1-q_j)q_j^b = \frac{1}{1-q_j} - 1 = \frac{1}{1-((n-2)/(n-1))^{w-j+1}} - 1,$$

and therefore,

$$\begin{aligned}
\mathrm{E}[A_i] &= -i + \sum_{j=1}^{i} \frac{1}{1-((n-2)/(n-1))^{w-j+1}} \\
&= -i + \sum_{j=1}^{i} \frac{1}{1-(1+(-(w-j+1)/(n-1))/(w-j+1))^{w-j+1}} \\
&\leqslant -i + \sum_{j=1}^{i} \frac{1}{1-\mathrm{e}^{-(w-j+1)/(n-1)}}, \text{ using (1)} \\
&= -i + \sum_{j=1}^{i} \frac{\mathrm{e}^{(w-j+1)/(n-1)}}{\mathrm{e}^{(w-j+1)/(n-1)}-1} \\
&= -i + \sum_{j=1}^{i} \frac{\mathrm{e}^{(w-j+1)/(n-1)}-1+1}{\mathrm{e}^{(w-j+1)/(n-1)}-1} \\
&= -i + \sum_{j=1}^{i} 1 + \frac{1}{\mathrm{e}^{(w-j+1)/(n-1)}-1} \\
&= \sum_{j=1}^{i} \frac{1}{\mathrm{e}^{(w-j+1)/(n-1)}-1} \\
&\leqslant \frac{i}{\mathrm{e}^{(w-i+1)/(n-1)}-1}, \text{ since } j \leqslant i \\
&\leqslant \frac{i(n-1)}{w-i+1}, \text{ using (2).} \qquad \square
\end{aligned}$$

Now we consider how long a message takes to get from u to v. Suppose we put the message on the first $\lambda \leqslant w$ distinct buses that arrive at u. Without loss of generality, we can assume these buses are numbered $1, 2, \ldots, \lambda$. As before, let A_i be the number of steps before i distinct buses arrive at u. Let D_i be the number of steps taken by the ith distinct bus in traveling from u to v. By the "expected delivery time" we mean $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right]$. (So this includes the time u must wait to put messages on buses and the time it takes for some bus to reach v.) We have

$$\begin{aligned}
\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] &\leqslant \mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_\lambda + D_i\}\right] \\
&= \mathrm{E}[A_\lambda] + \mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda} D_i\right].
\end{aligned}$$

We claim that $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda} D_i\right] = h_\lambda(u, v)$. To see this, suppose that the ith bus receives the message at u at time α_i. As above, let $\sigma^i = \{\sigma^i_j\}$ be the sequence of vertices visited by

the ith bus. Then $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda} D_i\right]$ is the minimum index X such that $\sigma^i_{\alpha_i+X} = v$ for some $i \in \{1, 2, \ldots, \lambda\}$. It is then clear that $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda} D_i\right]$ is the same as $h_\lambda(u, v)$. Hence, using Lemma 5.1, we have

$$\begin{aligned}\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] &\leqslant \mathrm{E}[A_\lambda] + h_\lambda(u, v)\\ &\leqslant \frac{\lambda(n-1)}{w-\lambda+1} + \frac{n-2}{\lambda} + 1. \end{aligned} \tag{3}$$

If we choose $\lambda \leqslant \sqrt{w}$, then

$$\frac{\lambda(n-1)}{w-\lambda+1} \leqslant \frac{\lambda(n-1)}{\lambda^2-\lambda+1},$$

so the expected delivery time is $\mathrm{O}(n/\lambda)$. Further note that

$$\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] \geqslant \mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda} D_i\right] = h_\lambda(u, v) \geqslant \frac{n-2}{\lambda}.$$

Therefore, for $\lambda \leqslant \sqrt{w}$ the expected delivery time is $\Theta(n/\lambda)$. So, in this case, linear speed-up in λ is achieved.

For $\lambda > \sqrt{w}$, an upper bound of $\mathrm{O}(n/\sqrt{w})$ holds, since, using (3), we have

$$\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] \leq \mathrm{E}\left[\min_{1\leqslant i\leqslant\lfloor\sqrt{w}\rfloor}\{A_i + D_i\}\right] = \mathrm{O}\left(\frac{n}{\sqrt{w}}\right).$$

We now develop a lower bound for the expected delivery time when $\lambda > \sqrt{w}$. Put $\theta = \lfloor\sqrt{w}\rfloor + 1$. Note that

$$\begin{aligned}\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] &= \mathrm{E}\left[\min\left\{\min_{1\leqslant i<\theta}\{A_i + D_i\},\ \min_{\theta\leqslant i\leqslant\lambda}\{A_i + D_i\}\right\}\right]\\ &\geqslant \mathrm{E}\left[\min\left\{\min_{1\leqslant i<\theta} D_i,\ A_\theta\right\}\right].\end{aligned}$$

Now, by Markov's Inequality, we have

$$\mathrm{E}\left[\min\left\{\min_{1\leqslant i<\theta} D_i,\ A_\theta\right\}\right] \geqslant \Pr\left[\min\left\{\min_{1\leqslant i<\theta} D_i,\ A_\theta\right\} \geqslant a\right]\cdot a$$

for any positive number a. Also, since the random variables $\min_{1\leqslant i<\theta} D_i$ and A_θ are independent, we have

$$\begin{aligned}\Pr\left[\min\left\{\min_{1\leqslant i<\theta} D_i,\ A_\theta\right\} \geqslant a\right] &= \Pr\left[\left(\min_{1\leqslant i<\theta} D_i \geqslant a\right)\cap A_\theta \geqslant a\right]\\ &= \Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant a\right]\Pr\left[A_\theta \geqslant a\right].\end{aligned}$$

Hence, we have

$$\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] \geqslant \Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant \frac{n}{c\sqrt{w}}\right]\Pr\left[A_\theta \geqslant \frac{n}{c\sqrt{w}}\right]\frac{n}{c\sqrt{w}},$$

for any positive number c. Our goal now is to show that for $c = 2e^6$ each of the probabilities in this last inequality are bounded away from 0. If we can do this, then the expected delivery time is $\Omega(n/\sqrt{w})$ for $\lambda > \sqrt{w}$.

We first establish a positive lower bound for the probability that $A_\theta \geqslant n/(2e^6\sqrt{w})$. Put $\vartheta = \lceil n/w \rceil$. Define Z_i to be 1 if $B_i \geqslant \vartheta$ and 0 otherwise. The variables $Z_1, \ldots, Z_\theta$ are independent Poisson trials with

$$\Pr[Z_i = 1] = \Pr[B_i \geqslant \vartheta] = \sum_{b=\vartheta}^{\infty} (1-q_i)q_i^b = q_i^\vartheta.$$

Define $Z = \sum_{i=1}^{\theta} Z_i$ and $\mu = \mathrm{E}[Z]$. Then the Chernoff bound [14, Theorem 4.2], tells us that

$$\Pr\left[Z < \tfrac{1}{2}\mu\right] < e^{-\mu/8}.$$

Hence, we have

$$\Pr\left[Z \geqslant \tfrac{1}{2}\mu\right] \geqslant 1 - \frac{1}{e^{\mu/8}}.$$

For $n \geqslant 3$ and $w \geqslant 1$, we have

$$\begin{aligned}\mu &= \sum_{i=1}^{\theta} q_i^\vartheta = \sum_{i=1}^{\theta} \left(\frac{n-2}{n-1}\right)^{\vartheta\cdot(w-i+1)} \geqslant \theta\left(\frac{n-2}{n-1}\right)^{\vartheta w} \\ &= \frac{\theta}{(1+(\vartheta w/(n-2))/\vartheta w)^{\vartheta w}} \geqslant \frac{\theta}{e^{\vartheta w/(n-2)}},\end{aligned}$$

using (1). Now, $\vartheta w < n + w \leqslant 2n$, so $\vartheta w/(n-2) < 6$ (since $n \geqslant 3$). Hence

$$\mu > \frac{\theta}{e^6} > \frac{\sqrt{w}}{e^6}.$$

Since $\theta \geqslant 2$ for $w \geqslant 1$, we find that

$$\Pr\left[Z \geqslant \frac{\sqrt{w}}{2e^6}\right] \geqslant \Pr[Z \geqslant \mu/2] \geqslant 1 - \frac{1}{e^{\mu/8}} \geqslant 1 - \frac{1}{e^{\theta/8e^6}} \geqslant 1 - \frac{1}{e^{1/4e^6}} > 0.0006.$$

We note that

$$A_\theta = \sum_{j=1}^{\theta} B_j \geqslant \vartheta \sum_{j=1}^{\theta} Z_j = \vartheta Z,$$

and therefore $1 - 1/e^{1/4e^6}$ is also a lower bound on the probability that $A_\theta \geqslant \vartheta\sqrt{w}/(2e^6) \geqslant n/(2e^6\sqrt{w})$.

The other probability is much simpler to bound. For any nonnegative integer v, we have

$$\begin{aligned}\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant v+1\right] &= \left(\frac{n-2}{n-1}\right)^{(\theta-1)v} \\ &= \frac{1}{((n-1)/(n-2))^{(\theta-1)v}},\end{aligned}$$

since at each of v steps each of $\theta-1$ buses goes to one of the $n-2$ vertices other than v and the vertex the bus was at. Writing

$$\left(\frac{n-1}{n-2}\right)^{(\theta-1)v} = \left(1+\frac{(\theta-1)v/(n-2)}{(\theta-1)v}\right)^{(\theta-1)v}$$

and applying (1), we get

$$\begin{aligned}\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant v+1\right] &\geqslant \frac{1}{\mathrm{e}^{(\theta-1)v/(n-2)}}\\ &\geqslant \frac{1}{\mathrm{e}^{v\sqrt{w}/(n-2)}}.\end{aligned}$$

For $v=\lfloor n/(2\mathrm{e}^6\sqrt{w})\rfloor$ and $n\geqslant 3$ (so $n/(n-2)\leqslant 3$), we have

$$\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant n/(2\mathrm{e}^6\sqrt{w})\right] \geqslant \Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant v+1\right] \geqslant \frac{1}{\mathrm{e}^{3/(2\mathrm{e}^6)}} > 0.99.$$

We have shown the following.

Theorem 5.1. *When G is complete, the expected delivery time of a message in the persistent multiple bus protocol is*

$$\Theta\left(\frac{n}{\sqrt{w}}+\frac{n}{\lambda}\right)$$

for all $1\leqslant\lambda\leqslant w\leqslant n$.

We now consider the question of congestion. Suppose the initial locations of the w buses are chosen independently and uniformly at random from $V(G)$. Then if we consider any fixed time step, bus positions are chosen independently and uniformly at random. The expected number of buses at a vertex is just w/n. What is the maximum number of buses at a vertex? This is modeled by the situation of throwing w balls in n bins independently and uniformly at random. Putting $k^*=\lceil(3w\ln n)/n\ln\ln n\rceil$, and reasoning as in [14, p. 44], we see that no bin has more than k^* balls in it with probability at least $1-1/n$, so the amount of congestion is reasonable.

6. Multiple persistent buses: directed collars

We now consider graphs that we will call directed collars. These graphs may be considered generalizations of a unidirectional ring, as considered in [8]. Also, results on complete bipartite graphs may be inferred from results on these directed collars, as we will note at the end of this section.

Definition 6.1. An r-partite directed collar is a directed graph G such that
(1) the vertex set V is the disjoint union of nonempty subsets $V_1, V_2, \ldots, V_r$.
(2) the arcs of G consist of arcs directed from each vertex in V_i to each vertex in V_{i+1} for $i=1,\ldots,r-1$ and from each vertex in V_r to each vertex in V_1.

For this section, G will denote an r-partite directed collar. Let $n_i = \#V_i$ for $i = 1, 2, \ldots, r$. Put $n = \sum_{i=1}^{r} n_i$ and put $n_* = \max_{1 \leqslant i \leqslant r} n_i$. We assume throughout this section that $w \leqslant n_*$. We also assume that r is a fixed number, and our goal is to bound the expected delivery time in terms of n_*.

Let u and v be two vertices of G. If $u \in V_i$ and $v \in V_j$, with $i \neq j$, then the distance from u to v (which is the length of the shortest $u - v$ path) is $j - i$ if $i < j$ and is $r - (i - j)$ if $i > j$. If u and v are distinct vertices in the same V_i, then the distance from u to v is k. We again consider w buses initially located at u that are randomly walking on G (only in the direction of the arcs, of course).

Lemma 6.1. *Let d denote the distance from u to v. Suppose $v \in V_m$. Then*

$$d + r\left(\frac{n_m - w - 1}{w}\right) \leqslant h_w(u, v) \leqslant d + r\left(\frac{n_m - 1}{w}\right).$$

Proof. As in the proof of Lemma 5.1, let $\sigma^i = \{\sigma^i_j\}$, for $1 \leqslant i \leqslant w$, be the sequence of vertices visited by the ith bus. Now, define a new sequence ξ by

$$\xi_j = \sigma^{(j \bmod w)+1}_{d+\lfloor j/w \rfloor r} \qquad \text{for } j \geqslant 0.$$

So,

$$\xi = \{\sigma^1_d, \ldots, \sigma^w_d, \sigma^1_{d+r}, \ldots, \sigma^w_{d+r}, \sigma^1_{d+2r}, \ldots\}.$$

If Y is the minimum index such that $\xi_Y = v$, then Y is a geometric random variable with parameter $1/n_m$ that may take the values $0, 1, \ldots$. Therefore, $\mathrm{E}[Y] = n_m - 1$. Note that here we have $h_w(u, v) = d + r \cdot \mathrm{E}[\lfloor Y/w \rfloor]$. Since

$$\frac{n_m - w - 1}{w} = \frac{n_m - 1}{w} - 1 = \mathrm{E}\left[\frac{Y}{w} - 1\right] \leqslant \mathrm{E}[\lfloor Y/w \rfloor] \leqslant \frac{\mathrm{E}[Y]}{w} = \frac{n_m - 1}{w},$$

it follows that

$$d + r\left(\frac{n_m - w - 1}{w}\right) \leqslant h_w(u, v) \leqslant d + r\left(\frac{n_m - 1}{w}\right). \qquad \square$$

We now consider how buses arrive at a node u in the case of a directed collar. Again, let A_i denote the number of time steps we have to wait before i distinct buses arrive at u, starting from some particular time.

Proposition 6.1. *Suppose $u \in V_l$. Then*

$$\mathrm{E}[A_i] \leqslant r\left(i + \frac{i n_l}{w - i + 1}\right) - i.$$

Proof. As in the proof of Proposition 5.1, we have $A_i = \sum_{j=1}^{i} B_j$, where B_j is the number of time steps between the arrival of the $(j-1)$st distinct bus and the arrival of the jth distinct

bus. However, we must make some modifications to the argument used in Proposition 5.1 in the case of a directed collar.

Define a new random variable $\overline{B}_j$ to be equal to s if $(s-1)r \leqslant B_j \leqslant sr-1$. Then $\overline{B}_j$ may take the values $1, 2, \ldots$. In the case of a directed collar, the probability that none of the $w-j+1$ buses that have yet to arrive at u will arrive at u *during the next* r steps is

$$\bar{q}_j = \left(\frac{n_l - 1}{n_l}\right)^{w-j+1}.$$

Then $\Pr[\overline{B}_j = s] = \bar{q}_j^{(s-1)}(1-\bar{q}_j)$. Thus $\overline{B}_j$ is a geometric random variable and

$$\mathrm{E}[\overline{B}_j] = \frac{1}{1-\bar{q}_j}.$$

Now, since $(\overline{B}_j - 1)r \leqslant B_j \leqslant \overline{B}_j r - 1$, we have

$$r\left(\frac{1}{1-\bar{q}_j} - 1\right) \leqslant \mathrm{E}[B_j] \leqslant r\left(\frac{1}{1-\bar{q}_j}\right) - 1.$$

Therefore,

$$\sum_{j=1}^{i} r\left(\frac{1}{1-\bar{q}_j} - 1\right) \leqslant \mathrm{E}[A_i] \leqslant \sum_{j=1}^{i} r\left(\frac{1}{1-\bar{q}_j}\right) - i.$$

Proceeding as in the proof of Proposition 5.1, we obtain

$$\mathrm{E}[A_i] \leqslant r\left(i + \frac{in_l}{w-i+1}\right) - i. \qquad \square$$

As in the previous section, let D_i be the number of steps taken by the ith distinct bus in traveling from u to v. We recall that the expected delivery time is $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right]$. Here, if $u \in V_l$ and $v \in V_m$, we have

$$\begin{aligned}\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] &\leqslant \mathrm{E}[A_\lambda] + h_\lambda(u,v)\\ &\leqslant r\left(\lambda + \frac{\lambda n_l}{w-\lambda+1}\right) - \lambda + d + r\left(\frac{n_m - 1}{\lambda}\right)\\ &\leqslant r\lambda + r\frac{\lambda n_*}{w-\lambda+1} - \lambda + r + r\left(\frac{n_* - 1}{\lambda}\right). \qquad (4)\end{aligned}$$

If we choose $\lambda \leqslant \sqrt{w}$, then notice that $\lambda^2 \leqslant n_*$, and so $\lambda \leqslant n_*/\lambda$. It follows then from (4) that the expected delivery time is $\mathrm{O}(rn_*/\lambda)$. For $\lambda > \sqrt{w}$, we have

$$\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i + D_i\}\right] \leqslant \mathrm{E}\left[\min_{1\leqslant i\leqslant\lfloor\sqrt{w}\rfloor}\{A_i + D_i\}\right] = \mathrm{O}\left(\frac{rn_*}{\sqrt{w}}\right).$$

Thus, we have shown

Theorem 6.1. *When G is an r-partite directed collar, the expected delivery time of a message in the persistent multiple bus protocol is*

$$\mathrm{O}\left(\frac{rn_*}{\sqrt{w}}+\frac{rn_*}{\lambda}\right)$$

for all $1\leqslant\lambda\leqslant w\leqslant n_*$.

A lower bound for delivery time for a directed collar would depend on the minimum of the number of vertices in each V_i. However, it is of more interest to find a lower bound for the maximum of the delivery time over all pairs of vertices. This maximum delivery time would occur when u and v both belong to a V_i with n_* vertices.

Theorem 6.2. *Let G be an r-partite directed collar. Suppose u and v are distinct vertices in V_l and $\#V_l=n_*$, where $n_*\geqslant 2$. Then the expected delivery time of a message from u to v is*

$$\Theta\left(\frac{rn_*}{\sqrt{w}}+\frac{rn_*}{\lambda}\right).$$

Proof. If $\lambda\leqslant\sqrt{w}$, then

$$\begin{aligned}\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\}\right]&\geqslant\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}D_i\right]\\&=h_\lambda(u,v)\geqslant r+r\left(\frac{n_*-\lambda-1}{\lambda}\right)=r\left(\frac{n_*-1}{\lambda}\right).\end{aligned}$$

It then suffices to show that the expected delivery time when $\lambda>\sqrt{w}$ is $\Omega(rn_*/\sqrt{w})$. To do that, we will proceed in a similar manner to the proof of the lower bound in Theorem 5.1. As in that proof, we have

$$\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\}\right]\geqslant\Pr\left[\min_{1\leqslant i<\theta}D_i\geqslant\frac{rn_*}{c\sqrt{w}}\right]\Pr\left[A_\theta\geqslant\frac{rn_*}{c\sqrt{w}}\right]\frac{rn_*}{c\sqrt{w}},$$

for any positive number c, where $\theta=\lfloor\sqrt{w}\rfloor+1$.

Put $\vartheta=\lceil n_*/w\rceil$. With $\overline{B}_j$ as in the proof of Proposition 6.1, define $\overline{Z}_j$ to be 1 if $\overline{B}_j-1\geqslant\vartheta$ and 0 otherwise. Then

$$\Pr[\overline{Z}_j=1]=\Pr[\overline{B}_j\geqslant\vartheta+1]=\sum_{s=\vartheta+1}^{\infty}\bar{q}_j^{(s-1)}(1-\bar{q}_j)=\bar{q}_j^{\vartheta}.$$

Put $\overline{Z}=\sum_{j=1}^{\theta}\overline{Z_j}$ and let $\bar{\mu}=\mathrm{E}[\overline{Z}]$. Then

$$\begin{aligned}\bar{\mu}&=\sum_{j=1}^{\theta}\bar{q}_j^{\vartheta}=\sum_{j=1}^{\theta}\left(\frac{n_*-1}{n_*}\right)^{\vartheta\cdot(w-j+1)}\geqslant\theta\left(\frac{n_*-1}{n_*}\right)^{\vartheta w}\\&=\frac{\theta}{(1+(\vartheta w/(n_*-1))/\vartheta w)^{\vartheta w}}\geqslant\frac{\theta}{\mathrm{e}^{\vartheta w/(n_*-1)}}.\end{aligned}$$

Now, $\vartheta w < n_* + w \leqslant 2n_*$. Hence, $\vartheta w/(n_* - 1) < 2n_*/(n_* - 1) \leqslant 4$ (since $n_* \geqslant 2$). Therefore, we have

$$\bar{\mu} > \frac{\theta}{e^4} > \frac{\sqrt{w}}{e^4}.$$

Applying the Chernoff bound as in the proof of Theorem 5.1, we obtain

$$\Pr\left[\overline{Z} \geqslant \frac{\sqrt{w}}{2e^4}\right] \geqslant \Pr\left[\overline{Z} \geqslant \mu/2\right] \geqslant 1 - \frac{1}{e^{\mu/8}} \geqslant 1 - \frac{1}{e^{\theta/8e^4}} \geqslant 1 - \frac{1}{e^{1/4e^4}} > 0.004.$$

Now, we have

$$A_\theta = \sum_{j=1}^{\theta} B_j \geqslant \sum_{j=1}^{\theta} (\overline{B}_j - 1)r \geqslant r\vartheta \sum_{j=1}^{\theta} \overline{Z}_j = r\vartheta\overline{Z}.$$

Thus,

$$\begin{aligned}\Pr\left[A_\theta \geqslant \frac{rn_*}{2e^4\sqrt{w}}\right] &\geqslant \Pr\left[A_\theta \geqslant \frac{r\vartheta\sqrt{w}}{2e^4}\right] \\ &\geqslant \Pr\left[r\vartheta\overline{Z} \geqslant \frac{r\vartheta\sqrt{w}}{2e^4}\right] = \Pr\left[\overline{Z} \geqslant \frac{\sqrt{w}}{2e^4}\right] > 0.004.\end{aligned}$$

Finally, we need to obtain a lower bound for $\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant \frac{rn_*}{2e^4\sqrt{w}}\right]$. Recall that D_i is the number of steps taken by the ith distinct bus in traveling from u to v. Since u and v are both in V_l, it follows that D_i will be a multiple of r. For any nonnegative integer v, we have

$$\begin{aligned}\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant r(v+1)\right] &= \left(\frac{n_* - 1}{n_*}\right)^{(\theta-1)v} \\ &= \frac{1}{(n_*/(n_* - 1))^{(\theta-1)v}},\end{aligned}$$

since at each of v trips completely around the collar (starting at u) each of $\theta - 1$ buses goes through one of the $n_* - 1$ vertices in V_l other than v. Writing

$$\left(\frac{n_*}{n_* - 1}\right)^{(\theta-1)v} = \left(1 + \frac{(\theta - 1)v/(n_* - 1)}{(\theta - 1)v}\right)^{(\theta-1)v}$$

and applying (1), we get

$$\begin{aligned}\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant r(v+1)\right] &\geqslant \frac{1}{e^{(\theta-1)v/(n_*-1)}} \\ &\geqslant \frac{1}{e^{v\sqrt{w}/(n_*-1)}}.\end{aligned}$$

For $v = \lfloor n_*/(2e^4\sqrt{w})\rfloor$ and $n_* \geqslant 2$ (so $n_*/(n_*-1) \leqslant 2$), we have

$$\Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant r n_*/(2e^4\sqrt{w})\right] \geqslant \Pr\left[\min_{1\leqslant i<\theta} D_i \geqslant r(v+1)\right] \geqslant \frac{1}{e^{1/e^4}} > 0.98. \qquad \square$$

We now consider the question of congestion in the directed collar case. Let w_i denote the number of buses initially at vertices in V_i. Then it is clear that there will never be more than $\max_{1\leqslant i\leqslant r} w_i$ buses at any vertex. In the case when $\#V_i = n/r$ for all i, by applying the balls into bins model as in the complete graph case to the "blocks" V_i and then to the vertices in each block, we see that if $k^* = \lceil (9w/n)(\ln r/\ln\ln r)(\ln(n/r)/\ln\ln(n/r))\rceil$, then no vertex will have more than k^* buses with probability at least $(1-1/r)(1-r/n)$.

Remark 3. A complete (undirected) bipartite graph may be viewed as a 2-partite directed collar with the arcs from V_1 to V_2 and from V_2 to V_1 being replaced by the edges of the complete bipartite graph. Hence, the results from this section, with $r = 2$, apply to the case of a complete bipartite graph.

7. Multiple perishable buses

We now return to the case of a complete graph and consider the situation where buses are perishable. This situation is significantly more complicated than the situation where buses are persistent. However, intuitively, if we choose k and ℓ large, then we expect that the performance of the system will be close to that where buses live forever. We show that for $k = \ell = n^{7/2}$, with high probability, the expected delivery time of a message is $O(\sqrt{n})$.

If we choose $k = \ell$, then after the first k steps, we will have exactly n buses in the system at all times (assuming buses do not crash and no node needs to be reset). This seems to be a nice choice, as the number of buses in the system is constant. (However, with so many buses, one may wish to consider using the naive system in which each node sends a real or nonsense message to every other node at each time step.) In this situation, as noted at the end of the previous section, the number of buses at any vertex is at most $\lceil 3\ln n/\ln\ln n\rceil$ with probability at least $1-1/n$, so the level of congestion is reasonable. So, for the remainder of this section we assume that $w = n$ and $\lambda = \sqrt{n}$.

In the situation where buses are not destroyed, a message placed on a bus is delivered with probability 1. If buses have lifetimes, this is no longer true. However, we are able to show that with high probability a message will be delivered.

Suppose we want to send a message at time t. Recall that each bus dies every ℓ time steps, and the times at which buses are initially created at nodes are uniformly distributed between 0 and $k-1 = \ell-1$. We calculate the probability that no bus is destroyed in the time period starting at t and ending at $t+n^{3/2}$. The probability that a given bus dies during this period is exactly $\max\{n^{3/2}/\ell, 1\}$. Assume from now on that $\ell \geqslant n^{3/2}$. Using Boole's Inequality, the probability that some bus dies during the previous mentioned time period is at most $wn^{3/2}/\ell = n^{5/2}/\ell$. So, if we pick $\ell = n^{7/2}$, then the probability that *no* bus dies

during this time period is at least $1-1/n$. Let Y denote the number of buses that die during this time period. Then, we've seen that $\Pr[Y=0] \geqslant 1-1/n$.

Markov's inequality (which holds for conditional probabilities and expectations) tells us that

$$\Pr\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\} < n^{3/2} \mid Y=0\right] \geqslant 1-\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\} \mid Y=0\right] \Big/ n^{3/2}$$
$$= 1-\mathrm{O}\left(\frac{1}{n}\right),$$

since, by the previous section, $\mathrm{E}\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\} \mid Y=0\right]$ is $\mathrm{O}(n/\lambda)=\mathrm{O}(\sqrt{n})$. Then, since

$$\Pr\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\} < n^{3/2}\right] \geqslant \Pr\left[\min_{1\leqslant i\leqslant\lambda}\{A_i+D_i\} < n^{3/2} \mid Y=0\right]\Pr[Y=0],$$

the probability that the message gets delivered during this time period is at least $1-\mathrm{O}(1/n)$. The expected delivery time given these circumstances is clearly $\mathrm{O}(\sqrt{n})$.

8. Conclusions

We have presented further analysis of the random walk anonymous message system proposed by Beimel and Dolev [2]. We have shown that by increasing the number of buses w in a system, delivery time is decreased. However, our lower bound of $\Omega(n/\sqrt{w})$ in the complete graph case shows that message delivery time does not decrease linearly with the number of buses. It would be quite interesting to obtain upper bound results for other types of communication network topologies besides complete graphs and directed collars.

Acknowledgements

We dedicate this work to the memory of our co-author Steve Seiden. We thank Nigel Gwee for helpful discussions and the referees for useful comments.

References

[1] D.J. Aldous, J.A. Fill, Reversible Markov Chains and Random Walks on Graphs, in preparation.
[2] A. Beimel, S. Dolev, Buses for anonymous message delivery, J. Cryptology 16 (2003) 25–39.
[3] M. Ben-Or, S. Goldwasser, A. Wigderson, Completeness theorems for non-cryptographic fault-tolerant distributed computation, in: Proc. 20th Annual ACM Symp. on the Theory of Computing, May 1988, pp. 1–10.
[4] R. Canetti, U. Feige, O. Goldreich, M. Naor, Adaptively secure multi-party computation, in: Proc. 28th Annual ACM Symp. on the Theory of Computing, May 1996, pp. 639–648.
[5] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, Commun. ACM 24 (2) (1981) 84–88.
[6] D. Chaum, The dining cryptographers problem: unconditional sender and recipient untraceability, J. Cryptology 1 (1988) 65–75.

[7] D. Chaum, C. Crépeau, I. Damgård, Multiparty unconditionally secure protocols, in: Proc. 20th Annual ACM Symp. on the Theory of Computing, May 1988, pp. 11–19.
[8] E.G. Coffman Jr., N. Kahale, F.T. Leighton, Processor-ring communication: a tight asymptotic bound on packet waiting time, SIAM J. Comput. 27 (1998) 1221–1236.
[9] D. Coppersmith, U. Feige, J. Shearer, Random walks on regular and irregular graphs, SIAM J. Discrete Math. 9 (1996) 301–308.
[10] S. Dolev, R. Ostrovsky, Xor-trees for efficient anonymous multicast and reception, ACM Trans. Inform. System Security 3 (2) (2000) 63–84.
[11] W. Feller, An Introduction to Probability Theory and its Applications, Vols. 1 & 2, Wiley, New York, 1957.
[12] O. Goldreich, S. Micali, A. Wigderson, How to play any mental game or a completeness theorem for protocols with honest majority, in: Proc. 19th Annual Symp. on the Theory of Computing, May 1987, pp. 218–229.
[13] L. Lovász, Random walks on graphs: a survey, in: Combinatorics, Paul Erdös is Eighty, Vol. 2, Keszthely, 1993, pp. 353–397 (Bolyai Society of Mathematical Studies, Vol. 2, János Bolyai Mathematical Society, Budapest, 1996).
[14] R. Motwani, P. Raghavan, Randomized Algorithms, Cambridge University Press, Cambridge, 1995.
[15] A. Pfitzmann, How to implement ISDNs without user observability - some remarks, Technical Report 14/85, Fakultät für Informatik, Universität Karlsruhe, 1985.
[16] C. Rackoff, D.R Simon, Cryptographic defense against traffic analysis, in: Proc. 25th Ann. ACM Symp. on the Theory of Computing, May 1993, pp. 672–681.
[17] P.F. Syverson, D.M. Goldschlag, M.G. Reed, Anonymous connections and onion routing, in: IEEE Symp. on Security and Privacy, May 1997, pp. 44–54.