Technical Report

UCAM-CL-TR-964
ISSN 1476-2986

Number 964

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Information dissemination via random walks

Hayk Saribekyan

November 2021

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

https://www.cl.cam.ac.uk/

This technical report is based on a dissertation submitted July 2021 by the author for the degree of Doctor of Philosophy to the University of Cambridge, St John's College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

https://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Abstract

Information dissemination via random walks

Hayk Saribekyan

Information dissemination is a fundamental task in distributed computing: How to deliver a piece of information from a node of a network to some or all other nodes? In the face of large and still growing modern networks, it is imperative that dissemination algorithms are decentralised and can operate under unreliable conditions. In the past decades, randomised rumour spreading algorithms have addressed these challenges. In these algorithms, a message is initially placed at a source node of a network, and, at regular intervals, each node contacts a randomly selected neighbour. A message may be transmitted in one or both directions during each of these communications, depending on the exact protocol. The main measure of performance for these algorithms is their broadcast time, which is the time until a message originating from a source node is disseminated to all nodes of the network. Apart from being extremely simple and robust to failures, randomised rumour spreading achieves theoretically optimal broadcast time in many common network topologies.

In this thesis, we propose an agent-based information dissemination algorithm, called VISIT-EXCHANGE. In our protocol, a number of agents perform independent random walks in the network. An agent becomes informed when it visits a node that has a message, and later informs all future nodes it visits. VISIT-EXCHANGE shares many of the properties of randomised rumour spreading, namely, it is very simple and uses the same amount of communication in a unit of time. Moreover, the protocol can be used as a simple model of non-recoverable epidemic processes.

We investigate the broadcast time of VISIT-EXCHANGE on a variety of network topologies, and compare it to traditional rumour spreading. On dense regular networks we show that the two types of protocols are equivalent, which means that in this setting the vast literature on randomised rumour spreading applies in our model as well. Since many networks of interest, including real-world ones, are very sparse, we also study agent-based broadcast for sparse networks. Our results include almost optimal or optimal bounds for sparse regular graphs, expanders, random regular graphs, balanced trees and grids. We establish that depending on the network topology, VISIT-EXCHANGE may be either slower or faster than traditional rumour spreading. In particular, in graphs consisting of hubs that are not well connected, broadcast using agents can be significantly faster. Our conclusion is that a combined broadcasting protocol that simultaneously uses both traditional rumour spreading and agent-based dissemination can be fast on a larger range of topologies than each of its components separately.

Acknowledgements

First and foremost, I would like to thank my supervisors Thomas Sauerwald and George Giakkoupis for their continuous support and patience during my PhD study. Without their guidance and feedback this work would not be.

I thank the members of our group, Nicolás Rivera, Leran Cai, John Sylvester and Luca Zanetti, as well as my co-authors, Frederik Mallmann-Trenn and Paul Bastide.

I am very grateful to the Gates Cambridge Trust and St John's College for their generous support of my work. I also thank Inria for hosting me in Rennes for research visits.

My interest in research was inspired by interactions with Nir Shavit and Yaron Meirovitch who helped me in my first steps. One of the most memorable moments in my life has happened during our work together, for which I am grateful to both. I also thank Garegin Papoian for his advice and insights into the scientific world.

I thank Cambridge University Gliding Club for opening the world of flying to me. At CUGC I was able to disconnect from my work and get a unique immersion in the UK culture by meeting the outstanding people that form the club. CUGC helped me appreciate the value of service and it was my privilege to contribute to the club.

I would not be the person I am without my wonderful friends, who supported and encouraged me at each step of this journey. I thank Varduhi, Vahe, Vahagn and Mike for all the motivating, inspiring and fun discussions we had. I look forward to many more.

I thank uncle Rubik, for always motivating me with his infinite energy and his excitement about science, and uncle Laert, for sharing his love and admiration for mathematics.

This journey would not be as meaningful without Nora. I thank her for the support and encouragement, for being excited about whatever I do and pushing me to do it better. Thank you, Nora, for dreaming with me.

There are no words that can express my gratitude towards my parents and my brother, who have always found the right words for me in difficult times. Their wisdom, their sacrifices and their unwavering belief in me got me to where I am now. Moving forward, I am left with a great responsibility that I will carry with pride, joy and honour.

To my grandmother,
Emma Buniatyan

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Networks are all around us, from logistics to social media to the world wide web. A core function of many of these networks is the dissemination of information between their nodes. The sheer size and complexity of modern networks mean that often the dissemination cannot be orchestrated by a central coordinator. Individual nodes must make independent decisions about their local communication without having a sense about the global state of the system [Eug+04]. The information flow in such decentralised systems is thus determined by these local decisions and the structure of the network. Understanding how these two factors affect information spread is one of the central topics of distributed computing and has been studied for decades. Many of the seminal works in the field have studied information dissemination from a pragmatic point of view of building computer networks capable of efficient data transfer and synchronisation [Dem+88; Ora01; VRB03]. A parallel and rich line of work studied information dissemination from a theoretical perspective and aimed to understand fundamental properties of networks that affect various dissemination methods. This viewpoint also helps to model and understand certain processes in naturally occurring networks, such as the spread of news in social networks or of infections in animal or human contact networks [HHL88; Sha07; Kar+00]. This thesis focuses on theoretical aspects of information dissemination.

We study *broadcasting* protocols (or algorithms) that are executed by the nodes of the network in parallel. A piece of information or a message originates at a source node of a network, and the goal of the protocol is to use message passing between neighbouring nodes to deliver the message to all other nodes (inform them). One simple broadcasting method is the flooding protocol, where each informed node continuously propagates the message it has received to all of its neighbouring nodes. While flooding is fast and robust to network changes, it is prohibitively expensive in terms of communication cost. In order to design fast, robust and communication-efficient protocols we turn to randomisation.

Randomised broadcasting protocols have been studied in the distributed computing community for many years. If chosen appropriately, they can broadcast a message quickly and without a large communication cost [Chi+18; BES14]. Additionally, often they are simpler than their deterministic counterparts and operate without requiring any memory at the nodes (this is impossible in the deterministic case as the nodes would repeat their actions). Another advantage of randomised protocols is that they better reflect communication patterns found in natural processes such as of the spread of a virus or a rumour in a population. An epidemic, in fact, is a physical example of a simple, randomised and fast spreading process that arguably uses little amount of "communication" and "coordination" between nodes, much like rumour spreading [Die67]. A simple protocol for information dissemination is PUSH, proposed in the 1980s: At every round, a node that has previously received a message, sends it to a neighbour picked uniformly at random [Dem+88]. PUSH-PULL is a slightly more complicated version of PUSH, whereby at every time step each node picks a random neighbour and if either of the two nodes is informed, then after the round both become informed. It turns out that with this small modification, PUSH-PULL can be significantly faster than PUSH (but, obviously, never the other way). The two protocols together are referred as *randomised rumour spreading* algorithms.

The main topic of this thesis is the study of broadcasting protocols that use *moving and interacting agents* (or particles) in the network. In most of the previous literature on broadcast or information dissemination, adjacent nodes of the network communicate directly like in randomised rumour spreading above. We propose an agent-based protocol, called VISIT-EXCHANGE, where the nodes communicate only using a number of agents. At the start of this protocol, each agent is placed at a random node of the network, chosen independently with probability proportional to the number of neighbours of the node. In other words, the agents are initially placed according to the stationary distribution of the network. Then, the agents perform an independent random walk, that is, at every round each agent selects a random neighbour of the node where it currently is located, and moves that that neighbour. If an agent visits an informed node, then it becomes informed, and after that informs all future nodes it visits.

Randomised rumour spreading, as the name suggests, can be used as a very simple model of rumour spreading in a friendship network, where the rumour is propagated by people randomly calling their friends and sharing information. The VISIT-EXCHANGE process can serve as a very simple model of an epidemic without recovery: The network represents a region with its nodes as cities and its edges as roads, the agents represent people moving between the cities, and the message represents a contagious virus that is spreading. Despite the fact that these models are perhaps too simple to be used in understanding of real-world phenomena, or in building of distributed systems, they provide initial insights in these directions. For example, the PUSH process belongs to a general class of gossiping protocols, which are characterised by the fact that nodes communicate with their neighbours one at a time, randomly. A more elaborate gossiping algorithm than PUSH is averaging, where nodes have a load (a real number) and at each step communicating nodes average their loads [Boy+06]. Averaging protocols have implications in machine learning and, in particular, in distributed stochastic gradient descent [Lia+18]. Similarly, we view the study of VISIT-EXCHANGE and its broadcast time as a first step towards understanding more complex agent-based algorithms, which will be used when networks that are composed of mobile entities become more abundant (e.g., drone networks, self-driving cars).

In this thesis we study agent-based broadcasting protocols in terms of the time it takes for broadcasting to complete, and make comparisons to randomised rumour spreading. We also conduct an experimental evaluation of the considered processes observing properties that were not obvious from the theoretical analysis. Our results show that agent-based methods have advantages in certain network types, including on some real-world networks. Additionally, we argue that combining agent-based and non-agent-based protocols can result in an algorithm that is efficient on a wider range of networks.

1.2 Overview of results

We consider the synchronised versions of the aforementioned PUSH and PUSH-PULL protocols, where the nodes execute them in parallel, taking steps simultaneously, in distinct rounds. Similarly, we study the synchronous VISIT-EXCHANGE process, where agents perform discrete random walks, taking steps at the same time. In all three protocols, before the process starts, a message is delivered to a source vertex which becomes informed. In our theoretical analysis, we assume that VISIT-EXCHANGE uses a linear number of agents in the size of the network, which also makes the comparison with PUSH or PUSH-PULL fair in terms of the amount of the communication used in one round.

The main measure of performance we study is the *broadcast time* of information dissemination protocols, the time it takes until all nodes of the network become informed, measured in the number of rounds. In general, the broadcast time of a protocol may depend on the source vertex. For VISIT-EXCHANGE and PUSH-PULL, however, it can be shown that the broadcast time is asymptotically the same for any source. In PUSH the source vertex matters, so in our bounds the worst possible vertex is considered, unless specifically mentioned otherwise. It should be noted here that since the protocols we consider are randomised, their broadcast time is a random variable. Most of the bounds we prove on the broadcast time hold with high probability,[1] while a few others hold in expectation only. This is usually a small price one has to pay to take advantage of the many benefits provided by randomised algorithms. The summary of our main results for the broadcast time of VISIT-EXCHANGE and its comparison to PUSH and PUSH-PULL can be found in Table 1.1.

Our first set of results compares randomised rumour spreading with the proposed agent-based protocol VISIT-EXCHANGE. We prove that in general graphs the two categories of processes are not comparable: There are instances where PUSH is significantly faster than VISIT-EXCHANGE, and instances where VISIT-EXCHANGE is significantly faster than PUSH-PULL. The networks in these instances are highly non-regular, that is, some nodes have significantly larger number of neighbours than others. In particular, it seems VISIT-EXCHANGE has an advantage over PUSH-PULL when the network consists of some number of hubs and many smaller nodes. This advantage can be attributed to the fact that VISIT-EXCHANGE uses all edges at the same frequency, while the other two protocols use edges connecting hubs less frequently.

On the other hand, we prove that in sufficiently dense *regular* graphs, where the number of neighbours of all vertices is the same and is at least logarithmic, the two types of protocols have the broadcast time asymptotically. This implies that for such graphs the vast literature on randomised rumour spreading also applies to VISIT-EXCHANGE, bounding its broadcast time in terms of the graph conductance, vertex expansion, diameter and degree [Chi+18; Gia14; Fei+90]. Intuitively, this is not surprising since in regular graphs in one round of VISIT-EXCHANGE a constant number of agents depart each vertex, in expectation, which should have the same effect as PUSH. The formalisation of this intuition and subsequent proofs are non-trivial. We use a coupling between VISIT-EXCHANGE and PUSH, which allows us to argue that if one of the processes makes progress along a path starting from the source, then the other process will also follow the same path in

[1] In the thesis, *with high probability* or *w.h.p.* means with probability at least $1 - n^{-c}$ for some constant $c > 0$, where n is the number of vertices of the graph under consideration. The constant can be adjusted at the expense of constant factors in the broadcast times of processes.

<table>
<tr><th rowspan="2">Graph</th><th rowspan="2">Conditions</th><th colspan="2">Broadcast times</th></tr>
<tr><th>VISIT-EXCHANGE</th><th>PUSH and PUSH-PULL</th></tr>
<tr><td rowspan="2">Regular</td><td>$d = \Omega(\log n)$</td><td colspan="2">Asymptotically the same broadcast time (Thm. 3.1.1)
$O(d \cdot (\mathrm{diam} + \log n))$ [Fei+90]</td></tr>
<tr><td>$d = O(\log n)$</td><td>$\widetilde{O}(d \cdot \mathrm{diam} + \log^3 n/d)$
Thm. 4.1.1</td><td>$O(d \cdot (\mathrm{diam} + \log n))$
[Fei+90]</td></tr>
<tr><td>Any</td><td>$d_{\min} = \Omega(d_{\mathrm{avg}})$</td><td>$O(d_{\mathrm{avg}} \log^2 n \cdot (\mathrm{diam} + \log n))$
Thm. 4.1.2</td><td>$O(d_{\max} \cdot (\mathrm{diam} + \log n))$
[Fei+90]</td></tr>
<tr><td rowspan="4">Regular expanders</td><td>$d = \Omega(\log n)$</td><td rowspan="3">$O(\log n)$
Thm. 5.1.1 to 5.1.3</td><td rowspan="4">$O(\log n)$
[Chi+18]</td></tr>
<tr><td>$d = O(1)$</td></tr>
<tr><td>Strong expander</td></tr>
<tr><td>Any d</td><td>$O(\log n \cdot \log \log n)$
Thm. 5.1.5</td></tr>
<tr><td>Random regular</td><td>Any d</td><td>$O(\log n)$
Thm. 5.1.4</td><td>$O(\log n)$
[FP10]</td></tr>
<tr><td rowspan="2">Balanced trees</td><td>b-ary tree of height $h \geqslant 2$</td><td>$O(h \log h + \log n)$
Thm. 6.1.1</td><td>$\Theta(b \cdot \log n)$
[Fei+90]</td></tr>
<tr><td>Binary tree</td><td>$\Theta(\log n \cdot \log \log n)$
Thm. 6.1.1</td><td>$\Theta(\log n)$
[Fei+90]</td></tr>
<tr><td>Grids, Tori</td><td>Constant dimensional</td><td>$\Theta(\mathrm{diam})$
Thm. 7.1.1</td><td>$\Theta(\mathrm{diam})$
[Fei+90]</td></tr>
</table>

Table 1.1 Summary of the main results in this thesis and their analogues for randomised rumour spreading for a graph with n vertices. An arbitrary source vertex is assumed. For the presented cases, PUSH and PUSH-PULL have the same broadcast time, w.h.p. The degree of regular graphs is denoted by d. The minimum, average and maximum degrees are denoted by $d_{\min}$, d_{avg} and $d_{\max}$, respectively, and the diameter is diam. The tilde notation hides factors of order at most $(\log \log n)^2$. All bounds hold w.h.p.

approximately the same number of rounds.

Given the earlier intuition for regular graphs, it is perhaps surprising that there are (sparse) regular graphs where the two protocols have slightly different broadcast times. Such examples are specifically constructed to create "node islands" which agents visit rarely. As a result, VISIT-EXCHANGE may inform almost all vertices of the graph, except the few remaining ones in the "island," increasing the broadcast time. This motivates the analysis of the partial broadcast time, which is the time until, for example, 90% of the vertices become informed. We have only done some basic experimental analysis on the partial broadcast, presented in Appendix A, and focus on the complete broadcast. In most sparse regular graphs, however, the equivalence between VISIT-EXCHANGE and randomised rumour spreading holds as we are able to give a tight upper bound on the broadcast time of VISIT-EXCHANGE on random regular graphs of any degree.

The next set of results studies VISIT-EXCHANGE on its own merit depending on the diameter, average degree and expansion of the graph. First, we prove an almost tight bound on broadcast time for arbitrary regular graphs, depending on the degree and the diameter of the graph. The bound is tight up to an additive poly-logarithmic term. Second, we present an asymptotically optimal bound for regular graphs with strong expansion properties. This result also gives us a tight bound on the broadcast time for random regular graphs. Analogous results also exist for PUSH [Chi+18; SS11; MS08], which is not surprising due to the earlier intuition on the equivalence of VISIT-EXCHANGE and PUSH on regular graphs. Third, we bound the broadcast time of VISIT-EXCHANGE with respect to the average degree and the diameter of the graph. The result most similar to this for PUSH depends on the maximum degree of the graph instead of the average [Fei+90]. This implies that on sparse graphs which have nodes of high degrees, VISIT-EXCHANGE can be significantly faster.

We also thoroughly analysed VISIT-EXCHANGE on balanced trees of arbitrary branching, proving both upper and lower bounds. One surprising result here is that PUSH and PUSH-PULL are slightly faster in low-degree balanced trees (e.g., the binary tree), even though these are almost regular graphs. However, as the degree increases beyond doubly-logarithmic in the size of the network (e.g., in the star graph, the balanced degree of height one), VISIT-EXCHANGE becomes significantly faster. This fact further reinforces the intuition that VISIT-EXCHANGE can be advantageous in graphs with hubs. Our results for the balanced trees are tight and also imply a lower bound on the cover time of many random walks in a balanced tree.

Our final theoretical study is of VISIT-EXCHANGE on grid graphs of any constant dimension. We prove that VISIT-EXCHANGE has an asymptotically optimal broadcast time. The technique we use adapts a beautiful line of papers by Kesten and Sidoravicius and works for grids of any constant dimension, including for dimension one, that is, the path graph [KS03; KS05]. The latter is a particularly challenging case as there are other techniques that are likely to be useful only if the dimension is at least two [GS18; Lam+12].

As mentioned in the introduction, our results indicate that combining agent-based and traditional rumour spreading protocols may result in an algorithm that is efficient in a wider range of networks. For this reason, we introduce a protocol called VX-PUSH-PULL. This protocol uses agents too, and in each round executes one step of VISIT-EXCHANGE and one step of PUSH-PULL, independently. Assuming that in VX-PUSH-PULL the number of agents is linear in the number of vertices, the upper bounds on the broadcast time of VISIT-EXCHANGE and PUSH-PULL apply also for this protocol. This results in asymptotically optimal protocol for all classes of graphs studied in the thesis. Beyond this simple conclusion, we have not evaluated VX-PUSH-PULL theoretically. However, we present preliminary experimental analysis of the VX-PUSH-PULL as well as of VISIT-EXCHANGE and PUSH-PULL separately, complementing the theoretical study. First, we observe that in our bounds the constants hidden in the asymptotic notation appear to be small. In order to compare VX-PUSH-PULL to the original processes fairly, we use half as many agents as in VISIT-EXCHANGE and let each communication in PUSH-PULL fail with probability 1/2. It appears that in a variety of networks, including in real-world networks, VX-PUSH-PULL is not slower than any of the other protocols and in some cases is positively faster, especially when considering partial broadcasting to only 90% of the vertices. These results indicate that agent-based information dissemination is not only a mathematically interesting process, but can also have practical implications.

1.3 Related work

The literature relevant to this thesis can be split in two categories. The first includes prior work on information dissemination in the field of distributed computing. The second is the random walk literature, particularly that considers processes with many random walks. We review these separately.

1.3.1 Information dissemination

The problem of information broadcast is central in distributed computing and there is a large volume of prior work. We focus on the review of theoretical works, and throughout this section assume that the results are on a connected, undirected graph $G = (V, E)$ with $|V| = n$ vertices.The first paper considering a randomised information dissemination is by Frieze and Grimmett, where a telephone call protocol is introduced that is equivalent to the PUSH process defined earlier [FG85]. In this work the authors consider the protocol on a complete graph of n vertices and prove that PUSH takes at most $\log_2 n + \ln n + o(\log n)$ rounds, w.h.p. Later, Pittel showed that the broadcast time is $\log_2 n + \ln n + O(1)$, almost surely. Later, Feige et al. considered the same process on topologies other than the complete graph [Fei+90]. They proved an optimal bound of $O(d_{\max} \cdot (\mathrm{diam} + \log n))$ for the broadcast time of PUSH on any graph G with maximum degree $d_{\max}$ and diameter diam. Although this bound is best possible for arbitrary graphs, if we narrow the class of graphs better bounds are possible. Indeed, [Fei+90] also provides better bounds for hypercube graphs and Erdös-Rényi random graphs.

On the distributed systems side, Demers et al. introduced the anti-entropy mechanism for database replication, in which a node periodically picks a random neighbour and the two resolve their differences (reducing the entropy of the system) [Dem+88]. A large volume of work has followed this paper in the context of databases, studying properties of anti-entropy with respect to robustness, correctness, database consistency. To study the efficiency of anti-entropy, Karp et al. considered the simplest version of anti-entropy that is PUSH-PULL [Kar+00]. Like earlier papers on PUSH, they also study the process on complete graphs but, in addition to the broadcast time, also put an emphasis on the number of messages used by the process. In particular, it is shown that PUSH-PULL completes in $\log_3 n + O(\ln \ln n)$ rounds, w.h.p., and the majority of the messages are exchanged in the last $O(\ln \ln n)$ rounds, meaning that the message-complexity of the protocol is $O(n \ln \ln n)$. Note that the simple PUSH requires $\Omega(n \log n)$ messages.

With the adoption of internet en masse and the subsequent rise of social networks, besides the earlier goal of algorithm design many studies on information dissemination also turned to modelling information flow in networks. For this reason PUSH and PUSH-PULL and similar protocols have been studied for a variety of classes of graphs such as expanders, regular graphs, hypercubes and random graphs. We describe such results next. In many of these, the diameter of the graph is logarithmic in n and it is shown that the broadcast time is also asymptotically logarithmic, i.e., the protocol is optimal.

Fountoulakis et al. consider the PUSH-PULL process on scale-free networks, which can be used to model certain real-world networks [FPS12]. (For a discussion on the abundance of scale-free networks, see [BC19].) They prove that under certain conditions on the scaling parameter, the broadcast to $1-\epsilon$ fraction of all vertices of the graph happens in $O(\log \log n)$ rounds, for any $\epsilon > 0$. The authors also consider the asynchronous version of PUSH-PULL, in which the nodes take steps independently at intervals determined by a Poisson clock rather than synchronously in a lock-step. They show that for the asynchronous PUSH-PULL the broadcast to $1-\epsilon$ fraction of vertices will finish in constant number of rounds. This is a surprising result given that the diameter of the graph is super-constant. Since the synchronous broadcast time is lower bounded by the diameter of a graph, the result also shows that PUSH-PULL and its asynchronous variant have different behaviours.

The above result contrasts with [Sau10], which proves that the asynchronous and synchronous PUSH processes have asymptotically the same runtime. As an intermediate step, the author uses the following fact about robustness of rumour spreading. Namely, if each communication of PUSH fails with a constant probability, then the asymptotic broadcast of the process remains unchanged [ES09]. Daknama, Panagiotou, and Reisser also consider the robustness of randomised rumour spreading [DPR21]. They focus on expander graphs where at each round a certain number of edges can disappear and bound the broadcast time of the processes in terms of failure parameters. These results reinforce our intuition that rumour spreading must be robust to failures due to local decision-making and randomness.

One possible drawback of randomised rumour spreading is its use of randomness at every round. In quasirandom rumour spreading, introduced by Doerr, Friedrich, and Sauerwald, each vertex creates a cyclic list of its neighbours, starting from a random one, and contacts its neighbours according to the list [DFS14]. The authors prove that quasirandom rumour spreading is at least as fast as regular rumour spreading for a large number of graph classes, such as expanders, k-ary trees, hypercubes, complete graphs, random graphs. [Doe+08] studies quasirandom rumour spreading empirically, complementing theoretical results. Furthermore, they show that in practice the quasirandom version is slightly faster than normal rumour spreading and its broadcast time is more concentrated around the mean.

A class of graphs for which rumour spreading is often studied is that of regular graphs. In regular graphs PUSH and PUSH-PULL use all graph edges at the same rate, i.e., the probability that any particular edge is used in a round is the same. (This is also the case for VISIT-EXCHANGE in all graphs, not just regular ones.) Giakkoupis et al. show that the two protocols have the same asymptotic runtime for any regular graph [GNW16]. Their result also holds in the asynchronous case. Fountoulakis et al. studied rumour spreading for random regular graphs and proved precise logarithmic bounds on the broadcast time [FP10]. One of the main results of this thesis is also the comparison of VISIT-EXCHANGE and PUSH for regular graphs.

Following many results studying rumour spreading for specific classes of graphs, a beautiful line of papers aimed to give general bounds on the broadcast time of randomised rumour spreading protocols in terms of the expansion parameters of the graph, such as its conductance ϕ and vertex expansion α, as it is natural to expect that high expansion implies fast broadcast. The first such result is due to Chierichetti, Lattanzi, and Panconesi, who showed that the broadcast time of PUSH-PULL is bounded by $O(\log^4 n/\phi^6)$, w.h.p., for any graph [CLP10]. They used a spectral sparsification of graphs by Spielman and Teng [ST11]. Furthermore, they also gave a lower bound of $\Omega(\log n/\phi)$ on the broadcast time. By a more direct approach, Giakkoupis showed that the matching upper bound of $O(\log n/\phi)$ on the broadcast time [Gia11]. These results appear in [Chi+18] and they close the analysis of broadcast time of PUSH-PULL with respect to graph conductance. Notice that the optimal bound generalises the earlier analyses proving logarithmic bounds for complete graphs, expanders or random regular graphs since for them the conductance ϕ is constant.

Similar to the results involving the conductance ϕ of a graph, several papers have studied bounds on the broadcast time of PUSH-PULL with respect to the vertex expansion α of a graph. The tight result appeared in [Gia14] which proved that the broadcast time is at most $O(\log n \cdot \log d_{\max}/\alpha)$, where $d_{\max}$ is the maximum degree of the graph.

A simple modification of randomised rumour spreading is when the vertices are equipped with a limited amount of memory. In particular, [BEF16; ES08] considered a variant of PUSH-PULL where each vertex stores the neighbours it contacted in the previous three rounds and does not contact them in the next round. The benefit of this modification is the reduction of message complexity from $O(n \log n)$ to $O(n \log \log n)$ for random regular graphs as well as Erdös-Rényi graphs, while still having logarithmic broadcast time. A similar variant was studied in [DFF11] where nodes can remember only the neighbour they contacted in the round before and do not contact it immediately after. It is shown that in Barabási-Albert preferential attachment graphs which model social networks [BA99], the broadcast time of the protocol is sub-logarithmic. They also show that in the original (memory-less) PUSH-PULL finishes in $O(\log n)$ rounds.

The results presented so far are on randomised rumour spreading protocols and related processes, which belong to a more general family of *gossiping* processes: Their distinctive property is that in each time step a vertex initiates a communication with only one neighbour. A trivial lower bound on the broadcast time for gossiping processes is the diameter of the graph G where they run. With respect to the diameter, the best bound for PUSH and related processes is that of [Fei+90]. Censor-Hillel et al. positively answered the natural question of whether a gossiping protocol exists that can broadcast in $O(\operatorname{diam}(G) + \operatorname{poly}\log n)$ rounds [Cen+17]. Note that there is no dependence on the conductance. In their protocol, PUSH-PULL is used as a subroutine to disseminate information in subgraphs of G that have high internal conductance, which is fast due to [Gia11]. Using carefully-designed rules it is also guaranteed that information is disseminated from one subgraph to another. Improving that line of work, Haeupler presented a deterministic algorithm for the problem [Hae15]. Using a distributed minimum spanning tree construction, [GK18] develop a gossiping protocol that uses messages composed of $O(\log n)$ bits and completes the broadcast in $O(\sqrt{\operatorname{diam}(G) \cdot n} \cdot \log n)$ rounds. The disadvantage of these results is that the protocols are no longer state-less and the steps taken in different rounds are not independent. At various stages, the vertices have to keep track which of their neighbours have received certain messages, and the transmitted messages per node in one round can be of linear size in n.

1.3.2 Moving particle processes

Many real-world processes can be modelled using moving and interacting particles or agents. For example, agent-based models for epidemic simulations have become common in recent years due to the availability of large computational resources [MN14]. Prior to that, compartmental models such as SIR were the standard which are not stochastic and often can be studied analytically [KMW27]. Analytic studies for agent-based processes are challenging since the state space of multi-agent systems is very large and complex. Nevertheless, by making simplifying assumptions on the processes and the agents, one can obtain theoretical results. The VISIT-EXCHANGE process is an attempt at such a simplification, where agents perform a simple random walk and execute a simple infectious process without recovery.

To our best knowledge, [ES09] is the only prior work where VISIT-EXCHANGE is mentioned. Their main result is that if each communication in PUSH fails with constant probability, then the asymptotic broadcast time of the process does not change. They also note that for sufficiently dense regular graphs VISIT-EXCHANGE behaves like PUSH with failures, and therefore, by their main result, VISIT-EXCHANGE has the same asymptotic broadcast time as the standard PUSH. We have proved this fact rigorously in Section 3.4, using a non-trivial technique that circumvents the dependencies that would arise in a direct analysis. Our proof also does not rely on the main claim of [ES09] that PUSH with failures is asymptotically equivalent to PUSH. An earlier result by the same authors considered a variant of VISIT-EXCHANGE process where at each step the agents are re-distributed randomly according to stationarity [ELS04]. This modification makes the rounds of the process independent from ane another, like in randomised rumour spreading, and hence the process is significantly easier to analyse.

More commonly, another similar multi-agent process is studied, which we call MEET-EXCHANGE here (as in [GMS19]). The difference of MEET-EXCHANGE and VISIT-EXCHANGE is that vertices do not become informed and agents pass information directly from one to another when they are at the same vertex. In this case, the broadcast time refers to the number of rounds until all agents become informed. The first result for MEET-EXCHANGE on finite graphs is by Dimitriou et al. [DNS06], who consider continuous walks and study the broadcast time of information *among agents*, that is, the time until all agents become informed. The main result in their paper is that the broadcast time is $O(t_{meet} \cdot \log m)$ in expectation, where t_{meet} is the maximum meeting time of two random walks on the underlying graph and m is the number of agents. On some graphs this bound is tight but for expanders and complete graphs tighter results were shown. Cooper et al. considered the MEET-EXCHANGE process on random d-regular graphs [CFR09]. They show that for m agents, the broadcast time is $\Theta(\frac{n \ln m}{m})$. In [GMS19] we have shown that MEET-EXCHANGE is not asymptotically faster than PUSH and VISIT-EXCHANGE for regular graphs of at least logarithmic degree (this result is not included in the thesis).

Of particular interest have been results on MEET-EXCHANGE (or its slight variations) on infinite or finite grids as well as torus graphs of dimension k. Kesten and Sidoravicius studied a continuous variant of MEET-EXCHANGE on infinite grids [KS05; KS08], where initially at each node the number of agents is a Poisson random variable with constant mean. The authors proved a theorem for the shape formed by the contour of informed agents and for the shape of the vertices that have been visited by informed agents in the limit. Roughly, they show that the shape grows linearly with time. We simplify and adopt their technique for (synchronous) VISIT-EXCHANGE on finite grids of any dimension. It is noteworthy that the multi-scale technique by the authors allows to prove tight bounds for the one dimensional grid, that is, for the path and cycle graphs, which are challenging instances due to the bad expansion properties of these graphs. For most other results $k \geqslant 2$ is required (e.g., the Lipschitz net framework in [GS18]).

A few results considered the broadcast time of MEET-EXCHANGE depending on the number of agents in the system. Pettarin et al. showed that for the 2-dimensional grid, $G_{2,n}$, the broadcast time is $\widetilde{\Theta}(n/\sqrt{m})$,[2] w.h.p., for m agents starting from stationarity [Pet+10]. Lam et al. studied the same problem for $k \geqslant 3$ dimensions, and showed that there is a phase transition depending on m: for large m the broadcast time is $\widetilde{\Theta}(n^{1+1/k}/\sqrt{m})$, while for small m it is $\widetilde{\Theta}(n/m)$ [Lam+12]. Furthermore, the authors show that there is no phase

[2]Here, the tilde asymptotic notation hides polylogarithmic factors in n.

transition for lower dimensions and the broadcast time is $\Theta(n/m)$ for $k=1$ and $\Theta(n/\sqrt{m})$ for $k=2$. Huq et al. studied the process for varying number of agents on path and cycle graphs [HP20]. Although they consider the whole possible range of m, the results are only tight up to logarithmic factors.

The *frog model* is another process of a similar flavour. In this model, a number of particles are placed in a graph and initially all particles but one are inactive. The active particles perform a random walk on a graph and they awake inactive particles that they encounter. Earlier results in this model considered an infinite graph such as the Cartesian graph $\mathbb{Z}^d$ and proved a shape theorem on the set of the visited vertices [AMP02; Pop03]. [Her18] also studies the frog model on finite b-ary trees.

Next, we give a brief overview of the literature that studies independent parallel random walks, without the particular task of information dissemination. In this setting, instead of considering a broadcast time, we analyse the cover time of the graph by many walks, that is, the time until each vertex is visited by at least one walk. The study of the cover time by many walks is motivated by the problem of s-t connectivity in a graph: Given a graph and two vertices s, t, the goal is to determine whether there is a path connecting the two. The standard depth-first-search algorithm starting from s is fastest possible to answer but uses large amount of memory. Now consider a random walk starting from s and execute it for long enough to either reach t or assert with high probability that there is no path connecting s and t. This algorithm uses very little amount of memory (to count the steps of the random walk) but is slow as it can take up to cubic time in n in general. Using many walks instead of one allows for a space-time trade-off, as proposed by Broder et al. [Bro+94]. To prove their main result, they showed that m walks, starting from stationarity, cover a graph in $O(|E|^2 \log^3 n/m^2)$ rounds.

Alon et al. also studied the cover time of many random walks [Alo+11]. They defined the *speed-up* for many random walks on a graph as the ratio of the cover time using m walks and the cover time using 1 walk. The m walks here are assumed to start from a single vertex (Note, that in VISIT-EXCHANGE we assume that the walks start from stationarity.) They analysed the speed-up for different classes of graphs, noticing that depending on the graph structure different speed-ups are possible, ranging from logarithmic to exponential. Later [ES11] improved and extended some of the bounds from [Alo+11]. Although the processes considered in these two papers are different from the ones in the thesis, some of our analyses are inspired by them. Moreover, these results give insight on how multiple walks spread in a graph.

The random walk literature is vast. Here we do not review results for a single random walk, however, in our proofs we use many standard results. These and other results on random walks can be found in [Lov93; LP17; AF02].

1.4 Thesis outline

The rest of the thesis is organised as follows. In Chapter 2 we present definitions and standard results in graph theory and probability. We also formally define the processes studied in this work. In Chapter 3 we present results that compare VISIT-EXCHANGE to randomised rumour spreading, showing that in dense regular graphs the broadcast times of the two processes are the same. Examples of sparse or non-regular graphs where the processes are significantly different are also given. Chapter 4 contains bounds on the broadcast time of VISIT-EXCHANGE in terms of the graph diameter, average and minimum degrees. Chapter 5 presents our results on expander graphs, which are also used to bound the broadcast time of VISIT-EXCHANGE on random regular graphs. Chapters 6 and 7 tightly analyse VISIT-EXCHANGE on balanced trees and grid graphs, respectively. Chapter 8 contains a summary of our results and some future directions on this line of work.

In Appendix A we evaluate the broadcasting protocols empirically.

The results from Chapter 3, except those in Section 3.2.4, were published at the 38th ACM Symposium on Principles of Distributed Computing (PODC'2019) [GMS19]. Most of the material from Chapters 4 to 7 was published at the 34th International Symposium on Distributed Computing (DISC'2020) [GSS20].[3]

Some passages have been quoted verbatim from these papers.

[3]The video presentation at DISC'2020 can be found at https://www.youtube.com/watch?v=3eKipiyDmnY.

Chapter 2

Preliminaries

This chapter provides the mathematical background required for the rest of the thesis, including formal definitions of the processes we study. We state standard results in graph theory, probability theory and random walks. For some claims, the precise version we need is not present in the literature. In such cases, we present their proofs, which, however, do not contain significant novelty.

2.1 Useful inequalities

The following lemma lists several inequalities frequently used throughout the thesis. They can be derived using elementary methods.

Lemma 2.1.1.

(a) *For any* $x \in \mathbb{R}$, $1 - x \leqslant e^{-x}$.

(b) *For any* $x \in (0,1)$ *and* $t \geqslant 0$, $(1-x)^t \leqslant \frac{1}{1+tx}$.

(c) *For any* $x \in (0,1)$ *and integers* m, n *with* $1 \leqslant m \leqslant n$, $x^m + x^n \leqslant x^{m-1} + x^{n+1}$.

(d) *For* $0 \leqslant x_1, \dots, x_n \leqslant 1$, $\prod_{i=1}^{n}(1-x_i) \geqslant 1 - \sum_{i=1}^{n} x_i$ *(Weierstrass' inequality).*

2.2 Graph theoretic preliminaries

Throughout the thesis we use $G = (V, E)$ to denote a connected, undirected, unweighted simple graph defined on the vertex set V and the edge set E. Typically, we use G to denote the graph on which the information dissemination process is being executed. We use $n = |V|$ to denote the number of vertices of the graph. For a vertex $u \in V$, we define $\Gamma(u)$ as the neighbourhood of u, that is the set of all vertices which have a common edge with u. The degree of vertex u is the number of neighbours of u, denoted by $\deg(u) = |\Gamma(u)|$.

For a subset of vertices $S \subseteq V$, the neighbourhood of S is denoted by ∂S and contains the vertices that are not in S but have a neighbour in S. The graph $G(S) = (S, E(S))$ with a vertex set S and edges $(u, v) \in E$ such that $u, v \in S$ is called the *induced* graph by the set S.

The minimum, maximum and average degrees of the graph are denoted by $d_{\min}, d_{\max}$ and d_{avg}, respectively. A graph is said to be *d-regular* if all vertices have degree d, i.e., $d_{\min} = d_{\max} = d$.

A *path* of length l in a graph is a sequence of vertices $\langle u_0, u_1, \dots, u_l \rangle$ such that for any $i \in \{1, \dots, l\}$, $(u_{i-1}, u_i) \in E$. For two vertices u and v, their distance is the length of the shortest path with u and v as its endpoints.

2.3 Probability

In this section we introduce some basic notions from probability theory that are used in the thesis. [MU17; GS01] cover the necessary material in more depth.

In the thesis we only consider discrete probability spaces $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the sample space, $\mathcal{F}$ is the set of allowable events (subsets of Ω), and $\mathbb{P}$ is the probability function. For a random variable X, we denote its expectation by $\mathbb{E}[X]$ and variance by $\mathbf{Var}[X]$. The probability of an event $\mathcal{E} \subseteq \mathcal{F}$ is denoted by $\mathbb{P}[\mathcal{E}]$. Throughout the thesis we will say that an event $\mathcal{E}$ holds *with high probability*, or w.h.p. in short, if $\mathbb{P}[\mathcal{E}] \geqslant 1 - n^{-c}$ for an arbitrary constant $c > 0$, where n will be the number of vertices of the graph we consider. Similarly, an event $\mathcal{E}$ holds with constant probability, or w.c.p., if $\mathbb{P}[\mathcal{E}]$ is lower bounded by positive constant.

2.3.1 Commonly used distributions

Next, we describe some distributions often used in the thesis. For a real $p \in [0, 1]$, a random variable X has the Bernoulli distribution with parameter p, or $X \sim \mathrm{Ber}(p)$, if $\mathbb{P}[X = 1] = p$ and $\mathbb{P}[X = 0] = 1 - p$. For an event $\mathcal{E}$, the random variable $\mathbb{1}_{\mathcal{E}}$ is a Bernoulli random variable that is equal to 1, when the event $\mathcal{E}$ holds. It is also called the indicator random variable of $\mathcal{E}$.

For $n \in \mathbb{N}$ and $p \in [0, 1]$, let $X_1, \dots, X_n \sim \mathrm{Ber}(p)$ be independent random variables and denote $X = \sum_{i=1}^{n} X_i$. Then, X is a binomial random variable, or $X \sim \mathrm{Bin}(n, p)$. The binomial random variable corresponds, say, to the number of heads one gets when tossing a coin n times, when each time it lands on heads with probability p. For $k = 0, \dots, n$, $\mathbb{P}[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}$.

For $p \in [0, 1]$, a geometric random variable $X \sim \mathrm{Geom}(p)$ corresponds to the number of tosses of a coin until the first heads appears, if in each toss it lands on heads with probability p. In other words, for $k \in \mathbb{N}$, $\mathbb{P}[X = k] = (1 - p)^{k-1} p$.

2.3.2 Probabilistic inequalities

Theorem 2.3.1 (Union bound. See, e.g., [MU17]). *For any finite or countably infinite sequence of events $\mathcal{E}_1, \mathcal{E}_2, \dots$,*

$$\mathbb{P}\left[\bigcup_{i \geqslant 1} \mathcal{E}_i\right] \leqslant \sum_{i \geqslant 1} \mathbb{P}[\mathcal{E}_i].$$

Theorem 2.3.2 (See, e.g., [MU17]). *For a non-negative discrete random variable X,*

$$\mathbb{E}[X] = \sum_{k=1}^{+\infty} \mathbb{P}[X \geqslant k].$$

The aim of concentration inequalities, presented next, is to bound how much a random variable deviates from its mean.

Theorem 2.3.3 (Markov's inequality. See, e.g. [MU17]). *For a positive random variable X and any $a > 0$,*

$$\mathbb{P}[X \geqslant a] \leqslant \frac{\mathbb{E}[X]}{a}.$$

Theorem 2.3.4 (Chernoff bounds for independent Bernoulli random variables. See, e.g., [MU17]). *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables with $\mathbb{P}[X_i = 1] = p_i$. Let $X = \sum_{i=1}^{n} X_i$ and denote $\mu = \mathbb{E}[X]$. Then,*

(a) *For any $\delta > 0$,*

$$\mathbb{P}[X \geqslant (1+\delta)\mu] \leqslant \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu},$$

(b) *For $0 < \delta < 1$,*

$$\mathbb{P}[X \geqslant (1+\delta)\mu] \leqslant e^{-\mu\delta^2/3},$$

(c) *For $c \geqslant 6$,*

$$\mathbb{P}[X \geqslant c \cdot \mu] \leqslant 2^{-c\mu},$$

(d) *For $0 < \delta < 1$,*

$$\mathbb{P}[X \leqslant (1-\delta)\mu] \leqslant \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu},$$

(e) *For $0 < \delta \leqslant 1$,*

$$\mathbb{P}[X \leqslant (1-\delta)\mu] \leqslant e^{-\mu\delta^2/2}.$$

In the dissertation, most often the simpler versions of Chernoff bounds in Theorem 2.3.4, (b), (c) and (e), are sufficient and we do not explicitly reference to them in the text. Next we prove two less standard versions of Chernoff bounds.

Lemma 2.3.5 (Chernoff bound for independent geometric random variables). *Let $F_1, \ldots, F_n$ be independent and identical geometrically distributed random variables with parameter p. Let $F = \sum_{i=1}^{n} F_i$. Then for any $k \geqslant 2 \cdot \mathbb{E}[F] = 2n/p$,*

$$\mathbb{P}[F \geqslant k] \leqslant \exp\left(-\frac{kp}{8}\right).$$

Proof. Follows trivially from [Jan17, Theorem 2.1]. □

Lemma 2.3.6. *If X is a sum of independent Bernoulli random variables, then for $b > 1$ and $x > b \cdot \mathbb{E}[X]$, $\mathbb{P}[X \geqslant x] \leqslant (b/e)^{-x}$.*

Proof. Let $\delta = \frac{x}{\mathbb{E}[X]} - 1$, then, by Theorem 2.3.4(a),

$$\begin{aligned}\mathbb{P}[X \geqslant x] &= \mathbb{P}[X \geqslant (1+\delta)\cdot\mathbb{E}[X]] \\ &\leqslant \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mathbb{E}[X]} \leqslant \left(\frac{e}{1+\delta}\right)^{(1+\delta)\mathbb{E}[X]} \\ &= \left(\frac{e\cdot\mathbb{E}[X]}{x}\right)^{x} \leqslant \left(\frac{b}{e}\right)^{-x}.\end{aligned}$$

□

Theorem 2.3.7 ([CL06, Theorem 3.7]). *Let X_i be independent random variables with $X_i \geqslant -M$, for $1 \leqslant i \leqslant n$. Let $X = \sum_{i=1}^{n} X_i$ and $||X|| = \sqrt{\sum_{i=1}^{n}\mathbb{E}[X_i^2]}$. Then, for any $\lambda \geqslant 0$,*

$$\mathbb{P}[X \leqslant \mathbb{E}[X] - \lambda] \leqslant \exp\left(-\frac{\lambda^2}{2(||X||^2 + M\lambda/3)}\right).$$

Theorem 2.3.8 (Method of bounded differences, [DP09, Corollary 5.2]). *Let $\mathbf{X} = (X_1, \ldots, X_n)$ be a vector of independent random variables. Suppose X_i is defined on the set $\mathcal{X}_i$ and $f(\mathbf{x})$ is a function defined on the space $\prod_{i=1}^{n}\mathcal{X}_i$. If there are constants d_i such that $|f(\mathbf{x}) - f(\mathbf{y})| \leqslant d_i$ when $\mathbf{x}$ and $\mathbf{y}$ differ only in the ith coordinate, then for any $\lambda \geqslant 0$,*

$$\mathbb{P}[|f(\mathbf{X}) - \mathbb{E}[f(\mathbf{X})]| \geqslant \lambda] \leqslant 2\exp\left(-\frac{2\lambda^2}{\sum_{i=1}^{n} d_i^2}\right).$$

The Chernoff bounds above also hold for a set of *negatively associated* random variables [DR96]. A set of random variables $X_1, \ldots, X_n$ is said to be negatively associated if for any two disjoint index sets $I, J \subset \{1, \ldots, n\}$, and two functions f, g both non-increasing or both non-decreasing

$$\mathbb{E}[f(\{X_i\}_{i\in I})\cdot g(\{X_j\}_{j\in J})] \leqslant \mathbb{E}[f(\{X_i\}_{i\in I})]\cdot\mathbb{E}[g(\{X_j\}_{j\in J})].$$

Intuitively, negative association means that if a set of these variables has large values then on another disjoint set should have lower values. For example, consider the balls-and-bins process, where n balls are uniformly randomly allocated to n bins. If X_i is the number of balls in the ith bin, then the variables X_i are negatively associated. Verifying negative association directly may be laborious but one can use the following closure properties for negative association:

(a) If $X_1, \ldots, X_n$ are negatively associated and $Y_1, \ldots Y_m$ are negatively associated, and $\{X_i\}_{1\leqslant i\leqslant n}$ and $\{Y_i\}_{1\leqslant i\leqslant m}$ are independent, then the union $X_1, \ldots, X_n, Y_1, \ldots, Y_m$ are negatively associated.

(b) Let $f_1, \ldots f_k$ be non-increasing or non-decreasing real functions, defined on R^n, such that the value of each function depends only on a subset of the variables, and these subsets are disjoint. Then, if $X_1, \ldots, X_n$ are negatively associated and $\mathbf{X} = (X_1, \ldots, X_n)$, we have $f_1(\mathbf{X}), \ldots, f_k(\mathbf{X})$ are negatively associated.

See [Waj17, Section 4] and [DR96, Section 2.2] for the details on negative association, and the worked out example on the balls-and-bins process.

The following lemma allows us to apply concentration bounds on dependent random variables.

Lemma 2.3.9. *Let $Z_1, \dots, Z_n$ be (dependent) random variables and $Z'_1, \dots, Z'_n$ be mutually independent random variables such that for any $1 \leqslant i \leqslant n$ and $z_1, \dots, z_i \in \mathbb{R}$,*

$$\mathbb{P}\left[Z_i \leqslant z \mid Z_1 = z_1, \dots, Z_{i-1} = z_{i-1}\right] \geqslant \mathbb{P}\left[Z'_i \leqslant z_i\right].$$

Then, for any $b \in \mathbb{R}$,

$$\mathbb{P}\left[\sum_{i=1}^{n} Z_i \leqslant b\right] \geqslant \mathbb{P}\left[\sum_{i=1}^{n} Z'_i \leqslant b\right].$$

Proof. Let $S_k = \sum_{i=1}^{k} Z_k$ and $S'_k = \sum_{i=1}^{k} Z'_k$. We can prove using induction on n. The statement is trivial for $n = 1$, so suppose that $\mathbb{P}\left[S_{k-1} \leqslant b\right] \geqslant \mathbb{P}\left[S'_{k-1} \leqslant b\right]$.

$$\begin{aligned}
\mathbb{P}\left[S_k \leqslant b\right] &= \mathbb{E}\left[\mathbb{1}_{S_k \leqslant b}\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{Z_k \leqslant b - S_{k-1}} \mid S_{k-1}\right]\right], \quad \text{by the tower property,} \\
&\geqslant \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{Z'_k \leqslant b - S_{k-1}} \mid S_{k-1}\right]\right], \quad \text{by the lemma condition,} \\
&= \mathbb{E}\left[\mathbb{1}_{Z'_k \leqslant b - S_{k-1}}\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{S_{k-1} \leqslant b - Z'_k} \mid Z'_k\right]\right] \\
&\geqslant \mathbb{E}\left[\mathbb{1}_{S'_{k-1} \leqslant b - Z'_k}\right], \quad \text{by the inductive hypothesis,} \\
&= \mathbb{P}\left[S'_k \leqslant b\right].
\end{aligned}$$

□

2.3.3 Couplings

For two random variables X and Y on real numbers, we say that X *stochastically dominates* Y, or $X \succcurlyeq Y$ if for any $r \in R$,

$$\mathbb{P}\left[X \leqslant r\right] \leqslant \mathbb{P}\left[Y \leqslant r\right].$$

We use $X \sim Y$ to denote the fact that X and Y have the same distribution. A *coupling* of random variables X and Y is a random variable $(\hat{X}, \hat{Y})$ such that $\hat{X} \sim X$ and $\hat{Y} \sim Y$. Normally, the variables $\hat{X}$ and $\hat{Y}$ are dependent random variables that allow us to deduce certain properties for the random variables X and Y. The following lemma is a powerful tool that allows one to compare the distributions of random variables X and Y, by constructing their coupling.

Lemma 2.3.10 ([GS01]). *A random variables X stochastically dominates another random variable Y, i.e., $X \succcurlyeq Y$ if and only if there is a coupling $(\hat{X}, \hat{Y})$ such that $\mathbb{P}\left[\hat{X} \geqslant \hat{Y}\right] = 1$.*

In the thesis we use couplings between the stochastic processes we study, in order to compare them. In that case, we represent the processes using a set of decisions they make. Then we let these decisions to be made jointly, and argue about the broadcast times of the processes under the coupling.

2.4 Random walks

A random walk is a stochastic process defined on the vertices of a graph $G = (V, E)$. It starts at some source vertex $s \in V$. At each next step the walk moves to a randomly selected neighbour of the current vertex. In this thesis we work with discrete random walks. Formally, we denote the position of a walk by a random variable $X(t) \in V$ for an integer round $t \geqslant 0$. Then $X(0) = s$ and for $t \geqslant 0$,

$$\mathbb{P}\left[X(t+1) = v \mid X(t) = u\right] = \begin{cases} 1/\deg(u), & \text{if } (u,v) \in E; \\ 0, & \text{otherwise,} \end{cases}$$

which is called the transition probability from vertex u to v, also denoted by $p_{u,v}$. $X(t)$ is referred as a *simple discrete random walk* on G. In parts of the analysis we will also consider an α-lazy walk for a constant $\alpha \in (0,1)$, in which case the walk stays at its current position with probability α and moves to a random neighbour with probability $1-\alpha$.

The transition probabilities $p_{u,v}$, when put in a matrix, form the transition matrix P of a simple random walk. Let the row vector $\mathbf{p}_t$ be the distribution of the walk $X(t)$ at round t, i.e., $\mathbf{p}_t(u) = \mathbb{P}\left[X(t) = u\right]$. Then,

$$\mathbf{p}_{t+1} = \mathbf{p}_t \cdot P = \mathbf{p}_0 \cdot P^t.$$

The distribution π over the vertices V for which $\pi = \pi \cdot P$ is called the *stationary distribution* of the walk. It is easy to see that $\pi(u) = \deg(u)/(2|E|)$. If G is not bipartite, then $\mathbf{p}_t$ converges to π from any starting distribution $\mathbf{p}_0$. The transition matrix of an α-lazy walk is given by $P_\alpha = \alpha I_n + (1-\alpha)P$, where I_n is the identity matrix of size n. For a lazy random walk, $\mathbf{p}_t$ always converges to the stationary distribution for bipartite graphs as well.

The time until the distribution $\mathbf{p}_t$ is sufficiently close to the stationary distribution π, starting from any vertex, is determined by the *mixing time*. We define and use the uniform mixing time t^{∞}_{mix}, for which we have that if $t \geqslant t^{\infty}_{\text{mix}}$,

$$\max_{u \in V} |\mathbf{p}_t(u) - \pi(u)| \leqslant \pi(u)/2.$$

For a simple random walk on a non-bipartite, connected graph t^{∞}_{mix} is finite. For lazy walks, the uniform mixing time is finite for bipartite graphs as well [LP17]. Thus, when using t^{∞}_{mix}, we switch to lazy walks and then show that the claim also holds for non-lazy walks. The mixing time can be bounded from above using spectral properties of the graph G.

In a connected undirected graph G, the random walk $X(t)$ will eventually visit all vertices of G, say in round $\tau_{\text{cov}}(s)$ when starting from vertex s. Then the *cover time* of a random walk is defined as

$$t_{\text{cov}} = \max_{s \in V} \mathbb{E}\left[\tau_{\text{cov}}(s)\right].$$

2.5 Protocol descriptions and notation

First, we formally define the VISIT-EXCHANGE process, the main protocol studied in this thesis. Consider a graph $G = (V, E)$ and let A be a set of mobile agents. Initially, in round $t = 0$, a piece of information is placed on a source vertex $s \in V$. Thus, s is *informed* in round 0 while all other vertices are not. Also, in round 0, each agent $g \in A$ is placed on a vertex of G chosen randomly according to the stationary distribution of the graph, i.e., the probability that g is at vertex u is $\deg(u)/(2|E|)$. Starting from round $t = 1$, the agents perform independent simple discrete random walks, taking steps in parallel. An agent learns the information the first time it visits some informed vertex (the vertex may

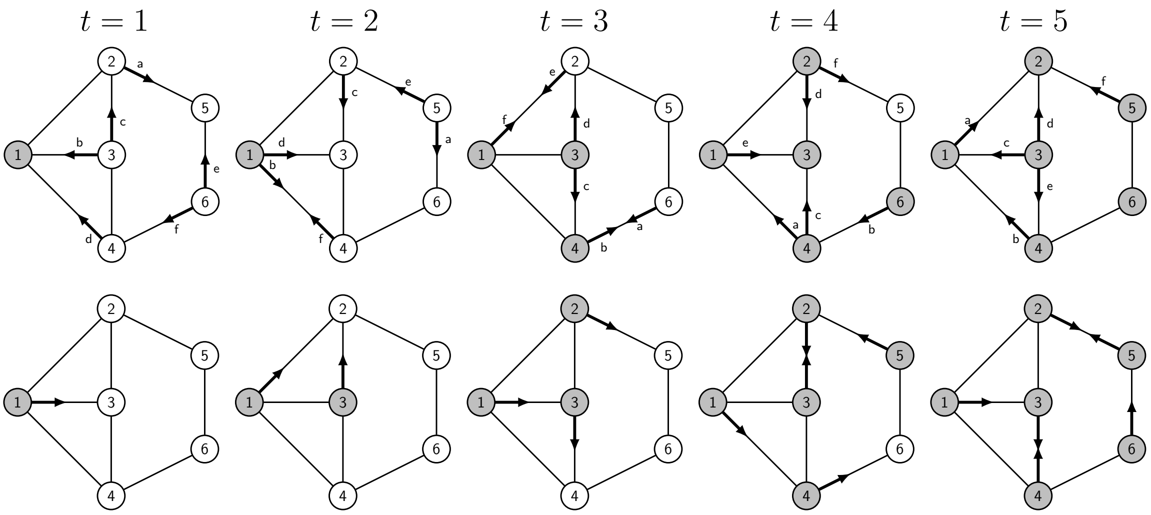

Figure 2.1 The first 5 rounds of example executions of VISIT-EXCHANGE (top row) and PUSH (bottom row). The arrow(s) departing from each vertex indicate the neighbour(s) chosen in the corresponding round by the vertex. The agents in VISIT-EXCHANGE are shown using letters. Informed vertices are shaded, and agents leaving these vertices are also informed.

have become informed on the same or any previous round). From that point on, every vertex the agent visits becomes informed.

The PUSH process also starts with an informed vertex s in round $t = 0$. In any round $t \geqslant 1$, if a vertex u is informed at the start of round t, then it picks a neighbour $v \in \Gamma(u)$ uniformly randomly and sends the message to v. Thus, at the end of the round v becomes informed.

Unlike PUSH, in PUSH-PULL uninformed vertices also initiate communication with their neighbours. The protocol proceeds as follows. In every round $t \geqslant 1$, every vertex u picks a uniformly random neighbour $v \in \Gamma(u)$. If either u or v are informed at the start of round t, then both of them become informed at the end of the round. Collectively, PUSH and PUSH-PULL are referred to as *randomised rumour spreading* algorithms due to the seminal paper by Karp et al. [Kar+00]. Fig. 2.1 illustrates the execution of PUSH and VISIT-EXCHANGE for 5 rounds.

Note that all three protocols are stateless, i.e., the nodes are not required to use any additional memory for the dissemination apart from a single bit indicating whether they are informed or not. As it was mentioned in the introduction, this property is difficult to maintain in deterministic protocols (except, perhaps, in simple flooding).

Since we consider connected graphs, all three protocols will eventually inform the whole graph. The *broadcast time from vertex s* is the number of rounds until every vertex is informed, if initially the source vertex s is informed. For the graph G and vertex s, we denote the broadcast time for VISIT-EXCHANGE, PUSH and PUSH-PULL by $T_{\text{visitx}}(G, s)$, $T_{\text{push}}(G, s)$ and $T_{\text{ppull}}(G, s)$, respectively. Where the graph or the source vertex are implicitly clear we may omit them from this notation.

In this thesis, we only consider the setting when the number of agents is linear in the number of vertices of the graph, i.e., $|A| = \Theta(|V|)$. We denote by α the ratio $|A|/|V| = \Theta(1)$. This assumption about the number of agents allows us to make a fair comparison with randomised rumour spreading in the following sense. In each round of randomised rumour spreading exactly $|V|$ edges are used for communication. In VISIT-EXCHANGE, assuming that each step of an agent costs one unit of communication, in one round the amount of communication is $|A|$. Thus, when $\alpha = 1$, VISIT-EXCHANGE and randomised rumour spreading use the same amount of communication. Our theoretical results also hold for an arbitrary constant value of α, at the cost of constants hidden under the asymptotic notation in the bounds for the broadcast time. We do not study the dependency of broadcast times from α theoretically, but in Appendix A we do a preliminary study of how modifying α may impact the VISIT-EXCHANGE process.

2.5.1 Symmetry of protocols

In general, the source vertex of broadcasting protocols may have an impact on the broadcast time. Consider a graph where $(n-1)$ vertices are fully connected and the remaining vertex u has a single neighbour. If u is the source vertex, then $T_{\text{push}} = O(\log n)$, w.h.p., while for any other source vertex $T_{\text{push}} = \Omega(n \log n)$, w.c.p., since u becomes informed very late. For PUSH-PULL and VISIT-EXCHANGE however, the following results allow us to assume that the source vertex is an arbitrary vertex in the graph.

Lemma 2.5.1. *For vertices u and v of a connected graph $G = (V, E)$, let $T_{u,v}$ be the number of rounds of* VISIT-EXCHANGE *until v is informed when the information originates at u. Then $T_{u,v} \sim T_{v,u}$, i.e., the two random variables have the same distribution.*

Proof. For round r, let Ω_r be the set of all possible executions in the first r rounds, i.e., every element $\omega \in \Omega_r$ is composed of the paths that are taken by each of the agents. Let $p(\omega)$ be the probability associated with the outcome ω. For an execution ω, let ω^* be the *reversal* of ω: If $X_g(t)$ is the walk taken by agent g then $X_g(t)(\omega^*) = X_g(r-t)(\omega)$ for $t \in \{0, \dots, r\}$. Clearly $\omega^* \in \Omega_r$.

Since g starts its walk from the stationary distribution π, for any path $\langle u_0, \dots, u_r \rangle$,

$$\mathbb{P}\left[X_g(t) = u_t \text{ for all } t\right] = \pi(u_0) \prod_{t=0}^{r-1} \frac{1}{\deg(u_t)} = \frac{1}{2|E|} \prod_{t=1}^{r-1} \frac{1}{\deg(u_t)} = \mathbb{P}\left[X_g(t) = u_{r-t} \text{ for all } t\right],$$

where the last equality holds since the product on its left hand side does not depend on the direction of the path. Applying this equality for every agent g, we get that $p(\omega) = p(\omega^*)$.

Additionally, notice that if in execution ω vertex v gets informed when u is the source, then u gets informed in ω^* when v is the source. Combining the two previous facts gives

$$\mathbb{P}\left[T_{u,v} \leqslant r\right] = \sum_{\omega \in \{T_{u,v} \leqslant r\}} p(\omega) = \sum_{\omega^* \in \{T_{v,u} \leqslant r\}} p(\omega^*) = \mathbb{P}\left[T_{v,u} \leqslant r\right]. \qquad \square$$

The following lemma is proved similarly.

Lemma 2.5.2 ([CLP10, Lemma 3.3])**.** *For vertices u and v of a connected graph $G = (V, E)$, let $T_{u,v}$ be the number of rounds of* PUSH-PULL *until v is informed when the information originates at u. Then for any round r, $\mathbb{P}\left[T_{u,v} \leqslant r\right] = \mathbb{P}\left[T_{v,u} \leqslant r\right]$.*

Corollary 2.5.3. *Consider either the* VISIT-EXCHANGE *or* PUSH-PULL *processes. Let u and v be any two vertices of a connected graph $G = (V, E)$, and let T_u denote the broadcast time when the information originates at u. If $\mathbb{P}\left[T_v \leqslant r\right] = 1 - \delta$ then $\mathbb{P}\left[T_u \leqslant 2r\right] \geqslant 1 - 2\delta$.*

Proof. We have

$$\begin{aligned}
\mathbb{P}\left[T_u \leqslant 2r\right] &\geqslant \mathbb{P}\left[T_u \leqslant 2r \mid T_{u,v} \leqslant r\right] \cdot \mathbb{P}\left[T_{u,v} \leqslant r\right] \\
&\geqslant \mathbb{P}\left[T_v \leqslant r\right] \cdot \mathbb{P}\left[T_{u,v} \leqslant r\right] \\
&\geqslant \mathbb{P}\left[T_v \leqslant r\right] \cdot \mathbb{P}\left[T_{v,u} \leqslant r\right], \quad \text{by Lemmas 2.5.1 and 2.5.2}, \\
&\geqslant \mathbb{P}\left[T_v \leqslant r\right]^2, \quad \text{since } T_{v,u} \leqslant T_v, \\
&\geqslant 1 - 2\delta, \text{by Weierstrass' inequality.}
\end{aligned}$$

□

Chapter 3

Comparison between Visit-Exchange and randomised rumour spreading

3.1 Introduction

As we mentioned in the introductory chapter, randomised rumour spreading algorithms have been studied extensively. In particular, tight bounds on their broadcast time are known in terms of various graph parameters, such as the conductance, the diameter, the degree distribution. In this chapter we consider VISIT-EXCHANGE with a linear number of agents and compare it to randomised rumour spreading. We answer two questions in this chapter. First, is VISIT-EXCHANGE always dominated by the other protocols or vice versa? We prove that this is not the case, and, in general, there are graphs where PUSH is faster than VISIT-EXCHANGE and VISIT-EXCHANGE is faster than PUSH-PULL. Given this, the second question then is whether there are graphs for which the protocols are comparable. The first step in this direction is our proof that in sufficiently dense regular graphs all three protocols have asymptotically the same broadcast time. Surprisingly, this is not the case in sparse regular graphs, as seen later in Section 3.2.4.

The three example graphs depicted in Fig. 3.1 show that the three processes can have significantly different broadcast times. Clearly, the PUSH process is never faster than PUSH-PULL and the three examples cover all possible orderings of VISIT-EXCHANGE, PUSH and PUSH-PULL in terms of their broadcast times.

The star graph in Fig. 3.1(a) is an example where PUSH is known to take $\Omega(n \log n)$ rounds, w.h.p., as the center must contact all leaves. The broadcast time of VISIT-EXCHANGE is $O(\log n)$ in the star graph, w.h.p. In the star, PUSH-PULL is also (extremely) fast. The next example, the double-star in Fig. 3.1(b), is a graph where PUSH-PULL (and thus also PUSH) is slow, whereas VISIT-EXCHANGE is still fast. This demonstrates the advantage of the fairness property of VISIT-EXCHANGE, that all edges are used at the same rate in the process. This is not the case in PUSH-PULL, which selects the edge between the two star centers only with probability $O(1/n)$ in each round. As a result the expected broadcast time of the protocol is $\Omega(n)$. In VISIT-EXCHANGE, on the other hand, the probability that some agent crosses any edge in a round is always constant in any graph, resulting in a logarithmic broadcast time in this case. Intuitively, this means that in graphs that have at least two hubs that are not well connected, VISIT-EXCHANGE is faster than the other processes.

Fig. 3.1(c) illustrates an example where PUSH and PUSH-PULL have an advantage over

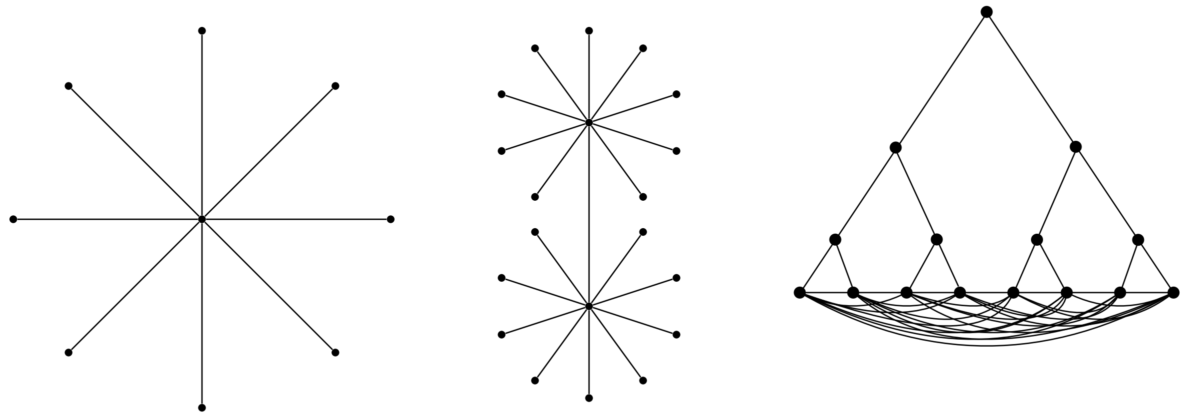

(a) (b) (c)

Figure 3.1 **(a)** Star S_n, on which VISIT-EXCHANGE is faster than PUSH but slower than PUSH-PULL: $\mathbb{E}[T_{\text{push}}] = \Omega(n \log n)$, $T_{\text{ppull}} = O(1)$ and $T_{\text{visitx}} = O(\log n)$, w.h.p. **(b)** Double-star S_n^2, on which VISIT-EXCHANGE is the fastest process: $\mathbb{E}[T_{\text{ppull}}] = \Omega(n)$, and $T_{\text{visitx}} = O(\log n)$ w.h.p. **(c)** Heavy binary tree B_n (leaves are connected to a clique), on which VISIT-EXCHANGE is the slowest process: $T_{\text{push}} = O(\log n)$ w.h.p., $\mathbb{E}[T_{\text{visitx}}] = \Omega(n)$.

VISIT-EXCHANGE. Here PUSH (and thus PUSH-PULL) has logarithmic broadcast time. However, for VISIT-EXCHANGE, at least linear time is needed. This is because the volume of the graph is concentrated on the leaves and it is likely that all agents are on the leaves at the start of the process, and then it takes linear number of rounds before the first walk reaches the root.

These results, which are proved rigorously later, suggest that in certain settings, agent-based information dissemination, separately or in combination with PUSH-PULL, may significantly improve the broadcast time. We stress that, even though the examples presented may seem contrived, they are intentionally simple to demonstrate the principle reasons that make the protocols perform differently, and we expect that similar result can be observed on a wide range of networks. In particular, the observations for the double-star example of Fig. 3.1(b) extend to more general tree-like topologies with high-degree internal nodes. Indeed, we will see in Chapter 6 theoretically and in Appendix A via experiments, that in balanced trees, VISIT-EXCHANGE becomes faster than PUSH-PULL as the branching of the tree increases. Additionally, Theorem 4.1.2 implies that if G is a tree then $T_{\text{visitx}} = \widetilde{O}(\text{diam}(G))$, w.h.p., while a similar bound for PUSH-PULL contains an additional factor of $d_{\max}$, the largest degree of the tree.

All examples of this chapter that we have discussed so far, involve highly non-regular graphs. Our main technical result concerns regular graphs, and can be stated somewhat informally as follows. (The formal, stronger statements are presented in Sections 3.3 and 3.4.)

Theorem 3.1.1. *For any d-regular graph on n vertices, where $d = \Omega(\log n)$, and any source vertex, the broadcast times of* PUSH *and* VISIT-EXCHANGE *with $\Theta(n)$ agents are asymptotically the same both in expectation and w.h.p., modulo constant multiplicative factors.*

Recall that PUSH and PUSH-PULL have asymptotically the same broadcast times on regular graphs [GNW16], so the result also applies for PUSH-PULL. Note also that their broadcast times on d-regular graphs can vary from logarithmic, e.g., in random d-regular graphs, to polynomial in n, e.g., in a path of d-cliques where the broadcast time is $\Omega(n)$. Since PUSH has been studied extensively, especially for regular graphs, this theorem immediately implies a variety of bounds for VISIT-EXCHANGE when G satisfies the requirements of the theorem. Namely, [Fei+90] implies that $T_{\text{visitx}}(G) = O(d \cdot (\text{diam}(G) + \log n))$, w.h.p., and by [Chi+18] we have that for a graph with conductance ϕ, $T_{\text{visitx}} = O(\log n/\phi)$, w.h.p. The latter result also implies that if G is a sufficiently dense expander, then $T_{\text{visitx}} = O(\log n)$, w.h.p. Sparse expanders are studied in Chapter 5.

The result that on regular graphs VISIT-EXCHANGE and PUSH have the same asymptotic runtime is not too surprising. Since the number of agents is linear in n, in any fixed round an informed vertex has a constant number of agents in expectation. It implies that it randomly picks a constant number of its neighbours in expectation and informs them. In PUSH too, every vertex contacts a constant number of agents in every round (one, to be precise), hence it is expected that any progress that a message makes along a path in PUSH will also happen in VISIT-EXCHANGE. The argument is not easy to formalise as the propagation of information depends on the location of the agents in each round, and thus, there are dependencies between rounds of VISIT-EXCHANGE. This is not the case in PUSH, where the neighbours chosen in each round are independent from previous rounds. The proof of Theorem 3.1.1 uses the fact that $d = \Omega(\log n)$ to remove some of the dependencies that occur in VISIT-EXCHANGE. We use a coupling argument which relates the random choices of vertices in PUSH, with the random walks in VISIT-EXCHANGE. Roughly speaking, for each node u, we consider the list of neighbours that u samples in PUSH, and the list of neighbours to which informed agents move to in their next step after visiting u in VISIT-EXCHANGE. Our coupling just sets the two lists to be identical for each u. Even though the coupling is straightforward, its analysis is not. On the one direction of the proof, showing that the broadcast time of PUSH is dominated by the broadcast time of VISIT-EXCHANGE, the main step is to bound the congestion, i.e., the maximum number of agents encountered along any path through which information travels.

The proof of the reverse direction is significantly simpler. We focus only on the fastest path through which information reaches each node in PUSH, and show that VISIT-EXCHANGE makes progress through the same path equally as fast. We use a slightly different coupling, and decide the agent destinations based on the PUSH process only every other round. This allows us to use the independence of the agents on the "non-coupled" rounds of VISIT-EXCHANGE, and argue that a constant number of agents arrive at each vertex at every other round, independently of the past. This proof implies that VISIT-EXCHANGE resembles a PUSH process with failures of constant probability, as noted in [ES09].

Given our earlier intuition about the equivalence of VISIT-EXCHANGE and PUSH on regular graphs, in terms of their broadcast times, a natural question is whether the condition $d = \Omega(\log n)$ is necessary in Theorem 3.1.1. Rather surprisingly, the answer to this question is negative.

Theorem 3.1.2. *There is a regular graph G, such that $\mathbb{E}[T_{\text{visitx}}(G)] = \omega(\mathbb{E}[T_{\text{push}}(G)])$.*

The example graph G we construct is 3-regular and has diameter $\Theta(\log n)$. Thus, by [Fei+90], $T_{\text{push}} = O(\text{diam}(G)) = O(\log n)$, w.h.p. We show that $T_{\text{visitx}} = \Omega(\log^2 n/\log\log n)$ in expectation.

We give the rough description of G here, also illustrated in Fig. 3.2. We start with a 3-regular graph R with $\Theta(n)$ vertices and diameter $\Theta(\log n)$ (e.g., a 3-regular expander).

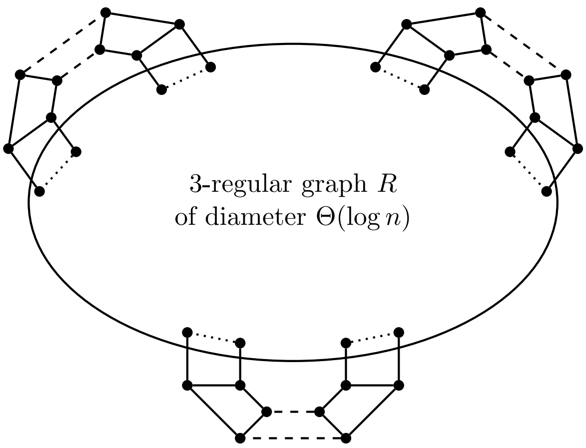

Figure 3.2 A 3-regular graph where $T_{\text{visitx}} = \omega(T_{\text{push}})$. $\Theta(\sqrt{n})$ copies of a ladder graph of length $O(\log n)$ are attached to a 3-regular graph R. The dashed lines represent the intermediate edges of the ladder graphs, not visible in the picture, and the dotted lines are the edges removed from R.

We then add $\Theta(\sqrt{n})$ "ladder" graphs to R, each of logarithmic length. The two vertices at the end of each ladder are identified with unique vertices in R that are connected by an edge, which is removed. The resulting graph is 3-regular, has $\Theta(n)$ vertices and a logarithmic diameter, so $T_{\text{push}} = O(\log n)$, w.h.p.

To show the lower bound on T_{visitx}, we argue that, with constant probability, at least one of the attached ladder graphs contains no agents initially. It will then take $\Omega(\log^2 n/\log\log n)$ rounds before all vertices of that ladder graph are reached by an agent. The precise construction of G and the detailed proof are in Section 3.2.4.

This example shows that on sparse regular graphs T_{visitx} and T_{push} may differ by a factor of $\widetilde{\Omega}(\log n)$. However, it may be that T_{visitx} and T_{push} simply differ be an *additive* $O(\log^2 n)$ term. We currently do not know which of the two cases holds in the general case.

A consequence of Theorem 3.1.2 is that known bounds for PUSH do not readily apply to VISIT-EXCHANGE for low-degree regular graphs, thus new bounds are needed. In view of that, in later chapters we consider a number of classes of sparse graphs and bound the broadcast time of VISIT-EXCHANGE for them.

3.1.1 Road-map

In Section 3.2 we analyse the broadcast times for the example graphs in Figs. 3.1 and 3.2. Then in Section 3.3 we prove one of the directions of Theorem 3.1.1, that PUSH is not slower than VISIT-EXCHANGE. Section 3.4 proves the opposite direction.

3.2 Examples where the two processes differ

In this section, we provide examples demonstrating that PUSH, PUSH-PULL and VISIT-EXCHANGE can have very different broadcast times on the same graph. In the first three examples we present the graphs are highly non-regular and the separation between T_{visitx} and either T_{push} or T_{ppull} is polynomial. In the last example, we present a regular but sparse graph for which the processes are not equivalent and their broadcast times differ by a logarithmic factor. In all examples, we assume that the number of agents is $|A| = \alpha n = \Theta(n)$ for some constant α.

3.2.1 Star graphs

Let S_n denote an *n-leaf star*, that is, a tree with one internal node (the center of the star), and n leaves. See Fig. 3.1(a) for an illustration. This is an example of a graph where PUSH is very slow, whereas all other processes are very fast.

Lemma 3.2.1. *For the graph S_n described above and any source vertex s,*

(a) $\mathbb{E}[T_{\text{push}}] = \Omega(n \log n)$,

(b) $T_{\text{ppull}} \leqslant 2$,

(c) $T_{\text{visitx}} = O(\log n)$, *w.h.p.*

Proof. **(a)** This bound is well-known. It follows from the observation that the center needs to sample each of the leaves (except possibly for one, if the source is a leaf) before all vertices are informed. The time for that at least the time needed to collect $n-2$ coupons in a coupon collector's problem with $n-1$ coupons, which is $\Theta(n \log n)$ in expectation [MU17, Example 2.4.1].

(b) This bound is also well-known (and trivial). It takes one round to inform all vertices if s is the source, and two rounds if s is a leaf.

(c) First, we show that a fixed vertex $u \in V$ and round t, u is visited by some agent by round $t + O(\log n)$. For any $v \in V$, the probability that an agent that is at v visits u in the next two rounds is at least $1/n$. Since the agents do independent walks, a standard Chernoff bound implies that, for any initial placement of the agents, one of the agents will visit u in $O(\log n)$ rounds, w.h.p.

By this observation, some agent visits the source s and becomes informed in the first $O(\log n)$ rounds, w.h.p. After at most two rounds all agents become informed, because agents visit the central vertex every other round. Finally, every leaf u gets informed in an additional $O(\log n)$ rounds, w.h.p, due to our first observation that every vertex is visited in logarithmically many rounds, w.h.p. □

3.2.2 Double star graphs

In the star example above only the PUSH version of randomised rumour spreading is slow, while PUSH-PULL is extremely fast. Next we present a graph where PUSH-PULL (and thus, PUSH) is slow, while VISIT-EXCHANGE is fast. Let S_n^2 denote a *double-star* graph: two star graphs with $n/2$ vertices with their centers connected by an edge, as can be seen in Fig. 3.1(b).

Lemma 3.2.2. *For the graph S_n^2 described above and any source vertex s,*

(a) $\mathbb{E}[T_{\text{ppull}}] = \Omega(n)$,

(b) $T_{\text{visitx}} = O(\log n)$, *w.h.p.*

Proof. **(a)** Let a, b be the centers of the two stars. For PUSH-PULL to complete, a must sample b or b must sample a, at least once. The probability of that happening in a given round is at most $2/(n/2)$. Thus, the expected number of rounds until PUSH-PULL completes is at least $(n/2)/2$.

(b) Let $\mathcal{E}_u(t)$ denote the event that at least $|A|/8$ agents are located at vertex $u \in \{a, b\}$ in round t. We consider the following modification to process VISIT-EXCHANGE.

> For any round $t \geqslant 0$ and $u \in \{a, b\}$, if the event $\mathcal{E}_u(t)$ does not hold, then before round $t+1$ we add a minimal number of new and informed agents to the graph, at node u, such that there are $|A|/8$ agents at u.

In VISIT-EXCHANGE, at any round t, the agents are distributed according to the stationary distribution of the graph. Hence, the expected number of agents that visit u is greater than $|A|/4$. It follows, $\mathbb{P}[\mathcal{E}_u(t)] \geqslant 1 - e^{-\Omega(|A|)} = 1 - e^{-\Omega(n)}$ by a Chernoff bound. By applying a union bound for each $u \in \{a, b\}$ and all rounds $t \leqslant \log^2 n$, we get that, with probability at least $1 - e^{-\Omega(n)}$, the modified process is identical to the original VISIT-EXCHANGE for the first $\log^2 n$ rounds. Since our goal is to prove that $T_{\text{visitx}} = O(\log n)$ w.h.p., it suffices to analyze the modified process.

Now suppose $s \notin \{a, b\}$ and s is adjacent, say, to a. In the modified process, since there are at least a linear number of agents at a before each round, it takes $O(\log n)$ rounds before one agent visits s and then a, thus, informing a in $O(\log n)$ rounds, w.h.p. (If $s = a$, then a is informed at round 0.) By a similar argument, it takes an additional $O(\log n)$ rounds until b becomes informed, and then another $O(\log n)$ rounds until all leaf vertices become informed, w.h.p. The total broadcast time is thus logarithmic, w.h.p. □

3.2.3 Heavy binary trees

Next we describe a graph where VISIT-EXCHANGE is slow, while the other processes are fast. Let B_n denote a *heavy binary tree*, which is constructed by adding an edge between every pair of leaves of a balanced binary tree with n vertices. Even though B_n is not a tree, we will refer to the leaves of the original binary tree as the leaves of B_n. The set of leaves of B_n induces a clique of $l = \lceil n/2 \rceil$ vertices. See Fig. 3.1(c) for an illustration.

Lemma 3.2.3. *For the graph B_n described above and any source vertex s,*

(a) $T_{\text{push}} = O(\log n)$, *w.h.p.,*

(b) $\mathbb{E}[T_{\text{visitx}}] = \Omega(n)$.

Proof. **(a)** First, we upper bound the number of rounds until some internal node is informed. This is zero if s is an internal node, so suppose s is a leaf. The number of rounds before all leaves are informed is $O(\log n)$ w.h.p. This follows from the well-known logarithmic bound on the PUSH broadcast time on a clique, and the fact that random failures of transmission with probability $1/l$ (corresponding to the case when a leaf samples its parent) do not change the broadcast time asymptotically [ES09]. Once all leaves are informed, it takes at most $O(\log n)$ additional rounds, w.h.p., until the first internal node is informed, because there are l leaves and, in each round, each leaf samples its parent with probability $1/l$. Once some internal node becomes informed, then all internal nodes become informed after at most $O(\log n)$ rounds w.h.p. This follows from the observation that the broadcast time of PUSH on B_n starting from an internal node is dominated by the broadcast time on a balanced binary tree with n vertices. Since the binary tree has bounded degree and logarithmic diameter, the broadcast time of PUSH is $O(\log n)$ w.h.p. [Fei+90]. Adding all these logarithmic bounds and applying a union bound completes the proof.

(b) At any fixed round t, the agents are distributed according to the stationary distribution of the graph. If ρ is the root of B_n, then the probability that a given agent is at the root some round t is $\deg(\rho)/(2|E|) \leqslant 8/n^2$. Recall that $\alpha = |A|/n$ and consider the first $\tau = n/(16\alpha)$ rounds of the process. Let X be the number of total number of visits by agents to the root in those rounds. Then, $\mathbb{E}[X] \leqslant |A| \cdot \tau \cdot 8/n^2 \leqslant 1/2$. By Markov's inequality, then $\mathbb{P}[X \geqslant 1] \leqslant 1/2$, which means that with probability at least $1/2$ no agent visits the root in any of the first $\tau = \Theta(n)$ rounds. This implies $\mathbb{E}[T_{\text{visitx}}] = \Omega(\log n)$. □

3.2.4 Sparse regular graphs

In this section, we prove Theorem 3.1.2 by constructing a graph $G = G_{n,\ell,m,k}$, illustrated in Fig. 3.2. Here n is the number of vertices of the graph, ℓ is an upper bound on the diameter of the graph R to which copies of the ladder graph H are added, m is the number of the such copies and $2k$ is the number of vertices in H. The formal definition of G follows. Let H be a ladder graph defined as the Cartesian product of a path graph of k vertices, and a path graph of two vertices (i.e., a single edge). Every vertex of H has degree 3, except for the 4 endpoints that have degree 2. Take a 3-regular graph R of $n - 2mk$ vertices and diameter at most ℓ, and remove an arbitrary set of $2m$ edges (u_i, v_i), $1 \leqslant i \leqslant 2m$. Create m copies of H, and denote the four endpoints of the ith copy by $x_{2i-1}, y_{2i-1}, x_{2i}, y_{2i}$, where x_j and y_j are connected by an edge. Then, join every copy of H with R, by adding edges (x_i, u_i) and (y_i, v_i), for all $1 \leqslant i \leqslant 2m$. The resulting graph is $G_{n,\ell,m,k}$. By construction, the graph is 3-regular with n vertices. Also,

$$\operatorname{diam}(G_{n,\ell,m,k}) \leqslant 3\ell + 2k + 2, \tag{3.1}$$

because for every edge (u_i, v_i) removed from R, a path $u_i x_i y_i v_i$ of length 3 is created, and the diameter of each copy of H is k.

Recall we assume that the total number of agents in VISIT-EXCHANGE is αn, for a constant α.

Lemma 3.2.4. *For $m = \lceil\sqrt{n}\rceil$, and $k = \lfloor \log n/(4\alpha)\rfloor$ and $\ell = O(\log n)$, $\mathbb{E}[T_{\text{visitx}}(G_{n,\ell,m,k})] = \Omega(\log^2 n/\log\log n)$.*

Proof. For $i \in \{1, \dots, m\}$, let U_i be the vertex set of the ith copy of H, and let S_i be the set of its endpoints. The expected number of unique agents that visit S_i in the first $r = \lceil \log^2 n/\log\log n\rceil$ rounds is $4\alpha r$, and since the agents move independently, by a Chernoff bound, with probability at least $1 - 1/n^2$ no more than $8\alpha r$ unique agents visit S_i during the first r rounds. We create a modified process M-VISIT-EXCHANGE, in which if for some $i \in \{1, \dots, m\}$ more than $8\alpha r$ agents visit S_i in the first r rounds, we remove the extra agents. By a union bound, VISIT-EXCHANGE and M-VISIT-EXCHANGE are identical with probability at least $1 - 1/n$, and for the rest of the proof we will consider M-VISIT-EXCHANGE.

A single agent starts its walk in some vertex of U_i with probability $|U_i|/n = 2k/n$, thus, U_i does not contain any agents at the start of VISIT-EXCHANGE with probability $(1 - 2k/n)^{\alpha n} \geqslant e^{-2\alpha k}/2 \geqslant 1/(2m)$. This implies that with a probability at least $1 - (1 - 1/(2m))^m \geqslant 1 - e^{-2}$ there is a set U_i that does not contain any agent at time 0. In the rest of the proof we condition on this event and assume that for some fixed i, set U_i does not contain any agents at time 0.

Consider a path graph P_h with vertices $0, \ldots, h = \lceil k/2 \rceil - 1$. For every agent g that visits P_i, we couple its movement to a new lazy random walk W_g on P_h, with holding probability $1/3$, that starts from vertex 0. While g is in U_i, the position of W_g in P_h is equal to the distance of g from S_i. When g leaves the set U_i, we freeze W_g at vertex 0 and activate it again when g returns.

By the construction of M-VISIT-EXCHANGE, at most $8\alpha r$ unique agents ever visit S_i, and therefore, at most that many walks exist in P_h. By [Alo+11], the expected number of steps for the walks W_g to cover P_h is $\Omega(h^2/\log(8\alpha r)) = \Omega(\log^2 n/\log\log n)$. On the other hand, by the time all vertices of U_i are visited by agents, the coupled walks must cover P_h. Combining that with the fact that VISIT-EXCHANGE and M-VISIT-EXCHANGE are identical w.h.p., and that a set U_i without agents at the start exists with constant probability, we get that $\mathbb{E}[T] = \Omega(\log^2 n/\log\log n)$. □

Choosing m and k as in Lemma 3.2.4, and also choosing the graph R in the construction of $G_{n,\ell,m,k}$ to have logarithmic diameter, i.e., $\ell = O(\log n)$, we obtain from (3.1), that $\operatorname{diam}(G_{n,\ell,m,k}) = O(\log n)$, and from Lemma 3.2.4, that $\mathbb{E}[T(G_{n,\ell,m,k})] = \Omega(\log^2 n/\log\log n)$. Also, by [Fei+90], $T_{\text{push}} = O(\operatorname{diam}(G_{n,\ell,m,k})) = O(\log n)$, w.h.p., since the graph is 3-regular. Thus, the graph $G = G_{n,\ell,m,k}$ establishes Theorem 3.1.2.

3.3 Upper bounding T_{push} in terms of T_{visitx}

In this section we prove one of the directions of Theorem 3.1.1, namely that PUSH does not broadcast slower than VISIT-EXCHANGE in sufficiently dense regular graphs. The claim is formalised by the following theorem.

Theorem 3.3.1. *For any constants $c, \alpha, \beta > 0$, there is a constant $\lambda > 0$, such that for any d-regular graph $G = (V, E)$ with $|V| = n$ and $d \geqslant \beta \log n$, and for any source vertex $s \in V$, the broadcast times of* PUSH *and* VISIT-EXCHANGE*, with $|A| = \alpha n$ agents, satisfy*

$$\mathbb{P}[T_{\text{push}} \leqslant \lambda k] \geqslant \mathbb{P}[T_{\text{visitx}} \leqslant k] - n^{-c},$$

for any $k \geqslant 0$.

3.3.1 Overview of the proof

The proof uses the following coupling of the PUSH and VISIT-EXCHANGE processes. For each vertex u, let $\langle \pi_u(1), \pi_u(2), \ldots \rangle$ be the sequence of neighbours that u samples in PUSH after getting informed. Similarly, for VISIT-EXCHANGE, consider all moves of informed agents from u to its neighbour vertices in chronological order, and let $\langle p_u(1), p_u(2), \ldots \rangle$ be the destination vertices in those moves (we order moves in the same round by, say, agent ID). We couple the two processes by setting $p_u(i) = \pi_u(i)$, for all $u \in V$ and $i \geqslant 1$. Fig. 3.3 illustrates this coupling for the source vertex.

To prove the theorem, it suffices to show that under our coupling, with probability at least $1 - n^{-c}$, if $T_{\text{visitx}} \leqslant k$ then $T_{\text{push}} \leqslant \lambda k$. We also assume that $k = \Omega(\log n)$. For $k = O(\log n)$, the theorem follows from the logarithmic lower bound on T_{visitx} that can be obtained using the proof of [ES11, Theorem 4.2].

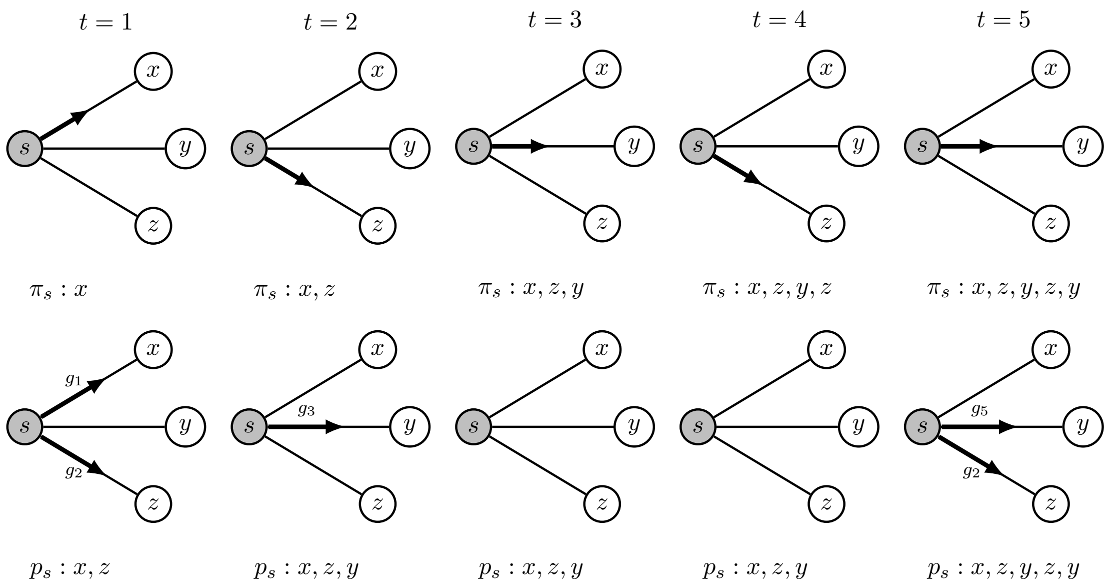

Figure 3.3 The first 5 rounds of the coupled processes used in the proof of Theorem 3.3.1 for the source vertex s. The top row corresponds to PUSH and the bottom row to VISIT-EXCHANGE. The agents g_i follow the choices made by the PUSH process. π_u is the list of neighbours that s contacts in the first five rounds of PUSH.

The intuition of this coupling is that it ensures that for each move of an informed agent from vertex u to its neighbour v, there is a corresponding round in PUSH when u samples the same neighbour v. Thus, if there were a constant upper bound λ on the *actual* number of visits to each vertex on each round, then the coupling would immediately yield $T_{\text{push}} \leqslant \lambda \cdot T_{\text{push}}$. In reality, however, such a bound exists only *in expectation* and a super-constant number of agents may visit a vertex in certain rounds. Moreover, the actual number of visits depends on the past history of the process.

The main idea we use to tackle the dependencies uses the lower bound of $\Omega(\log n)$ on the degree of the graph. We introduce a modified version of VISIT-EXCHANGE, called M-VISIT-EXCHANGE. The only difference between these two processes is that by removing some agents arbitrarily in M-VISIT-EXCHANGE, we ensure that the neighbourhood of any vertex contains $O(d)$ agents in any round. Due to the fact that $d = \Omega(\log n)$ and $|A| = O(n)$, we can show that the two processes are identical in the first polynomially many rounds of VISIT-EXCHANGE, w.h.p. It implies that we can use M-VISIT-EXCHANGE for the main part of the proof. This property implies that for each round t, the number of agents visiting a vertex at round $t+1$ is upper bounded by binomial random variable $\text{Bin}(\Theta(d), 1/d)$, independently from the past execution of M-VISIT-EXCHANGE.

To prove that $T_{\text{visitx}} \leqslant k$ implies $T_{\text{push}} \leqslant \lambda k$, w.h.p., we consider *all* possible paths of length k through which information travels in VISIT-EXCHANGE.[1] We follow each path, moving one vertex further each round, and count the total number of (non-distinct) agents encountered along this path. This number is called the congestion of the path. Formally, we use the notion of a *canonical walk* θ, which is represented by a sequence of vertices $\theta = \langle \theta_0, \theta_1, \ldots, \theta_k \rangle$ starting from $\theta_0 = s$: In each round $1 \leqslant t \leqslant k$, the walk either stays

[1] It is natural to consider a single path via which VISIT-EXCHANGE makes progress towards a particular vertex u and prove that PUSH follows the same path. While this technique works for the opposite direction in Theorem 3.4.1, fixing a path in VISIT-EXCHANGE introduces dependencies from the future and we can no longer argue that agents move independently.

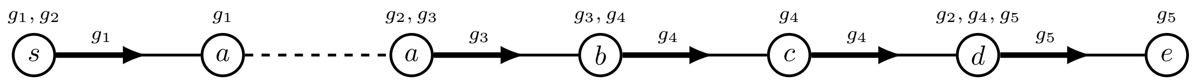

Figure 3.4 A labelled canonical walk $\theta = \langle s, g_1, a, \perp, a, g_3, b, g_4, c, g_4, d, g_5, e\rangle$ of length 6. The congestion of the walk is $Q(\theta) = 11$. Note that in round 2, even though agent g_1 is at vertex a, the walk stays put. A PUSH process coupled to VISIT-EXCHANGE, would take at most 11 rounds to pass information along the same path.

put and $\theta_t = \theta_{t-1}$, or it follows one of the agents g that leave θ_{t-1} in round t, and, in that case, θ_t is the new vertex that g moves to. For any round t, we count the agents that are in θ_t. The sum of these counts, for $0 \leqslant t < k$ is the *congestion* $Q(\theta)$ of the walk θ. Note that the number of agents in θ_k at round k does not contribute to the congestion. Fig. 3.4 illustrates the definition of the congestion for some canonical walk.

The congestion of a canonical walk is used to bound the time needed for information to travel along the *same* path in the coupled PUSH process. Intuitively, larger congestion implies longer travel time for PUSH, for the following reason. Suppose there are m agents in u at some round after it is informed by VISIT-EXCHANGE. The coupled PUSH process, using the same random decisions for the choice of neighbours as VISIT-EXCHANGE, will take m rounds to "go through" these m agents.

To formalise the relation between the congestion of canonical walks and the time it takes for information to spread in PUSH, we introduce *C-counters*: For each vertex u, we maintain a counter $C_u(t)$ for round t. The counter is initialised in the round t_u in which u becomes informed in VISIT-EXCHANGE. Its initial value is the value of the C-counter of the neighbour from which the first informed agent arrived to u. In each subsequent round $t > t_u$, C_u increases by the number of agents that visited u in round $t-1$. C-counters have the following two properties: If τ_u is the round when u gets informed in PUSH then $\tau_u \leqslant C_u(t_u)$; and for any $t \geqslant t_u$, there is a canonical walk θ of length t such that $C_u(t) = Q(\theta)$. Therefore, to show that w.h.p. $T_{\text{visitx}} \leqslant k$ implies $T_{\text{push}} \leqslant \lambda k$, it suffices to show that the maximum congestion of *all* canonical walks of length k is at most λk, w.h.p.

We can bound the congestion of a *single* canonical walk of length k using the property of M-VISIT-EXCHANGE that the number of agents at a node is bounded by a binomial distribution with constant mean. This results in the desired bound of λk for a single walk with probability at least $1 - a^{-k}$, for some constant $a > 1$. We would like to take a union bound over all canonical walks, which would complete the proof. For this to work, however, we should also bound the total number of canonical walks of length k by at most a^k/n^c, which does not work using trivial methods.

We bound the number of canonical walks of length k by introducing a set of *descriptors* for these walks. A descriptor is represented by a matrix, which, together with a given execution of VISIT-EXCHANGE, uniquely defines a canonical walk. Additionally, the set of descriptors suffices to encode all canonical walks, and therefore, it is at least as large as the set of all walks. Thus, we can use a bound on the number of descriptors that can be computed by a simple combinatorial argument involving the number of elements used in the matrix, and the values they can take. A naive construction of descriptors, however, is too wasteful giving us a much larger bound than the a^k/n^c we need. A key idea here is that the majority of the descriptors represent walks only in executions that happen with low probability. So, we construct a set of *concise descriptors* that can describe all canonical walks in a random execution, w.h.p. We show that the size of the set of concise descriptors can be bounded by a^k/n^c, as desired. The theorem then follows by switching back from the M-VISIT-EXCHANGE process to the original VISIT-EXCHANGE.

3.3.2 Notation and definition of the coupling

For each vertex $u \in V$, we denote by τ_u the round when u gets informed in PUSH. For $i \geqslant 1$, let $\pi_u(i)$ be the ith vertex that u samples, i.e., the vertex it samples in round $\tau_u + i$. Note that $\tau_{\pi_u(i)} \leqslant \tau_u + i$. In VISIT-EXCHANGE, we denote by t_u the round when vertex u gets informed. For any agent $g \in A$ and $t \geqslant 0$, we denote by $x_g(t)$, the vertex that g visits in round t. Thus, $\{x_g(t)\}_{t\geqslant 0}$ is a random walk on G. Let $Z_u(t)$ be the set of all agents that visit u in round t, i.e.,

$$Z_u(t) = \{g \in A \mid x_g(t) = u\}.$$

Thus, $Z_u(t)$ is also the set of agents that depart from u in round $t+1$. Consider all visits to u in rounds $t \geqslant t_u$, in chronological order, ordering visits in the same round with respect to a predefined total order over agents. For each $i \geqslant 1$, consider the agent g that does the ith such visit, and let $p_u(i)$ be the vertex that g visits next. Formally, let $W_u = \{(t, g) \mid t \geqslant t_u, x_g(t) = u\}$, and order its elements such that $(t, g) < (t', g')$ if $t < t'$, or $t = t'$ and $g < g'$. If (t, g) is the ith smallest element in W_u, then $p_u(i) = x_g(t+1)$.

Coupling. We couple processes PUSH and VISIT-EXCHANGE by setting $\pi_u(i) = p_u(i)$. Formally, let $\{w_u(i)\}_{u\in V, i\geqslant 1}$, be a collection of independent random variables, where $w_u(i)$ takes a uniformly random value from the set $\Gamma(u)$ of u's neighbours. Then, for every $u \in V$ and $i \geqslant 1$, we set $\pi_u(i) = p_u(i) = w_u(i)$. See Fig. 3.3 for an illustration of the coupling.

3.3.3 A modified Visit-Exchange process

We will use the next simple bound on the number of agents that visit a given set S of vertices in some round t of VISIT-EXCHANGE. The proof is by a simple Chernoff bound, and relies on the assumption that agents execute independent walks starting from stationarity.

Lemma 3.3.2. *For any $S \subseteq V$, $t \geqslant 0$, and $\gamma \geqslant 2e \cdot |A|/n$,*

$$\mathbb{P}\left[\sum_{v\in S} |Z_v(t)| \leqslant \gamma \cdot |S|\right] \geqslant 1 - 2^{-\gamma\,|S|}.$$

Proof. Since each random walk starts from stationarity, and G is a regular graph, it follows that for any agent $g \in A$, $\mathbb{P}[x_g(t) \in S] = |S|/n$. Thus, the expected number of agents that visit S in round t is $|A| \cdot |S|/n \leqslant \gamma \cdot |S|/(2e)$. Then, by the independence of the random walks, we can use a standard Chernoff bound to show that the number of agents that visit S at t is at most $\gamma \cdot |S|$ with probability at least $1 - 2^{-\gamma\cdot|S|}$. □

We remark that Lemma 3.3.2 holds also in the case where $|A| = n$ and exactly one walk starts from each vertex. This implies that Theorem 3.3.1 holds in the above case as well, because the rest of the proof does not require any assumptions about the initial distribution of agents.

Next we define a modified variant of VISIT-EXCHANGE, called M-VISIT-EXCHANGE, defined as follows. Let

$$\gamma \geqslant 2e \cdot |A|/n \tag{3.2}$$

be a (sufficiently large) constant to be specified later. If in some round $t \geqslant 0$, there is a vertex $u \in V$ for which the following condition does *not* hold:

$$\sum_{v \in \Gamma(u)} |Z_v(t)| \leqslant \gamma \cdot d, \tag{3.3}$$

then before round $t+1$, we remove a minimal set of agents from the graph in such a way that the above condition holds for all vertices u, when counting just the remaining agents.

It follows from Lemma 3.3.2 that if γ is large enough, and $d = \Omega(\log n)$, then w.h.p. the modified process is identical to the original in the first polynomial number of rounds.

Lemma 3.3.3. *The* VISIT-EXCHANGE *and* M-VISIT-EXCHANGE *processes are identical for the first k rounds of their execution, with probability at least $1 - kn \cdot 2^{-\gamma d}$.*

Proof. The claim follows by applying Lemma 3.3.2, for each $0 \leqslant t < k$ and each pair u, S, where $u \in V$ and $S = \Gamma(u)$, and then combining the results using a union bound over all vertices $u \in V$ and rounds t up to k. □

This lemma allows us to use the M-VISIT-EXCHANGE process in the main part of the proof instead of VISIT-EXCHANGE. We use the same notations for both processes.

3.3.4 Canonical walks

Let $\theta = \langle \theta_0, \theta_1, \ldots, \theta_k \rangle$, where $\theta_0 = s$ and $\theta_i \in \Gamma(\theta_{i-1}) \cup \{\theta_{i-1}\}$ for $1 \leqslant i \leqslant k$, be a walk on G constructed from VISIT-EXCHANGE as follows. We start from vertex $\theta_0 = s$ in round zero, and in each round $1 \leqslant t \leqslant k$, we either stay put, in which case $\theta_t = \theta_{t-1}$, or we choose one of the agents $g \in Z_{\theta_{i-1}}(t-1)$, which visited θ_{i-1} in the previous round, and move to the same vertex as g in round t, i.e., $\theta_t = x_t(g)$. We call θ a *canonical walk* of length k. A *labelled canonical walk* is a canonical walk that specifies also the agent g_t that the walk follows in each step t, if $\theta_t \neq \theta_{t-1}$. Formally, a labelled canonical walk corresponding to θ is $\eta = \langle \theta_0, g_1, \theta_1, g_2, \ldots, g_k, \theta_k \rangle$, where $g_t \in Z_{\theta_{t-1}}(t-1) \cap Z_{\theta_t}(t)$ if $\theta_t \neq \theta_{t-1}$, and $g_t = \perp$ if $\theta_t = \theta_{t-1}$. Note that different labelled canonical walks may correspond to the same (unlabelled) canonical walk.

Concise descriptors of canonical walks

In this section, we bound the number of distinct labelled canonical walks of a given length k. For that, we present a concise description for such walks, and bound the total number of the walks by the total number of different possible descriptions.

We start with a rather wasteful way to describe labelled canonical walks, which we then refine in two steps. Let $\mathcal{A}_k$ denote the set of all $\alpha n \times k$ matrices $A_k = [a_{i,j}]$, where $a_{i,j} \in \{0, \ldots, i\}$. Let us fix the first k rounds of VISIT-EXCHANGE, and consider a labelled canonical walk $\eta = \langle \theta_0 = s, g_1, \theta_1, \ldots, g_k, \theta_k \rangle$. For each $1 \leqslant t \leqslant k$, let

$$\delta_t = |Z_{\theta_{t-1}}(t-1)|$$

be the number of agents that visit θ_{t-1} in round $t-1$, and thus also the number of agents that depart from θ_{t-1} in round t. Let $\rho_t = 0$ if $g_t = \perp$, otherwise, ρ_t is equal to the *rank* of g_t in set $Z_{\theta_{t-1}}(t-1)$, i.e., $\rho_t = |\{g \in Z_{\theta_{t-1}}(t-1) \mid g \leqslant g_t\}|$. We describe walk η by a matrix $A_k \in \mathcal{A}_k$ with the following entries: For each $1 \leqslant t \leqslant k$, if $\delta_t > 0$, then $a_{\delta_t, j} = \rho_t$, for

$\delta_t : 2, 1, 2, 2, 1, 3$
$\rho_t : 1, 0, 2, 2, 1, 3$

(a)

0	1				
1	2	2			
3					

(b)

Figure 3.5 **(a)** The sequences δ_t and ρ_t for $t \geqslant 1$ of the canonical walk presented in Fig. 3.4. **(b)** The non-concise descriptor corresponding to the walk (here we assume that there are 5 agents in total by using 5 rows). The missing elements of the table can have arbitrary values.

$j = |\{t' \leqslant t \mid \delta_{t'} = \delta_t\}|$, i.e., value ρ_t is stored in the first unused entry of row $A_k[\delta_t, \cdot]$. At most k of the entries of A_k are specified that way; the remaining entries can have arbitrary values. We call A_k a *non-concise descriptor* of η. An illustration of this construction can be seen in Fig. 3.5.

For any given realisation of VISIT-EXCHANGE, each $A_k \in \mathcal{A}_k$ describes exactly one labelled canonical walk of length k. To construct such a canonical walk from A_k, we start from s and add its elements consecutively. For each round, suppose u is the most recently added vertex to the walk. If u contains δ agents, then we consider the next unused element of the row $A_k[\delta, \cdot]$, say ρ. If $\rho = 0$, then the walk stays put, i.e., we append $\langle \perp, u \rangle$ to it. Otherwise the walk follows the agent of rank ρ from among the δ agents at u at that round. If that agent is g and in the next round it visits the neighbour v of u, then we append $\langle g, v \rangle$ to the walk.

The total number of different non-concise descriptors is $|\mathcal{A}_k| = \prod_{1 \leqslant i \leqslant \alpha n} (i+1)^k$, which is too large for our purposes. A simple improvement is to use only entries in rows $A_k[i, \cdot]$ for which i is a power of 2 (we assume w.l.o.g. that αn is also a power of 2). Roughly speaking, if δ_t is between $2^{\ell-1}$ and 2^ℓ then ρ_t is stored in row $A_k[2^\ell, \cdot]$. Formally, let b be a (large enough) constant, to be specified later, which is a power of 2. The matrix $A_k \in \mathcal{A}_k$ we use to describe η has the following entries. For each $1 \leqslant t \leqslant k$:

1. If $2^{\ell-1} < \delta_t \leqslant 2^\ell$, where $\ell \in \{1 + \log b, \ldots, \log(\alpha n)\}$, and $|\{t' \leqslant t \mid 2^{\ell-1} < \delta_{t'} \leqslant 2^\ell\}| = j$, then
   (a) if $\rho_t \neq 0$, we have $a_{2^\ell, j} = \rho_t$,
   (b) if $\rho_t = 0$, $a_{2^\ell, j}$ can take any value in $\{0\} \cup \{\delta_t + 1, \ldots, 2^\ell\}$.
2. If $0 \leqslant \delta_t \leqslant b$ and $|\{t' \leqslant t \mid 0 < \delta_{t'} \leqslant b\}| = j$, then
   (a) if $\rho_t \neq 0$, we have $a_{b,j} = \rho_t$,
   (b) if $\rho_t = 0$, $a_{b,j}$ can take any value in $\{0\} \cup \{\delta_t + 1, \ldots, b\}$.

The purpose of subcases (b) is to maintain the property that *every* A_k describes a labelled canonical walk, which would not be the case if we just set $a_{2^\ell, j} = 0$ or $a_{b,j} = 0$, since values greater than δ_t would not correspond to a walk. We call the matrix A_k above a *semi-concise descriptor* of η.

A second modification we make is based on the observation that, even in the logarithmic number of rows used in the above scheme, most entries are still very unlikely to be used. For each row $i = 2^\ell$, we specify a threshold index $k_i \leqslant k$, such that the first k_i entries in each row $A_k[i,\cdot]$ suffice w.h.p. to describe all labelled canonical walks of length k, in a random realisation of VISIT-EXCHANGE. Let $\mathcal{B}_k$ be a subset of $\mathcal{A}_k$ defined as follows. Let

$$k_i = b \cdot k/i,$$

and recall that b is a constant power of 2. The set $\mathcal{B}_k$ consists of all $A_k = [a_{i,j}] \in \mathcal{A}_k$ such that

$$\begin{aligned} &a_{i,j} \in \{0,\ldots,i\}, && \text{if } i \in \{2^\ell \mid \log b \leqslant \ell \leqslant \log(\alpha n)\} \text{ and } j \leqslant k_i \\ &a_{i,j} = 0, && \text{otherwise.} \end{aligned}$$

A *concise descriptor* of a labelled canonical walk η of length k is any semi-concise descriptor A_k of η that belongs to set $\mathcal{B}_k$.

Next we compute an upper bound on the number of all possible concise descriptors of length k.

Lemma 3.3.4. $|\mathcal{B}_k| \leqslant (4b)^{2k}$.

Proof. From the definition of $\mathcal{B}_k$, we have

$$\begin{aligned} |\mathcal{B}_k| &\leqslant \prod_{\log b \leqslant \ell \leqslant \log(\alpha n)} (2^\ell + 1)^{bk/2^\ell} \\ &= \prod_{\log b \leqslant \ell \leqslant \log(\alpha n)} 2^{\ell bk/2^\ell} \cdot \prod_{\log b \leqslant \ell \leqslant \log(\alpha n)} (1+2^{-\ell})^{bk/2^\ell} \\ &\leqslant \frac{\prod_{\ell \geqslant 1} 2^{\ell bk/2^\ell}}{\prod_{\ell \leqslant \log b - 1} 2^{\ell bk/2^\ell}} \cdot \prod_{\ell \geqslant \log b} e^{bk/4^\ell} \\ &= \frac{2^{2bk}}{2^{(2(b-\log b-1)k}} \cdot e^{(4/3)k/b} \\ &\leqslant 2^{2(\log b+2)k}, \end{aligned}$$

where in the second-last line we used $\sum_{\ell \geqslant 1} \ell/2^\ell = 2$, $\sum_{\ell \leqslant y} \ell/2^\ell = 2^{-y}(2^{y+1} - y - 2)$, and $\sum_{\ell \geqslant 0} 1/4^\ell = 4/3$; and in the last line we used that $e^{(4/3)} < 4$. $\square$

Concise descriptors encode all canonical walks

For any realisation of VISIT-EXCHANGE, each $A_k \in \mathcal{B}_k$ is a concise descriptor of some labelled canonical walk of length k. However it is not always the case that a labelled canonical walk has a concise descriptor. The next lemma shows that w.h.p. all labelled canonical walks of length k have concise descriptors for an appropriate choice of constant parameter b. Note that the lemma assumes the M-VISIT-EXCHANGE process.

Lemma 3.3.5. *If $b \geqslant \max\{2\gamma e^2, 64\}$ then, with probability at least $1 - 2^{-bk/4}\log(\alpha n)$, all labelled canonical walks of length k in a random realisation of* M-VISIT-EXCHANGE *have concise descriptors.*

First, we bound the number of steps t in which more than i agents are encountered in a canonical walk of length k.

Lemma 3.3.6. *Fix any $A_k \in \mathcal{A}_k$, and let $\eta = \langle \theta_0, g_1, \theta_1, \ldots, g_k, \theta_k \rangle$ be the labelled canonical walk with semi-concise (or non-concise) descriptor A_k in* M-VISIT-EXCHANGE. *For any $i \geqslant e^2\gamma$ and $\epsilon \geqslant e^2\gamma$,*

$$\mathbb{P}\left[|\{t \in \{1, \ldots, k\} \mid \delta_t > i\}| \geqslant \epsilon k/i\right] \leqslant 2^{-\epsilon k}.$$

Proof. Recall that $\delta_t = |Z_{\theta_{t-1}}(t-1)|$ is the number of agents that visit vertex θ_{t-1} in round $t-1$, and thus also the number of agents that depart from θ_{t-1} in round t. We argue that for any $t \geqslant 1$, conditioned on $\delta_1, \ldots, \delta_t$, variable δ_{t+1} is stochastically dominated by the binomial random variable $\mathrm{Bin}(\gamma d, 1/d) + 1$: From (3.3), applied for vertex θ_t and round $t-1$, we get

$$\sum_{v \in \Gamma(\theta_t)} |Z_v(t-1)| \leqslant \gamma \cdot d.$$

Thus, there are at most γd agents in the neighbourhood of θ_t before round t. If $\theta_t = \theta_{t-1}$, then each one of those at most γd agents will visit θ_t in round t independently with probability $1/d$. If $\theta_t \neq \theta_{t-1}$ (thus $g_t \in Z_{\theta_{t-1}}(t-1) \cap Z_{\theta_t}(t)$), then each of the at most γd agents will visit θ_t in round t independently with probability $1/d$, *except* for agent g_t who visits θ_t with probability 1. In both cases, the number δ_{t+1} of agents that visit θ_t is dominated by $\mathrm{Bin}(\gamma d, 1/d) + 1$. It follows that for any $t \geqslant 1$ and $i \geqslant 1$,

$$\begin{aligned}\mathbb{P}\left[\delta_{t+1} > i \mid \delta_1, \ldots, \delta_t\right] &\leqslant \mathbb{P}\left[\mathrm{Bin}(\gamma d, 1/d) + 1 > i\right] = \mathbb{P}\left[\mathrm{Bin}(\gamma d, 1/d) \geqslant i\right] \\ &\leqslant \binom{\gamma d}{i} \cdot \frac{1}{d^i} \leqslant \left(\frac{e\gamma d}{i}\right)^i \cdot \frac{1}{d^i} = \left(\frac{e\gamma}{i}\right)^i.\end{aligned}$$

Similarly, for δ_1 we have

$$\mathbb{P}\left[\delta_1 \geqslant i\right] = \mathbb{P}\left[\mathrm{Bin}(\alpha n, 1/n) \geqslant i\right] \leqslant \left(\frac{e\alpha}{i}\right)^i < \left(\frac{e\gamma}{i}\right)^i.$$

Let $p_i = \left(\frac{e\gamma}{i}\right)^i$. It follows from the above that for any $\ell \geqslant 1$,

$$\mathbb{P}\left[|\{t \in \{1, \ldots, k\} \mid \delta_t > i\}| \geqslant \ell\right] \leqslant \mathbb{P}\left[\mathrm{Bin}(k, p_i) \geqslant \ell\right] \leqslant \binom{k}{\ell} \cdot p_i^\ell \leqslant \left(\frac{ekp_i}{\ell}\right)^\ell. \tag{3.4}$$

For $\ell \geqslant \epsilon k/i$ and $i \geqslant e^2\gamma$,

$$\begin{aligned}\left(\frac{ekp_i}{\ell}\right)^\ell &\leqslant \left(\frac{ek(e\gamma/i)^i}{\epsilon k/i}\right)^\ell, \qquad \text{by } p_i = \left(\frac{e\gamma}{i}\right)^i \text{ and } \ell \geqslant \epsilon k/i \\ &= \left(\frac{e^2\gamma}{\epsilon} \cdot \left(\frac{e\gamma}{i}\right)^{i-1}\right)^\ell \leqslant \left(\frac{e\gamma}{i}\right)^{(i-1)\ell}, \qquad \text{by } \epsilon \geqslant e^2\gamma \\ &\leqslant \left(\frac{e\gamma}{i}\right)^{(1-1/i)\epsilon k}, \qquad \text{by } \ell \geqslant \epsilon k/i \\ &\leqslant \left(\frac{1}{e}\right)^{(1-1/e^2)\epsilon k}, \qquad \text{by } i \geqslant e^2\gamma \geqslant e^2 \\ &\leqslant 2^{-\epsilon k}.\end{aligned}$$

Substituting that to (3.4) completes the proof of Lemma 3.3.6. ∎

We proceed now to the proof of Lemma 3.3.5.

Proof of Lemma 3.3.5. For any $A_k \in \mathcal{A}_k$, and for $\eta = \langle \theta_0, g_1, \theta_1, \ldots, \theta_k \rangle$ being the labelled canonical walk with semi-concise descriptor A_k, let $\mathcal{E}_{A_k}$ denote the event:

$$|\{t \in \{1, \ldots, k\} \mid 2^{\ell-1} < \delta_t \leqslant 2^\ell\}| \leqslant k_{2^\ell}, \text{ for all } \ell \in \{\log b + 1, \ldots, \log(\alpha n)\}.$$

Applying Lemma 3.3.6, for $i = 2^{\ell-1}$ and $\epsilon = b/2$, for each $\ell \in \{\log b + 1, \ldots, \log(\alpha n)\}$, and then using a union bound, we obtain

$$\mathbb{P}[\mathcal{E}_{A_k}] \geqslant 1 - 2^{-bk/2} \log(\alpha n).$$

By another union bound and Lemma 3.3.4,

$$\begin{aligned} \mathbb{P}\left[\bigcap_{A_k \in \mathcal{B}_k} \mathcal{E}_{A_k}\right] &\geqslant 1 - |\mathcal{B}_k| \cdot 2^{-bk/2} \log(\alpha n) \geqslant 1 - (4b)^{2k} \cdot 2^{-bk/2} \log(\alpha n) \\ &\geqslant 1 - 2^{-bk/4} \log(\alpha n), \end{aligned} \tag{3.5}$$

where the last inequality holds if $b \geqslant 64$. Next we show that the event $\bigcap_{A_k \in \mathcal{B}_k} \mathcal{E}_{A_k}$ implies that every labelled canonical walk η has a concise descriptor $A_k \in \mathcal{B}_k$. From this and (3.5), the lemma follows.

Fix a realisation of M-VISIT-EXCHANGE conditioned on the event $\bigcap_{A_k \in \mathcal{B}_k} \mathcal{E}_{A_k}$. Suppose, for contradiction, that there is some labelled canonical walk $\eta' = \langle \theta'_0, g'_1, \theta'_1, \ldots, g'_k, \theta_k \rangle$ that does *not* have a concise descriptor. Let $\eta = \langle \theta_0, g_1, \theta_1, \ldots, g_k, \theta_k \rangle$ be a labelled canonical walk that *does* have a concise descriptor $A_k \in \mathcal{B}_k$, and shares a *maximal common prefix* with η'. Consider the first element where η' and η are different. We first argue that this element is not a vertex: Suppose, for contradiction, that $\langle \theta'_0, \ldots, g'_i \rangle = \langle \theta_0, \ldots, g_i \rangle$ and $\theta'_i \neq \theta_i$, for some $0 \leqslant i \leqslant k$. Then $i \neq 0$, as $\theta'_0 = s = \theta_0$. Moreover, if $i > 0$, then by definition, $\langle \theta'_0, \ldots, g'_i \rangle = \langle \theta_0, \ldots, g_i \rangle$ implies $\theta'_i = \theta_i$, contradicting our assumption. Thus, the first element where η' and η are different must be an agent. Suppose $\eta'' = \langle \theta'_0, g'_1, \ldots, \theta'_{i-1} \rangle = \langle \theta_0, g_1, \ldots, \theta_{i-1} \rangle$ and $g'_i \neq g_i$, for some $1 \leqslant i \leqslant k$. Then, by the maximal prefix assumption, the labelled canonical walk $\langle \theta_0, \ldots, \theta_{i-1}, g'_i, \theta'_i, \bot, \theta'_i, \bot, \ldots, \bot, \theta'_i \rangle$, which stays put at vertex θ'_i in rounds $i+1$ up to k, has no concise descriptor. This can only be true if $|\{t \in \{1, \ldots, i-1\} \mid 2^{\ell-1} < \delta_t \leqslant 2^\ell\}| > k_{2^\ell}$, for some $\ell \in \{\log b + 1, \ldots, \log n\}$, because the descriptor of η'' is the same as A_k except some rows end with more number of 0's. But this contradicts event $\mathcal{E}_{A_k}$. Therefore, there exists no labelled canonical walk η' of length k such that η' has no concise descriptor. □

3.3.5 Congestion of canonical walks

For a canonical walk $\theta = \langle \theta_0, \ldots, \theta_k \rangle$ we define its *congestion* $Q(\theta)$ as the total number of agents encountered along the walk,[2] not counting the last step, i.e.,

$$Q(\theta) = \sum_{0 \leqslant t < k} |Z_{\theta_t}(t)|.$$

The congestion of a *labelled* canonical walk is the same as the congestion of the corresponding unlabelled walk. Fig. 3.4 illustrates the congestion for some path starting from the source vertex.

[2]The same agent is counted more than once if encountered in multiple rounds.

C-Counters

We now introduce C-counters, which allow us to bound the round at which a vertex u becomes informed by the congestion of a canonical walk to u. Recall that t_u is the round when vertex u gets informed in VISIT-EXCHANGE. If $u \neq s$, this is the first round when some informed agent visits u. We are interested in the neighbour v of u from which that agent arrived. Note that $t_v < t_u$. Note also that there may be more than one such neighbours v, if more than one informed agent visits u at round t_u. For each $u \in V$, let

$$S_u = \{v \in \Gamma(u) \mid t_v < t_u,\ Z_v(t_u - 1) \cap Z_u(t_u) \neq \emptyset\},$$

i.e., S_u contains all neighbours v of u for which some informed agent moved from v to u in round t_u. Next, for each $t \geqslant 0$, we define the counter variable

$$C_u(t) = \begin{cases} 0, & \text{if } t < t_u \text{ or } t = t_u = 0 \\ \min_{v \in S_u} C_v(t), & \text{if } t = t_u > 0 \\ C_u(t-1) + |Z_u(t-1)|, & \text{if } t > t_u. \end{cases} \tag{3.6}$$

That is, C_u is initialised in round t_u to the minimum counter value of the neighbours in S_u (or to zero if $u = s$), and $C_u(t) - C_u(t_u)$ is the number of visits to u from round t_u until round $t-1$, or equivalently, the number of departures of agents from u in rounds $t_u + 1$ up to t.

The next two lemmas imply that if the congestion of all canonical walks to vertex u of length t_u is at most ct_u then $\tau_u \leqslant \lambda t_u$.

Lemma 3.3.7. *For any $u \in V$, $\tau_u \leqslant C_u(t_u)$.*

Proof. Consider the following path through which information reaches u in VISIT-EXCHANGE. The path is $\langle v_0, v_1, \ldots, v_k\rangle$, where $v_0 = s$, $v_k = u$, and for each $0 < j \leqslant k$, we have $v_{j-1} \in S_{v_j}$ and $C_{v_{j-1}}(t_{v_j}) = \min_{v \in S_{v_j}} C_v(t_{v_j}) = C_{v_j}(t_{v_j})$. We prove by induction on $0 \leqslant j \leqslant k$ that

$$\tau_{v_j} \leqslant C_{v_j}(t_{v_j}). \tag{3.7}$$

This holds for $j = 0$, because $v_0 = s$, $t_s = 0$, and $\tau_s = 0 = C_s(0)$. Let $0 < j \leqslant k$, and suppose that $\tau_{v_{j-1}} \leqslant C_{v_{j-1}}(t_{v_{j-1}})$; we will show that $\tau_{v_j} \leqslant C_{v_j}(t_{v_j})$. We have

$$\begin{aligned} C_{v_j}(t_{v_j}) &= C_{v_{j-1}}(t_{v_j}), \quad \text{by the path property} \\ &= C_{v_{j-1}}(t_{v_{j-1}}) + \sum_{t_{v_{j-1}} \leqslant t < t_{v_j}} |Z_{v_{j-1}}(t)|, \quad \text{by recursive application of (3.6)} \\ &\geqslant \tau_{v_{j-1}} + \sum_{t_{v_{j-1}} \leqslant t < t_{v_j}} |Z_{v_{j-1}}(t)|, \quad \text{by the induction hypothesis.} \end{aligned}$$

Let $\ell = \min\{i \mid p_{v_{j-1}}(i) = v_j\}$, let g be the agent that does the ℓth visit to v_{j-1} since round $t_{v_{j-1}}$, and let r be the round when that visit takes place, thus $x_g(r) = v_{j-1}$ and $x_g(r+1) = v_j$. By the minimality of ℓ, $r+1$ is the first round when some informed agent moves to v_j from v_{j-1}. Since $v_{j-1} \in S_{v_j}$, it follows that $r + 1 = t_{v_j}$. Then

$$\ell \leqslant \sum_{t_{v_{j-1}} \leqslant t \leqslant r} |Z_{v_{j-1}}(t)| = \sum_{t_{v_{j-1}} \leqslant t < t_{v_j}} |Z_{v_{j-1}}(t)|.$$

Also, from the coupling, $\pi_{v_{j-1}}(\ell) = p_{v_{j-1}}(\ell) = v_j$, which implies

$$\tau_{v_j} \leqslant \tau_{v_{j-1}} + \ell.$$

Combining all the above we obtain $C_{v_j}(t_{v_j}) \geqslant \tau_{v_{j-1}} + \ell \geqslant \tau_{v_j}$, completing the inductive proof of (3.7). Applying (3.7) for $j = k$, we obtain $\tau_u \leqslant C_u(t_u)$. $\square$

Lemma 3.3.8. *For any $u \in V$ and $t \geqslant t_u$, there is a canonical walk θ of length t with $Q(\theta) = C_u(t)$.*

Proof. We consider the same path $\langle v_0, v_1, \ldots, v_k \rangle$ as in the proof of Lemma 3.3.7, where $v_0 = s$, $v_k = u$, and for each $0 < j \leqslant k$, $v_{j-1} \in S_{v_j}$ and $C_{v_j}(t_{v_j}) = C_{v_{j-1}}(t_{v_j})$. Consider the canonical walk θ obtained from this path by adding between each pair of consecutive vertices v_{j-1} and v_j, $t_{v_j} - t_{v_{j-1}} - 1$ copies of v_{j-1}, and also appending after v_k a number of $t - t_{v_k}$ copies of v_k. It is then easy to show by induction that $Q(\theta) = C_u(t)$. $\square$

Recall that $\mathcal{B}_k$ is the set of concise descriptors of canonical walks of length k. The next lemma gives un upper bound on the congestion of a single canonical walk of length k given its concise descriptor.

Lemma 3.3.9. *Fix any $A_k \in \mathcal{B}_k$, and let η be the labelled canonical walk with concise descriptor A_k in* M-VISIT-EXCHANGE. *Then, for any $\lambda \geqslant 2e\gamma + 1$, $\mathbb{P}[Q(\eta) \leqslant \lambda k] \geqslant 1 - 2^{-(\lambda-1)k}$.*

Proof. Let $\eta = \langle \theta_0, g_1, \theta_1, \ldots, g_k, \theta_k \rangle$. Then $Q(\eta) = \sum_{1 \leqslant t \leqslant k} \delta_t$, where $\delta_t = |Z_{\theta_{t-1}}(t-1)|$. By the same reasoning as in the proof of Lemma 3.3.6, $Q(\eta)$ is stochastically dominated by $k + \sum_{1 \leqslant t \leqslant k} B_t$, where $B_1, \ldots, B_k$ are independent binomial random variables, such that $B_1 \sim \text{Bin}(\gamma n, 1/n)$ and, for $t > 1$, $B_t \sim \text{Bin}(\gamma d, 1/d)$. It follows that $\mathbb{E}[\,Q(\eta) - k\,] \leqslant k\gamma$, and

$$\mathbb{P}[Q(\eta) \geqslant \lambda k] = \mathbb{P}[Q(\eta) - k \geqslant (\lambda - 1)k] \leqslant 2^{-(f-1)k},$$

by a Chernoff bound, since $(\lambda - 1)k \geqslant 2e \cdot \mathbb{E}[\,Q(\eta) - k\,]$. $\square$

3.3.6 Proof of Theorem 3.3.1

Recall that our goal is to bound $\mathbb{P}[T_{\text{push}} \leqslant \lambda k]$ in terms of $\mathbb{P}[T_{\text{visitx}} \leqslant k]$. First we consider the case where k is at most logarithmic. By [ES11, Theorem 4.2] the expected cover time of n random walks, all starting from a fixed arbitrary vertex, is at least $\Omega(\log n)$. The exact same proof also implies that the cover time of n random walks starting from stationarity is at least $\Omega(\log n)$, w.h.p. This implies that $T_{\text{visitx}} = \Omega(\log n)$ w.h.p. Thus, there is some constant $\epsilon > 0$ such that if $k \leqslant \epsilon \log n$, $\mathbb{P}[T_{\text{visitx}} \leqslant k] \leqslant n^{-c}$. From this, the theorem's statement follows for $k \leqslant \epsilon \log n$. In the rest of the proof, we assume that $k \geqslant \epsilon \log n$.

We have $T_{\text{push}} = \max_{u \in V} \tau_u$, and from Lemma 3.3.7,

$$T_{\text{push}} \leqslant \max_{u \in V} C_u(t_u).$$

Since for any fixed realisation of VISIT-EXCHANGE and any $u \in V$, $C_u(t)$ is a non-decreasing function of t, and since $t_u \leqslant T_{\text{visitx}}$, it follows that

$$T_{\text{push}} \leqslant \max_{u \in V} C_u(T_{\text{visitx}}).$$

By Lemma 3.3.8, for any $u \in V$, there is a canonical walk θ of length $t = T_{\text{visitx}}$ with congestion $Q(\theta) = C_u(T_{\text{visitx}})$. Thus, there is also a *labelled* canonical walk η of length T_{visitx} with $Q(\eta) = Q(\theta) = C_u(T_{\text{visitx}})$. It follows that

$$T_{\text{push}} \leqslant \max_{\eta \in \mathcal{H}(T_{\text{visitx}})} Q(\eta), \tag{3.8}$$

where $\mathcal{H}(t)$ denotes the set of all labelled canonical walks of length t in VISIT-EXCHANGE.

Next we bound $\max_{\eta \in \mathcal{H}(k)} Q(\eta)$. Consider M-VISIT-EXCHANGE, and for any $A_k \in \mathcal{B}_k$, let η_{A_k} be the labelled canonical walk with concise descriptor A_k in M-VISIT-EXCHANGE. From Lemma 3.3.9, for any $A_k \in \mathcal{B}_k$ and $\lambda \geqslant 2e\gamma + 1$, $\mathbb{P}\left[Q(\eta_{A_k}) \leqslant \lambda k\right] \geqslant 1 - 2^{-(\lambda-1)k}$. Then

$$\mathbb{P}\left[\max_{A_k \in \mathcal{B}_k} Q(\eta_{A_k}) \leqslant \lambda k\right] \geqslant 1 - 2^{-(\lambda-1)k} \cdot |\mathcal{B}_k| \geqslant 1 - 2^{-(\lambda-1)k} \cdot (4b)^{2k},$$

by Lemma 3.3.4. Choosing the constant λ large enough so that $(\lambda - 1)/2 \geqslant 2\log(4b)$, yields

$$\mathbb{P}\left[\max_{A_k \in \mathcal{B}_k} Q(\eta_{A_k}) \leqslant \lambda k\right] \geqslant 1 - 2^{-(\lambda-1)k/2}.$$

From Lemma 3.3.5, the probability that all labelled canonical walks of length k have concise descriptors is at least $1 - 2^{-bk/4}\log(\alpha n)$, if $b \geqslant \max\{2\gamma e^2, 64\}$. It follows that

$$\mathbb{P}\left[\max_{A_k \in \mathcal{B}_k} Q(\eta_{A_k}) = \max_{\eta \in \mathcal{H}^*(k)} Q(\eta)\right] \geqslant 1 - 2^{-bk/4}\log(\alpha n),$$

where $\mathcal{H}^*(t)$ is the set of all labelled canonical walks of length t in M-VISIT-EXCHANGE. By Lemma 3.3.3, however, we can couple VISIT-EXCHANGE and M-VISIT-EXCHANGE, by using the same collection of random walks for both, such that the two processes are identical until round k with probability at least $1 - kn \cdot 2^{-ad}$. Thus

$$\mathbb{P}\left[\mathcal{H}(k) = \mathcal{H}^*(k)\right] \geqslant 1 - kn \cdot 2^{-\gamma d}.$$

Combining the last three inequalities above, we obtain

$$\mathbb{P}\left[\max_{\eta \in \mathcal{H}(k)} Q(\eta) \leqslant \lambda k\right] \geqslant 1 - 2^{-(\lambda-1)k/2} - 2^{-bk/4}\log(\alpha n) - kn \cdot e^{-\gamma d}.$$

Since $k \geqslant \epsilon \log n$ and $d \geqslant \beta \log n$, for any given constant $c > 0$ we can choose constants λ, b, γ large enough such that

$$\mathbb{P}\left[\max_{\eta \in \mathcal{H}(k)} Q(\eta) \leqslant \lambda k\right] \geqslant 1 - n^{-c}. \tag{3.9}$$

From (3.8) and (3.9), we obtain

$$\begin{aligned}
\mathbb{P}\left[T_{\text{push}} \leqslant \lambda k\right] &\geqslant \mathbb{P}\left[\max_{\eta \in \mathcal{H}(T_{\text{visitx}})} Q(\eta) \leqslant \lambda k\right], \qquad \text{by (3.8)} \\
&\geqslant \mathbb{P}\left[\{T_{\text{visitx}} \leqslant k\} \cap \left\{\max_{\eta \in \mathcal{H}(k)} Q(\eta) \leqslant \lambda k\right\}\right] \\
&\geqslant \mathbb{P}\left[T_{\text{visitx}} \leqslant k\right] - \mathbb{P}\left[\max_{\eta \in \mathcal{H}(k)} Q(\eta) > \lambda k\right] \\
&\geqslant \mathbb{P}\left[T_{\text{visitx}} \leqslant k\right] - n^{-c}, \qquad \text{by (3.9)}.
\end{aligned}$$

This completes the proof of Theorem 3.3.1.

3.4 Upper bounding T_{visitx} in terms of T_{push}

The following theorem formally states the other direction of Theorem 3.1.1, which is the upper bound on the broadcast time of VISIT-EXCHANGE by the broadcast time of PUSH in regular graphs.

Theorem 3.4.1. *Let $c, \alpha, \beta > 0$ be constants with $\alpha\beta$ sufficiently large (depending on c). There is a constant $\lambda > 0$, such that for any d-regular graph $G = (V, E)$ with $|V| = n$ and $d \geqslant \beta \log n$, and for any source $s \in V$, the broadcast times of* PUSH *and* VISIT-EXCHANGE*, with $|A| = \alpha n$ agents, satisfy*

$$\mathbb{P}\left[T_{\text{visitx}} \leqslant \lambda k\right] \geqslant \mathbb{P}\left[T_{\text{push}} \leqslant k\right] - n^{-c},$$

for any $k \geqslant 0$.

Theorem 3.4.1 is the only result in the thesis where α is not an arbitrary constant. Instead there is an interplay between the number of agents in VISIT-EXCHANGE (determined by α) and the degree of the graph (determined by β). From the proof, (3.12) in particular, it follows that $\alpha\beta \geqslant 8c+40$ suffices for the theorem to hold (note, that these constants have not been optimised). This relation is only required when $d = \Omega(\log n)$. If $d = \omega(\log n)$, then the theorem holds for arbitrary values of α.

From Theorem 3.4.1, it is immediate that if $T_{\text{push}} \leqslant T$ w.h.p., then $T_{\text{visitx}} = O(T)$ w.h.p. We can also bound $\mathbb{E}\left[T_{\text{visitx}}\right]$ in terms of $\mathbb{E}\left[T_{\text{push}}\right]$. Fix a constant $c > 3$ and choose the constant λ determined by Theorem 3.4.1. By [LP17, Theorem 11.6], the expected cover time of a random walk on a regular graph is at most $O(n^2)$, thus, by Markov's inequality, $T_{\text{visitx}} = O(n^2)$ with probability at least $1/2$. By repeated application, we have that $\mathbb{P}\left[T_{\text{visitx}} \geqslant T\right] \leqslant n^{-c}$, for some $T = O(n^2 \log n)$, and moreover, $\mathbb{P}\left[T_{\text{visitx}} \geqslant \ell T\right] \leqslant n^{-c\ell}$, for any integer $l \geqslant 1$.

$$\begin{aligned}
\mathbb{E}\left[T_{\text{visitx}}\right] &\leqslant \sum_{k=0}^{+\infty} \mathbb{P}\left[T_{\text{visitx}} \geqslant k\right], \text{ by Theorem 2.3.2,}\\
&\leqslant \sum_{k=0}^{T-1} \mathbb{P}\left[T_{\text{visitx}} \geqslant \lfloor k/\lambda \rfloor \cdot \lambda\right] + \sum_{k=T}^{+\infty} \mathbb{P}\left[T_{\text{visitx}} \geqslant k\right]\\
&\leqslant \lambda \sum_{\ell=0}^{\lfloor T/\lambda \rfloor} \left(\mathbb{P}\left[T_{\text{push}} \geqslant \ell\right] + n^{-c}\right) + T \sum_{\ell=1}^{+\infty} \mathbb{P}\left[T_{\text{visitx}} \geqslant \ell T\right], \quad \text{by Theorem 3.4.1,}\\
&\leqslant \lambda \sum_{\ell=0}^{+\infty} \mathbb{P}\left[T_{\text{push}} \geqslant \ell\right] + 3Tn^{-c}\\
&= O(\mathbb{E}\left[T_{\text{push}}\right]), \quad \text{by Theorem 2.3.2.}
\end{aligned}$$

3.4.1 Overview of the proof

We use a coupling which is similar to that in the proof of the converse result, stated in Theorem 3.3.1, but with a small change (which we describe momentarily). Unlike in the proof of Theorem 3.3.1, where we essentially consider all possible paths through which information travels, here we consider the first path which PUSH uses to inform each vertex and prove that VISIT-EXCHANGE progresses along the same path in a comparable number of rounds. Let $P = \langle u_0 = s, u_1, \ldots, u_k = u\rangle$ be one such path for vertex u of the graph G, i.e., each vertex u_i in the path gets informed by u_{i-1}. Let δ_i be the number of rounds it takes for u_{i-1} to sample (and inform) u_i in PUSH. We consider the same path in VISIT-EXCHANGE, and compare δ_i to the number D_i of rounds until some informed agent moves from u_{i-1} to u_i, counting from the round when u_{i-1} becomes informed. Note that $\sum_i \delta_i$ is precisely the round when u is informed in PUSH, while $\sum_i D_i$ is an upper bound on the round when u is informed in VISIT-EXCHANGE. It follows that we have to compare these two sums.

The coupling used for the proof of Theorem 3.3.1 seems suitable for this setup. Recall that there we let the lists of neighbours sampled by each vertex u once it is informed to be identical in PUSH and VISIT-EXCHANGE. A similar intuition applies here. Namely, on average each vertex is visited by $\Theta(1)$ agents per round, therefore, D_i and δ_i should be comparable. We can apply a similar technique and introduce a process M-VISIT-EXCHANGE to avoid some dependencies. This process is similar to VISIT-EXCHANGE except that it ensures that the number of agents in the neighbourhood of each vertex in each round is at least $\Omega(d)$. This means that the number of agents visiting a vertex in a single round of VISIT-EXCHANGE is bounded from below by $\mathrm{Bin}(\Omega(d), 1/d)$. We show that M-VISIT-EXCHANGE and VISIT-EXCHANGE are identical in the first polynomially many rounds, and therefore, we can consider M-VISIT-EXCHANGE in our proof.

There is, however, a problem with this proof plan. By fixing the first path P informing u in PUSH, we introduce *dependencies from the future*, i.e., we condition on the fact that P is indeed the first such path. The following idea modifies the coupling slightly and allows us to overcome this problem. We only consider the *odd rounds* of VISIT-EXCHANGE in the coupling, i.e., we match the list of neighbours that a vertex v samples in PUSH (in all rounds), to the list of neighbours that informed agents visit in round $2k+1$ after visiting u in round $2k$, for all $k \geqslant 0$. In even rounds, agents take steps *independently* of the coupled PUSH process. Fig. 3.6 illustrates the first five rounds of the coupling for the source vertex. The only difference for the other vertices is that the coupling does not start at the same round in both processes.

To summarise, under this coupling, the proof proceeds as follows. We introduce M-VISIT-EXCHANGE and show that it is identical to VISIT-EXCHANGE for the purpose of most of the proof. Then we fix *all* random choices made by PUSH, and thus, the information path P from the source to a fixed vertex u. Suppose u_i is the most recently informed vertex in P. Then, for each even round of VISIT-EXCHANGE, the vertex u_i is visited by at least one agent with constant probability, independently of the past and of the fixed choices in future odd rounds. If indeed some agent visits u_i, then in the subsequent odd round it visits a vertex dictated by the coupling. By a Chernoff bound on geometric random variables (indicating the number of rounds between consecutive visits to vertices in P in even rounds), we can show that $\sum_i D_i \leqslant \lambda \cdot (\sum_i \delta_i + \log n)$, w.h.p. We get rid of the $\log n$ term in the final bound by using the fact that $T_{\text{push}} = \Omega(\log n)$.

3.4.2 Definition of another coupling

We recall the notation from Section 3.3.2. For each vertex $u \in V$, τ_u is the round when u first gets informed in PUSH. For $i \geqslant 1$, let $\pi_u(i)$ be the ith vertex that u samples after being informed, i.e., the vertex it samples in round $\tau_u + i$. In VISIT-EXCHANGE, t_u is the round when vertex u gets informed. For any agent $g \in A$ and $t \geqslant 0$, we denote by $x_g(t)$

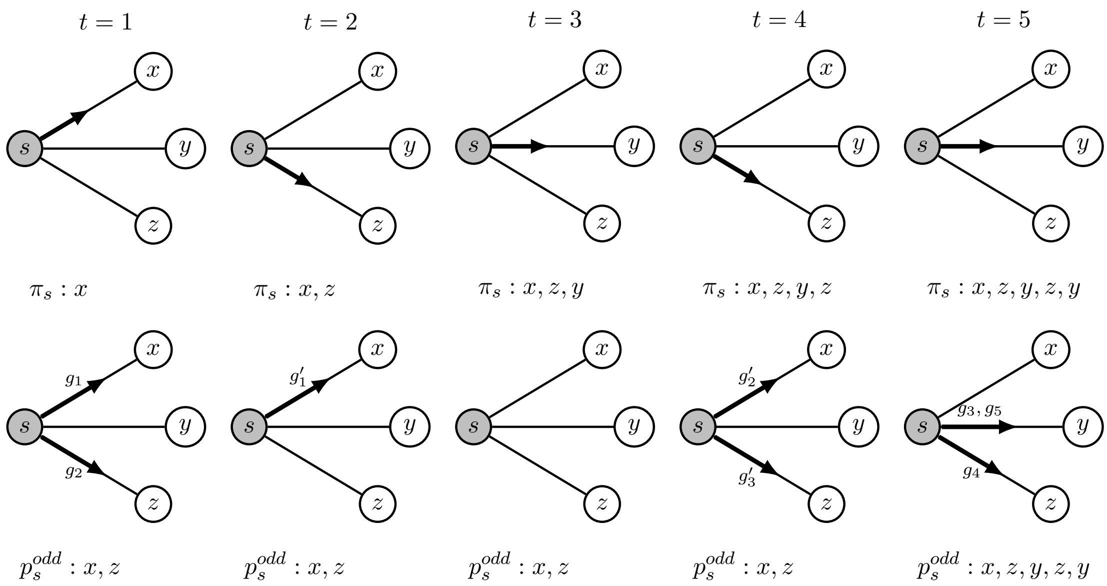

Figure 3.6 The first 5 rounds of the coupled processes used in the proof of Theorem 3.4.1 for the source vertex s. The top row corresponds to PUSH and the bottom row to VISIT-EXCHANGE. π_u is the list of neighbours that s contacts in the first five rounds of PUSH. p_u^{odd} is the list of neighbours of s where agents are sent in odd rounds. Agents g_i follow the choices made by PUSH, while agents g'_i proceed independently of PUSH. Note that in round $t = 3$, no agents are present in s, hence p_u^{odd} is remains unchanged.

the random walk performed by agent g. Thus, $x_g(t)$ is the vertex that g visits in round t. Recall that $Z_u(t)$ is the set of all agents that visit u in round t.

For vertex $u \in V$, consider all visits to u in *even* rounds $t \geqslant t_u$, in chronological order, ordering visits in the same round with respect to a predefined but arbitrary total order over agents. We call these visits *even visits* to vertex u. For each $i \geqslant 1$, consider the agent g that performs the ith even visit and let $p_u^{odd}(i)$ be the vertex that g visits in the next (*odd*) round. Formally, let

$$W_u^{even} = \{(t, g) \mid t \geqslant t_u, t \in \mathbb{Z}_+^{even}, x_g(t) = u\},$$

where $\mathbb{Z}_+^{even}$ is the set of non-negative even integers. We order the elements of W_u^{even} such that $(t, g) < (t', g')$ if $t < t'$, or $t = t'$ and $g < g'$. If (t, g) is the ith smallest element in W_u^{even}, then $p_u^{odd}(i) = x_g(t + 1)$.

Coupling. We couple processes PUSH and VISIT-EXCHANGE by setting $\pi_u(i) = p_u^{odd}(i)$, for all $i \geqslant 1$. Formally, let $\{w_u(i)\}_{u \in V, i \geqslant 1}$, be a collection of independent random variables each taking a uniformly random value from the set $\Gamma(u)$ of u's neighbours in G. For all $u \in V$ and $i \geqslant 1$, we set

$$\pi_u(i) = p_u^{odd}(i) = w_u(i).$$

See Fig. 3.6 for an illustration of the coupling for the source vertex s.

3.4.3 A modified Visit-Exchange process

Here we introduce the M-VISIT-EXCHANGE process, which is identical to VISIT-EXCHANGE except for the following simple modification. Recall that d is the degree of the regular graph G that we consider. If in some odd round $t \geqslant 0$, there is a vertex $u \in V$ for which the next condition is *not* true, i.e.,

$$\sum_{v \in \Gamma(u)} |Z_v(t)| \geqslant \frac{|A|}{2n} \cdot d = \alpha d/2, \tag{3.10}$$

then before round $t+1$, we add a minimal set of new agents to the graph such that the above condition holds for all vertices u. An agent g added to vertex u adopts the state (informed or non-informed) of u at the end of round t.

Recall that $|A| = \alpha n$. The following lemma allows us to consider the M-VISIT-EXCHANGE process in the rest of the proof, and argue that the results also hold for VISIT-EXCHANGE.

Lemma 3.4.2. *The probability that the condition (3.10) holds simultaneously for all $u \in V$ and $0 \leqslant t < k$ is at least $1 - kn \cdot 2^{-\alpha d/8}$.*

Proof. We first fix a vertex $u \in V$ and a round t such that $0 \leqslant t < k$, and prove that (3.10) holds for u at round t, and then apply a union bound.

For an agent $g \in A$, let X_g be an indicator random variable that g is in the neighbourhood $\Gamma(u)$ of vertex u in round t. Then, $X = \sum_{g \in A} X_g$ is the number of agents in the neighbourhood of u in round t. We have that $\mathbb{E}[X] = |A| \cdot |\Gamma(u)|/n = \alpha d$ and since the random variables X_g are independent, it follows by a Chernoff bound that

$$\mathbb{P}[X \geqslant \alpha d/2] \geqslant 1 - e^{-\mathbb{E}[X]/8} = 1 - e^{-\alpha d/8}.$$

The lemma follows after applying a union bound for each $0 \leqslant t < k$ and each $u \in V$. □

3.4.4 Proof of Theorem 3.4.1

We first compare the times until a given vertex u gets informed in PUSH and in M-VISIT-EXCHANGE.

Lemma 3.4.3. *The coupling described in Section 3.4.2, when applied to* PUSH *and* M-VISIT-EXCHANGE, *yields the following property. For any constant $c > 0$, there is a constant $\lambda > 0$ such that for any $u \in V$,*

$$\mathbb{P}[t'_u \geqslant \lambda(\tau_u + \log n)] \leqslant n^{-(c+2)},$$

where τ_u and t'_u are the rounds when u is informed in the coupled processes PUSH *and in* M-VISIT-EXCHANGE, *respectively.*

Proof. In this proof, we will use the same notation for M-VISIT-EXCHANGE as defined for VISIT-EXCHANGE. (We used t'_u instead of t_u in the lemma's statement to avoid confusion when we apply the lemma, but in the proof there is no such issue, because only M-VISIT-EXCHANGE is used.)

As described in the proof overview, we consider a path from the source s to vertex u that PUSH uses to inform u, and count the number of rounds VISIT-EXCHANGE takes to traverse the same path. First, we consider a single edge (v, w) such that w is informed by v in a realisation of PUSH that we fix. We also fix the first t_v rounds of M-VISIT-EXCHANGE, i.e., until v becomes informed. Let $\delta_{v,w} = \tau_w - \tau_v$ be the number of rounds that the PUSH process takes to inform w, counting from when v gets informed. Similarly, we define $D_{v,w} = t_w - t_v$ for M-VISIT-EXCHANGE. We will bound $D_{v,w}$ in terms of $\delta_{v,w}$.

Recall that we have defined a natural total order over the set W_v^{even} of even visits to vertex v. For $j \geqslant 1$, let (t, g) be the jth element of W_v^{even} in that order. By the coupling, at the odd round $t + 1$, agent g will move to the neighbour of v that is sampled by PUSH in round $\pi_v(j) = \tau_v + j$. In particular, since $\pi_v(j) = w$ for $j = \delta_{v,w}$, vertex w becomes informed after $\delta_{v,w}$ even visits to v in M-VISIT-EXCHANGE (possibly earlier).

Formally, let $B_v^{(j)}$ be the number of M-VISIT-EXCHANGE rounds between even visits $j - 1$ and j (when $j = 1$, $B_v^{(j)}$ is the number of rounds until the first even visit since t_v). $B_v^{(j)}$ can be 0, if two agents visit v at the same even round. With this definition,

$$D_{v,w} \leqslant \sum_{j=1}^{\delta_{v,w}} B_v^{(j)}. \tag{3.11}$$

By condition (3.10), there are at least $\alpha \cdot d/2$ agents in the neighbourhood of v at any round of M-VISIT-EXCHANGE. Let $p = 1 - e^{-\alpha/2}$ and recall that, for an even $t > 0$, the agents move independently from PUSH, and therefore, some agent visits v in round t with probability at least $1 - (1 - 1/d)^{\alpha d/2} \geqslant p$. Also, for $t = 0$, when agents are placed according to the stationary distribution, some agent is placed at v with probability $1-(1-1/n)^{\alpha n} \geqslant 1-e^{-\alpha} \geqslant p$. It follows that the number of rounds between two even visits to v, namely $B_v^{(j)}$ for $1 \leqslant j \leqslant \delta_{v,w}$, is stochastically dominated by $2 \cdot F_v^{(j)}$, where $\{F_v^{(j)}\}_{j\geqslant 1}$ is a collection of independent geometric random variables with success probability p. The coefficient 2 appears because we have to take into account both odd and even rounds. In other words, for any $b \geqslant 0$ and $1 \leqslant j \leqslant \delta_{v,w}$,

$$\mathbb{P}\left[B_v^{(j)} \leqslant b \mid B_v^{(1)}, \ldots, B_v^{(j-1)}\right] \geqslant \mathbb{P}\left[2 \cdot F_v^{(j)} \leqslant b\right].$$

Using Lemma 2.3.9, we get that, given v is informed, $D_{v,w}$ is stochastically dominated by $2 \cdot \sum_{j=1}^{\delta_{v,w}} F_v^{(j)}$:

$$\mathbb{P}\left[D_{v,w} \leqslant b \mid t_v\right] \geqslant \mathbb{P}\left[\sum_{j=1}^{\delta_{v,w}} B_v^{(j)} \leqslant b \mid t_v\right] \geqslant \mathbb{P}\left[2 \cdot \sum_{j=1}^{\delta_{v,w}} F_v^{(j)} \leqslant b\right].$$

We apply the above result to all edges on the path from s to u through which PUSH informed u. Let $P_u = \langle s = u_0, u_1, \ldots, u_k = u\rangle$ be a path in G such that, in PUSH, u_i is informed from u_{i-1}, for all $1 \leqslant i \leqslant k$. By definition of τ_u, u_{i-1} samples its neighbour u_i at round τ_{u_i}. Define $\delta_i = \tau_{u_i} - \tau_{u_{i-1}}$ and $D_i = t_{u_i} - t_{u_{i-1}}$ for $1 \leqslant i \leqslant k$. From our result above for a single edge it follows that

$$\mathbb{P}\left[D_i \leqslant b \mid D_1, \ldots, D_{i-1}\right] \geqslant \mathbb{P}\left[2 \cdot \sum_{j=1}^{\delta_i} F_{u_i}^{(j)} \leqslant b\right].$$

Once again, by Lemma 2.3.9 and the fact that $t_u = t_{u_k} = \sum_{i=1}^{k} D_i$, we have that t_u is stochastically dominated by $2F = 2 \cdot \sum_{i=1}^{k} \sum_{j=1}^{\delta_i} F_{u_{i-1}}^{(j)}$, i.e., for any $b \geqslant 0$,

$$\mathbb{P}\left[t_u \leqslant b\right] \geqslant \mathbb{P}\left[2F \leqslant b\right].$$

The random variable F is a sum of exactly τ_u independent and identical geometrically distributed random variables with mean $1/p$, hence, $\mathbb{E}[F] = \tau_u/p$. Thus, for any constant $\lambda \geqslant 4/p$, by Lemma 2.3.5,

$$\begin{aligned}\mathbb{P}[t_u \geqslant \lambda(\tau_u + \log n)] &\leqslant \mathbb{P}\left[F \geqslant \frac{\lambda}{2}(\tau_u + \log n)\right] \\ &\leqslant \exp\left(-\frac{\lambda(\tau_u + \log n) \cdot p}{16}\right) \\ &\leqslant n^{-\lambda p/16}.\end{aligned}$$

Choosing λ large enough so that $\lambda p/16 \geqslant c + 2$, completes the proof. □

We can now complete the proof of our main result, where we use the previous lemma relating M-VISIT-EXCHANGE and PUSH together with the earlier equivalence of M-VISIT-EXCHANGE and VISIT-EXCHANGE.

Theorem 3.4.1. *Let $c, \alpha, \beta > 0$ be constants with $\alpha\beta$ sufficiently large (depending on c). There is a constant $\lambda > 0$, such that for any d-regular graph $G = (V, E)$ with $|V| = n$ and $d \geqslant \beta \log n$, and for any source $s \in V$, the broadcast times of* PUSH *and* VISIT-EXCHANGE*, with $|A| = \alpha n$ agents, satisfy*

$$\mathbb{P}[T_{\text{visitx}} \leqslant \lambda k] \geqslant \mathbb{P}[T_{\text{push}} \leqslant k] - n^{-c},$$

for any $k \geqslant 0$.

Proof. Recall that τ_u, t_u and t'_u are the rounds when vertex u gets informed in PUSH, VISIT-EXCHANGE, and M-VISIT-EXCHANGE, respectively. From Lemma 3.4.3, and a union bound over all vertices, we obtain that for any constant $c > 0$, there is a constant $\lambda > 0$ such that

$$\mathbb{P}[\forall\, u \in V \mid t'_u \leqslant \lambda(\tau_u + \log n)] \geqslant 1 - n \cdot n^{-(c+2)}.$$

Thus,

$$\mathbb{P}\left[\max_{u \in V} t'_u \leqslant \lambda\left(\max_{u \in V} \tau_u + \log n\right)\right] \geqslant 1 - n^{-(c+1)}.$$

It follows that for any $k \geqslant 0$,

$$\begin{aligned}\mathbb{P}\left[\max_{u \in V} t'_u \leqslant \lambda\,(k + \log n)\right] &\geqslant \mathbb{P}\left[\max_{u \in V} t'_u \leqslant \lambda\left(\max_{u \in V} \tau_u + \log n\right) \cap \max_{u \in V} \tau_u \leqslant k\right] \\ &\geqslant \mathbb{P}\left[\max_{u \in V} \tau_u \leqslant k\right] - n^{-(c+1)}.\end{aligned}$$

From Lemma 3.4.2, it follows that

$$\mathbb{P}\left[\max_{u \in V} t'_u \leqslant \lambda\,(k + \log n)\right] - \mathbb{P}\left[\max_{u \in V} t_u \leqslant \lambda\,(k + \log n)\right] \leqslant \lambda(k + \log n) \cdot n \cdot e^{-\alpha d/8}.$$

Combining the last two inequalities above we obtain

$$\mathbb{P}\left[\max_{u \in V} t_u \leqslant \lambda\,(k + \log n)\right] \geqslant \mathbb{P}\left[\max_{u \in V} \tau_u \leqslant k\right] - n^{-(c+1)} - \lambda(k + \log n) \cdot n \cdot e^{-\alpha d/8}.$$

Substituting $T_{\text{visitx}} = \max_{u \in V} t_u$ and $T_{\text{push}} = \max_{u \in V} \tau_u$, and using $d \geqslant \beta \log n$, yields

$$\mathbb{P}\left[T_{\text{visitx}} \leqslant \lambda\left(k + \log n\right)\right] \geqslant \mathbb{P}\left[T_{\text{push}} \leqslant k\right] - n^{-(c+1)} - \lambda(k + \log n) \cdot n^{1-\alpha\beta/8}. \tag{3.12}$$

This implies the theorem for $\log n \leqslant k \leqslant \text{poly}(n)$. For larger k, the theorem follows from the known polynomial upper bound on the cover time on regular graphs. For smaller k, it follows from the fact that $T_{\text{push}} = \Omega(\log n)$, w.h.p. □

Chapter 4

General bounds in terms of node degrees and diameter

4.1 Introduction

In this chapter we present two results on the broadcast time of VISIT-EXCHANGE. The first one is for sparse regular graphs, bounding the broadcast time in terms of the diameter and the degree of the graph. The second bound applies to general graphs, and is in terms of the average degree and the diameter of the graph. The two results use a similar technique that involves the return probability of a random walk.

Theorem 4.1.1. *For any d-regular graph G with $d = O(\log n)$ and any source vertex, $T_{\text{visitx}} = \widetilde{O}(d \cdot \text{diam}(G) + \log^3 n/d)$, w.h.p., where the tilde notation hides factors of order at most* $(\log \log n)^2$.

In the above bound, the dependence on the diameter is best possible (e.g., the broadcast time along a cycle of d-cliques is proportional to the path length multiplied by d). An additive term is also needed when the diameter is sub-logarithmic, but it is not clear whether the term $\log^3 n/d$ is tight. Recall that the corresponding upper bound for T_{push} in [Fei+90] is $O(d \cdot (\text{diam}(G) + \log n))$. Thus, it would be reasonable to guess that the right additive term is $d \cdot \log n$. However, the example in Theorem 3.1.2 shows that the term must be at least $\widetilde{\Omega}(\log^2 n)$. We conjecture that the tight bound is $\widetilde{O}(d \cdot \text{diam}(G) + \log^2 n)$.

The proof of Theorem 4.1.1 bounds the time that the information takes to spread along a given (shortest) path in the graph. We divide time into phases of length $r = \Theta(\log^2 n)$ rounds, and in each phase, we lower bound the probability that the information spreads along a sub-path of length $\widetilde{\Omega}(\log^2 n/d)$. For $d = \omega(\log \log n)$, we show this probability to be $1 - e^{-\Omega(d)}$. Moreover, we ensure that this probability bound holds, essentially, *independently* of previous phases, by considering every other phase. We prove the bound by showing a concentration result on the number of agents at the neighbourhood of each individual vertex in the sub-path, at each round of the phase, and then applying a union bound. To boost the above probability to $1 - e^{-\Omega(\log n)}$, we need $\log n/d$ phases, which yields the $\log^3 n/d$ term of the bound. For the case of $d = O(\log \log n)$, we use a similar approach, but give a lower bound on the number of agents visiting a vertex over an interval of multiple rounds (rather than considering the neighbourhood in each round).

For non-regular graphs, a similar analysis as for Theorem 4.1.1 yields the following result.

Theorem 4.1.2. *Let $d_{\min}$ and d_{avg} be the minimum and average degrees of a graph G, respectively. If $d_{\min} = \Omega(d_{\text{avg}})$, then $T_{\text{visitx}} = O(d_{\text{avg}} \cdot \log^2 n \cdot (\text{diam}(G) + \log n))$, w.h.p., for any source vertex.*

Even though this bound is likely not tight, it is interesting because it does not have an analogue for randomised rumour spreading. The corresponding upper bound for PUSH, shown in [Fei+90], is $O(d_{\max} \cdot (\text{diam}(G) + \log n))$, where $d_{\max}$ is the maximum degree. Note also that many common networks, including preferential attachment graphs, have a constant average degree, a logarithmic diameter, but a polynomial maximum degree. In such cases the bound in Theorem 4.1.2 represents an exponential improvement over the corresponding bound from [Fei+90]. The graph consisting of two stars with their centers connected by an edge, illustrated in Fig. 3.1(b), shows this difference on a concrete example. It should be noted that when d_{avg} is constant, the bound from Theorem 4.1.2 is tight up to logarithmic factors since the diameter of the graph is a trivial lower bound for the broadcast time.

Finally, for graphs with constant average degree, our result bears some resemblance to the result of [Hae15] (see also [Cen+17]), which is an algorithm that spreads an information in $O(\log^2 n \cdot \text{diam}(G))$. While our process obviously benefits from a more efficient bandwidth utilisation (i.e., higher degree vertices tend to initiate more connections in each round), it is simpler in the sense that it requires no memory about previously used edges as opposed to [Hae15; Cen+17]. Moreover, these works require messages of size linear in n.

The main technical tool we use in this chapter is an upper bound on the return probability of a random walk from [OP19], which holds for arbitrary graphs and gives rise the length $r = O(\log^2 n)$ of a phase in our analysis. For other graph topologies, where better bounds on return probability hold, r can be smaller, resulting in a tighter analysis of VISIT-EXCHANGE. This is the case for expanders, where it is possible to set $r = O(\log \log n)$ (Section 5.5).

4.1.1 Road-map

In Section 4.2 we use a bound on return probability of a single walk by [OP19] to prove a technical result that is the necessary building block for the two theorems proved in this chapter. In Sections 4.3 and 4.4 we present the proofs of Theorems 4.1.1 and 4.1.2, respectively.

4.2 Preliminaries

Let $G = (V, E)$ be any graph (not necessarily a regular one), and let A be the set of agents in VISIT-EXCHANGE. We denote the ratio $|A|/n$ by α, which is a constant since we assume that $|A| = \Theta(n)$. The agents in A start their walks from the stationary distribution π. For a vertex u, let $N_u(t)$ be the number of agents that are at vertex u at round t. For an integer $r > 0$ and round t, let

$$\hat{N}_u(t, r) = \mathbb{E}\left[N_u(t + r) \mid N_v(t), \text{ for all } v \in V\right] = \sum_{v \in V} p^r_{v,u} \cdot N_v(t),$$

where $p^r_{v,u}$ is the probability that a random walk starting from v is at u after exactly r rounds. Note that $\hat{N}_u(t, r)$ is a random variable that depends on the positions of the agents in round t, and $\mathbb{E}\left[\hat{N}_t(t,r)\right] = \mathbb{E}\left[N_u(t+r)\right] = \sum_{u\in V} p^r_{v,u} \cdot \mathbb{E}\left[N_v(t)\right]$. The following key lemma allows us to analyse VISIT-EXCHANGE by splitting the process into phases of r rounds and argue that progress is made in a phase, independently of the past phases.

Lemma 4.2.1. *For any vertex u, round t, and integer r,*

$$\mathbb{P}\left[\hat{N}_u(t,r) \leqslant |A| \cdot \pi(u)/2\right] \leqslant \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right).$$

Proof. Let $X^t_{v,g}$ be an indicator random variable, which is 1 when agent g is at vertex v at round t. Then, $N_v(t) = \sum_{g\in A} X^t_{v,g}$, which implies

$$\hat{N}_u(t,r) = \sum_{v\in V} p^r_{v,u} \sum_{g\in A} X^t_{v,g} = \sum_{g\in A}\sum_{v\in V} p^r_{v,u} \cdot X^t_{v,g} = \sum_{g\in A} Y_g,$$

where Y_g is the internal sum above for agent g. The random variables Y_g, $g \in A$, are independent, since the agents perform independent random walks. We compute the second moment of random variables Y_g to argue about the concentration of $\hat{N}_u(t,r)$ around its mean.

$$\begin{aligned}
\mathbb{E}\left[Y_g^2\right] &= \mathbb{E}\left[\sum_{v_1,v_2\in V} p^r_{v_1,u} p^r_{v_2,u} \cdot X^t_{v_1,g} \cdot X^t_{v_2,g}\right] \\
&= \sum_{v\in V} \left(p^r_{v,u}\right)^2 \cdot \mathbb{E}\left[X^t_{v,g}\right], \quad \text{as } g \text{ cannot be in two vertices simultaneously,} \\
&= \sum_{v\in V} p^r_{v,u} \cdot \left(p^r_{v,u} \cdot \pi(v)\right), \quad \text{since } g \text{ is placed according to } \pi, \\
&= \sum_{v\in V} p^r_{u,v} \cdot \left(\pi(u) \cdot p^r_{u,v}\right), \quad \text{by reversibility,} \\
&= \pi(u) \cdot \sum_{v\in V} p^r_{u,v} \cdot p^r_{v,u} \\
&= \pi(u) \cdot p^{2r}_{u,u}.
\end{aligned}$$

Also, since the agents are initially distributed according to the stationary distribution π,

$$\mathbb{E}\left[\hat{N}_u(t,r)\right] = \mathbb{E}\left[N_u(t+r)\right] = |A| \cdot \pi(u).$$

We apply Theorem 2.3.7, setting $\lambda = \mathbb{E}\left[\hat{N}_u(t,r)\right]/2$ and $M = 0$, to obtain

$$\begin{aligned}
\mathbb{P}\left[\hat{N}_u(t,r) \leqslant |A| \cdot \pi(u)/2\right] &\leqslant \exp\left(-\frac{\lambda^2}{2 \cdot \sum_{g\in A} \mathbb{E}\left[Y_g^2\right]}\right) \\
&\leqslant \exp\left(-\frac{(|A| \cdot \pi(u))^2}{8 \cdot \sum_{g\in A} \pi(u) \cdot p^{2r}_{u,u}}\right) = \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right). \quad \square
\end{aligned}$$

Next we present bounds on the return probability of a random walk that will be applied in Lemma 4.2.1 in the main proofs.

Lemma 4.2.2 ([OP19, Theorem 1.2])**.** *For a lazy random walk $X(t)$ that starts at vertex $u \in V$ of graph $G = (V, E)$, and has holding probability $1/2$,*

$$\mathbb{P}[X(t) = u] - \pi(u) \leqslant \frac{10 \cdot \deg(u)}{d_{\min}\sqrt{t+1}} \cdot \min\left\{1, \sqrt{\frac{t_{rel}}{t+1}}\right\},$$

where d_{min} is the minimal degree of the graph and t_{rel} is the relaxation time of the random walk.[1]

Lemma 4.2.2 concerns a lazy random walk, while in this thesis we consider VISIT-EXCHANGE for non-lazy random walks. Thus, we have to extend it to the non-lazy case. First, we prove the following auxiliary lemma.

Lemma 4.2.3. *For a simple random walk $X(t)$ on a graph $G = (V, E)$ and any vertex u, $\mathbb{P}[X(2t) = u \mid X(0) = u]$ is non-increasing for any integer $t \geqslant 0$, i.e., the return probability does not increase at even steps.*[2]

Proof. In the proof we use the results from [LP17, Chapter 12]. Let P be the transition matrix of the walk $X(t)$ and π be the stationary distribution, i.e, $\pi(u) = \deg(u)/(2|E|)$, P is a reversible Markov chain, i.e., for any $u, v \in V$

$$\pi(u) \cdot P_{u,v} = \pi(v) \cdot P_{v,u}.$$

(This fact is also easy to see by simply computing the transition probabilities, given the degrees of the vertices.) By [LP17, Lemmas 12.1-12.2], P has real-valued eigenfunctions f_i corresponding to real eigenvalues λ_i, for $i \in \{1, \dots, |V|\}$ and $|\lambda_i| \leqslant 1$. Moreover, by the same lemmas, for any $t \geqslant 0$, by we can write

$$P^{2t}(u, u) = \pi(u) \cdot \sum_{i=1}^{|V|} f_i(u)^2 \lambda_i^{2t},$$

which is non-increasing in t. □

Lemma 4.2.4. *Let $X(t)$ be a non-lazy random walk on a graph $G = (V, E)$, and $X'(t)$ be a lazy walk with holding probability $1/2$. If both walks start from some vertex $u \in V$, then for any even $t \geqslant 0$,*

(a) $\mathbb{P}[X(t) = u] \leqslant 2 \cdot \mathbb{P}[X'(t) = u]$,

(b) $\mathbb{P}[X(t) = u] \leqslant \frac{\deg(u)}{|E|} + \frac{20 \cdot \deg(u)}{d_{\min}\sqrt{t+1}}$.

Proof. Let L_t be the number of times $X'(t)$ stays put in its first t rounds. Then,

$$\mathbb{P}[X'(t) = u] \geqslant \sum_{t'=0}^{t/2} \mathbb{P}[X'(t) = u \mid L_t = 2t'] \cdot \mathbb{P}[L_t = 2t']$$

[1]The relaxation time is the reciprocal of the spectral gap of the transition matrix P of the random walk. In the later proofs, we do not use the bound involving the relaxation time.

[2]This in fact is true for any reversible Markov chain using the same proof, but here we are concerned with simple random walks.

$$= \sum_{t'=0}^{t/2} \mathbb{P}\left[X(t-2t')=u\right] \cdot \binom{t}{2t'} \cdot 2^{-t}$$
$$\geqslant \sum_{t'=0}^{t/2} \mathbb{P}\left[X(t)=u\right] \cdot \binom{t}{2t'} \cdot 2^{-t}, \quad \text{by Lemma 4.2.3, since } t-2t' \text{ is even.}$$

By expanding $0 = (1+(-1))^t$ using the binomial theorem, we get that $\sum_{t'=0}^{t/2} \binom{t}{2t'} = \sum_{t'=0}^{t/2} \binom{t}{2t'+1} = 2^{t-1}$, since the sum of the two sides of this equality is 2^t. Therefore,

$$\mathbb{P}\left[X'(t)=u\right] \geqslant \frac{1}{2} \cdot \mathbb{P}\left[X(t)=u\right],$$

which completes the proof of part (a). Part (b) of the lemma follows by an application of Lemma 4.2.2 to upper bound $\mathbb{P}\left[X'(t)=u\right]$ and substituting $\pi(u) = \deg(u)/(2|E|)$. □

4.3 Upper bound for regular graphs

In this section we prove Theorem 4.1.1. Suppose that $G = (V, E)$ is a d-regular graph with $d = O(\log n)$, thus $\pi(u) = 1/n$ for any $u \in V$. For a constant $\rho > 0$ define $r = r(\rho)$ as the smallest even integer such that

$$r \geqslant \max\{\rho \cdot \log^2 n,\ 256d \cdot \log n/\alpha\} = \Theta(\log^2 n). \tag{4.1}$$

We modify the VISIT-EXCHANGE process to create a new process called M-VISIT-EXCHANGE$_r$, as follows: At the end of each round $t \geqslant 0$, we add a minimal set of agents to the process to make sure that $\hat{N}_u(t,r) \geqslant |A| \cdot \pi(u)/2 = \alpha/2$, for every vertex u. Next we prove that, in the first polynomially many rounds M-VISIT-EXCHANGE$_r$ and VISIT-EXCHANGE are equivalent, w.h.p. Therefore, the results that we prove for M-VISIT-EXCHANGE$_r$, also hold for VISIT-EXCHANGE, w.h.p. This technique allows us to avoid dealing with dependencies of the random walks, which would arise if we directly analysed VISIT-EXCHANGE conditioned on $\hat{N}_u(t,r) \geqslant \alpha/2$ for all u and t. (Similar modified processes are also used in Chapters 3 and 6, as well as in Theorem 4.1.2.)

Lemma 4.3.1. *For any constant $c > 0$, there is a constant ρ, such that for $r = r(\rho)$ VISIT-EXCHANGE and M-VISIT-EXCHANGE$_r$ are identical for the first T' rounds of their execution with probability at least $1 - T' \cdot n^{-(c+2)}$.*

Proof. By Lemma 4.2.4, $p_{u,u}^{2r} \leqslant \frac{2}{n} + \frac{20}{\sqrt{2r+1}} \leqslant \frac{20}{\sqrt{r}}$, since $r = O(\log^2 n)$. For $t < T'$, we substitute the above inequality into Lemma 4.2.1, and use the fact that $|A| \cdot \pi(u) = \alpha$, to get that

$$\mathbb{P}\left[\hat{N}_u(t,r) \leqslant \alpha/2\right] \leqslant \exp\left(-\frac{\alpha}{8 \cdot p_{u,u}^{2r}}\right) \leqslant \exp\left(-\frac{\alpha}{160} \cdot \sqrt{r}\right) \leqslant n^{-(c+3)},$$

for a sufficiently large constant ρ. By applying a union bound over all vertices u and rounds $t < T'$, we complete the proof. □

Consider two vertices u and v with distance $O(r/\max\{d, \log^2 \log n\})$, and assume u is informed at some round t_0. The next key lemma provides a lower bound for the probability that v becomes informed $O(r)$ rounds after t_0. The lemma holds for any execution prefix of M-VISIT-EXCHANGE$_r$ up to round t_0, which means we can apply it repeatedly to prove Theorem 4.1.1. Let $\mathcal{K}_t$ be the σ-field that determines the execution of M-VISIT-EXCHANGE$_r$ until round t.

Lemma 4.3.2. *Let $h = \max\{d, \log\log n\}$, and $k_{\max}(\gamma) = \frac{\gamma \cdot r}{\max\{d, (\log\log n)^2\}}$. There are constants $\gamma, \beta > 0$, such that the following holds for any t_0 and $u, v \in V$ with $\operatorname{dist}(u,v) \leqslant k_{\max}(\gamma)$: Given $\mathcal{K}_{t_0}$ and that u is informed at round t_0, vertex v is informed at round $t_0 + 2r$ with probability at least $1 - e^{-\beta \cdot h}$.*

Proof. Case $d = \omega(\log\log n)$. To simplify the presentation, we assume $t_0 = 0$ and omit the conditioning on $\mathcal{K}_{t_0}$ throughout the proof. Fix the constant γ such that $k_{\max}(\gamma) \leqslant \frac{\alpha r}{256d}$. Consider two vertices u, v such that a shortest path between them is $\langle u = u_0, \ldots, u_k = v \rangle$, where $k = \operatorname{dist}(u,v) \leqslant k_{\max}(\gamma)$. For a round $t \geqslant r$ and $i \in \{0, \ldots, k-1\}$, let $Z_{i,t}$ be the number of agents in the neighbourhood $\Gamma(u_i)$ of vertex u_i at round t. Then, by definition of M-VISIT-EXCHANGE$_r$,

$$\mathbb{E}[Z_{i,t}] = \sum_{w \in \Gamma(u_i)} \mathbb{E}[N_{u_i}(t)] = \sum_{w \in \Gamma(u_i)} \mathbb{E}\left[\hat{N}_{u_i}(t-r, r)\right] \geqslant \alpha \cdot d/2.$$

Since the agents make independent random walks, by a Chernoff bound we get that

$$\mathbb{P}[Z_{i,t} \geqslant \alpha \cdot d/4] \geqslant 1 - e^{-\alpha \cdot d/16}.$$

If $\mathcal{E}$ is the event that $Z_{i,t} \geqslant \alpha \cdot d/4$ for all $i \in \{0, \ldots k-1\}$ and $t \in \{r, \ldots 2r\}$ simultaneously, then, by a union bound,

$$\mathbb{P}[\mathcal{E}] \geqslant 1 - k \cdot r \cdot e^{-\alpha \cdot d/16} \geqslant 1 - e^{-\beta d}/2,$$

for a small enough constant β, because $kr = O(\operatorname{poly}(\log n))$ and $d = \omega(\log\log n)$.

We modify M-VISIT-EXCHANGE$_r$ as follows: If $\mathcal{E}$ does not hold, then we add a minimum number of agents to the process so that $\mathcal{E}$ holds. We call the new process R-VISIT-EXCHANGE$_r$, and observe that M-VISIT-EXCHANGE$_r$ and R-VISIT-EXCHANGE$_r$ are identical with probability at least $1 - e^{-\beta d}/2$.

We divide the rounds $r, \ldots, 2r-1$ of R-VISIT-EXCHANGE$_r$ into $r/2$ phases of 2 rounds each. For each $0 \leqslant i < r/2$, let $\mathcal{K}'_i$ be the σ-algebra which determines the execution prefix of R-VISIT-EXCHANGE$_r$ until round $r + 2i \leqslant 2r$. Let p_i be the largest integer, between 0 and k, such that vertex $w = u_{p_i}$ is informed at round $r + 2i$. If $p_i < k$, then each agent that is in the neighbourhood of w in round $r + 2i$, informs vertex u_{p_i+1} after two rounds, with probability $1/d^2$, by going through w. Define a Bernoulli random variable X_i, such that $X_i = 1$ if $p_i < k$ and u_{p_i+1} is informed in round $r + 2(i+1)$, i.e., the ith phase is successful. For technical convenience, we also define $X_i = 1$ if $p_i = k$, i.e., v is already informed in that phase. Then,

$$\mathbb{P}[X_i = 1 \mid \mathcal{K}'_i] \geqslant 1 - \left(1 - d^{-2}\right)^{\alpha \cdot d/4} \geqslant 1 - e^{-\alpha/(4d)} \geqslant \frac{\alpha}{8d}. \tag{4.2}$$

Define $Y = \sum_{i=0}^{r/2-1} Y_i$, where Y_i are independent Bernoulli random variables with success probability $\alpha/8d$. By our choice of γ and (4.1),

$$\mathbb{E}[Y] = \frac{\alpha r}{16d} \geqslant 8(k_{\max}(\gamma) + \log n) \geqslant 8(k + \log n),$$

and, by a Chernoff bound,

$$\mathbb{P}\left[Y \geqslant k\right] \geqslant \mathbb{P}\left[Y \geqslant \mathbb{E}\left[Y\right]/2\right] \geqslant 1 - e^{-\mathbb{E}[Y]/8} \geqslant 1 - 1/n \geqslant 1 - e^{-\beta d}/2,$$

since $d = O(\log n)$ and by choosing constant β smaller if necessary. On the other hand, for $X = \sum_{i=1}^{r/2-1} X_i$, (4.2) implies that X stochastically dominates Y, in particular,

$$\mathbb{P}\left[X \geqslant k\right] \geqslant \mathbb{P}\left[Y \geqslant k\right] \geqslant 1 - e^{-\beta d}/2.$$

Note, $X \geqslant k$ implies that v is informed in R-VISIT-EXCHANGE$_r$ at round $2r$. Since R-VISIT-EXCHANGE$_r$ and M-VISIT-EXCHANGE$_r$ are identical with probability $1 - e^{-\beta d}/2$, vertex v must be informed in M-VISIT-EXCHANGE$_r$ at round $2r$ with probability at least $1 - e^{-\beta d} = 1 - e^{-\beta h}$.

Case $d = O(\log\log n)$. As in the previous case, we assume $t_0 = 0$ and consider the spread of information along a shortest path from u to v, namely, $\langle u = u_0, \ldots, u_k = v\rangle$. Fix a round $t \geqslant r$ and some $i \in \{0, \ldots, k-1\}$. Let $l = (\eta \log\log n)^2$ for some constant η that will be specified later. For an agent g define R_g as the number of times agent g visits u_i in rounds $t, \ldots, t+l-1$. If $X_g(t')$ is the position of the agent g at round t', then $R_g = \sum_{t'=t}^{t+l-1} \mathbb{1}_{X_g(t')=u_i}$, so by Lemma 4.2.4,

$$\mathbb{E}\left[R_g \mid X_g > 0\right] = \sum_{t'=t}^{t+l-1} \mathbb{P}\left[X_g(t') = u_i \mid R_g > 0\right] \leqslant 1 + \sum_{t'=t}^{t+l-1} \left(\frac{1}{n} + \frac{20}{\sqrt{t-t'+1}}\right) \leqslant 50 \cdot \sqrt{l}.$$

Let $Z_{i,t}$ be the number of unique agents that visit u_i in rounds $t, \ldots, t+l-1$.

$$\begin{aligned}\mathbb{E}\left[Z_{i,t}\right] &= \sum_{g \in A} \mathbb{P}\left[R_g > 0\right] = \sum_{g\in A} \frac{\mathbb{E}\left[R_g\right]}{\mathbb{E}\left[R_g \mid R_g > 0\right]} \\ &\geqslant \frac{\sum_{g\in A} \mathbb{E}\left[R_g\right]}{50 \cdot \sqrt{l}} = \frac{\sum_{t'=t}^{t+l-1} \mathbb{E}\left[N_{u_i}(t')\right]}{50 \cdot \sqrt{l}} = \frac{\sum_{t'=t}^{t+l-1} \mathbb{E}\left[\hat{N}_{u_i}(t'-r, r)\right]}{50 \cdot \sqrt{l}} \\ &\geqslant \frac{l \cdot \alpha/2}{50 \cdot \sqrt{l}} = \frac{\alpha \cdot \sqrt{l}}{100}.\end{aligned}$$

Since the agents are performing independent random walks, then by a Chernoff bound,

$$\mathbb{P}\left[Z_{i,t} \geqslant \alpha \cdot \sqrt{l}/200\right] \geqslant 1 - \exp\left(-\frac{\alpha\eta}{800} \cdot \log\log n\right) \geqslant 1 - 1/\log^5 n,$$

for a suitable choice of η. We now let $\mathcal{E}$ be the event $Z_{i,t} \geqslant \alpha\cdot\sqrt{l}/200$ for all $i \in \{0, \ldots, k-1\}$ and $t \in \{r, \ldots, 2r\}$, simultaneously. As before, we create R-VISIT-EXCHANGE$_r$ by adding minimum number of agents to M-VISIT-EXCHANGE$_r$ to ensure that $\mathcal{E}$ holds. Since $rk = O(\log^4 n)$, by a union bound, there is a constant β such that $\mathbb{P}\left[\mathcal{E}\right] \geqslant 1 - e^{-\beta h}/2$.

The rest of the proof follows the same line of logic as in the case of $d = \omega(\log\log n)$. The only difference is that instead of phases of 2 rounds, we consider phases of l rounds. $\mathcal{E}$ implies that after each phase R-VISIT-EXCHANGE$_r$ informs the next vertex on the path with a constant probability since $\sqrt{l} = \Omega(d)$. Therefore, as long as $k \leqslant \gamma \cdot r/l$ for a sufficiently small γ, vertex v becomes informed at round $2r$ of R-VISIT-EXCHANGE$_r$ w.h.p., which completes the proof. □

We are now ready to prove Theorem 4.1.1, by considering each vertex separately and then using a union bound. For each vertex, we divide the path from the source to that vertex into phases of $2r$ rounds each and show that sufficient progress is made in each phase with constant probability.

Proof of Theorem 4.1.1. First, we consider the M-VISIT-EXCHANGE$_r$ process for a constant ρ chosen by Lemma 4.3.1 such that M-VISIT-EXCHANGE$_r$ is identical to VISIT-EXCHANGE in the first n^2 rounds of its execution, with probability at least $1-n^{-2}$. Consider a shortest path $\langle s = u_0, \dots, u_m = u\rangle$ from source vertex s to vertex u. Let $k = k_{max}(\gamma)$ be the upper bound on the distance from Lemma 4.3.2, and as before $h = \max\{d, \log\log n\}$. We divide the execution of M-VISIT-EXCHANGE$_r$ into phases of $2r$ rounds each. If vertex u_i is informed at the end of a phase, then by Lemma 4.3.2, the vertex $u_{\min\{m,i+k\}}$ will be informed in the next phase of $2r$ rounds with probability at least $1-e^{-\beta h}$, independently from the past.

For some constant $\eta \in (0,1)$, let $l = \lceil m/k + \log n/h\rceil/(1-\eta)$. For $i \in \{1,\dots,l\}$, let X_i be a Bernoulli random variable that is 0 if in the ith phase of M-VISIT-EXCHANGE$_r$ either k new vertices along the specified path become informed, or vertex u becomes informed, i.e., the phase is successful. For $X = \sum_{i=1}^{l} X_i$, if $X < l - \lceil m/k\rceil$ then vertex u is informed at the end of the lth phase, because at least $\lceil m/k\rceil$ phases were successful. By a stochastic dominance argument as in Lemma 4.3.2 we upper bound $\mathbb{P}[X < l - \lceil m/k\rceil]$.

Let $\{Y_i\}_{1\leqslant i\leqslant l}$ be a collection of independent Bernoulli random variables $\mathbb{P}[Y_i = 1] = e^{-\beta h}$. By Lemma 4.3.2, $\mathbb{P}[X_i = 1 \mid X_1,\dots,X_{i-1}] \leqslant \mathbb{P}[Y_i = 1]$, and therefore, for $Y = \sum_{i=1}^{l} Y_i$,

$$\begin{aligned}\mathbb{P}[X > l - \lceil m/k\rceil] &\leqslant \mathbb{P}[Y > l - \lceil m/k\rceil] \leqslant \mathbb{P}[Y \geqslant l - \lceil m/k + \log n/h\rceil]\\ &= \mathbb{P}[Y \geqslant \eta\cdot l] = \mathbb{P}\left[Y \geqslant \eta\cdot e^{\beta h}\cdot \mathbb{E}[Y]\right]\\ &\leqslant \left(\eta\cdot e^{\beta h-1}\right)^{-\eta\cdot l} \leqslant n^{-3},\end{aligned}$$

by a Chernoff bound (Lemma 2.3.6) and by taking a value of η that is sufficiently close to 1. Thus, after $l\cdot 2r$ rounds of M-VISIT-EXCHANGE$_r$ vertex u is informed with probability $1-n^{-3}$. By a union bound over all vertices, and the fact that M-VISIT-EXCHANGE$_r$ and VISIT-EXCHANGE are identical in the first n^2 rounds we get that $T \leqslant l\cdot 2r$ w.h.p. Since $k = O(r/\max\{d,(\log\log n)^2\})$, and $m \leqslant \operatorname{diam}(G)$, and $h = \max\{d, \log\log n\}$, we finally get that, w.h.p.,

$$T = O\left(\max\{d,(\log\log n)^2\}\cdot \operatorname{diam}(G) + \frac{\log^3 n}{h}\right) = \widetilde{O}\left(d\cdot\operatorname{diam}(G) + \frac{\log^3 n}{d}\right). \qquad \square$$

4.4 Upper bound in terms of average degree

In this section we prove Theorem 4.1.2. Recall that $G = (V,E)$ is a graph with average degree d_{avg} and minimum degree $d_{\min} = \Omega(d_{\text{avg}})$ and A is the set of agents in VISIT-EXCHANGE. The set of agents of VISIT-EXCHANGE is A, and $|A| = \alpha\cdot n$ for a constant $\alpha > 0$. The agents in A start their walks from the stationary distribution π. Let $\epsilon = d_{\min}/d_{\text{avg}} = \Omega(1)$. Then, for every vertex $u \in V$,

$$\pi(u) = \frac{\deg(u)}{2|E|} = \frac{\deg(u)}{n\cdot d_{\text{avg}}} \geqslant \frac{\epsilon}{n}. \tag{4.3}$$

We define $N_u(t)$, $\hat{N}_u(t,r)$ and $p^t_{v,u}$ as in Section 4.3.

As in Section 4.3, we modify the VISIT-EXCHANGE process to create a new process called M-VISIT-EXCHANGE$_r$, that depends on a parameter $r = \Theta(\log^2 n)$: For all rounds t and vertices u, we add a minimal set of agents to the process to make sure that $\hat{N}_u(t,r) \geqslant |A| \cdot \pi(u)/2$.

Lemma 4.4.1. *For any constant $c > 0$, there is a parameter $r = O(\log^2 n)$ such that* VISIT-EXCHANGE *and* M-VISIT-EXCHANGE$_r$ *are identical for the first T' rounds of their execution with probability at least $1 - T' \cdot n^{-(c+2)}$.*

Proof. By Lemma 4.2.4 and condition (4.3),

$$p^{2r}_{u,u} \leqslant 2\pi(u)\left(1 + \frac{20 \cdot |E|}{d_{\min} \cdot \sqrt{2r+1}}\right) \leqslant 2\pi(u)\left(1 + \frac{40n}{\epsilon\sqrt{2r+1}}\right) \leqslant \frac{100n \cdot \pi(u)}{\epsilon\sqrt{2r+1}},$$

where the last inequality holds assuming a large value of n and $r = O(\log^2 n)$. Substituting in Lemma 4.2.1, gives

$$\begin{aligned}\mathbb{P}\left[\hat{N}_u(t,r) \leqslant |A| \cdot \pi(u)/2\right] &\leqslant \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right) \leqslant \exp\left(-\frac{|A| \cdot \epsilon\sqrt{2r+1}}{100n}\right)\\ &= \exp\left(-\frac{\alpha\epsilon}{100} \cdot \sqrt{2r+1}\right) \leqslant n^{-(c+3)},\end{aligned}$$

for $r = \eta \log^2 n$ for a sufficiently large constant η. By applying a union bound over all vertices u and rounds $t < T'$, we complete the proof. $\square$

Lemma 4.4.2. *Let k be the length of a shortest path from the source vertex s to vertex u. For any constant $c > 0$ and integer r, vertex u becomes informed in at most $O(r \cdot d_{\text{avg}} \cdot (k + \log n))$ rounds of* M-VISIT-EXCHANGE$_r$*, with probability at least $1 - n^{-(c+1)}$.*

Proof. Let $\langle s = u_0, u_1, \ldots, u_k = u\rangle$ be a shortest path from vertex s to u. We divide the execution of M-VISIT-EXCHANGE$_r$ into phases of $(r+1)$ rounds each. For each $i \geqslant 0$, let $\mathcal{K}_i$ be the σ-algebra fixing the execution (prefix) of M-VISIT-EXCHANGE$_r$ up to round $i(r+1)$. Let p_i be the largest integer, between 0 and k, such that $w = u_{p_i}$ is informed at round $i(r+1) - 1$. By the definition of M-VISIT-EXCHANGE,

$$\mathbb{E}\left[N_w\left((i+1)(r+1) - 1\right) \mid \mathcal{K}_i\right] = \hat{N}_w(i(r+1), r) \geqslant \frac{|A| \cdot \pi(w)}{2}.$$

Since the agents move independently, by a Chernoff bound we have that for an event $\mathcal{E}_i = \left\{N_w\left((i+1)(r+1) - 1\right) \geqslant \frac{|A| \cdot \pi(w)}{4}\right\}$,

$$\mathbb{P}\left[\mathcal{E}_i \mid \mathcal{K}_i\right] \geqslant 1 - \exp\left(-\frac{|A| \cdot \pi(w)}{16}\right) \geqslant 1 - e^{-\alpha\epsilon/16}.$$

Notice that, if $p_i < k$, then each agent that visits u in round $(i+1)(r+1) - 1$, informs vertex u_{p_i+1} with probability $1/\deg(w)$ at the next round. Define $Y_i = 1$ if either u_{p_i+1} is informed in round $(i+1)(r+1)$, or $p_i = k$. Then,

$$\mathbb{P}\left[Y_i = 1 \mid \mathcal{E}_i; \mathcal{K}_i\right] \geqslant 1 - \left(1 - \frac{1}{\deg(w)}\right)^{|A| \cdot \pi(w)/4} \geqslant 1 - \exp\left(-\frac{|A| \cdot \pi(w)}{4 \cdot \deg(w)}\right)$$

$$\geqslant 1 - \exp\left(-\frac{\alpha}{4} \cdot \frac{n}{2|E|}\right) = 1 - e^{-\frac{\alpha}{4d_{\mathrm{avg}}}} \geqslant \min\left\{\frac{1}{2}, \frac{\alpha}{8d_{\mathrm{avg}}}\right\}.$$

Then,

$$\begin{aligned}
\mathbb{P}\left[Y_i = 1 \mid \mathcal{K}_i\right] &\geqslant \mathbb{P}\left[Y_i = 1 \mid \mathcal{K}_i; \mathcal{E}_i\right] \cdot \mathbb{P}\left[\mathcal{E}_i \mid \mathcal{K}_i\right] \\
&\geqslant \min\left\{\frac{1}{2}, \frac{\alpha}{8d_{\mathrm{avg}}}\right\} \cdot \left(1 - e^{-\alpha\epsilon/16}\right) \\
&= \eta/d_{\mathrm{avg}}, \qquad (4.4)
\end{aligned}$$

where η is a constant that could depend on d_{avg} if $d_{\mathrm{avg}} = O(1)$. As in Lemma 4.3.2, (4.4) implies that after at most $O(d_{\mathrm{avg}}(k + \log n))$ phases, vertex u must become informed w.h.p. Since each phase lasts $r + 1$ rounds, we complete the proof. □

Proof of Theorem 4.1.2. Let $c > 0$ be any fixed constant and $r = O(\log^2 n)$ be as determined from Lemma 4.4.2. Let k be the length of a shortest path from the source vertex to a fixed vertex $u \in V$. By Lemma 4.4.2, in at most $O(d_{\mathrm{avg}} \cdot r \cdot (k + \log n)) = O(d_{\mathrm{avg}} \cdot \log^2 n \cdot (\mathrm{diam}(G) + \log n))$ rounds u becomes informed with probability at least $1 - n^{-(c+1)}$. Applying a union bound for all n vertices, it follows that M-VISIT-EXCHANGE$_r$ informs all vertices in $T' = O(d_{\mathrm{avg}} \cdot \log^2 n \cdot (\mathrm{diam}(G) + \log n))$ rounds, with probability at least $1 - n^{-c}$. Finally, by Lemma 4.4.1, M-VISIT-EXCHANGE$_r$ and VISIT-EXCHANGE are identical in the first $T' \leqslant n^2$ rounds of their executions, with probability at least $1 - n^{-c}$, and therefore, VISIT-EXCHANGE informs all vertices of G in T' rounds with probability at least $1 - 2n^{-c}$. □

Chapter 5

Bounds for expanders

5.1 Introduction

In this chapter we study VISIT-EXCHANGE for d-regular expander graphs. Expanders are graphs that have strong connectivity properties, while possibly being very sparse. This property of expander graphs makes them naturally appealing in the context of the design and analysis of communication networks and information dissemination. Expanders have also proved to be extremely useful in other areas of theoretical computer science and mathematics such as error correcting codes, de-randomisation, analysis of algorithms. One way to quantify the well-connectedness of graphs is using their conductance ϕ. A graph is said to be an expander if ϕ is constant. We give precise definitions in Section 5.2 and also refer the reader to [HLW06] for a review on expanders.

Due to the strong connectivity properties of expander graphs, one expects fast broadcasting for all protocols discussed in this thesis. Indeed, [Chi+18] proved that in a graph with conductance ϕ, the broadcast time of PUSH-PULL is $O(\log n/\phi)$, w.h.p. Therefore, on expanders, where $\phi > 0$ is constant, we have $T_{\text{ppull}} = O(\log n)$, w.h.p., which in general is optimal. By the equivalence of PUSH and VISIT-EXCHANGE for sufficiently dense regular graphs, proved in Chapter 3 (Theorem 3.4.1 in particular), we immediately get the following result:

Theorem 5.1.1. *Let $c, \alpha, \beta > 0$ be constants with $\alpha\beta$ sufficiently large (depending on c). Consider the* VISIT-EXCHANGE *process with $|A| = \alpha n$ agents on any d-regular expander G such that $d \geqslant \beta \log n$. For any source vertex, $T_{\text{visitx}} = O(\log n)$, with probability at least $1 - n^{-c}$.*

We therefore focus on sparse regular expanders with $d = O(\log n)$. Our first result proves an optimal bound for constant degree expanders.

Theorem 5.1.2. *For any d-regular expander G with $d \geqslant 3$ constant, and any source vertex, $T_{\text{visitx}} = O(\log n)$, w.h.p.*

Note that the only 2-regular connected graphs are cycles, which are not expanders. See Chapter 7 for the analysis of VISIT-EXCHANGE on cycles.

The proof of Theorem 5.1.2 uses a method different from the ones presented in earlier chapters. Instead of arguing about VISIT-EXCHANGE informing individual vertices by following a certain progress path, to prove Theorem 5.1.2 we argue that the set of all informed vertices grows exponentially during the execution. This method is similar to those used for proving bounds for randomised rumour spreading, e.g., in [Chi+18].

For round t, denote by I_t the set of vertices that were informed in any round up to and including round t. By the expansion property, it is easy to see that at least a constant fraction of vertices in I_t have an uninformed neighbour. Denote this subset of I_t by S. We claim that a constant fraction of vertices in S are visited by some agent between rounds t and $t+r-1$, w.h.p., for any t and large enough constant r. Since d is constant, this implies that the number of informed vertices increases by a constant factor every r rounds, w.h.p. The key technical argument in the proof is that the probability a given agent visits S between rounds t and $t+r-1$ is proportional to $|S|$ and r. Thus, S is not visited by sufficiently many agents in these r rounds with probability decreasing exponentially in $r \cdot |S|$ or, equivalently, in rk where $k = |I_t|$. This fact holds for a fixed round but we need it to hold for all rounds. Since I_t is connected, the number of its different instantiations of is bounded by $d^{\Theta(k)}$, which does not depend on r. It implies that by taking a sufficiently large constant r, we can apply a union bound over all possible instantiations of I_t, and argue that S is visited by sufficiently many agents in all rounds, w.h.p. As a result, we are able to prove that the number of newly informed vertices between rounds t and $t+r$ is at least a constant fraction of $|S|$. This implies the exponential increase of the size of the informed set. It should be noted that this expansion only holds if there are at least $\Omega(\log n)$ and at most $n/2$ informed vertices. For these other (extremal) cases, a separate and simpler analysis is used, thus, the whole proof is done in three phases with the middle being the main one.

We are currently not able to extend Theorem 5.1.2 to arbitrary d-regular expanders for $\omega(1) \leqslant d \leqslant O(\log n)$ (for $d = \Omega(\log n)$, the result follows from Chapter 3, as mentioned earlier). However, if in addition to having a constant conductance the graph G has strong vertex-expanding properties, then we can prove the optimal logarithmic bound for G. We call such graphs *strong* expanders. Roughly, a graph $G = (V, E)$ is a strong expander, if for every set $S \subset V$ with $|S| = O(n/d)$ the neighbourhood of S contains at least $\Omega(d|S|)$ vertices. The precise definition is given in Section 5.2.

Theorem 5.1.3. *For any d-regular strong expander G, and any source vertex, $T_{\text{visitx}} = O(\log n)$, w.h.p.*

We note immediately that random regular graphs are strong expanders and therefore we have the following corollary.

Theorem 5.1.4. *For any integers n and d such that nd is even, if G is a random d-regular graph of n vertices, then $T_{\text{visitx}} = O(\log n)$, w.h.p., for any source vertex.*

The proof of Theorem 5.1.3 has many similarities to that of Theorem 5.1.2. In particular, the first and third phases of the analyses are exactly the same since they do not require the additional properties of strong expanders. For the middle phase where the set of informed vertices increases exponentially, we cannot directly apply the proof for Theorem 5.1.2. This is because the number of possible instantiations of I_t of size k, that is $d^{\Theta(k)}$, is too large when $d = \omega(1)$ and a union bound does not work any more. To overcome this challenge, we can use the strong expansion property that the neighbourhood of any set S contains at least $\Omega(d|S|)$ vertices. Unlike in the proof of Theorem 5.1.2, where we wait for some $r = O(1)$ rounds until agents arrive to the set S and use them to inform new vertices, here we simply consider the agents arriving to S from its neighbourhood in round $t+1$.

The neighbourhood of S is large enough that we can lower bound the number of these agents with sufficiently high probability that the union bound over $d^{\Theta(k)}$ sets works. This method only works when $d = \omega(1)$, but since all constant degree expanders are also strong expanders, we have covered the $d = O(1)$ case in Theorem 5.1.2.

Finally, we are able to prove a sub-optimal bound on the broadcast time of VISIT-EXCHANGE for any d-regular expander combining ideas from the previous two chapters.

Theorem 5.1.5. *For any d-regular expander graph G and any source vertex, $T_{\text{visitx}} = O(\log n \cdot \log \log n)$, w.h.p.*

We prove this theorem by bounding T_{visitx} in terms of T_{push}, via a coupling like in the proof of Theorem 3.4.1. Unlike in Theorem 3.4.1, where the processes are coupled every other round, here they are coupled only every $(r+1)$th round for some $r = O(\log \log n)$. We use Lemma 4.2.1 to argue that, given a configuration of agents, in r rounds each vertex receives an agent, w.c.p. These two facts together allow us to fix a path that PUSH uses to inform a fixed vertex u and argue that VISIT-EXCHANGE makes progress via the same path every r rounds. Since for regular graphs $T_{\text{push}} = O(T_{\text{ppull}}) = O(\log n)$, w.h.p., by [Chi+18], the theorem follows. We believe that this bound is not tight and $T_{\text{visitx}} = O(\log n)$ for general regular expanders, but we do not have a proof for that.

5.1.1 Road-map

In Section 5.2 we give the definition of expanders and strong expanders precisely. We also prove necessary results on random walks and on graphs. Section 5.3 contains the proofs of the main claims of this chapter, Theorems 5.1.2 and 5.1.3. Section 5.4 and Section 5.5 contain the proofs of Theorems 5.1.4 and 5.1.5, respectively.

5.2 Preliminaries

First we define graph expansion parameters. Let $G = (V, E)$ be an n-vertex graph. For a set $S \subset V$, let $E(S, V \setminus S)$ be the set of edges (u, v) such that $u \in S$ and $v \notin S$. The *volume* of a set S is defined as $\text{vol}(S) = \sum_{u \in S} \deg(u)$. The *conductance* of a set $S \subset V$ with $S \neq \emptyset, V$ is defined as

$$\phi(S) = \frac{|E(S, V \setminus S|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}}.$$

The conductance ϕ of the graph is the minimum possible conductance achieved by any non-trivial subset of V:

$$\phi = \min_{\emptyset \subsetneq S \subsetneq V} \phi(S).$$

The conductance is always between 0 and 1. Fig. 5.1 shows an example graph and its conductance. The graph G is said to be an (edge)-*expander*, if $\phi = \Omega(1)$.

If G is d-regular, the definition of an expander can be simplified as follows: For any set S such that $0 < |S| \leqslant n/2$, we have that

$$|E(S, V \setminus S)| \geqslant \phi \cdot d|S|.$$

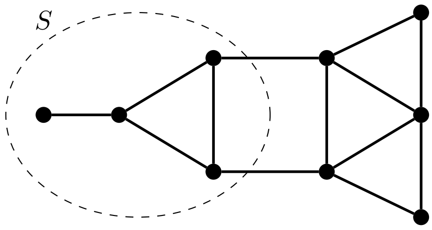

Figure 5.1 For the subset S of vertices of the depicted graph, $\mathrm{vol}(S) = 10$ and $\mathrm{vol}(V \setminus S) = 16$, while $|E(S, V \setminus S)| = 2$, hence $\phi(S) = 2/10 = 1/5$. It is not hard to see that S is the set with the smallest conductance, therefore, $\phi(G) = 1/5$.

Next we define strong expansion. For any set $S \subset V$, let ∂S be the neighbourhood of S, i.e.,

$$\partial S = \{u \in V \setminus S \mid (u, v) \in E \text{ for some } v \in S\}.$$

An expander graph with conductance ϕ is a *strong expander*, if there are constants $\epsilon, \delta > 0$ such that for any set $S \subset V$ if $1 \leqslant |S| \leqslant \delta n/d$ then $|\partial S| \geqslant \epsilon d|S|$. We will say that G has strong expansion parameters (ϕ, ϵ, δ).

Note that if G is an expander and $d = O(1)$, then for any set S with $|S| \leqslant n/2$, $|\partial S| \geqslant |E(S, V \setminus S)|/d \geqslant \phi|S|$, thus, G is also a strong expander with parameters $(\phi, \phi, d/2)$.

It is also possible to give an equivalent spectral definition of expanders. For a connected graph $G = (V, E)$ let P be the transition matrix of a random walk on G. We denote by λ the largest non-trivial eigenvalue of P, that is $\lambda = \lambda_2$ where $1 = \lambda_1 > \lambda_2 \geqslant \ldots \geqslant \lambda_n \geqslant -1$ are the n real-valued eigenvalues of P. The value $1 - \lambda$ is called a spectral gap of P. G is a *spectral* expander the spectral gap of P is constant, i.e., $1 - \lambda = \Omega(1)$.

By Cheeger's inequality (e.g., [LP17, Inequality (13.6)]),

$$\frac{\phi^2}{2} \leqslant 1 - \lambda \leqslant 2\phi,$$

which means that a graph is an edge-expander if and only if it is a spectral expander. This allows us to use both definitions of expanders depending on convenience.

The benefit of the spectral definition of expanders is that it allows to reason about the distribution of a random walk after a number of rounds. Recall from Chapter 2 that the uniform mixing time t^{∞}_{mix} of a random walk is the smallest round for which if $t \geqslant t^{\infty}_{\text{mix}}$, then

$$||\mathbf{p}_t - \pi||_{\infty} = \max_{u \in V} |\mathbf{p}_t(u) - \pi(u)| \leqslant \pi(u)/2,$$

where $\mathbf{p}_t$ is the distribution of the walk at round t. For a lazy random walk on an expander graph we have that $t^{\infty}_{\text{mix}} = O(\log n/(1 - \lambda)) = O(\log n)$, by [LP17, Inequality (12.11)]. Unfortunately, such bound does not exist for non-lazy walks, because a random walk does not mix in bipartite graphs. Thus, for some claims which require the mixing time bound we assume that the agents perform lazy walks walks and then argue that the claim holds for non-lazy walks as well.

The following result on probability amplification for expanders uses tools from [AF02].

Lemma 5.2.1. *Let $X(t)$ be a random walk on a d-regular graph $G = (V, E)$. starting from the (uniform) stationary distribution, with a transition matrix P and a second largest* *eigenvalue of λ. Then, for a set $S \subset V$ and integer $t \leqslant 2n/|S|$, if τ_S is the first time when the walk $X(t)$ visits any vertex in S*

$$\mathbb{P}\left[\tau_S \leqslant t\right] \geqslant \begin{cases} \frac{t(1-\lambda)|S|}{2n}, & \textit{if } X(t) \textit{ is a } 1/2\textit{-lazy walk;} \\ \frac{t(1-\lambda)|S|}{4n}, & \textit{if } X(t) \textit{ is a non-lazy walk.} \end{cases}$$

Proof. First we consider the case when $X(t)$ is lazy. Let Q be the transition matrix P that is restricted to the set $\overline{S} = V \setminus S$, i.e., $Q_{u,v} = P_{u,v}$ for $u, v \notin S$, and 0 otherwise. If a walk does not reach the set S, it must start in $\overline{S}$ and only follow edges that are present in Q, up until round t. Thus,

$$\mathbb{P}\left[\tau_S > t\right] = \sum_{u \notin S} \pi(u) \sum_{v \notin S} (Q^t)_{u,v} = \frac{1}{n} \sum_{u,v \in \overline{S}} (Q^t)_{u,v}.$$

The sum of the elements of the matrix Q^t above can be written as $||Q^t x^T||_1$, where x is a row vector taking values 1 on the set $\overline{S}$, and 0 otherwise. Let λ_Q be the largest non-trivial eigenvalue of Q. Then,

$$\mathbb{P}\left[\tau_S > t\right] = \frac{1}{n} \cdot ||Q^t x^T||_1 \leqslant \frac{1}{n} \cdot |\overline{S}| \cdot ||Q^t x^T||_\infty \leqslant \frac{|\overline{S}|}{n} \lambda_Q^t \leqslant \lambda_Q^t.$$

By [AF02, Corollary 3.34] and using the definition of relaxation time of chain P (which is $1/(1-\lambda)$), we get that for the chain Q,

$$\frac{1}{1-\lambda_Q} \leqslant \frac{1}{1-\lambda} \cdot \frac{n}{|S|},$$

which implies that

$$\lambda_Q \leqslant 1 - \frac{(1-\lambda)|S|}{n},$$

and thus,

$$\begin{aligned} \mathbb{P}\left[\tau_S \leqslant t\right] &\geqslant 1 - \left(1 - \frac{(1-\lambda)|S|}{n}\right)^t \\ &\geqslant 1 - \frac{1}{1 + t(1-\lambda)|S|/n}, \quad \text{by Lemma 2.1.1(b)}, \\ &\geqslant \frac{t(1-\lambda)|S|}{2n}, \end{aligned}$$

where the last inequality holds because $t \leqslant 2n/|S|$. This completes the proof when the walk $X(t)$ is lazy with holding probability $1/2$.

When $X(t)$ is not lazy, we consider a 1/2-lazy walk $X'(t)$ on G. If τ'_S is the first time when $X'(t)$ visits a vertex in S then, by a simple coupling, and the result for lazy walks above,

$$\mathbb{P}\left[\tau_S \leqslant t\right] \geqslant \mathbb{P}\left[\tau'_S \leqslant t\right] \geqslant \frac{t(1-\lambda')|S|}{2n}.$$

The transition matrix of $X'(t)$ is $P' = (P+I)/2$, where I is an identity matrix. Therefore, the spectral gap of P' is $1 - \lambda' = (1-\lambda)/2$. We complete the proof after a substituting λ instead of λ'. □

We also present this combinatorial fact about connected graphs, which will be useful for a union bound over all informed subsets of vertices.

Lemma 5.2.2. *Let u be any vertex of a graph H with largest degree Δ. For any integer $k > 0$, there are at most Δ^{2k} connected subgraphs of H that contain u and have k vertices.*

Proof. Consider a (not necessarily simple) path in H that has length $2k$, starts at u and contains at most k unique vertices. The subgraph induced by the vertices on the path is connected and contains u. The number of such paths is at most Δ^{2k} because we can construct them starting from u, choosing one of the at most Δ neighbours at each step. Therefore, it suffices to show that each connected subgraph H' of H, that contains u can be traversed by a path of length at most $2k$. Consider a spanning tree R of H'. A depth first search traversal path of R starting from u uses at most $2(k-1)$ edges, therefore, it satisfies our requirement. □

5.3 Optimal bound for strong expanders

In this section we prove Theorems 5.1.2 and 5.1.3. Throughout the proofs we assume that $G = (V, E)$ is a d-regular graph with $d = O(\log n)$ (for larger d, the theorems follow from Chapter 3). We assume that G has strong expansion parameters (ϕ, ϵ, δ), and the spectral gap is $1-\lambda$. We denote by I_t the set of informed vertices after round t of VISIT-EXCHANGE, so $I_0 = \{s\}$. The proof proceeds in three phases and we show that each takes at most $O(\log n)$ rounds. In the first phase at least $\Omega(\log n)$ vertices become informed, after the second one the number of informed agents is at least $\Omega(n/d)$, and, finally, in the third phase all vertices become informed. We present the analysis of the second phase separately for the case when $d = O(1)$ and $d = \omega(1)$, corresponding to Theorems 5.1.2 and 5.1.3, respectively.

5.3.1 Phase 1: $\Omega(\log n)$ informed vertices

We prove that if G is a d-regular expander, then in at most $O(\log n)$ rounds there are at least $\Omega(\log n)$ informed vertices, w.h.p. In this section we do not need the strong expansion of G.

Lemma 5.3.1. *For any $b, c > 0$, there is a round $\tau_1 = O(\log n)$ such that $|I_{\tau_1}| \geqslant b \log n$ with probability at least $1 - n^{-c}$.*

Proof. For an agent g, let X_g be the indicator variable that g visits the source vertex in the first $\tau' = \eta' \log n$ rounds. We apply Lemma 5.2.1 for the singleton set $S = \{s\}$ containing only the source vertex:

$$\mathbb{P}[X_g = 1] \geqslant \frac{(1-\lambda)\tau'}{4n}.$$

Thus, if A' is the set of agents that have visited the source in the first τ' rounds, then

$$\mathbb{E}[|A'|] \geqslant \frac{\alpha(1-\lambda)\tau'}{4}.$$

Furthermore, since the agents perform independent random walks, by a Chernoff bound,

$$\mathbb{P}\left[|A'| \geqslant \frac{\alpha(1-\lambda)\eta'}{8} \cdot \log n\right] \geqslant 1 - \exp\left(-\frac{\alpha(1-\lambda)\eta'}{32} \cdot \log n\right).$$

Thus, for any constant a', we can take a large enough η' such that

$$\mathbb{P}\left[|A'| \geqslant a' \log n\right] \geqslant 1 - n^{-c}/2. \tag{5.1}$$

Recall that t_{mix}^{∞} is the uniform mixing time of a lazy random walk, for now assume that the agents perform lazy random walks. We set $\tau = \tau' + t_{\text{mix}}^{\infty} = O(\log n)$. Let U be the number of vertices that contain an agent from A' at round τ. By the property of mixing, for a vertex u, the probability that a given agent is at u is at least $1/(2n)$. Thus, for $N = |U|$, we have $\mathbb{E}\left[N \mid A'\right] \geqslant |A'|/2$.

Furthermore, N is a function of the independent walks performed by the agents in A', and changing the walk of one of the agents can change N by at most 1. Therefore, by the method of bounded differences (Theorem 2.3.8),

$$\mathbb{P}\left[N \geqslant \frac{\mathbb{E}\left[N\right]}{2} \;\middle|\; A'\right] \geqslant 1 - e^{-\frac{2\cdot(\mathbb{E}[N]/2)^2}{|A'|}} = 1 - 2e^{-\frac{|A'|}{8}},$$

which implies that

$$\mathbb{P}\left[N \geqslant \frac{a' \cdot \log n}{2} \;\middle|\; |A'| \geqslant a' \cdot \log n\right] \geqslant 1 - 2n^{-\frac{a'}{8}}. \tag{5.2}$$

We take $a' \geqslant 2b$ and also large enough so that $2n^{-\frac{a'}{8}} \leqslant n^{-c}/2$.

Now consider the non-lazy agents again and couple their executions to corresponding lazy walks in a natural way, i.e., each non-lazy walk simply skips the holding steps of a lazy walk. For each vertex in $u \in U$, there is a non-lazy walk from A' that is either at u or has passed it earlier, by the coupling. Thus, the vertices in U are informed and, by applying union bound for (5.1) and (5.2), we get

$$\mathbb{P}\left[|I_\tau| \geqslant b \cdot \log n\right] \geqslant \mathbb{P}\left[N \geqslant b \cdot \log n\right] \geqslant 1 - n^{-c}. \qquad \square$$

5.3.2 Phase 2: $\Omega(n/d)$ informed vertices

In the previous section we proved that in $O(\log n)$ rounds of VISIT-EXCHANGE, at least $\Omega(\log n)$ vertices of a d-regular expander graph G become informed, w.h.p., regardless of the value of d. In this section our goal is to show that until at least $\Omega(n/d)$ vertices of G become informed, the set of informed vertices I_t grows exponentially, w.h.p. Hence, in $O(\log n)$ rounds at least $\Omega(n/d)$ vertices will become informed, w.h.p. It will remain to show that in further $O(\log n)$ rounds all vertices become informed, which is the goal of the next section.

The analysis of this section is the main part of the proofs of Theorems 5.1.2 and 5.1.3. We treat the cases $d = O(1)$ and $d = \omega(1)$ separately. While the two proofs have similarities, the separate presentation benefits the readability.

Constant degree expanders

We start by describing a modification of the VISIT-EXCHANGE process, that depends on two constants b and r, which are fixed later. Recall that ϕ is the conductance of the graph and let $\epsilon = \phi/(2d)$. For a set $S \subset V$, let $G(S)$ be the subgraph of G induced by S. Define the following set of subsets of V:

$$\mathcal{S}(b) = \{S \subset V \mid s \in S,\ G(S) \text{ is connected, and } b \cdot \log n \leqslant |S| \leqslant n/d\}.$$

For a round $t \geqslant 0$, we say that a set $S \in \mathcal{S}(b)$ is (t, r)-good, if at least $(1-\epsilon)|S|$ vertices in S are visited by some agent in rounds $t, \ldots, t+r-1$.

If for some round t, some set $S \in \mathcal{S}(b)$ is not (t, r)-good, then we add a minimal set of agents in S immediately after round $t+r-1$ to turn S into a (t, r)-good set. The modified process is called M-VISIT-EXCHANGE$_{b,r}$, for which we will use the same notation as for VISIT-EXCHANGE. We show that M-VISIT-EXCHANGE$_{b,r}$ and VISIT-EXCHANGE are identical in the first polynomial number of rounds, w.h.p.

Lemma 5.3.2. *Let T' be any positive integer. For any constants $b, c > 0$, there is a positive integer r, such that every set $S \in \mathcal{S}(b)$ is (t, r)-good for every round $t \leqslant T'$ of* VISIT-EXCHANGE, *with probability at least $1 - T' \cdot n^{-(c+1)}$.*

Proof. Consider a fixed round $t \leqslant T'$ and a set $S' \subseteq S$ for some $S \in \mathcal{S}(b)$, such that $|S'| \geqslant \epsilon|S|$. For an integer r to be fixed later in the proof, let X_g be an indicator random variable that agent g visits some vertex in S' between rounds t and $t+r-1$. By Lemma 5.2.1, $\mathbb{P}[X_g = 1] \geqslant \frac{(1-\lambda)r|S'|}{4n}$ and thus, if $N_{S'}$ is the number of unique agents visiting S' between rounds t and $t+r-1$, then

$$\mathbb{E}[N_{S'}] \geqslant |A| \cdot \frac{(1-\lambda)r}{4n} \cdot |S'| \geqslant \frac{\epsilon\alpha(1-\lambda)r}{4} \cdot |S|.$$

By an application of a Chernoff bound, and setting $\eta = \frac{\epsilon\alpha(1-\lambda)}{32}$ for conciseness, we get

$$\mathbb{P}[N_{S'} \geqslant 1] \geqslant \mathbb{P}\left[N_{S'} \geqslant \frac{\mathbb{E}[N_{S'}]}{2}\right] \geqslant 1 - e^{-\eta r|S|}.$$

Next, notice that S is (t, r)-good if and only if for all $S' \subset S$ with $|S'| \geqslant \epsilon|S|$, $N_{S'} \geqslant 1$. Thus,

$$\mathbb{P}[S \text{ is } (t,r)\text{-good}] = \mathbb{P}\left[\bigcap_{\substack{S' \subset S \\ |S'| \geqslant \epsilon|S|}} \{N_{S'} \geqslant 1\}\right] \geqslant 1 - 2^{|S|} \cdot e^{-\eta r|S|} \geqslant 1 - e^{-(\eta r - 1)\cdot|S|}.$$

Next we apply another union bound for all sets S:

$$\begin{aligned}
\mathbb{P}[S \text{ is } (t,r)\text{-good for all } S \in \mathcal{S}(b)] &\geqslant 1 - \sum_{S \in \mathcal{S}(b)} e^{-(\eta r-1)\cdot|S|} \\
&\geqslant 1 - \sum_{k \geqslant b\log n} \sum_{\substack{S \in \mathcal{S}(b) \\ |S| = k}} e^{-(\eta r - 1)\cdot k} \\
&\geqslant 1 - \sum_{k \geqslant b \log n} \left(d^2 \cdot e^{-\eta r + 1}\right)^k, \quad \text{by Lemma 5.2.2,}
\end{aligned}$$

$$\geqslant 1 - 2\left(e^{-\eta r + 2\ln d + 1}\right)^{b \log n}$$
$$\geqslant 1 - n^{-(c+1)},$$

for a sufficiently large constant r, that depends on b and c. Applying another union bound over all rounds $t \leqslant T'$ completes the proof. $\square$

Lemma 5.3.3. *For any constant $c > 0$, there is a constant b such that for any r, if $I_t \in \mathcal{S}(b)$ in* M-VISIT-EXCHANGE$_{b,r}$*, then $|I_{t+r}| \geqslant (1+\psi)|I_t|$, with probability at least $1 - n^{-(c+1)}$, where $\psi = \epsilon/(2d^2)$.*

Proof. Let S be the set of vertices $u \in I_t$ that have an uninformed neighbour. Then,

$$|E(I_t, V \setminus I_t)| = \sum_{u \in S} \deg_{V \setminus I_t}(u) \leqslant d \cdot |S|.$$

On the other hand, since G is an expander with conductance ϕ,

$$|E(I_t, V \setminus I_t)| \geqslant \phi \cdot |I_t|,$$

since $|I_t| \leqslant n/2$ by the condition of the lemma. Combining the two inequalities above gives

$$|S| \geqslant \phi \cdot |I_t|/d = 2\epsilon |I_t|.$$

Since we are considering the M-VISIT-EXCHANGE$_{b,r}$ process, the set I_t is (t,r)-good and therefore, at most $\epsilon|I_t|$ vertices in I_t are not visited by any agent in rounds $t, \ldots, t+r-1$. By the fact that $S \subseteq I_t$ has at least $2\epsilon|I_t|$ vertices, we conclude that at least $\epsilon|I_t|$ vertices in S are visited by some agent in rounds $t, \ldots, t+r-1$. Suppose the set of these vertices is S'.

For $u \in S'$, let X_u be an indicator random variable that the first agent that visits u in rounds $t, \ldots, t+r-1$ (there must be one), visits a vertex in $V \setminus I_t$ in the next round. Then, $\mathbb{P}[X_u = 1 \mid I_t, S'] \geqslant 1/d$ and for $N = \sum_{u \in S'} X_u$,

$$\mathbb{E}[N \mid I_t, S'] \geqslant \frac{|S'|}{d} \geqslant \frac{\epsilon |I_t|}{d} = 2\psi d \cdot |I_t|.$$

Furthermore, since the variables X_u are independent, we can apply a Chernoff bound:

$$\mathbb{P}[N \geqslant \psi d \cdot |I_t| \mid I_t, S'] \geqslant 1 - e^{-\epsilon|I_t|/(8d)}.$$

On the other hand, $|I_{t+r}| \geqslant |I_t| + N/d$, because every vertex in $I_{t+r} \setminus I_t$ has at most d neighbours in I_t. Therefore,

$$\mathbb{P}[|I_{t+r}| \geqslant (1+\psi) \cdot |I_t| \mid I_t] \geqslant 1 - e^{-\epsilon|I_t|/(8d)} \geqslant 1 - n^{-b\epsilon/(8d)}.$$

By taking $b \geqslant 8(c+1)d/\epsilon$, we complete the proof of the lemma. Note that b depends on c but not r. $\square$

Strong expanders with $d = \omega(1)$

Here we assume that $d = \omega(1)$. Recall that for constants δ and ϵ, if $|S| \leqslant \delta n/d$ for $S \subseteq V$, then $|\partial S| \geqslant \epsilon d|S|$. As in the previous case here too we introduce the M-VISIT-EXCHANGE$_b$ process, a modification of VISIT-EXCHANGE, parametrised by the constant b, to be fixed later. Let

$$\mathcal{S}(b) = \{S \subset V \mid s \in S, G(S) \text{ is connected and } b \cdot \log n \leqslant |S| \leqslant \delta n/d\}.$$

We say that a set $S \in \mathcal{S}(b)$ is *t-good* if there are at least $\alpha\epsilon d|S|/2$ vertices in ∂S that contain an agent in round t. If for some set $S \in \mathcal{S}(b)$ and some round t the set S is not t-good, then we add a minimal number of agents in S in round t to make it t-good. As in proofs before, the next lemma allows us to use the modified process instead of VISIT-EXCHANGE throughout the proof and avoid dealing with dependencies of agents.

Lemma 5.3.4. *Let T' be any positive integer. For any $c > 0$, there is a constant b such that* M-VISIT-EXCHANGE$_b$ *and* VISIT-EXCHANGE *are identical in the first T' rounds of their executions with probability at least $1 - T' \cdot n^{-(c+1)}$.*

Proof. For a constant b to be determined later, fix a set $S \in \mathcal{S}(b)$. For a node $v \in \partial S$, let X_v be an indicator random variable that is 1 if v has an agent at round t. Since the agents in VISIT-EXCHANGE are initially distributed according to the (uniform) stationary distribution and make independent walks, $\mathbb{E}[X_v] = \alpha$. The collection of random variables $\{X_v \mid v \in \partial S\}$ are negatively associated by exactly the same argument as that of for the loads in the balls-and-bins process [Waj17; DR96]. Thus, we can use a Chernoff bound for the number of vertices in ∂S that contain an agent in round t, denoted by $n_t(\partial S) = \sum_{v \in \partial S} X_v$. In particular, we have $\mathbb{E}[n_t(\partial S)] = \alpha|\partial S| \geqslant \alpha\epsilon d|S|$, by strong expansion, and thus,

$$\mathbb{P}[n_t(\partial S) \geqslant \alpha\epsilon d|S|/2] \geqslant 1 - e^{-\alpha\epsilon d|S|/8}.$$

Next, we can take a union bound over all rounds $t \leqslant T'$ and sets $S \in \mathcal{S}(b)$:

$$\begin{aligned}
\mathbb{P}\left[\bigcap_{\substack{S \in \mathcal{S}(b) \\ 1 \leqslant t \leqslant T'}} \{S \text{ is } t\text{-good}\}\right] &= \mathbb{P}\left[\bigcap_{\substack{S \in \mathcal{S}(b) \\ 1 \leqslant t \leqslant T}} \{n_t(\partial S) \geqslant \alpha\epsilon d|S|/2\}\right] \\
&\geqslant 1 - \sum_{t=1}^{T'} \sum_{k=\lceil b \cdot \log n \rceil}^{\lfloor \delta n/d \rfloor} \sum_{S \in \mathcal{S}(b)\,:\,|S|=k} e^{-\alpha\epsilon d|S|/8} \\
&\geqslant 1 - T' \cdot \sum_{k=\lceil b \cdot \log n \rceil}^{\lfloor \delta n/d \rfloor} d^{2k} \cdot e^{-\alpha\epsilon dk/8}, \quad \text{by Lemma 5.2.2}, \\
&\geqslant 1 - T' \cdot \sum_{k=\lceil b \cdot \log n \rceil}^{n} \left(d^2 \cdot e^{-\alpha\epsilon d/8}\right)^k, \\
&\geqslant 1 - T' \cdot \sum_{k=\lceil b \cdot \log n \rceil}^{n} (1/2)^k, \quad \text{since } d = \omega(1), \\
&\geqslant 1 - T' \cdot n^{-b+1}.
\end{aligned}$$

By taking $b = c + 2$, we complete the proof. □

Lemma 5.3.5. *For any $c > 0$, there are constants $b, \psi > 0$ such that in* M-VISIT-EXCHANGE$_b$, *if $I_t \in \mathcal{S}(b)$ then*

$$\mathbb{P}\left[|I_{t+2}| \geqslant (1+\psi) \cdot |I_t| \mid I_t\right] \geqslant 1 - n^{-(c+1)}.$$

Proof. Let W be the set of vertices in ∂I_t that contain at least one agent in round t. By the construction of M-VISIT-EXCHANGE$_b$,

$$|W| = n_t(\partial I_t) \geqslant \alpha\epsilon d|I_t|/2. \tag{5.3}$$

We arbitrarily partition W into disjoint sets $\{W_u\}_{u\in I_t}$ in such a way that for every $v \in W_u$, $(v, u) \in E$. For $u \in I_t$, define X_u as the indicator variable that in round $t+1$ there is an agent at u that was in W_u at round t. For each agent, this event happens with probability exactly $1/d$, and since the agents move independently,

$$\mathbb{P}\left[X_u = 1 \mid I_t\right] = 1 - (1 - 1/d)^{|W_u|} \geqslant 1 - e^{-|W_u|/d}.$$

From now on, throughout the proof, we will omit the conditioning on the informed set at round t to make the proof more readable. If $S \subseteq I_t$ is set of vertices u that receive an agent from W_u at round $t+1$, then

$$\mathbb{E}\left[|S|\right] \geqslant \sum_{u\in I_t}(1 - e^{-|W_u|/d}) = |I_t| - \sum_{u\in I_t}\left(e^{-1/d}\right)^{|W_u|}. \tag{5.4}$$

To maximise the sum in (5.4), we have to set $|W_u| = d$ for as many vertices as possible, because if for some vertices u and v, $1 \leqslant |W_u| \leqslant |W_v| < d$, then we can increase the sum by reducing $|W_u|$ by one and increasing $|W_v|$ by one, due to Lemma 2.1.1(c). The number of vertices u such that $|W_u| = d$ can be at most $k = \lfloor |W|/d \rfloor$, and assuming for the other u, $|W_u| = 0$, we get

$$\mathbb{E}\left[|S|\right] \geqslant |I_t| - k \cdot e^{-1} - (|I_t| - k) = k(1 - e^{-1}) \geqslant \frac{\alpha\epsilon}{4} \cdot |I_t|,$$

by (5.3). Since the random variables X_u are independent, and $|S| = \sum_{u\in I_t} X_u$, by a Chernoff bound,

$$\mathbb{P}\left[|S| \geqslant \frac{\alpha\epsilon}{8} \cdot |I_t|\right] \geqslant 1 - \exp\left(-\frac{\alpha\epsilon}{32} \cdot |I_t|\right) \geqslant 1 - n^{-\alpha\epsilon b/32} \geqslant 1/2. \tag{5.5}$$

For $v \in \partial I_t$, let Y_v be an indicator variable that an agent that was in ∂I_t at round t, visits I_t at round $t+1$ and then visits v at round $t+2$. Then,

$$\mathbb{P}\left[Y_v = 1\right] \geqslant \mathbb{P}\left[v \in \partial S\right]/d,$$

Clearly $Y_v = 1$ implies that v becomes informed, hence we would like to lower bound the sum $Y = \sum_{v\in\partial I_t} Y_v$. First,

$$\mathbb{E}\left[Y\right] \geqslant \frac{1}{d} \cdot \sum_{v\in\partial I_t} \mathbb{E}\left[\mathbb{1}_{v\in\partial S}\right] = \frac{1}{d} \cdot \mathbb{E}\left[|\partial S \cap \partial I_t|\right] = \frac{1}{d} \cdot \mathbb{E}\left[|\partial S \setminus I_t|\right].$$

Using conditional expectation with (5.5) and the fact that $|\partial S| \geqslant \epsilon d|S|$, we get

$$\mathbb{E}\left[Y\right] \geqslant \frac{1}{d} \cdot \frac{1}{2}\left(\frac{\alpha\epsilon^2 d}{8} - 1\right) \cdot |I_t| \geqslant \frac{\alpha\epsilon^2}{32} \cdot |I_t|.$$

Similar to the example on the balls-and-bins-process [Waj17; DR96], the variables Y_v are negatively associated. Thus, we can apply Chernoff bound,

$$\mathbb{P}\left[Y \geqslant \frac{\alpha\epsilon^2}{64} \cdot |I_t|\right] \geqslant 1 - \exp\left(-\frac{\alpha\epsilon^2}{256} \cdot |I_t|\right) \geqslant 1 - n^{-\alpha\epsilon^2 b/256} \geqslant 1 - n^{-(c+1)},$$

if b is taken sufficiently large. Thus, for $\psi = \alpha\epsilon^2/64$ we have the desired result. $\square$

5.3.3 Phase 3: n informed vertices

In this section we present the final phase of the analysis of VISIT-EXCHANGE on a d-regular expander G. Recall that we assume $d = O(\log n)$. We show that if the number of informed vertices is at least $\Omega(n/d)$ at round t_0 of VISIT-EXCHANGE, then after at most $O(\log n)$ rounds all vertices of G become informed, w.h.p. The analysis does not require strong expansion properties even when $d = \omega(1)$.

Lemma 5.3.6. *For any constants $c, \delta > 0$, if for some round t_0, $|I_{t_0}| \geqslant \delta n/\log n$, then there is a round $\tau_3 = t_0 + O(\log n)$ such that $|I_{\tau_3}| = n$, with probability at least $1 - n^{-c}$.*

Proof. As in the proof of Lemma 5.3.1, here too we consider lazy random walks as we need to use a bound on the mixing time of a random walk. The result holds by the same coupling argument, which we do not repeat here.[1] Let $t_1 = t_0 + t^{\infty}_{\text{mix}}$. By definition of uniform mixing, for any agent $g \in A$ and vertex $u \in V$, if $X_g(t)$ is the location of g at round t, then

$$\mathbb{P}\left[X_g(t_1) = u \mid I_{t_0}\right] \geqslant 1/(2n) \tag{5.6}$$

Consider a lazy random walk $X'(t)$ that starts from stationarity in round t_1. Let τ_g and τ' be the first round after t_1 when the walks $X_g(t)$ and $X'(t)$ visit I_{t_0}, respectively. Denote by $t_2 = t_1 + \eta_1 \cdot \log n$, for $\eta_1 = (\delta(1-\lambda))^{-1}$. Omitting I_{t_0} from the conditionals for readability, we have that

$$\begin{aligned}
\mathbb{P}\left[\tau_g \leqslant t_2\right] &= \sum_{u\in V} \mathbb{P}\left[\tau_g \leqslant t_2 \mid X_g(t_1) = u\right] \cdot \mathbb{P}\left[X_g(t_1) = u\right] \\
&\geqslant \sum_{u\in V} \mathbb{P}\left[\tau_g \leqslant t_2 \mid X_g(t_1) = u\right] \cdot \frac{1}{2n}, \quad \text{by (5.6)}, \\
&= \frac{1}{2} \cdot \sum_{u\in V} \mathbb{P}\left[\tau' \leqslant t_2 \mid X'(t_1) = u\right] \cdot \mathbb{P}\left[X'(t_1) = u\right] \\
&= \frac{1}{2} \cdot \mathbb{P}\left[\tau' \leqslant t_2\right] \\
&\geqslant \frac{\eta_1(1-\lambda)\log n}{2n} \cdot |I_{t_0}|, \quad \text{by Lemma 5.2.1} \\
&\geqslant 1/2,
\end{aligned} \tag{5.7}$$

where the last inequality holds due to the condition on I_{t_0} and the choice of η_1. This implies that if A' is the set of informed agents at round t_2, then $\mathbb{E}\left[|A'|\right] \geqslant \alpha n/2$ and by a Chernoff bound,

$$\mathbb{P}\left[|A'| \geqslant \alpha n/4 \mid I_t\right] \geqslant 1 - e^{-\alpha n/16}. \tag{5.8}$$

Finally, we argue that in $O(\log n)$ rounds all vertices will be visited by one of the agents in A', using similar arguments as above. Finally, let $\tau_3 = t_2 + t^{\infty}_{\text{mix}} + \eta_2 \cdot \log n$ for some constant $\eta_2 > 0$. For a vertex $u \in V$ and agent $g \in A'$ we denote by $\tau_{u,g}$ the first round

[1] Note that in VISIT-EXCHANGE it is not always possible to transform results for lazy walks to non-lazy ones as easily. This is because, in general, the order in which agents arrive at a vertex matters, but not in these two lemmas.

after $t_2 + t_{\text{mix}}^{\infty}$ when g visits u. By the derivation similar to (5.7) and applying Lemma 5.2.1 for the set $S = \{u\}$, we get that

$$\mathbb{P}\left[\tau_{u,g} \leqslant \tau_3\right] \geqslant \frac{\eta_2(1-\lambda)\log n}{4n}.$$

Let $\mathcal{E}_u$ be the event that u is informed at round τ_3, which will happen if at least one of the agents in A' visits u at or before round τ_3. Since the walks are independent,

$$\begin{aligned}
\mathbb{P}\left[\mathcal{E}_u\right] &\geqslant \mathbb{P}\left[\mathcal{E}_u \mid |A'| \geqslant \alpha n/4\right] \cdot \mathbb{P}\left[|A'| \geqslant \alpha n/4\right] \\
&\geqslant 1 - \left(1 - \frac{\eta_2(1-\lambda)\log n}{4n}\right)^{\alpha n/4} - e^{-\alpha n/16}, \quad \text{by (5.8)}, \\
&\geqslant 1 - n^{-\eta\alpha(1-\lambda)/16} - e^{-\alpha n/16} \\
&\geqslant 1 - n^{-(c+1)},
\end{aligned}$$

for a sufficiently large constant η_2. Using a union bound for all $u \in V$ completes the proof since $\tau_3 = t_0 + O(\log n)$. □

5.3.4 Putting pieces together

Proof of Theorem 5.1.2. Fix a constant $c > 0$. If $d = O(1)$, then we choose b and ψ by Lemma 5.3.3 and then set r by Lemma 5.3.2. If $d = \omega(1)$, then we choose b and ψ by Lemma 5.3.5 and set $r = 2$. We consider the corresponding M-VISIT-EXCHANGE$_b$ process. Lemma 5.3.1 implies that for some $\tau_1 = O(\log n)$, $|I_{\tau_1}| \geqslant b \cdot \log n$, with probability at least $1 - n^{-c}$. Next, for $i \in \{1, \ldots, \lceil \log_{1+\psi} n \rceil\}$, consider rounds $\tau_1 + i \cdot r$. By Lemma 5.3.3 and Lemma 5.3.5, for each i, given $I_{\tau_1+(i-1)\cdot r}$, either $|I_{\tau_1+i\cdot r}| \geqslant \delta n/d$ or $|I_{\tau_1+(i+1)\cdot r}| \geqslant (1+\psi) \cdot |I_{\tau_1+i\cdot r}|$, with probability at least $1 - n^{-(c+1)}$. By a union bound over all values of i, we have that $|I_{\tau_1+\tau_2}| \geqslant \delta n/d$, for $\tau_2 = O(\log n)$, with probability at least $1 - 2n^{-c}$. Next, we apply Lemma 5.3.6 and union bound again, showing that M-VISIT-EXCHANGE$_b$ informs all vertices of the graph in $\tau = \tau_1 + \tau_2 + \tau_3 = O(\log n)$ rounds, with probability at least $1 - 3n^{-c}$. Finally, by Lemmas 5.3.2 and 5.3.4, M-VISIT-EXCHANGE$_b$ and VISIT-EXCHANGE are identical in the first $\tau = O(\log n)$ rounds, with probability at least $1 - n^{-c}$. Thus, $T = O(\log n)$, w.h.p. □

5.4 Random regular graphs

We prove Theorem 5.1.4 that for random regular graphs, $T_{\text{visitx}} = O(\log n)$, w.h.p.

Proof of Theorem 5.1.4. Fix a constant $c > 0$. By [Bro+98, Lemma 18], the spectral gap of a random d-regular graph is at least $1 - \gamma/\sqrt{d}$ for some constant γ, with probability at least $1 - n^{-c}$. On the other hand, it is not hard to verify [DFS14, Theorem 4.12], if a d-regular graph's spectral gap is at least $1 - O(1/\sqrt{d})$, then it is a strong expander.[2] By Theorem 5.1.3, $T_{\text{visitx}} = O(\log n)$ for strong expanders with probability at least $1 - n^{-c}$. Combining these facts we have that on random regular graph the broadcast time of VISIT-EXCHANGE is $O(\log n)$ with probability at least $1 - 2n^{-c}$. □

[2] In the reference, the term "expanding graph" is used for what we call "strong expander" here.

5.5 General expanders

In this section we prove Theorem 5.1.5, that in d-regular expander graphs VISIT-EXCHANGE using $|A| = \alpha \cdot n$ agents has a broadcast time of $O(\log n \cdot \log\log n)$. As explained in the introduction we prove the bound T_{visitx} in terms of T_{push}, via the same stages as in the proof of Theorem 3.4.1. The difference of the two proofs is that the processes are coupled every $(r+1)$th round for $r = O(\log\log n)$ instead of every other round, and we use Lemma 4.2.1 to argue that sufficiently many agents arrive at a vertex to make progress along the path via which PUSH informed a vertex.

5.5.1 Coupling

We define the coupling of VISIT-EXCHANGE and PUSH here that depends on an integer $r \geqslant 1$. When $r = 1$, the coupling is identical to that of in Section 3.4.2, so we use the notation from that section. For a vertex u, denote by τ_u and t_u the rounds when u becomes informed in PUSH and VISIT-EXCHANGE, respectively. For an integer $i \geqslant 1$, we denote by $\pi_u(i)$ the ith vertex sampled by u after being informed in PUSH.

In VISIT-EXCHANGE, we denote by $Z_u(t)$ the set of agents that visit u in round t, which is also the set of agents departing u in round $t+1$. For an integer $\ell \geqslant 0$, consider all visits to u in round $\ell(r+1)$, in chronological order, ordering the visits in the same round with respect to a predefined but arbitrary total order over the agents. We call these visits as r-visits (in Section 3.4.2 they are called even visits). For each $i \geqslant 1$, consider the agent that performs the ith r-visit to u. We denote by $p_u(i)$ the neighbour of u where the agent moves in round $\ell(r+1)$. Formally, let

$$W_u = \{(\ell, g) \mid \ell(r+1) \geqslant t_u, j \in \mathbb{N}, x_g(\ell(r+1)) = u\},$$

where $x_g(t)$ is the random walk performed by the agent $g \in A$. We order the elements of W_u such that $(\ell, g) < (\ell', g')$ if $\ell < \ell'$, or $\ell = \ell'$ and $g < g'$. Then, if (ℓ, g) is the ith smallest element of W_u, we set $p_u(i) = x_g(\ell(r+1)+1)$.

The coupling is defined by setting $\pi_u(i) = p_u(i)$, for all $i \geqslant 1$.

5.5.2 A modified Visit-Exchange process

We use the notation from Section 4.2. For a vertex u, $N_u(t)$ is the number of agents that are at vertex u at round t. For an integer $r > 0$ and round t, let

$$\hat{N}_u(t, r) = \mathbb{E}\left[N_u(t+r) \mid N_v(t),\ \text{for all } v \in V\right] = \sum_{v \in V} p^r_{v,u} \cdot N_v(t),$$

where $p^r_{v,u}$ is the probability that a random walk starting from v is at u after exactly r rounds.

The M-VISIT-EXCHANGE$_r$ process, parametrised by an integer r, is defined as follows. If for some round t and vertex u, the following condition does *not* hold:

$$\hat{N}_u(t, r) \geqslant \alpha/2, \tag{5.9}$$

then we add a minimal set of agents in the graph immediately after round t, so that the condition holds. The following lemma allows us to bound the broadcast time of M-VISIT-EXCHANGE and argue that the bound also holds for VISIT-EXCHANGE.

Lemma 5.5.1. *For any constant $c > 0$, there is an integer $r = O(\log\log n)$, such that* VISIT-EXCHANGE *and* M-VISIT-EXCHANGE$_r$ *are identical for the first T' rounds of their execution with probability at least $1 - T' \cdot n^{-(c+1)}$.*

Proof. Our goal is to apply Lemma 4.2.1 to bound $\hat{N}_r(t, r)$, hence we start by considering the return probability $p^{2r}_{u,u}$ for some vertex u. Let ℓ be an even integer. If $X'(\ell)$ is a lazy walk starting at vertex u, then by Lemma 4.2.4a, we have that

$$p^{\ell}_{uu} \leqslant 2 \cdot \mathbb{P}\left[X'(\ell) = u\right].$$

Let $1 = \lambda'_1 \geqslant \ldots \geqslant \lambda'_n$ be the eigenvalues of the transition matrix P' of the walk $X'(\ell)$, with corresponding eigenfunctions f_i. By [LP17, Lemma 12.2], f_i can be chosen to be orthonormal, in which case,

$$\mathbb{P}\left[X'(\ell) = u\right] = \pi(u) + \pi(u) \cdot \sum_{i=2}^{n} f_i(u)^2 \cdot \lambda_i^{\ell}.$$

Since $X'(\ell)$ is a lazy walk, $0 \leqslant \lambda'_i \leqslant (1 + \lambda)/2 < 1$. We also have that $\pi(u) = 1/n$, and by orthonormality of f_i, $\sum_{i=2}^{n} f_i(u)^2 \leqslant n$. Combining these facts together we get

$$p^{\ell}_{u,u} \leqslant 2 \cdot \mathbb{P}\left[X'(\ell) = u\right] \leqslant \frac{2}{n} + 2 \cdot \left(\frac{1+\lambda}{2}\right)^{\ell}. \tag{5.10}$$

Now we return to the main claim of the lemma. Since $\lambda < 1$, then by (5.10), for any constant $\eta > 0$, there is an $r = \Theta(\log\log n)$, such that

$$p^{2r}_{u,u} \leqslant \frac{\eta}{\log n}.$$

Thus, by substituting $\pi(u) = 1/n$ and $|A| = \alpha n$ in Lemma 4.2.1, we get that

$$\mathbb{P}\left[\hat{N}_u(t, r) \leqslant \alpha/2\right] \leqslant \exp\left(-\frac{\alpha \log n}{8\eta}\right) = n^{-\alpha/(8\eta)}.$$

We can take $\eta \geqslant 8/\alpha + c + 2$, and take a union bound for all vertices u and all rounds t up to T' to complete the proof. □

5.5.3 Proof outline of Theorem 5.1.5

The construction of M-VISIT-EXCHANGE and the coupling of the processes have the necessary components to apply the arguments from the proof of Lemma 3.4.3 presented in Section 3.4.4. Thus, we do not go into all details of the proof of the following analogous lemma and use the same notation as in the earlier proof.

Lemma 5.5.2. *For any constant $c > 0$, there is an integer $r = O(\log\log n)$ and such that with the coupling presented in Section 5.5.1, for any $u \in V$,*

$$\mathbb{P}\left[t'_u \geqslant (r+1)(\tau_u + \log n)\right] \leqslant n^{-(c+1)},$$

where τ_u and t'_u are the rounds when u is informed in the coupled processes PUSH *and* M-VISIT-EXCHANGE$_r$*, respectively.*

Proof outline. In the proof we do not deal with VISIT-EXCHANGE, so we use t_u instead of t'_u. As in Theorem 3.4.1, we compare the number of rounds M-VISIT-EXCHANGE$_r$ and PUSH take to inform a vertex via a specific edge (v, w). More formally, denote by $\delta_{v,w}$ and $\Delta_{v,w}$ the number of rounds until w becomes informed since the round when v is informed, in PUSH and M-VISIT-EXCHANGE$_r$, respectively. We would like to bound $\Delta_{v,w}$ in terms of $\delta_{v,w}$.

Let $B_v^{(j)}$ be the number of M-VISIT-EXCHANGE$_r$ rounds between r-visits $j-1$ and j to v. Then, analogously to (3.11),

$$D_{v,w} \leqslant \sum_{j=1}^{\delta_{v,w}} B_v^{(j)},$$

by the coupling of PUSH and M-VISIT-EXCHANGE.

For an agent g, let Y_g be the indicator random variable that g visits v in round $t = \ell(r+1) \geqslant t_v + r$ for the some $\ell \geqslant 0$. If $Y = \sum_{g \in A} Y_g$, then by (5.9), $\mathbb{E}[Y] = \mathbb{E}\left[\hat{N}_v(t-r, r)\right] \geqslant \alpha/2$. Since the agents perform independent walks, by a Chernoff bound we have that

$$\mathbb{P}[Y \geqslant 1] \geqslant \mathbb{P}[Y \geqslant \alpha/4] \geqslant 1 - e^{-\alpha/16}.$$

Thus, for $p = 1 - e^{-\alpha/16}$, in rounds $t = \ell(r+1)$ an agent visits v with probability p, independently from the execution up until round $t - r$. Thus, the number of rounds between two r-visits to v, namely $B_v^{(j)}$ for $1 \leqslant j \leqslant \delta_{v,w}$, is stochastically dominated by $(r+1)F_v^{(j)}$, where $\{F_v^{(j)}\}_{j \geqslant 1}$ is a collection of independent geometric random variables with success probability p.

The rest of the proof follows Lemma 3.4.3 exactly, except that instead of a factor of 2 corresponding to every odd coupled round, there is a factor of $(r+1)$ since here we only couple every $(r+1)$th round of M-VISIT-EXCHANGE$_r$ process. □

Proof of Theorem 5.1.5. Fix a constant $c > 0$. By [Chi+18], there is an integer $T' = O(\log n)$ such that $T_{\text{push}} \leqslant T'$ with probability at least $1 - n^{-(c+1)}$. Let $r = O(\log\log n)$ be determined from Lemma 5.5.1 given the constant c, and consider the M-VISIT-EXCHANGE$_r$ process. For a fixed vertex u, we have that

$$\begin{aligned}\mathbb{P}[t'_u \geqslant (r+1)(T' + \log n)] &\leqslant \mathbb{P}[t'_u \geqslant (r+1)(\tau_u + \log n) \mid \tau_u \leqslant T'] \cdot \mathbb{P}[\tau_u \leqslant T'] + \mathbb{P}[\tau_u > T'] \\ &\leqslant 2n^{-(c+1)},\end{aligned}$$

by Lemma 5.5.2. Thus, by taking a union bound over all vertices u, the broadcast time of M-VISIT-EXCHANGE is at most $(r+1)(T' + \log n) = O(\log n \cdot \log\log n)$, with probability $1 - 2n^{-c}$. Since M-VISIT-EXCHANGE and VISIT-EXCHANGE are identical with probability at least $1 - T'n^{-(c+1)}$, we have that $T_{\text{visitx}} = O(\log n \cdot \log\log n)$, w.h.p. □

Chapter 6

Bounds for balanced trees

6.1 Introduction

In this section, we analyse the broadcast time of VISIT-EXCHANGE on balanced trees, as an example of a hierarchical communication network. We denote by $R_{b,h}$ a rooted b-ary tree where each vertex is at a distance at most h from the root. The total number of vertices is $n = (b^{h+1} - 1)/(b - 1)$. The following is the main result of the chapter.

Theorem 6.1.1. *For any b-ary tree $R_{b,h}$ with $b \geqslant 2$ and any source vertex, $T_{\text{visitx}} = O(h \log h + \log n)$, w.h.p. Furthermore, for the binary tree $R_{2,h}$, $T_{\text{visitx}} = \Omega(h \log h) = \Omega(\log n \cdot \log \log n)$, w.h.p.*

Note that the broadcast time of randomised rumour spreading on $R_{b,h}$ is $\Theta(b \log n)$, w.h.p. (For PUSH-PULL this holds only if $h \geqslant 2$ to exclude the special case of a star graph.) Thus, VISIT-EXCHANGE is slower than PUSH for small b, and faster than PUSH-PULL for larger b. An interesting implication of the upper bound of Theorem 6.1.1 is that the cover time of the tree by n random walks starting from stationarity has a super-linear speed-up, compared to the cover time for a single random walk, which is $\Omega(n \log^2 n)$ [LP17, Section 13.3.1].

We give now an overview of the proof of Theorem 6.1.1, for the binary tree; the case of $b > 2$ is similar. To prove the upper bound, we first show the following: For a fixed vertex u of distance at most $h - \log h$ from the root ρ of the tree, information spreads from u to ρ in $O(\log n)$ rounds, w.h.p. For every vertex v on the path from u to ρ, we identify a subset S_v of the descendants of v at distance $m = \Theta(\log h)$ from v. We show that agents that were in S_v visit v and then its parent at a constant rate (so if v were informed, information progresses towards the root). Furthermore, this holds for all vertices v *independently* from one another since we can construct the sets S_v such that they are disjoint. This means that we can apply a simple concentration bound to prove that ρ becomes informed in $O(\log n)$ rounds.

To show a constant rate of visits to v and then its parents from S_v, instead of a simpler rate of $1/\Theta(\log h)$, a careful pipeline argument is used, which essentially allows us to pretend that a completely new set of agents leaves S_v in every round by "reusing" the agents that return to S_v. In the pipeline argument, we are only concerned with visits of agents to a vertex v from its descendants. Thus, it suffices to only work with the distance Y_t of the agents from vertex v. Y_t is biased random walk on a line, so we study a new process called LUCKY-GAMBLER in which a set of gamblers performs a biased random walk on a path of length m. Gamblers originate at one endpoint of the path and we prove that they arrive at the other endpoint (against the bias) at a constant rate. By coupling the gamblers to agents of VISIT-EXCHANGE we arrive at the required claim about visits to each vertex v on the path from ρ to u. This completes the proof for the "top" of the tree.

We use a different method for the dissemination in the lower $\Theta(\log n)$ levels of the tree. First, we use another coupling of agents with gamblers from LUCKY-GAMBLER to show that at least h agents arrive at the roots of the lower subtrees of height $\Theta(\log h)$. Then we simply bound the cover time of a tree of height $\Theta(\log h)$ by h walks starting from the root. The cover time is $O(h \log h)$ steps w.h.p. (in n), completing the proof of the upper bound in Theorem 6.1.1.

Finally, to show the lower bound of Theorem 6.1.1, we bound from below the cover time of the tree by n random walks starting from stationarity. We prove that with high probability, there are subtrees of height $\Theta(\log \log n)$ that are not visited by any agent in the first $\Omega(\log n \cdot \log \log n)$ rounds. This lower bound is similar (but weaker) to the one for the sparse regular graph for which VISIT-EXCHANGE is slower than PUSH, presented in Section 3.2.4. In both cases, there are logarithmically many nodes which are not visited by any agents for a long period of time.

In Appendix A we complement our analysis by an experimental evaluation of VISIT-EXCHANGE and randomised rumour spreading processes for balanced trees. Fig. A.3 validates our results showing that for small values of b PUSH-PULL is faster but as b increases VISIT-EXCHANGE becomes faster. This also indicates that the constants hidden in the asymptotic notation of our bounds do not appear to be very large.

6.1.1 Notation

The graph $R_{b,h}$ is a rooted b-ary tree, where each vertex at distance less than h from the root has b children, and all leaves are at distance h from the root; thus h is the *height* of the tree. The total number of vertices is $n = (b^{h+1} - 1)/(b - 1)$. The set of children of vertex u is denoted by C_u. The set of descendants of u is denoted D_u; precisely, D_u contains the vertices in the subtree rooted at u, including u itself. The height of that subtree is denoted by h_u.

We define the set $B_{u,l} = \{v \in D_u \mid h_v = h_u - l\}$, which contains all descendants of v at distance l from u. Finally, as before, $Z_u(t)$ denotes the set of agents at vertex u at round t, and $Z_S(t) = \bigcup_{u \in S} Z(t)$ is the set of agents in the set $S \subseteq V$ at that round.

6.2 The Lucky-Gambler process

In this section we define an auxiliary process, called LUCKY-GAMBLER, which will be used in the analysis. The analysis of LUCKY-GAMBLER is interesting in its own right. The process has three parameters: two integers $m, k > 0$, and a probability $p < 1/2$. Consider a path graph P_m of length m, with vertices 0 up to m. For every integer $s \geqslant 0$, at round s exactly k *gamblers* appear on vertex 1 and make a biased random walk: for $0 < i < m$, the probability of moving from vertex i to $(i+1)$ and $(i-1)$ is $p_{i,i+1} = p$ and $p_{i,i-1} = 1 - p = q$, respectively. When the gambler reaches vertex 0 or m, it stops, i.e., $p_{0,0} = p_{m,m} = 1$ (states $0, m$ are absorbing). We will write LUCKY-GAMBLER(m, p, k) to explicitly state the parameters of the process.

For a vertex v of $R_{b,h}$, where $h_v \geqslant m$, we are going to couple the movement of the agents in part of the subtree of v, with the gamblers in LUCKY-GAMBLER. Using the coupling and the next lemmas, we argue that v receives agents at a constant rate. By carefully selecting the agents that are coupled, we can claim that agents arrive at constant rate to every vertex v on a given path to the root, independently for each vertex.

Lemma 6.2.1. *If $p = 1/(b+1)$ and $k \geqslant \epsilon \cdot b^{m-1}$, for some constant $\epsilon > 0$, then there is a constant $\beta < 1$ such that for any round $t \geqslant 4m$ and positive integer Δ the probability that no gambler reaches vertex m during any round in $\gamma_0 = \{t, \dots, t+\Delta-1\}$ is at most $(1-\beta)^{\Delta}$.*

Lemma 6.2.2. *If $p = 1/(b+1)$ and $k \geqslant \kappa \cdot b^{m-1}$, for some integer κ, then there is a constant γ, such that for any integer $\tau \geqslant 8m$, at least $\gamma\kappa\tau$ gamblers reach vertex m in the first τ rounds, with probability at least $1 - e^{-\gamma\kappa\tau/4}$.*

To prove Lemmas 6.2.1 and 6.2.2 we will use the next two results for a single gambler g making a biased random walk on P_m starting at round 0. Let $X_g(t)$ be the position of gambler g at round t and let $\tau_g(i) = \min\{t \mid X_g(t) = i\}$ be the hitting time of vertex i of g. We denote the event that $\tau_g(m) < \tau_g(0)$ as $\mathcal{L}_g$, and we will say that g is *lucky* if it occurs.

Lemma 6.2.3 ([Fel68, Chapter 14])**.** *If $p \neq q$, then for $0 < i < m$, $\mathbb{P}[\mathcal{L}_g \mid X_g(0) = i] = \frac{(q/p)^i - 1}{(q/p)^m - 1}$.*

Lemma 6.2.4. *If $p < q$, then for $0 < i < m$, $\mathbb{E}[\tau_g(m) \mid \mathcal{L}_g, X_g(0) = i] \leqslant \frac{m-i}{q-p}$.*

Proof of Lemma 6.2.4. By Bayes' theorem and Lemma 6.2.3, we can explicitly find the transition probabilities of the Markov chain of g, conditioned on $\mathcal{L}_g$. For $0 < j < m$,

$$\begin{aligned}
\mathbb{P}[X_g(1) = j+1 \mid \mathcal{L}_g, X_g(0) = j] &= \frac{\mathbb{P}[X_g(1) = j+1, \mathcal{L}_g \mid X(0) = j]}{\mathbb{P}[\mathcal{L}_g \mid X_g(0) = j]} \\
&= p_{j,j+1} \cdot \frac{\mathbb{P}[\mathcal{L}_g \mid X_g(0) = j+1]}{\mathbb{P}[\mathcal{L}_g \mid X_g(0) = j]} \\
&= p \cdot \frac{(q/p)^{j+1} - 1}{(q/p)^j - 1} = p \cdot \left(\frac{q}{p} + \frac{(q/p) - 1}{(q/p)^j - 1}\right) \geqslant q.
\end{aligned}$$

Consider a new random walk $X'(t)$ on $\mathbb{Z}$ with transition probabilities $p'_{j,j+1} = q$ and $p'_{j+1,j} = p$, for any $j \in \mathbb{Z}$. Let $\tau'(j)$ be the hitting time of $X'(t)$ of vertex j. The inequality above implies that $\tau(j)$ conditioned on $\mathcal{L}_g$ is stochastically dominated by $\tau'(j)$, when both start from the same vertex i, and thus,

$$\mathbb{E}[\tau_g(m) \mid \mathcal{L}_g, X_g(0) = i] \leqslant \mathbb{E}[\tau'(m) \mid X'(0) = i] = \frac{m-i}{q-p}. \qquad \square$$

We are ready to prove the main claims of this section.

Proof of Lemma 6.2.1. For $s \geqslant 0$ and $1 \leqslant i \leqslant k$, let $g_{s,i}$ be the ith gambler that starts its walk at round s at vertex 1. Let $\tau_{s,i} = \tau_{g_{s,i}}$ be defined as for the single gambler g above. Clearly, $\tau_{s,i}(j) - s$ and $\tau_g(j)$ are identically distributed, if $X_g(0) = 1$. We also extend the definition of γ_0, letting $\gamma_s = \{t-s, \dots, t+\Delta-s-1\}$.

We would like to study the number of lucky gamblers that reach m at rounds in γ_0. Consider first a "toy" example, which assumes that for each s, exactly one gambler is lucky among the k gamblers that start their walk at round s. Suppose that g'_s is that lucky gambler. We study the expected number of these agents that reach m during the rounds in γ_0:

$$\mathbb{E}\left[\sum_{s\geqslant 0} \mathbf{1}_{\{\tau_{g'_s}(m)\in\gamma_0\}} \mid \mathcal{L}_{g'_s} \text{ for } s\geqslant 0\right] = \sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_{g'_s}(m)\in\gamma_0 \mid \mathcal{L}_{g'_s}\right] = \sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right].$$

The setup in the "toy" example is unlikely to occur, however, we use it as a motivation to lower bound the last quantity, which will be used in the main part of the proof.

$$\begin{aligned}
\sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right] &= \sum_{l=0}^{\Delta-1} \sum_{\substack{0\leqslant s\leqslant t+\Delta \\ s\equiv l \pmod{\Delta}}} \mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right], \\
&\qquad \text{the inner sum is over every } \Delta\text{th summand}, \\
&\geqslant \sum_{l=0}^{\Delta-1} \mathbb{P}\left[\tau_g(m)<t \mid \mathcal{L}_g\right], \quad \text{by union of disjoint events}, \\
&= \Delta\cdot\mathbb{P}\left[\tau_g(m)<t \mid \mathcal{L}_g\right] \\
&\geqslant \Delta\cdot\left(1-\frac{\mathbb{E}\left[\tau_g(m)\mid\mathcal{L}_g\right]}{t}\right), \quad \text{by Markov's inequality}, \\
&\geqslant \Delta\cdot\left(1-\frac{m\cdot(b+1)}{t\cdot(b-1)}\right), \quad \text{by Lemma 6.2.4 as } q-p=\frac{b-1}{b+1}, \\
&\geqslant \Delta\cdot\left(1-\frac{b+1}{4(b-1)}\right), \quad \text{since } t\geqslant 4m, \\
&\geqslant \Delta/4.
\end{aligned}$$

We can now bound the probability that no agent visits vertex m between rounds t and $t+\Delta$:

$$\begin{aligned}
\mathbb{P}\left[\bigcap_{\substack{0\leqslant s\leqslant t+\Delta \\ 1\leqslant i\leqslant k}} \{\tau_{s,i}(m)\notin\gamma_0\}\right] &= \prod_{s=0}^{t+\Delta}\left(\mathbb{P}\left[\tau_{s,i}(m)\notin\gamma_0\right]\right)^k, \quad \text{by independence of the walks}, \\
&= \prod_{s=0}^{t+\Delta}\left(\mathbb{P}\left[\tau_g(m)\notin\gamma_s\right]\right)^k \\
&= \prod_{s=0}^{t+\Delta}\left(1-\mathbb{P}\left[\mathcal{L}_g\right]\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s\mid\mathcal{L}_g\right]\right)^k \\
&= \prod_{s=0}^{t+\Delta}\left(1-\frac{b-1}{b^m-1}\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s\mid\mathcal{L}_g\right]\right)^k, \text{ by Lemma 6.2.3}, \\
&\leqslant \prod_{s=0}^{t+\Delta}\exp\left(-\frac{k\cdot(b-1)}{b^m-1}\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s\mid\mathcal{L}_g\right]\right) \\
&\leqslant \exp\left(-\epsilon\cdot\frac{b^{m-1}(b-1)}{b^m-1}\cdot\sum_{s=0}^{t+\Delta}\mathbb{P}\left[\tau_g(m)\in\gamma_s\mid\mathcal{L}_g\right]\right)
\end{aligned}$$

$$\leqslant \exp\left(-\frac{\epsilon\Delta}{8}\right), \quad \text{by the analysis of the toy example.} \qquad \square$$

Proof of Lemma 6.2.2. For $i \in \{1, \ldots, k\}$, consider a gambler $g_{s,i}$ that starts its walk at round $s \leqslant t/2$. Let $X_{s,i} = 1$ if $g_{s,i}$ is lucky and reaches vertex m before round t, i.e., $\tau_{s,i}(m) \leqslant t$. Since $\tau_{s,i}(m) - s$ and $\tau_g(m)$ are identically distributed,

$$\begin{aligned}
\mathbb{P}[X_{s,i} = 1] &= \mathbb{P}\left[\tau_{s,i}(m) \leqslant t \mid \mathcal{L}_{g_{s,i}}\right] \cdot \mathbb{P}\left[\mathcal{L}_{g_{s,i}}\right] \\
&= \mathbb{P}\left[\tau_g(m) + s \leqslant t \mid \mathcal{L}_g\right] \cdot \mathbb{P}\left[\mathcal{L}_g\right] \\
&\geqslant \mathbb{P}\left[\tau_g(m) \leqslant t/2 \mid \mathcal{L}_g\right] \cdot \mathbb{P}\left[\mathcal{L}_g\right], \quad \text{since } s \leqslant t/2, \\
&\geqslant \left(1 - \frac{\mathbb{E}\left[\tau_g(m) \mid \mathcal{L}_g\right]}{t/2}\right) \cdot \mathbb{P}\left[\mathcal{L}_g\right], \quad \text{by Markov's inequality,} \\
&\geqslant \left(1 - \frac{2m(b+1)}{t(b-1)}\right) \cdot \frac{b-1}{b^m - 1}, \quad \text{by Lemmas 6.2.3 and 6.2.4,} \\
&\geqslant \frac{1}{8 \cdot b^{m-1}}.
\end{aligned}$$

If N is the number of gamblers that arrive at vertex m at before round t, then

$$\mathbb{E}[N] \geqslant \sum_{s=0}^{t/2} \sum_{i=1}^{k} \mathbb{E}[X_{s,i}] \geqslant \frac{\kappa b^{m-1} t}{2} \cdot \frac{1}{8b^{m-1}} = \frac{\kappa t}{16}.$$

Since the variables $X_{s,i}$ are independent, we can prove the lemma by an application of Chernoff bound. $\square$

6.3 A modified Visit-Exchange process

We define another auxiliary process, called M-VISIT-EXCHANGE, which is a slight modification of the original VISIT-EXCHANGE process. We use M-VISIT-EXCHANGE in most of the analysis and then use its equivalence (w.h.p.) with VISIT-EXCHANGE to prove the main theorem for the latter.

6.3.1 Process definition

Let m be the *smallest* integer such that $b^m \geqslant \mu \cdot \ln n$ for a constant μ to be defined later, and let $k = \lceil \alpha \cdot b^m/8 \rceil$. Recall that $B_{u,m}$ is the set of descendants of u at distance m and for a set $S \subset V$, $Z_S(t)$ is the set of agents in S in round t. Consider a vertex u of the tree, such that $h_u \geqslant m$. Let v be one of the children of u and define $Z'_{u,v}(t)$ be the set of agents that are in $B_{u,m-1}$ at round t and were in $B_{u,m} \setminus B_{v,m-1}$ the round before, i.e.,

$$Z'_{u,v}(t) = Z_{B_{u,m-1}}(t) \cap Z_{B_{u,m} \setminus B_{v,m-1}}(t-1).$$

For a round $t \geqslant 0$ let $q_{u,v}(t)$ be the smallest non-negative integer for which

$$|Z'_{u,v}(t)| + q_{u,v}(t) \geqslant \left\lceil \frac{\alpha}{8} \cdot |B_{u,m}| \right\rceil = \left\lceil \frac{\alpha}{8} \cdot b^m \right\rceil = k.$$

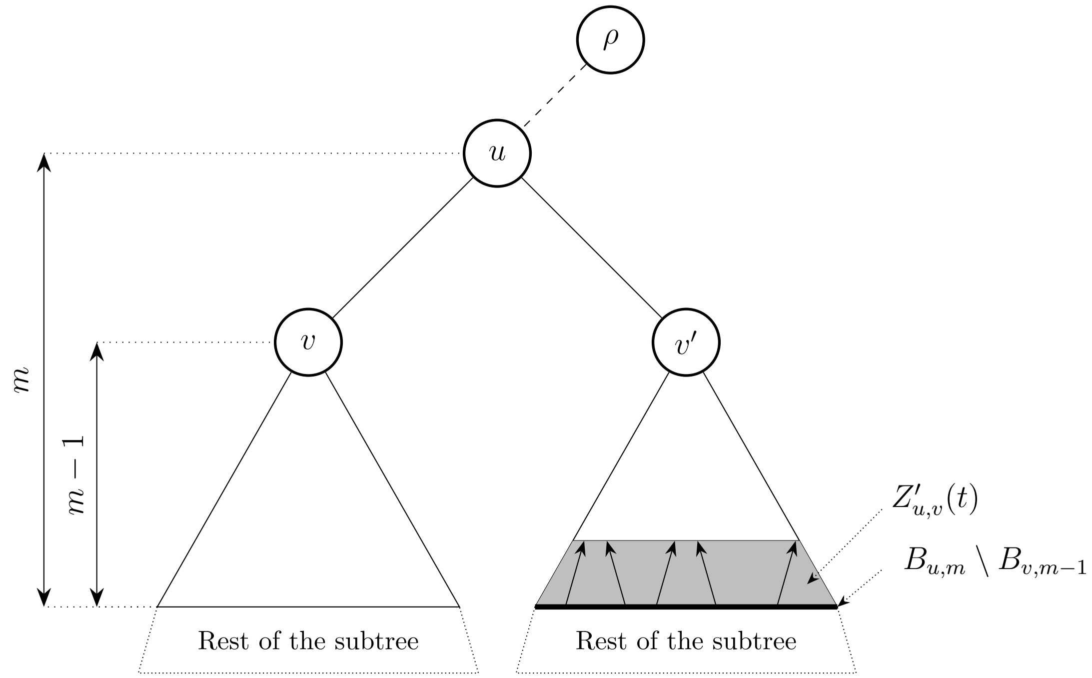

Figure 6.1 The construction of M-VISIT-EXCHANGE process for vertex u of a binary tree with a root ρ. The shaded area indicates the agents in set $Z'_{u,v}(t)$ that move from $B_{u,m}$ closer to u in round t, and are *not* in the subtree of v. In M-VISIT-EXCHANGE, $|Z'_{u,v}(t)| \geqslant \alpha b^m/8$.

To construct M-VISIT-EXCHANGE we add exactly $q_{u,v}(t)$ agents in $B_{u,m-1}$ at round t (it is not important to which vertices in $B_{u,m-1}$ these agents are added). See Fig. 6.1 for an illustration of these definitions for binary trees, where $B_{u,m} \setminus B_{v,m-1} = B_{v',m-1}$ if v' is the sibling of v.

To motivate the construction of M-VISIT-EXCHANGE, consider a vertex u and its child v, such that $m \leqslant h_u < h$. In round t of M-VISIT-EXCHANGE, there are at least k agents at vertices in $B_{u,m} \setminus B_{v,m-1}$ (of height $h_u - m$) that move closer to u in the next round. This allows us to couple these agents to that of gamblers in a LUCKY-GAMBLER$(m{+}1, 1/(b{+}1), k)$ process, and use our results from Section 6.2 to show that agents arrive at the parent of u at a constant rate. A key insight is that by not considering agents that are in descendants of v, the same argument can be made for vertex v, independently of u, if $h_v \geqslant m$ too. By repeating this argument, we show that in $O(\log n)$ rounds all vertices of height at least m are informed once one such vertex is informed. M-VISIT-EXCHANGE and LUCKY-GAMBLER are also used to analyse the spread of the message in the vertices of height at most m.

6.3.2 Equivalence to Visit-Exchange

Using a Chernoff bound we can show that M-VISIT-EXCHANGE and VISIT-EXCHANGE are equivalent in the first polynomially many rounds, w.h.p.

Lemma 6.3.1. *The probability that no agent is added in the* M-VISIT-EXCHANGE *process in the first r rounds is at least* $1 - r \cdot n^{-\frac{\alpha \cdot \mu}{32}+1}$.

Proof. Fix a vertex u with $h_u \geqslant m$ and a round $t \geqslant 0$ of VISIT-EXCHANGE. For an agent $g \in A$, let X_g be the indicator random variable that $g \in Z'_u(t)$. Since the agents are distributed by the stationary distribution, the probability that in any round, agent g traverses an edge in a particular direction is exactly $1/(2 \cdot |E|) = 1/(2 \cdot (n-1))$. Thus,

$$\mathbb{E}\left[|Z'_u(t)|\right] = \sum_{g \in A} \mathbb{P}\left[X_g = 1\right] = \frac{\alpha n \cdot |B_{u,m} \setminus B_{v,m-1}|}{2 \cdot (n-1)} \geqslant \frac{\alpha}{2} \cdot (b-1) \cdot b^{m-1} \geqslant \frac{\alpha}{4} \cdot b^m.$$

By an application of a Chernoff bound we get that

$$\mathbb{P}\left[b_{u,v}(t) > 0\right] = \mathbb{P}\left[|Z'_u(t)| < \frac{\alpha}{8} \cdot b^m\right] \leqslant \exp\left(-\frac{\alpha}{32} \cdot b^m\right) \leqslant n^{-\frac{\alpha \cdot \mu}{32}}.$$

By taking a union bound over all rounds t up until r and edges uv, we complete the proof. □

We will use the same notation for M-VISIT-EXCHANGE and VISIT-EXCHANGE processes.

6.4 Proof of Theorem 6.1.1: the upper bound

Lemma 6.4.1. *Let u be any vertex of the tree $R_{b,h}$ such that $h_u \geqslant m$. For any constant $c > 0$, if u is informed, then after $O(\log n)$ rounds of* M-VISIT-EXCHANGE *the root ρ of $R_{b,h}$ gets informed, with probability at least $1 - n^{-c}$.*

Proof. Consider the path $\langle u = u_1, \ldots, u_l = \rho \rangle$ from u to the root of the tree. Due to the symmetry of the tree, we can assume that the path is the "leftmost" path of the tree, i.e., for any $i \geqslant 1$, u_{i-1} is the leftmost child of u_i (for consistency, we let u_0 be the leftmost child of u_1). Roughly speaking, we show that for any i, the number of rounds between two consecutive visits to u_i (by a certain subset of agent) follows a geometric distribution, independently of the other $u_{i'}$. To that end, we couple the movement of agents of M-VISIT-EXCHANGE to $l-1$ independent instances of process LUCKY-GAMBLER$(m+1, 1/(b+1), k)$, one corresponding to each of the vertices u_i for $1 \leqslant i < l$.

Next we give some definitions and describe the coupling for a fixed i. For simplicity, define $B_i = B_{u_i,m}$ and $B'_i = B_i \setminus B_{u_{i-1},m-1} = \bigcup_{v \in C_{u_i} \setminus \{u_{i-1}\}} B_{v,m-1}$. I.e., B_i is the set of descendants of u_i at distance m from it, and to get B'_i we remove the descendants of u_{i-1} from B_i. Let $g_1, \ldots, g_{z_{i,t}}$ be the agents in M-VISIT-EXCHANGE that were at B'_i in round $t-1$ and moved closer to the root in the next round. By definition of M-VISIT-EXCHANGE, there are at least $k = \lceil \alpha \cdot b^m/8 \rceil$ such agents.

In the LUCKY-GAMBLER$(m+1, 1/(b+1), k)$ process that corresponds to vertex u_i, we start k gamblers in round t, denoted $g'_1, \ldots, g'_k$. For each $1 \leqslant j \leqslant k$, and for each round $t' \geqslant t$ until g_j reaches u_{i+1} or any vertex in B_i, the walks g_j and g'_j are coupled: if g_j moves closer to the root then g'_j moves to the right on the path, and if g_j moves away from the root, g'_j moves left. If g_j is at u_{i+1} or in $B_{u_i,m}$, then by the coupling, g'_j has finished its walk at one of the endpoints of the path. Before this happens we say that g_j is *i-coupled*.

Let $t_1 = 4 \cdot (m+1)$, and let t_{i+1} be the first round after t_i when u_{i+1} receives an i-coupled agent from u_i. Now, notice that by construction no agent can be i-coupled and i'-coupled at the same time for $i' \neq i$. It implies that the rounds when u_{i+1} receives i-coupled agents are independent from the walks of i'-coupled agents. On the other hand the walks of i-coupled agents are coupled with an independent LUCKY-GAMBLER process thus, Lemma 6.2.1 implies

$$\mathbb{P}\left[t_{i+1} - t_i \leqslant s \mid t_1, \ldots, t_i\right] = (1-\beta)^s = \mathbb{P}\left[F_i \geqslant s\right],$$

where $F_i \sim Geom(\beta)$, $1 \leqslant i < l$, are a collection of independent geometric random variables with success probability β. If τ_ρ is the round when the root is informed then $\tau_\rho \leqslant t_l = t_1 + \sum_{i=1}^{l-1}(t_{i+1} - t_i)$. It follows that $(\tau_\rho - t_1)$ is stochastically dominated by $F = \sum_{i=1}^{l-1} F_i$, and from a Chernoff bound for the sum of independent geometric random variables (Lemma 2.3.5),

$$\mathbb{P}\left[\tau_\rho \geqslant f + t_1\right] \leqslant \mathbb{P}\left[F \geqslant f\right] \leqslant e^{-f \cdot \beta/8},$$

for any $f \geqslant 2h/\beta$. Since $t_1 = O(h)$, we can take a large enough $f = O(\log n)$, completing the proof. □

Next we prove that if vertex u of height $h_u = m$ is informed, then after at most $O(m \ln n)$ rounds a given leaf v in u's subtree becomes informed, w.h.p. For that, we first show that there are at least $\Theta(m \ln n)$ visits to u in those rounds (possibly multiple times by the same agent). Using a lower bound on the probability that an agent that is at u visits v before returning to u, we can show that one of these agents will visit v in $O(m \ln n)$ rounds, w.h.p.

Lemma 6.4.2. *Let u be such that $h_u = m$. For any constant $c > 0$, there is a round $\tau = O(m \ln n)$ such that in the first τ rounds of* M-VISIT-EXCHANGE, *u is visited at least $c \cdot mb \cdot \ln n$ times, with probability at least $1 - n^{-cmb}$.*

Proof. For a round t, let $g_1, \ldots, g_{z_{u,t}}$ be the agents that are in $B_{u,m-1}$ at round t, and have also been at the leaf vertices $B_{u,m}$ in the previous round. By the definition of M-VISIT-EXCHANGE, $z_{u,t} \geqslant k$, where $k = \lceil \alpha b^m/8 \rceil$. We construct an instance of LUCKY-GAMBLER$(m, 1/(b+1), k)$ as follows. If $g'_1, \ldots, g'_k$ are the gamblers that started their walk at round t, then for each $1 \leqslant j \leqslant k$, the walk of agent g_j is coupled with the walk of the gambler g'_k: If g_j moves closer to the root of the tree, then g'_j moves right on the path and left otherwise. The coupling ends when g'_j arrives at either vertex 0 or m of its path. That corresponds to g_j either visiting a leaf vertex in $B_{u,m}$ or visiting vertex u.

Consider the first τ rounds of M-VISIT-EXCHANGE. Since $k \geqslant \alpha b^m/8$, we can apply Lemma 6.2.2 with parameter $\kappa = \alpha b/8$ to the coupled LUCKY-GAMBLER process. Let γ be the constant guaranteed by the lemma and let $\tau = \frac{8c}{\alpha\gamma} \cdot m \ln n$. Lemma 6.2.2 implies that in the first τ rounds of LUCKY-GAMBLER there are at least $\gamma\kappa\tau = c \cdot mb \cdot \ln n$ lucky gamblers, with probability at least $1 - e^{-\gamma\kappa\tau/4} = 1 - e^{-cmb \ln n} = 1 - n^{-cmb}$. Since each lucky gambler corresponds to a single visit to u by some agent, we complete the proof. □

Lemma 6.4.3. *Let u be such that $h_u = m$ and let v be a leaf in the subtree of u. For any constant $c_l > 0$, if vertex u is informed then after at most $O(m \ln n)$ rounds of* M-VISIT-EXCHANGE, *vertex v is informed with probability at least $1 - n^{-c_l}$.*

Proof. Let τ be the round guaranteed by Lemma 6.4.2 for a constant $c > 0$. If after the first τ rounds of M-VISIT-EXCHANGE, there have been fewer than $cmb \ln n$ visits to u, then we add a minimal number of agents to u at round τ to have at least $cmb \ln n$ agents there. We call the resulting process M-VISIT-EXCHANGE$_u$. By Lemma 6.4.2 and an application of union bound over the first $\log^2 n = \omega(m \ln n)$ rounds, M-VISIT-EXCHANGE$_u$ and M-VISIT-EXCHANGE are identical in the first $\Theta(m \ln n)$ rounds of execution with probability at least $1 - n^{-cmb} \log^2 n$. We therefore analyse M-VISIT-EXCHANGE$_u$.

For a round $t \leqslant \tau$, consider an agent g that visits u at round t. Let $\mathcal{D}_{g,t}$ be the event that g moves to one of u's children at round $t+1$. Let also $\mathcal{E}_{g,t}$ be the event that g visits

v before returning to u, and before round $\tau' = \tau + 8mb^{m-1}$. Clearly, $\mathcal{E}_{g,t}$ implies $\mathcal{D}_{g,t}$, and $\mathbb{P}[\mathcal{D}_{g,t}] = \frac{b}{b+1}$. Also, we can show that $\mathbb{P}[\mathcal{E}_{g,t} \mid \mathcal{D}_{g,t}] \geqslant 1/(12mb)$, by analysing a single random walk in $R_{b,m}$ that starts in the root of the tree (Lemma 6.6.2). Therefore,

$$\mathbb{P}[\mathcal{E}_{g,t}] = \mathbb{P}[\mathcal{E}_{g,t} \cap \mathcal{D}_{g,t}] = \mathbb{P}[\mathcal{D}_{g,t}] \cdot \mathbb{P}[\mathcal{E}_{g,t} \mid \mathcal{D}_{g,t}] \geqslant \frac{b}{b+1} \cdot \frac{1}{12mb} \geqslant \frac{1}{18mb}.$$

The probability that v is not visited by any informed agent before round τ' is at most

$$\mathbb{P}\Bigg[\bigcap_{t \leqslant \tau,\, g \in Z_u(t)} \neg\mathcal{E}_{g,t}\Bigg] \leqslant \left(1 - \frac{1}{18mb}\right)^{cmb\ln n} \leqslant e^{-c\ln n/18} \leqslant n^{-c/18} \leqslant n^{-c_l-1},$$

for a large enough constant c. Notice that $\tau' = \tau + 8mb^{m-1} = O(m\ln n)$ by the definition of m. Since M-VISIT-EXCHANGE and M-VISIT-EXCHANGE$_u$ are identical in the first $\log^2 n$ rounds with probability at least $1 - n^{-cmb}\log^2 n$, v will be informed in $O(m\ln n)$ rounds in M-VISIT-EXCHANGE, with probability at least $1 - n^{-c_l-1} - n^{-cmb}\log^2 n \geqslant 1 - n^{-c_l}$. □

Proof of the upper bound of Theorem 6.1.1. We will use the following simple symmetry lemma, which holds for any graph (Lemma 2.5.1): If $T_{u,v}$ is the number of rounds of VISIT-EXCHANGE until vertex v is informed when the information originates at u, then the random variables $T_{u,v}$ and $T_{v,u}$ have the same distribution.

Consider the M-VISIT-EXCHANGE process, and suppose that the source of the information is vertex u with $h_u = m$, for m as defined at the beginning of Section 6.3. By Lemma 6.4.1, for an arbitrary constant c, there is $T_1 = O(\log n)$ such that the root ρ is informed by time T_1, with probability at least $1 - n^{-c}$. Lemma 6.3.1 then implies that the same bound T_1 holds for the VISIT-EXCHANGE process, with probability $p \geqslant 1 - n^{-c} - n^{-\alpha\mu/32}$, for an arbitrary large μ. From the symmetry lemma above, it follows that if ρ is the initial source of the information instead, then u becomes informed within T_1 rounds of VISIT-EXCHANGE with the same probability $p \geqslant 1 - n^{-c} - n^{-\alpha\mu/32}$.

Suppose again that information originates at some u with $h_u = m$, and let v be any leaf that is a descendant of u. From Lemma 6.4.3 and Lemma 6.3.1, for an arbitrary constant c, there is some $T_2 = O(m\log n)$, such that v gets informed after at most T_2 rounds of VISIT-EXCHANGE, with probability at least $1 - n^{-c} - n^{-\alpha\mu/32}$.

Combining the above we obtain that if ρ is the source of the information, then any given leaf v is informed after at most $T_1 + T_2$ rounds of VISIT-EXCHANGE, with probability at least $1 - 2n^{-c} - 2n^{-\alpha\mu/32}$. And by a union bound, all leaves (and thus all vertices) are informed within $T_1 + T_2$ rounds with probability at least $1 - 2n^{-c+1} - 2n^{-\alpha\mu/32+1}$.

Finally, by employing the symmetry argument above again, we obtain that for any source vertex (not just ρ), all vertices are informed within $2(T_1 + T_2)$ rounds with probability at least $1 - 4n^{-c+1} - 4n^{-\alpha\mu/32+1}$. Since $T_1 + T_2 = O(\log n + m\log n) = O(\log n + \log_b \log n \cdot \log n) = O(\log n + h\log h)$, the theorem follows. □

6.5 Proof of Theorem 6.1.1: the lower bound

We prove the lower bound part of Theorem 6.1.1, i.e., the spreading time of VISIT-EXCHANGE on a binary tree is $\Omega(\log n \cdot \log\log n)$, w.h.p.

Proof of the lower bound of Theorem 6.1.1. We show that there is a leaf vertex that is not visited by any agent in the first $\tau = c\ln n \cdot \ln\ln n$ rounds of VISIT-EXCHANGE, w.h.p., where c is a small enough constant, to be determined later. For convenience, we assume that τ is even. For a fixed leaf vertex v and an agent g, let $N_g(v)$ be the number of times g visits v in the rounds $0, \ldots, \tau - 1$ of its walk. Since g starts from a stationary distribution,

$$\mathbb{E}\left[N_g(v)\right] = \pi(v) \cdot \tau = \tau/(2(n-1)).$$

Let τ_v be the first time when g visits v. Then,

$$\begin{aligned}\mathbb{E}\left[N_g(v) \mid \tau_v < \tau\right] &\geqslant \mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] \cdot \mathbb{P}\left[\tau_v < \tau/2 \mid \tau_v < \tau\right] \\ &\geqslant \mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] \cdot (1/2),\end{aligned}$$

where the second inequality holds because $\mathbb{P}\left[\tau_v < \tau\right] \leqslant 2\mathbb{P}\left[\tau_v < \tau/2\right]$, as g is equally likely to visit v in the intervals $0, \ldots, \tau/2 - 1$ and $\tau/2, \ldots, \tau - 1$ for g starts its walk from stationarity. From Lemma 6.6.4, if $X_g(t)$ denotes the position of the random walk of g at t,

$$\begin{aligned}\mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] \geqslant \mathbb{E}\left[N_g(v) \mid \tau_v = \tau/2 - 1\right] &\geqslant \sum_{t=\tau/2}^{\tau-1} \mathbb{P}\left[X_g(t) = v \mid \tau_v = \tau/2 - 1\right] \\ &\geqslant \sum_{t=\tau/2}^{\tau-1} \frac{1}{32 \cdot (t - (\tau/2 - 1))} \geqslant \frac{\ln(\tau/2)}{32} \geqslant \frac{\ln \ln n}{32},\end{aligned}$$

for n sufficiently large. It follows

$$\mathbb{P}\left[\tau_v < \tau\right] = \mathbb{P}\left[N_g(v) \geqslant 1\right] = \frac{\mathbb{E}\left[N_g(v)\right]}{\mathbb{E}\left[N_g(v) \mid N_g(v) \geqslant 1\right]} \leqslant \frac{\tau/(2(n-1))}{(1/2) \cdot (\ln \ln n/32)} = \frac{32c \cdot \ln n}{n-1}.$$

Returning to the case of n agents, for the leaf v,

$$\mathbb{P}\left[v \text{ not visited by any agent}\right] \geqslant \left(1 - \frac{32c \cdot \ln n}{n-1}\right)^n \geqslant \frac{1}{2} \cdot e^{-32c \cdot \ln n} = \frac{1}{2} \cdot n^{-32c}.$$

Thus, if X is the number of leaves that are *not* visited by any agent, then $\mathbb{E}\left[X\right] \geqslant n^{1-32c}/4$, as there are at least $n/2$ leaves. We can now use the method of bounded differences [DP09, Sec. 5.4] to give a lower bound on the probability that at least one vertex is not visited by any agent. If L_g is the set of leaves that g visits, then X can be written as a function of independent variables L_g. Notice that changing L_g can change X by at most τ. Thus,

$$\mathbb{P}\left[X \leqslant \frac{\mathbb{E}\left[X\right]}{2}\right] \leqslant \exp\left(-\frac{\mathbb{E}\left[X\right]^2}{2 \cdot n \cdot \tau^2}\right) \leqslant \exp\left(-\frac{n^{1-64c}}{64 \cdot \tau^2}\right).$$

By taking $c < 1/64$ we get that, w.h.p., there is at least one leaf that has not been visited by any agent. □

6.6 Auxiliary lemmas for random walks on trees

Consider a random walk on a b-ary tree $R_{b,h}$ of height h, with $n = (b^{h+1} - 1)/(b - 1)$ vertices. Denote the number of rounds the walk takes to reach from vertex u to v as $\tau_{u,v}$. The return time τ_u^+ is the number of rounds it takes for the walk starting from u to return to u.

Lemma 6.6.1. *For a balanced tree $R_{b,h}$ with any $b \geqslant 2$ and $h \geqslant 1$, let ρ be its root and v any leaf vertex. Then,* **(a)** $\mathbb{E}\left[\tau_\rho^+\right] \geqslant 2b^{h-1}$, **(b)** $\mathbb{E}\left[\tau_{v,\rho}\right] \geqslant b^{h-1}$, **(c)** $\mathbb{E}\left[\tau_{v,\rho}\right] \leqslant 8b^{h-1}$, **(d)** $\mathbb{E}\left[\tau_{\rho,v}\right] \leqslant 4hb^h$.

Proof. **(a)** Let π be the stationary distribution of $R_{b,h}$, which has $|E| = n-1 = (b^{h+1} - b)/(b-1)$ edges. Then, by [LP17, Proposition 1.19],

$$\mathbb{E}\left[\tau_\rho^+\right] = \frac{1}{\pi(\rho)} = \frac{\deg(\rho)}{2|E|} = \frac{2(b^h-1)}{b-1} \geqslant 2b^{h-1}.$$

(b–c) We follow the derivation for a binary tree in [LP17, Example 10.17]. For a vertex u, we denote by $s(u)$ the number edges in the subtree of u. Consider the path $\langle v = v_0, \ldots, v_h = \rho\rangle$. Then $s(v_i) = \sum_{j=1}^{i} b^j = (b^{i+1} - b)/(b-1)$. By [LP17, (10.24)]

$$\mathbb{E}\left[\tau_{v_{i-1},v_i}\right] = 2(s(v_{i-1}) + 1) - 1 = 2s(v_{i-1}) + 1.$$

Therefore,

$$\begin{aligned}
\mathbb{E}\left[\tau_{v_0,\rho}\right] &= \sum_{i=1}^{h} \mathbb{E}\left[\tau_{v_{i-1},v_i}\right] = h + \frac{2b}{b-1} \cdot \sum_{i=1}^{h} (b^{i-1} - 1) = 2b \cdot \frac{b^h - 1}{(b-1)^2} - h \cdot \frac{b+1}{b-1} \\
&= b^{h-1} \cdot \frac{2b^2}{(b-1)^2} - \frac{2b}{(b-1)^2} - h \cdot \frac{b+1}{b-1} \qquad (6.1)\\
&\geqslant b^{h-1} \cdot \frac{2b^2}{(b-1)^2} - 4 - 3h, \quad \text{since } b \geqslant 2, \\
&\geqslant b^{h-1}.
\end{aligned}$$

To see why the last inequality holds, notice that it is true for $h = 1$ and that it becomes stronger as h increases. This completes the proof of part (b). Part (c) follows from (6.1) and the fact that $2b^2/(b-1)^2 \leqslant 8$ since $b \geqslant 2$.

(d) We use the same notation and technique as in the previous parts. We have

$$\mathbb{E}\left[\tau_{v_i,v_{i-1}}\right] = 2(|E| - s(v_{i-1})) - 1 = \frac{2b}{b-1} \cdot (b^h - b^{i-1}) - 1,$$

which implies that

$$\mathbb{E}\left[\tau_{\rho,v_0}\right] = \sum_{i=h}^{1} \mathbb{E}\left[\tau_{v_i,v_{i-1}}\right] = \frac{2b}{b-1} \cdot \left(hb^h - \frac{b^h - 1}{b-1}\right) - h \leqslant 4hb^h,$$

since $2b/(b-1) \leqslant 4$. □

Lemma 6.6.2. *For the root ρ of tree $R_{b,h}$ and any leaf vertex v,*

$$\mathbb{P}\left[\tau_{\rho,v} < \min\{\tau_\rho^+, 8hb^{h-1}\}\right] \geqslant \frac{1}{12hb}.$$

Proof. Let $\mathcal{E}$ denote the event $\{\tau_{\rho,v} < \tau_\rho^+\}$, that a walk starting from ρ hits v before returning to ρ. We have

$$\mathbb{E}\left[\tau_{\rho,v}\right] \geqslant \mathbb{E}\left[\tau_{\rho,v} \mid \neg\mathcal{E}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right] = \left(\mathbb{E}\left[\tau_\rho^+ \mid \neg\mathcal{E}\right] + \mathbb{E}\left[\tau_{\rho,v}\right]\right) \cdot \mathbb{P}\left[\neg\mathcal{E}\right].$$

Also

$$\mathbb{E}\left[\tau_\rho^+ \mid \neg\mathcal{E}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right] = \mathbb{E}\left[\tau_\rho^+\right] - \mathbb{E}\left[\tau_\rho^+ \mid \mathcal{E}\right] \cdot \mathbb{P}\left[\mathcal{E}\right],$$

and

$$\mathbb{E}\left[\tau_\rho^+ \mid \mathcal{E}\right] = \mathbb{E}\left[\tau_{\rho,v} \mid \mathcal{E}\right] + \mathbb{E}\left[\tau_{v,\rho}\right] \leqslant \mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right].$$

Combining these three inequalities we obtain

$$\mathbb{E}\left[\tau_{\rho,v}\right] \geqslant \mathbb{E}\left[\tau_\rho^+\right] - \left(\mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right]\right) \cdot \mathbb{P}\left[\mathcal{E}\right] + \mathbb{E}\left[\tau_{\rho,v}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right].$$

Substituting $\mathbb{P}\left[\neg\mathcal{E}\right] = 1 - \mathbb{P}\left[\mathcal{E}\right]$, solving for $\mathbb{P}\left[\mathcal{E}\right]$, and using Lemma 6.6.1, yields

$$\mathbb{P}\left[\mathcal{E}\right] \geqslant \frac{\mathbb{E}\left[\tau_\rho^+\right]}{2\mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right]} \geqslant \frac{2b^{h-1}}{8hb^h + 8b^{h-1}} = \frac{1}{4(hb+1)} \geqslant \frac{1}{6hb}.$$

Next, we bound $\mathbb{P}\left[\tau_{\rho,v} < 4hb^{h-1} \mid \mathcal{E}\right]$. Let u be the child of ρ that is also an ancestor of v. Then, given $\mathcal{E}$, until the walk returns to ρ, the walk is a restricted to the subtree of u, which is a b-ary tree of height $h-1$. In particular, in the first step, the walk visits u. Therefore, by Lemma 6.6.1(d), we have that $\mathbb{E}\left[\tau_{\rho,v} - 1 \mid \mathcal{E}\right] \leqslant 4(h-1)b^{h-1} < 4hb^{h-1}$. Then, by Markov's inequality,

$$\mathbb{P}\left[\tau_{\rho,v} < 8hb^{h-1} \mid \mathcal{E}\right] \geqslant 1/2.$$

Finally,

$$\begin{aligned}\mathbb{P}\left[\tau_{\rho,v} < \min\{\tau_\rho^+, 8hb^{h-1}\}\right] &= \mathbb{P}\left[\mathcal{E} \cap \{\tau_{\rho,v} < 8hb^{h-1}\}\right] \\ &= \mathbb{P}\left[\tau_{\rho,v} < 8hb^{h-1} \mid \mathcal{E}\right] \cdot \mathbb{P}\left[\mathcal{E}\right] \geqslant \frac{1}{12hb}.\end{aligned}$$

□

Lemma 6.6.3. *Let v be a leaf of $R_{b,h}$, and u be its ancestor of height $x \geqslant 1$. For any $\epsilon > 0$,*

$$\mathbb{P}\left[\tau_{v,u} \geqslant \epsilon \cdot b^{x-1}\right] \geqslant 1 - \epsilon.$$

Proof. For brevity denote $k' = \epsilon \cdot b^{x-1}$ and $p = \mathbb{P}\left[\tau_{v,u} \geqslant k'\right]$. Then, for an integer $i \geqslant 1$,

$$\begin{aligned}\mathbb{P}\left[\tau_{v,u} \geqslant i \cdot k'\right] &= \mathbb{P}\left[\tau_{v,u} \geqslant i \cdot k' \mid \tau_{v,u} \geqslant (i-1) \cdot k'\right] \cdot \mathbb{P}\left[\tau_{v,u} \geqslant (i-1) \cdot k'\right] \\ &\leqslant p \cdot \mathbb{P}\left[\tau_{v,u} \geqslant (i-1) \cdot k'\right], \\ &\qquad \text{by Markov property and because } \mathbb{P}\left[\tau_{v',u} \geqslant k'\right] \leqslant p \text{ for any } v', \\ &\leqslant p^i, \ \text{ by iterating the argument.}\end{aligned}$$

This implies that

$$\mathbb{E}\left[\tau_{v,u}\right] = \sum_{t \geqslant 1} \mathbb{P}\left[\tau_{v,u} \geqslant t\right] \leqslant \sum_{i \geqslant 0} k' \cdot \mathbb{P}\left[\tau_{v,u} \geqslant i \cdot k'\right] \leqslant \frac{k'}{1-p} = \frac{\epsilon \cdot b^{x-1}}{1-p}.$$

By Lemma 6.6.1(b), $\mathbb{E}\left[\tau_{v,u}\right] \geqslant b^{x-1}$. Combining the last two inequalities, we get $p \geqslant 1 - \epsilon$. □

Lemma 6.6.4. *Let $X(t)$ be the location of a simple random walk that starts at a leaf v of $R_{b,h}$ at round 0. Then, for any even integer $t > 0$,*

$$\mathbb{P}\left[X(t) = v\right] \geqslant \max\left(\frac{1}{16bt}, \frac{1}{2n}\right).$$

Proof. For even rounds t, $\mathbb{P}[X(t) = v]$ monotonically decreases towards the stationary distribution at v. Thus,

$$\mathbb{P}[X(t) = v] \geqslant \frac{1}{2(n-1)} \geqslant \frac{1}{2n}. \tag{6.2}$$

It implies that we have to show the inequality in the case when the first term under the max is larger, i.e., when $t \leqslant n/(4b)$. Let $x = 1 + \lceil \log_b(2t) \rceil$. First, we prove that v has an ancestor of height x, i.e., that $x \leqslant h$.

$$\begin{aligned} x &= 1 + \lceil \log_b(2t) \rceil \\ &\leqslant 1 + \lceil \log_b(n/(2b)) \rceil \\ &\leqslant 1 + \lceil \log_b(2 \cdot b^h/(2b) \rceil, \text{ because } n \leqslant 2 \cdot b^h, \\ &\leqslant 1 + \lceil \log_b(b^{h-1}) \rceil \\ &= h. \end{aligned}$$

Thus, we can define u as the ancestor of v of height x. Let $\mathcal{R} = \{\tau_{v,u} > t\}$ be the event that the random walk $X(t)$ does not visit u in the first t rounds. Since $t \leqslant b^{x-1}/2$ by construction, then

$$\mathbb{P}[\mathcal{R}] \geqslant \mathbb{P}\left[\tau_{v,u} \geqslant b^{x-1}/2\right] \geqslant \frac{1}{2},$$

by Lemma 6.6.3. Then,

$$\mathbb{P}[X(t) = v] \geqslant \mathbb{P}[X(t) = v \mid \mathcal{R}] \cdot \mathbb{P}[\mathcal{R}] \geqslant \frac{1}{2} \cdot \mathbb{P}[X(t) = v \mid \mathcal{R}].$$

Let u' be the child of u that is also an ancestor of v. Let k' be the number of vertices in the subtree of u'. Given the event $\mathcal{R}$, the random walk can only visit vertices in the subtree of u', thus, for an even t, as in (6.2),

$$\mathbb{P}[X(t) = v \mid \mathcal{R}] \geqslant \frac{1}{2k'} = \frac{1}{2(1 + b + \cdots + b^{x-1})} \geqslant \frac{1}{4 \cdot b^{x-1}} \geqslant \frac{1}{16bt},$$

where the last inequality holds because $t > b^{x-2}/2$ by the choice of x. Combining with the previous inequality we finish the proof. □

Chapter 7

Bounds for grid graphs

7.1 Introduction

In this chapter we give a tight analysis of the broadcast time of VISIT-EXCHANGE on k-dimensional grid graphs and torus graphs where k is constant. For integers k and n, where $n^{1/k}$ is an integer, let $G_{k,n}$ and $\widehat{G}_{k,n}$ be k-dimensional grid and torus graphs with n vertices, respectively. Note that $G_{1,n}$ and $\widehat{G}_{1,n}$ are simply path and cycle graphs, respectively.

Theorem 7.1.1. *If $G = G_{k,n}$ or $G = \widehat{G}_{k,n}$ for a constant number of dimensions $k \geqslant 1$, then for any source vertex, $T_{\text{visitx}}(G) = O(\text{diam}(G))$, w.h.p.*

Each vertex of a k-dimensional grid graph has a degree between k and $2k$ (i.e., the grid is almost regular). A torus graph is $2k$-regular. The relatively simple analysis of [Fei+90] implies that the broadcast time of randomised rumour spreading for both $G = G_{k,n}$ and $G = \widehat{G}_{k,n}$ graphs is at most $O(\text{diam}(G))$, w.h.p. Since torus graphs are regular, Theorem 4.1.1 gives us a weaker bound of $T_{\text{visitx}} = O(\text{diam}(\widehat{G}_{k,n}) \cdot \text{poly}(\log\log n))$, w.h.p. The same bound also holds for grid graphs $G_{k,n}$ by adapting the proof of Theorem 4.1.1 for almost-regular graphs, which is possible since the key technical lemma needed for the theorem (Lemma 4.2.2) also holds for non-regular graphs. To remove the additional $\log\log n$ factors, a fine-grained analysis is required that fully exploits the structure of the graph.

To motivate our analysis technique, which is closely related to that of [KS03; KS05], we consider a path graph and assume that the source vertex is at one end. At any round, the agents are located according to the stationary distribution, therefore, w.h.p., there is a sub-path of logarithmic length without any agents. When the most recently informed vertex belongs to such a sub-path, the progress of informing new vertices is delayed, hence we have to argue that such situations are rare. We will show that for the majority of the rounds up to $T^* = O(\text{diam}(G))$, there is an agent at most constant steps away from the most recently informed vertex. With constant probability this agent will visit the informed vertices and move closer to the endpoint of the path that is not informed, making progress.

We tessellate the space-time into square blocks of constant side length Δ_1, both in space and time. Given an execution of VISIT-EXCHANGE up to a certain round, a *good* block is one which is likely to be densely populated by agents after Δ_1 rounds. When the most recently informed vertex is in a good block, then in the subsequent $\Delta_1 = \Theta(1)$ rounds information will be transmitted to a new vertex along the path, with constant probability. Therefore, we aim to show that throughout the runtime of the process, a constant fraction of the blocks that contain the most recently informed vertex are good.

For the blocks that are close to one another (in time or space), there are dependencies between the goodness of blocks, hence the above argument cannot be made directly. To tackle the dependencies we build a hierarchy of R tessellations of the space-time into square blocks of increasing sizes Δ_r, for a scale parameter $r \in \{1, \ldots, R\}$. The coarsest tessellations uses blocks of size $\Delta_R = O(\text{poly} \log n)$, and it is easy to argue that they are all good, w.h.p. Then we show that with a sufficiently high probability a good block of size Δ_r does not contain any bad blocks in a finer tessellation that uses blocks of size Δ_{r-1}. Furthermore, if two finer blocks are contained in larger blocks that are sufficiently far apart, then the finer ones satisfy the property of being good independently from one another. This allows us to recursively bound the number of bad blocks in each tessellation, starting from the coarsest one, concluding that at most a constant fraction of all blocks in the finest tessellation are bad. We note that various aspects of our proof are simpler that in the original proof of Kesten and Sidoravicius, mainly because our process stores the information at vertices, resulting in information paths that are easier to analyse.

The key ideas of the proof are the same for grids and tori of dimension $k > 1$, but there is the following difference. In path graphs, if the source is at one endpoint of the path, we simply argue that information progresses towards the other endpoint, which is equivalent to completing a broadcast. In higher dimensions, however, there is no "direction of progress." Instead, we fix a target vertex and prove that it becomes informed in $O(\text{diam}(G))$ rounds, w.h.p., by showing that the distance from the target vertex to the most recently informed vertex decreases. Then, we take a union bound over all vertices and complete the proof.

The above multi-scale argument works for the "central" area of the path that only contains vertices at least $\Theta(\log n)$ away from the endpoints of the path. For grids, we use well-known cover-time arguments to prove that the remaining vertices become informed quickly [LP17]. For tori, as there is no "central" area, we can choose the central vertices based on the target vertex, therefore, there is no need to analyse the edge cases separately.

Lastly, in Appendix A we have evaluated VISIT-EXCHANGE on 1-dimensional grid graphs, i.e., paths. From Fig. A.1 we can observe that VISIT-EXCHANGE and PUSH have very close broadcast times. This also indicates that despite the fact that we are generous with constants in the theoretical analysis, in reality the constant appears to be very close to 2.

7.2 Notation and definitions

First we define grid and torus graphs formally. Let k and n be two integers such that $\ell = n^{1/k} - 1 \geqslant 1$ is also an integer. Let V be the set of k-dimensional vectors $\mathbf{x} = (x_1, \ldots, x_k) \in \mathbb{Z}^k$ such that $x_j \in \{0, \ldots, \ell\}$ for $1 \leqslant j \leqslant k$. The grid graph $G_{k,n} = (V, E)$ has vertex set V and edge set E, where $(\mathbf{x}, \mathbf{y}) \in E$ if $|x_{j'} - y_{j'}| = 1$ for some j', and $x_j = y_j$ for the remaining $k-1$ coordinates j. In other words, $(\mathbf{x}, \mathbf{y})$ is an edge if $\mathbf{x}$ and $\mathbf{y}$ differ by 1 in one coordinate and are the same in the remaining $k-1$ coordinates. The torus graph $\widehat{G}_{k,n} = (V, \widetilde{E})$ also has the same vertex set V, but has some additional edges. Precisely, $(\mathbf{x}, \mathbf{y}) \in \widetilde{E}$ if $|x_{j'} - y_{j'}| \in \{1, \ell\}$ for some j', and $x_j = y_j$ for the remaining $k-1$ values of j.

We consider the case when k is a constant. Thus, for $G = G_{k,n}$ or $G = \widehat{G}_{k,n}$, $\text{diam}(G) = \Theta(\ell) = \Theta(n^{1/k})$. For $\mathbf{x} = (x_1, ..., x_k) \in \mathbb{Z}^k$ and an integer z, we denote $\mathbf{x} + z = (x_1 + z, \ldots, x_k + z)$. Also, for $\mathbf{y} = (y_1, ..., y_k) \in \mathbb{Z}^k$, we write $\mathbf{x} < \mathbf{y}$ to denote that

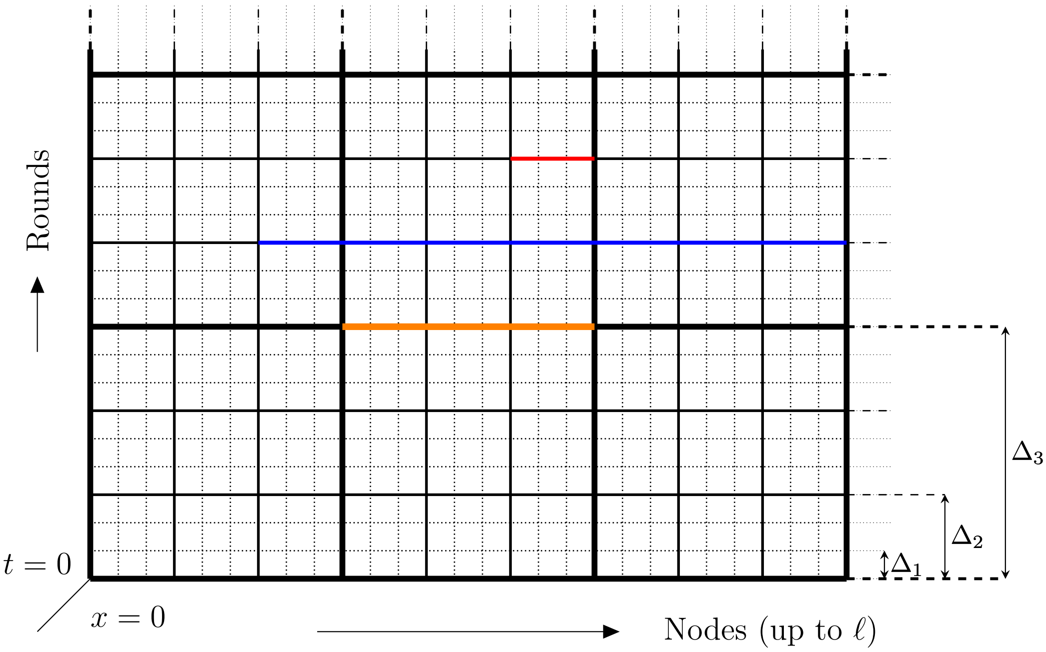

Figure 7.1 An illustration of 3 consecutive tessellations of space-time for $k = 1$, i.e., the path graph. Each tessellation is 3 times finer than the previous one. The set $\widetilde{S}_2(5,5)$ is shown in red, its base $\widetilde{B}_2(5,4)$ in blue and its parent $\widetilde{S}_3(1,1)$ in orange.

$x_i < y_i$, for all $1 \leqslant i \leqslant k$; and similarly define $\mathbf{x} \leqslant \mathbf{y}$.

For convenience we will define $\alpha = |A|/n = \Theta(1)$, where A is the set of agents in VISIT-EXCHANGE.

We emphasise that in this chapter, as in the rest of this thesis, $\mathrm{dist}(\mathbf{x}, \mathbf{y})$ refers to the distance between vertices $\mathbf{x}$ and $\mathbf{y}$ in the underlying graph. In other words, $\mathrm{dist}(\mathbf{x}, \mathbf{y})$ is the Manhattan distance between $\mathbf{x}$ and $\mathbf{y}$, not the Euclidean distance.

7.2.1 Space-time tessellation

Let $C > 0$ be a constant even integer constant to be defined later. We prove our result using a multi-scale argument, with a scaling parameter $r \in \{1, 2, \ldots, R\}$, for some $R = \Theta(\log \log n)$ that depends on C and will be defined precisely later in Lemma 7.3.4. For each r, we define $\Delta_r = C^{4kr}$. During our analysis, for each scale r, we only consider rounds $s \cdot \Delta_r$ for an integer $s \geqslant 0$. For a vector $\mathbf{i} \in \mathbb{Z}^k$ and an integer s, we define the following sets in space and space-time, respectively:

$$S_r(\mathbf{i}) = \{\mathbf{x} \in \mathbb{Z}^k \mid \mathbf{i} \cdot \Delta_r \leqslant \mathbf{x} < (\mathbf{i} + \mathbf{1}) \cdot \Delta_r\}, \quad \widetilde{S}_r(\mathbf{i}, s) = S_r(\mathbf{i}) \times \{s \cdot \Delta_r\}.$$

The collection of sets $\{S_r(\mathbf{i})\}_{\mathbf{i} \in \mathbb{Z}^k}$ partitions $\mathbb{Z}^k$. Additionally, we define extended versions of these sets:

$$B_r(\mathbf{i}) = \{\mathbf{x} \in \mathbb{Z}^k \mid (\mathbf{i} - \mathbf{3}) \cdot \Delta_r \leqslant \mathbf{x} < (\mathbf{i} + \mathbf{4}) \cdot \Delta_r\}, \quad \widetilde{B}_r(\mathbf{i}, s) = B_r(\mathbf{i}) \times \{s \cdot \Delta_r\}.$$

The second definitions, in space-time, are used as a shorthand, so instead of saying "agents in $S_r(\mathbf{i})$ at round $s \cdot \Delta_r$," we can say "agents at $\widetilde{S}_r(\mathbf{i}, s)$."

We call $\widetilde{B}_r(\mathbf{i}, s-1)$ the *base* of $\widetilde{S}_r(\mathbf{i}, s)$. If an agent is at $\widetilde{S}_r(\mathbf{i}, s)$, then it must have also been at its base. The *parent* of $\widetilde{S}_r(\mathbf{i}, s)$ is the set $\widetilde{S}_{r+1}(\mathbf{j}, l)$ corresponding to the unique pair $(\mathbf{j}, l)$ such that $S_r(\mathbf{i}) \subset S_{r+1}(\mathbf{j})$ and $s \cdot \Delta_r \in [l \cdot \Delta_{r+1}, (l+1) \cdot \Delta_{r+1})$. Correspondingly, $\widetilde{S}_r(\mathbf{i}, s)$ is one of the *children* of $\widetilde{S}_{r+1}(\mathbf{j}, l)$. Fig. 7.1 illustrates space-time tessellation for the one-dimensional case.[1]

Let $V' = \prod_{j=1}^{k}\{6\Delta_R, \ldots, \ell - 6\Delta_R\}$ be the "central" part of V. We will only consider sets $\widetilde{S}_r(\mathbf{i}, s)$ for which it holds $\widetilde{S}_r(\mathbf{i}, s) \subset V' \times [\Delta_R, T^* - \Delta_R]$, where $T^* = O(\ell)$. If P is the set of pairs $(\mathbf{i}, s)$ that satisfy the last relation, then for any $(\mathbf{i}, s) \in P$, we have $\widetilde{B}_r(\mathbf{i}, s) \subset V \times [0, T^*]$, and moreover if $r < R$, then $\widetilde{S}_r(\mathbf{i}, s)$ has a parent $\widetilde{S}_{r+1}(\mathbf{j}, l)$ with $(\mathbf{j}, l) \in P$. Note that the definition of a central part is arbitrary for the torus graphs.

7.2.2 Good and bad sets

Let $\gamma_R = \alpha/2$. For $1 \leqslant r < R$, define $\gamma_r = \gamma_{r+1} \cdot (1 - C^{-(r+1)/8})$. Then,

$$\begin{aligned}
\gamma_1 &= \gamma_R \cdot \prod_{j=1}^{R-1}(1 - C^{-(j+1)/8}) \\
&\geqslant \gamma_R \cdot \left(1 - \sum_{j=1}^{R-1} C^{-(j+1)/8}\right), \text{ by Weierstrass' inequality,} \\
&\geqslant \gamma_R \cdot \left(1 - \frac{C^{-1/4}}{1 - C^{-1/8}}\right) \\
&\geqslant \gamma_R/2,
\end{aligned}$$

for $C \geqslant 256$. Since $\gamma_r \geqslant \gamma_{r-1}$ for any $r \geqslant 2$, we have that for any r, $\gamma_r \in [\alpha/4, \alpha/2]$.

For any set of vertices S, let $N(S, t)$ be the number of agents in S at round t, and for $\mathbf{x} \in V$, we write $N(\mathbf{x}, t) = N(\{\mathbf{x}\}, t)$. For a space-time set $\widetilde{S}$, $N(\widetilde{S}) = \sum_{(\mathbf{x},t)\in\widetilde{S}} N(\mathbf{x}, t)$. Next, for any $\mathbf{x} \in \mathbb{Z}^k$ and integer s define

$$Q_r(\mathbf{x}) = \{\mathbf{y} \in \mathbb{Z}^k \mid \mathbf{x} \leqslant \mathbf{y} < \mathbf{x} + C^r\}, \quad \widetilde{Q}_r(\mathbf{x}, s) = Q_r(\mathbf{x}) \times \{s \cdot \Delta_r\}.$$

We say that $\widetilde{S}_r(\mathbf{i}, s)$ is *good*, if for every $\mathbf{x}$ such that $Q_r(\mathbf{x}) \subset B_r(\mathbf{i})$,

$$N(\widetilde{Q}_r(\mathbf{x}, s-1)) \geqslant \gamma_r \cdot |Q_r(\mathbf{x})| = \gamma_r \cdot C^{kr}.$$

Otherwise, $\widetilde{S}_r(\mathbf{i}, s)$ is *bad*. Informally, we have that for a good set $\widetilde{S}_r(\mathbf{i}, s)$, its base $\widetilde{B}_r(\mathbf{i}, s-1)$ contains agents in a way that all blocks of size C^{kr} in it are sufficiently densely populated. Since any agent visiting the set $\widetilde{S}_{r-1}(\mathbf{j}, l)$ must have been at the base of its parent $\widetilde{B}_r(\mathbf{i}, s-1)$, we can argue that all children of a good set are good, w.h.p.

7.3 The multi-scale analysis

The goal of this section is to prove that if two vertices are central, i.e., are in V', and one of them is informed, then after at most $T^* = O(\ell)$ rounds the other one will also become informed, w.h.p.

Consider a fastest path via which information progresses to the target vertex. If a set $\widetilde{S}_1(\mathbf{i}, s)$ of scale $r = 1$ through which this path passes is good, then with constant

[1] The ratio Δ_{r+1}/Δ_r must be even in our proof. In the figure it is 3 solely for illustrative purposes.

probability the path moves closer to the target while it is in the set. Thus, it suffices to prove that at least a constant fraction of the sets of scale $r = 1$ that the path intersects are good (we assume that in the bad sets, the path stays put). However, a priori, we do not know the path VISIT-EXCHANGE will take to deliver the information to the target vertex. Instead, we are able to prove that the desired property holds for *all* possible information paths, with high probability. This is done by a recursive argument starting from $r = R$ down to $r = 1$. At the scale $r = R$, it is easy to see that every set $\widetilde{S}_r(\mathbf{i}, s)$ is good. The key argument gives upper bounds the probability that a good set of scale $r \geqslant 2$ has a bad child of scale $r - 1$. Using this, we can bound the number of bad sets that any possible information path passes through at each scale. The main claim then follows.

7.3.1 Probability that a good set has a bad child

In this section we consider a set $\widetilde{S}_{r+1}(\mathbf{i}, s)$ for a scale $r \geqslant 1$ and an integer $s \geqslant 1$. Our goal is to show that if a set $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is a good set, then all its children are also good with high enough probability. To achieve that, we first show that in expectation there are sufficiently many agents in each set $Q_r(\mathbf{y}, t')$ that is contained in a base of some child of $\widetilde{S}_{r+1}(\mathbf{i}, s)$. Then, by the independence of the walks and an application of a Chernoff bound, we can lower bound the desired probability. This result holds given any execution of the walks until round $(s-1)\Delta_{r+1}$, denoted by a σ-field $\mathcal{K}_{r+1}(s-1)$, which allows us to apply it for a number of (space- and time-separated) sets at once.

Lemma 7.3.1. *There is a constant $C_1 > 0$, such that for any $C \geqslant C_1$ and $1 \leqslant r < R$, if $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is good, then for any even integer $u \in [\Delta_{r+1} - \Delta_r, 2\Delta_{r+1})$ and any vertex $\mathbf{y}$ with $Q_r(\mathbf{y}) \subset \{\mathbf{x} \mid \mathbf{i} \cdot \Delta_{r+1} - 3 \cdot \Delta_r \leqslant \mathbf{x} < (\mathbf{i} + \mathbf{1}) \cdot \Delta_{r+1} + 3 \cdot \Delta_r\}$,*

$$\mathbb{E}\left[N(Q_r(\mathbf{y}), t + u)\right] \geqslant \gamma_{r+1} \cdot C^{kr} \cdot \left(1 - C^{-(r+1)/2}\right),$$

where $t = (s-1)\Delta_{r+1}$.

Proof. Notice that only agents that are at $\widetilde{B}_{r+1}(\mathbf{i}, s-1)$ can be at $Q_r(\mathbf{y})$ at round $t + u$. For $\mathbf{j} \in \mathbb{Z}^k$, define $\mathbf{x}_\mathbf{j} = \mathbf{y} + \mathbf{j} \cdot C^{r+1}$. Construct a partition of $\mathbb{Z}^k$ into a grid of blocks $M_\mathbf{j}$ which have corners at vertices $\mathbf{x}_\mathbf{j}$:

$$M_\mathbf{j} = \{\mathbf{x} \mid \mathbf{x}_\mathbf{j} \leqslant \mathbf{x} < \mathbf{x}_{\mathbf{j}+1}\}.$$

Each set $M_\mathbf{j}$ contains $C^{k(r+1)}$ vertices. Notice that $\mathbf{y}$ is also a corner for some of the blocks ($2k$ of them). We will only consider the set J of indices $\mathbf{j}$ such that $M_\mathbf{j} \subset B_{r+1}(\mathbf{i})$. Since $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is good, $N(M_\mathbf{j}, t) \geqslant \gamma_{r+1} C^{k(r+1)}$ by definition.

Let $W(u)$ be the position of a random walk in $\mathbb{Z}^k$ at round u, assuming $W(0) = \mathbf{0}$. For a vertex $\mathbf{x} \in V$, let $W_\mathbf{x}(u)$ be the position at round u of a random walk in G that starts at $\mathbf{x}$. Note that since we only consider $\mathbf{x} \in Q_r(\mathbf{y})$, the vertex $\mathbf{x}$ is sufficiently far from the edges of G, hence $W_\mathbf{x}(u)$ has the same distribution as $W(u) + \mathbf{x}$. We have that,

$$\begin{aligned}\mathbb{E}\left[N(Q_r(\mathbf{y}), t + u)\right] &\geqslant \sum_{\mathbf{j} \in J} \sum_{\mathbf{x} \in M_\mathbf{j}} N(\mathbf{x}, t) \cdot \mathbb{P}\left[W_\mathbf{x}(u) \in Q_r(\mathbf{y})\right] \\ &\geqslant \sum_{\mathbf{j} \in J} \sum_{\mathbf{x} \in M_\mathbf{j}} N(\mathbf{x}, t) \cdot \min_{\mathbf{x}' \in M_\mathbf{j}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right]\end{aligned}$$

$$= \sum_{\mathbf{j}\in J} N(M_{\mathbf{j}}, t) \cdot \min_{\mathbf{x}'\in M_{\mathbf{j}}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right]$$
$$\geqslant \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}\in J} \min_{\mathbf{x}'\in M_{\mathbf{j}}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right].$$

The probability that $W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})$ is minimised when $\mathbf{x}'$ is the farthest possible vertex from $Q_r(\mathbf{y})$ in $M_{\mathbf{j}}$ due to Lemma 7.5.2. Thus, it will be minimised at one of the corners of $M_{\mathbf{j}}$, suppose $\mathbf{x}'_{\mathbf{j}}$. First, notice that $\mathbf{x}'_{\mathbf{j}}$ cannot share a coordinate with $\mathbf{y}$ because, otherwise, we could change that coordinate and get farther from $Q_r(\mathbf{y})$. Additionally, for $\mathbf{j_1} \neq \mathbf{j_2}$, $\mathbf{x}'_{\mathbf{j}_1} \neq \mathbf{x}'_{\mathbf{j}_2}$. Thus, the collection of $\mathbf{x}'_{\mathbf{j}}$ is precisely the set of corners $\mathbf{x}_{\mathbf{j}'}$ which do not share a coordinate with $\mathbf{y}$, i.e., when $\mathbf{j}'$ does not have a 0 coordinate. So we define

$$J' = \{\mathbf{j}' \in J \mid \mathbf{j}' \text{ does not have coordinate } 0\},$$

and, continuing our bound,

$$\begin{aligned}
\mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] &\geqslant \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}\in J} \mathbb{P}\left[W_{\mathbf{x}'_{\mathbf{j}}}(u) \in Q_r(\mathbf{y})\right] \\
&= \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}'\in J'} \mathbb{P}\left[W_{\mathbf{x}_{\mathbf{j}'}}(u) \in Q_r(\mathbf{y})\right] \\
&\geqslant \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}'\in J'} \mathbb{P}\left[W(u) + \mathbf{x}_{\mathbf{j}'} \in Q_r(\mathbf{y})\right] \\
&= \gamma_{r+1} \cdot C^{k(r+1)} \cdot \Bigg(\sum_{\mathbf{j}\in J} \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0})\right] - \\
&\qquad \sum_{\mathbf{j}\in J\setminus J'} \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0})\right] \Bigg). \qquad (7.1)
\end{aligned}$$

Next, we bound the sums above separately. Let $W'(u)$ be a random walk on a k-dimensional torus $\widehat{H} = \widehat{G}_{k,C^{k(r+1)}}$, with vertex set $V(\widehat{H}) = \{0, \ldots, C^{r+1}-1\}^k$. We assume that W' has started vertex $\mathbf{0}$ in round 0. Then,

$$\begin{aligned}
\sum_{\mathbf{j}\in J} \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0})\right] &= \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0}) \text{ for some } \mathbf{j} \in J\right] \\
&= \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0}) \text{ for some } \mathbf{j} \in \mathbb{Z}^k\right] \\
&= \mathbb{P}\left[W'(u) \in Q_r(\mathbf{0})\right],
\end{aligned}$$

where the second equality holds because in $u \leqslant 2\Delta_{r+1}$ steps, the walk $W(u)$ cannot reach a vertex $\mathbf{z} - \mathbf{j}\cdot C^{r+1}$ for $\mathbf{z} \in Q_r(\mathbf{0})$ and $j \notin J$. We have that $W'(0) \in Q_r(\mathbf{0})$, so as $u' \geqslant 0$ increases, $\mathbb{P}\left[W'(2u') \in Q_r(\mathbf{0})\right]$ decreases monotonically to its stationary value. Thus, $\mathbb{P}\left[W'(u) \in Q_r(\mathbf{0})\right] \geqslant |Q_r(\mathbf{0})/|V(H_k)| = C^{-k}$, which bounds the first sum in (7.1).

For the second sum in (7.1), we consider the cases when each component l is 0 separately. For $l \in \{1, \ldots, k\}$, let $J_l = \{(j_1, \ldots, j_k) \in J \mid j_l = 0\}$. For an integer $h \geqslant 0$, let $L_l^h = \{\mathbf{x} \in Q_r(\mathbf{0}) \mid \mathbf{x}_l = h\}$, which partition $Q_r(\mathbf{0})$ into disjoint sets of size $C^{(k-1)r}$. For $\mathbf{j} \in J_l$ the probability that a walk starting at $\mathbf{j}\cdot C^{r+1}$ is in L_l^h is greatest for $h = 0$ by Lemma 7.5.2. Therefore,

$$\sum_{\mathbf{j}\in J\setminus J'} \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0})\right] \leqslant \sum_{l=1}^{k} \sum_{\mathbf{j}\in J_l} \mathbb{P}\left[W(u) + \mathbf{j}\cdot C^{r+1} \in Q_r(\mathbf{0})\right]$$

$$
\begin{aligned}
&= \sum_{l=1}^{k} \sum_{\mathbf{j} \in J_l} \sum_{h=0}^{C^r-1} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in L_l^h\right] \\
&\leqslant C^r \cdot \sum_{l=1}^{k} \sum_{\mathbf{j} \in J_l} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in L_l^0\right] \\
&= C^r \cdot \sum_{l=1}^{k} P_l, \quad \text{where } P_l \text{ is the internal sum above.}
\end{aligned}
$$

Consider a fixed $l \in \{1, \ldots, k\}$. For any $\mathbf{j} \in J_l$ and $\mathbf{x} \in L_l^0$, the $\mathbf{x} - \mathbf{j} \cdot C^{r+1}$ are all unique vectors, and have 0 as their lth coordinate. Therefore, if $x(u)$ is $W(u)$'s lth coordinate, then $P_l \leqslant \mathbb{P}[x(u) = 0]$. Notice, however, that $x(u)$ is a lazy random walk on $\mathbb{Z}$, starting at 0, with holding probability $1 - 1/k$. Therefore, by Lemma 7.5.1 there is a constant η, such that

$$
P_l \leqslant \mathbb{P}[x(u) = 0] \leqslant \frac{\eta}{\sqrt{u}} \leqslant \frac{\eta}{\sqrt{\Delta_{r+1} - \Delta_r}} \leqslant \frac{\eta}{\sqrt{\Delta_{r+1}/2}} \leqslant \eta \cdot \sqrt{2} C^{-2k(r+1)}.
$$

We substitute these bounds in (7.1):

$$
\begin{aligned}
\mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] &\geqslant \gamma_{r+1} \cdot C^{k(r+1)} \cdot \left(C^{-k} - \sqrt{2}\eta k \cdot C^r \cdot C^{-2k(r+1)}\right) \\
&\geqslant \gamma_{r+1} \cdot C^{kr}\left(1 - \sqrt{2}\eta k \cdot C^{-(r+1)}\right) \geqslant \gamma_{r+1} \cdot C^{kr}\left(1 - C^{-(r+1)/2}\right),
\end{aligned}
$$

for $C \geqslant C_1 = \sqrt{2}\eta k$. □

Lemma 7.3.2. *There is a constant $C_2 > 0$, such that if $C \geqslant C_2$, then given $\mathcal{K}_{r+1}(s-1)$ and the event that $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is good, the probability that all the children of $\widetilde{S}_{r+1}(\mathbf{i}, s)$ are good is at least $1 - \rho_{r+1}$, where*

$$
\rho_{r+1} = C^{8k^2(r+1)} \cdot \exp\left(-\frac{\gamma_{r+1}}{4} \cdot C^{k(r+1)/4}\right). \tag{7.2}
$$

Proof. For convenience denote $t = (s-1) \cdot \Delta_{r+1}$. Suppose the child $\widetilde{S}_r(\mathbf{j}, m)$ of $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is bad for some $\mathbf{j}$ and m. Then there is a set $Q_r(\mathbf{y}) \subset B_r(\mathbf{j})$ such that $N(\widetilde{Q}_r(\mathbf{y}, m-1)) < \gamma_r \cdot C^{kr}$. We fix such $\mathbf{y}$ and m, then bound the probability of $N(\widetilde{Q}_r(\mathbf{y}, m-1)) < \gamma_r \cdot C^{kr}$, and take a union bound over all such pairs.

All agents at $\widetilde{Q}_r(\mathbf{y}, m-1)$ must have been at $\widetilde{B}_{r+1}(\mathbf{i}, s-1)$ before. For an agent g in $\widetilde{B}_{r+1}(\mathbf{i}, s-1)$, let X_g be the indicator random variable that g is at $\widetilde{Q}_r(\mathbf{y}, m-1)$. Note that g has to travel for $u = (m-1) \cdot \Delta_r - t$ rounds to be at $\widetilde{Q}_r(\mathbf{y}, m-1)$. By definition of being a child, $u \in [\Delta_{r+1} - \Delta_r, 2\Delta_{r+1})$. Then

$$
N(\widetilde{Q}_r(\mathbf{y}, m-1)) = N(Q_r(\mathbf{y}), t+u) = \sum_{g \text{ at } \widetilde{B}_{r+1}(\mathbf{i}, s-1)} X_g.
$$

We take $C_2 = \max\{6, C_1\}$, where C_1 is determined in Lemma 7.3.1. Since the agents move independently after round t, we can apply Chernoff bound:

$$
\mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \gamma_r \cdot C^{kr}\right] = \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \gamma_{r+1} \cdot C^{kr} \cdot (1 - C^{-(r+1)/8})\right]
$$

$$
\begin{aligned}
&\leqslant \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] \cdot \frac{1-C^{-(r+1)/8}}{1-C^{-(r+1)/2}}\right] \\
&= \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] \cdot \left(1 - \frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1-C^{-(r+1)/2}}\right)\right] \\
&\leqslant \exp\left(-\frac{1}{2} \cdot \left(\frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1-C^{-(r+1)/2}}\right)^2 \cdot \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right]\right) \\
&\leqslant \exp\left(-\frac{1}{2} \cdot \left(\frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1-C^{-(r+1)/2}}\right)^2 \cdot \gamma_{r+1} \cdot C^{kr} \cdot \left(1 - C^{-(r+1)/2}\right)\right) \\
&= \exp\left(-\frac{\gamma_{r+1} \cdot C^{kr}}{2} \cdot C^{-(r+1)/4} \cdot \frac{\left(1-C^{-3(r+1)/8}\right)^2}{1-C^{-(r+1)/2}}\right) \\
&\leqslant \exp\left(-\frac{\gamma_{r+1}}{4} \cdot C^{k(r+1)/4}\right),
\end{aligned}
$$

where it is easy to verify that the last inequality holds since $C \geqslant 6$. The number of pairs $(\mathbf{y}, m)$, which may contain less than $\gamma_r \cdot C^{kr}$ agents and render the child $\widetilde{S}_r(\mathbf{j}, m)$ of $\widetilde{S}_{r+1}(\mathbf{i}, s)$ bad, is at most $(3 \cdot \Delta_{r+1})^k \cdot (\Delta_{r+1}/\Delta_r) = 3^k \cdot C^{4k^2(r+1)+4k} \leqslant C^{8k^2(r+1)}$. Thus, the proof is complete after an application of a union bound. $\square$

7.3.2 Bound on the number of bad sets

An *information path* $\overline{\mathbf{x}}$ is defined as a sequence $\mathbf{x}_t$ of vertices in V', such that for any $t \geqslant 0$, either $\mathbf{x}_{t+1} = \mathbf{x}_t$ or $\operatorname{dist}(\mathbf{x}_{t+1}, \mathbf{x}_t) = 1$. Let Θ be the set of all information paths of length exactly $T^* = c\ell$, for a constant $c > 0$. We say that an information path $\overline{\mathbf{x}} \in \Theta$ intersects set $\widetilde{S}_r(\mathbf{i}, s)$ if $\mathbf{x}_{s \cdot \Delta_r} \in S_r(\mathbf{i})$. Let $\phi_r(\overline{\mathbf{x}})$ be the number of bad sets $\widetilde{S}_r(\mathbf{i}, s)$ that $\overline{\mathbf{x}}$ intersects, and

$$\Phi_r = \max_{\overline{\mathbf{x}} \in \Theta} \phi_r(\overline{\mathbf{x}}).$$

For $r > 1$, we also define $\psi_r(\overline{\mathbf{x}})$ as the number of good sets $\widetilde{S}_r(\mathbf{i}, s)$ that have a bad child and intersect $\overline{\mathbf{x}}$. We define

$$\Psi_r = \max_{\overline{\mathbf{x}} \in \Theta} \psi_r(\overline{\mathbf{x}}).$$

In this section, we prove an upper bound on the maximum number of bad sets at scale r that intersect an information path in Θ. In particular, our final lemma, bounding Φ_1, argues that at least a constant fraction of sets $\widetilde{S}_1(\mathbf{i}, s)$ that intersect any given information path are good, w.h.p. This allows us to argue, roughly, that if we split time into phases of $\Delta_1 = O(1)$, then in a constant fraction of those phases, progress is made toward a target vertex with constant probability.

Lemma 7.3.3. *For any scale $r \in \{1, \dots, R-1\}$,*

$$\Phi_r \leqslant (\Phi_{r+1} + \Psi_{r+1}) \cdot (\Delta_{r+1}/\Delta_r) = (\Phi_{r+1} + \Psi_{r+1}) \cdot C^{4k}.$$

Proof. If $\overline{\mathbf{x}} \in \Theta$ intersects the set $\widetilde{S}_{r+1}(\mathbf{i}, s)$, then it can also intersect at most Δ_{r+1}/Δ_r of its children, since it can only intersect a child at rounds $s \cdot \Delta_{r+1} + i \cdot \Delta_r$ for an integer $0 \leqslant i < \Delta_{r+1}/\Delta_r = C^{4k}$. If $\widetilde{S}_{r+1}(\mathbf{i}, s)$ is either bad, or has a bad child, we assume that all its children that $\overline{\mathbf{x}}$ intersects are bad. This gives us an upper bound:

$$\phi_r(\overline{\mathbf{x}}) \leqslant \phi_{r+1}(\overline{\mathbf{x}}) \cdot C^{4k} + \psi_{r+1}(\overline{\mathbf{x}}) \cdot C^{4k}.$$

The proof is completed by taking a maximum on both sides of the inequality with respect to all paths $\overline{\mathbf{x}} \in \Theta$. $\square$

Lemma 7.3.4. *For any constants $c, C > 0$ and $\kappa > 0$, there is a value $R = \Theta(\log\log n)$ such that $\mathbb{P}[\Phi_R = 0] \geqslant 1 - n^{-\kappa}$.*

Proof. Let $R = \lceil \log_C(\eta \ln n)/k \rceil$ for a constant $\eta > 0$. Consider some space-time set $\widetilde{Q}_R(\mathbf{x}, s) \subset V \times [0, T^*)$. By definition $|Q_R(\mathbf{x})| = C^{kR} \geqslant \eta \cdot \ln n$. By a Chernoff bound,

$$\mathbb{P}\left[N(\widetilde{Q}_R(\mathbf{x}, s) \geqslant \frac{\alpha}{2} \cdot |Q_R(\mathbf{x})|\right] \geqslant 1 - e^{-\frac{\alpha}{8} \cdot |Q_R(\mathbf{x})|} \geqslant 1 - n^{-\frac{\alpha \cdot \eta}{8}}.$$

If $\Phi_R > 0$, then $N(\widetilde{Q}_R(\mathbf{x}, s)) < \gamma_R \cdot C^{kR} = \alpha \cdot C^{kR}/2$ for some vertex $\mathbf{x}$ and integer s. Since the number of such sets $\widetilde{Q}_R(\mathbf{x}, s)$ is at most $(T^*/\Delta_R) \cdot \ell^k < c \cdot \ell^{k+1}$, by a union bound,

$$\mathbb{P}[\Phi_R = 0] \geqslant 1 - c \cdot \ell^{k+1} \cdot n^{-\frac{\alpha \cdot \eta}{8}} \geqslant 1 - n^{-\kappa},$$

for a constant η large enough. $\square$

Lemma 7.3.5. *There is a constant $C_3 > 0$, such that for any $C \geqslant C_3$, if $r \geqslant 2$, then*

$$\mathbb{P}\left[\Psi_r \geqslant e^{-r} \cdot \frac{T^*}{\Delta_r}\right] \leqslant e^{-T^*/\Delta_r}.$$

Proof. Recall that V' is the set of central vertices of V. Let P_r denote the set of pairs $(\mathbf{i}, s)$ that satisfy $\widetilde{S}_r(\mathbf{i}, s) \subset V' \times [\Delta_R, T^* - \Delta_R]$. We partition P_r into $m = 2 \cdot 7^k$ disjoint sets, which are defined using an integer $v \in \{0, 1\}$ and a vector $\mathbf{h} \in H = \{0, \ldots, 6\}^k$, as follows:

$$P_r(\mathbf{h}, v) = \{(\mathbf{i}, s) \in P_r \mid \mathbf{i} \equiv \mathbf{h} \pmod 7,\ s \equiv v \pmod 2\}.$$

Let $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ be the number of good sets $\widetilde{S}_r(\mathbf{i}, s)$, for $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$, that have a bad child and are intersected by $\overline{\mathbf{x}} \in \Theta$. Let also

$$\Psi_r(\mathbf{h}, v) = \max_{\overline{\mathbf{x}} \in \Theta} \psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}).$$

To prove the lemma we first bound $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ for a fixed pair $\mathbf{h}, v$ and path $\overline{\mathbf{x}}$. We then use union bound twice: first to bound $\Psi_r(\mathbf{h}, v)$, and then to bound Ψ_r, as the latter is the sum of all $\Psi_r(\mathbf{h}, v)$.

Consider a fixed pair $\mathbf{h}, v$. For $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$, define $Y(\mathbf{i}, s)$ as the indicator random variable that is 1 if $\widetilde{S}_r(\mathbf{i}, s)$ is good, but has a bad child. From Lemma 7.3.2,

$$\begin{aligned}\mathbb{P}[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1)] = \mathbb{P}&\left[\widetilde{S}_r(\mathbf{i}, s) \text{ has a bad child} \mid \mathcal{K}_r(\mathbf{i}, s-1); \widetilde{S}_r(\mathbf{i}, s) \text{ is good}\right] \cdot \\ &\mathbb{P}\left[\widetilde{S}_r(\mathbf{i}, s) \text{ is good} \mid \mathcal{K}_r(\mathbf{i}, s-1)\right] \\ \leqslant \rho_r.&\end{aligned}$$

Consider the following ordering of elements of $P_r(\mathbf{h}, v)$: $(\mathbf{i}', s') \prec (\mathbf{i}, s)$ if $s' < s$, or $s' = s$ and $\mathbf{i}'$ is lexicographically smaller than $\mathbf{i}$ (this decision is arbitrary). Due to the space and time separation of the sets $\widetilde{S}_r(\mathbf{i}, s)$ for $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$,

$$\mathbb{P}[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1); Y(\mathbf{i}', s') \text{ for all } (\mathbf{i}', s') \prec (\mathbf{i}, s)]$$

$$= \mathbb{P}\left[Y(\mathbf{i},s)=1 \mid \mathcal{K}_r(\mathbf{i},s-1)\right] \leqslant \rho_r.$$

For pairs $(\mathbf{i},s) \in P_r(\mathbf{h},v)$, let $\{Z(\mathbf{i},s)\}$ be a collection of independent Bernoulli random variables with success probability ρ_r. From the above it follows that

$$\mathbb{P}\left[Y(\mathbf{i},s)=1 \mid \mathcal{K}_r(\mathbf{i},s-1);\ Y(\mathbf{i}',s') \text{ for all } (\mathbf{i}',s') \prec (\mathbf{i},s)\right] \leqslant \mathbb{P}\left[Z(\mathbf{i},s)=1\right]. \tag{7.3}$$

Notice that to bound $\Psi_r(\mathbf{h},v)$ it is wasteful to take a union bound over all paths in Θ, because many information paths intersect exactly the same collection of sets $\widetilde{S}_r(\mathbf{i},s)$, for $(\mathbf{i},s) \in P_r(\mathbf{h},v)$. Thus, we can group them into such equivalence classes based on sets they intersect, reducing the number of objects we need to take a union bound over. For an information path $\overline{\mathbf{x}}$, define

$$I_{r,\mathbf{h},v}(\overline{\mathbf{x}}) = \{(\mathbf{i},s) \in P_r(\mathbf{h},v) \mid \overline{\mathbf{x}} \text{ intersects } \widetilde{S}_r(\mathbf{i},s)\}.$$

Then,

$$\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) = \sum_{(\mathbf{i},s)\in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})} Y(\mathbf{i},s) \leqslant |I_{r,\mathbf{h},v}|.$$

Next, we bound the probability that $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) \geqslant \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}$, where $m = 2 \cdot 7^k$ (recall, m is the number of sets $P_r(\mathbf{h},v)$ to which P_r is partitioned). Let $Z = \sum_{(\mathbf{i},s)\in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})} Z(\mathbf{i},s)$, and $b = \frac{2e^{-r}}{m \cdot \rho_r}$. Then,

$$b \cdot \mathbb{E}\left[Z\right] = b \cdot |I_{r,\mathbf{h},v}(\overline{\mathbf{x}})| \cdot \rho_r \leqslant b \cdot \frac{T^*}{2\Delta_r} \cdot \rho_r = \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}. \tag{7.4}$$

Let C_2 be as in Lemma 7.3.2, and let C_3' be the smallest constant such that for any $C \geqslant C_3'$,

$$\frac{\alpha}{32} \cdot C^{kr/4} \geqslant r(1+8k^2 \ln C), \quad \text{and} \quad C^{kr/8} \cdot \frac{e^{-r}}{m} \cdot \frac{\alpha}{64} \geqslant 1, \text{ for any } r \geqslant 1.$$

For $C \geqslant \max\{256, C_2, C_3'\}$, by substituting for ρ_r from (7.2) and using $\gamma_r \geqslant \alpha/4$,

$$\ln b \geqslant \ln \frac{2}{m} - r(1+8k^2 \ln C) + \frac{\gamma_r}{4} \cdot C^{kr/4} \geqslant \frac{\alpha}{32} \cdot C^{kr/4}. \tag{7.5}$$

By Lemma 2.3.9 and (7.3), $Z = \sum_{(\mathbf{i},s)\in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})} Z(\mathbf{i},s)$ stochastically dominates $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$, so

$$\begin{aligned}
\mathbb{P}\left[\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) \geqslant \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}\right] &\leqslant \mathbb{P}\left[Z \geqslant \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}\right] \\
&\leqslant \left(\frac{b}{e}\right)^{-\frac{e^{-r}}{m}\cdot\frac{T^*}{\Delta_r}}, \quad \text{by Lemma 2.3.6 and (7.4)}, \\
&\leqslant \exp\left(-\frac{T^*}{\Delta_r} \cdot \frac{e^{-r}}{m} \cdot \left(\frac{\alpha}{32} \cdot C^{kr/4} - 1\right)\right), \quad \text{by (7.5)}, \\
&\leqslant \exp\left(-\frac{T^*}{\Delta_r} \cdot \frac{e^{-r}}{m} \cdot \frac{\alpha}{64} \cdot C^{kr/4}\right) \\
&\leqslant \exp\left(-\frac{T^*}{\Delta_r} \cdot C^{kr/8}\right),
\end{aligned}$$

where the last two inequalities hold since $C \geqslant C_3'$.

Next we upper bound the number of distinct values that the set $I_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ takes, for all $\overline{\mathbf{x}} \in \Theta$, i.e., the cardinality of set $\bigcup_{\overline{\mathbf{x}}\in\Theta}\{I_{r,\mathbf{h},v}(\overline{\mathbf{x}})\}$. It suffices to bound instead the cardinality of $I_r = \bigcup_{\overline{\mathbf{x}}\in\Theta}\{I_r(\overline{\mathbf{x}})\}$, where $I_r(\overline{\mathbf{x}}) = \bigcup_{\mathbf{h}\in H, v\in\{0,1\}} I_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ is the set of all $\widetilde{S}_r(\mathbf{i}, s)$ that $\overline{\mathbf{x}}$ intersects. We do that by looking at how many possible sets $\widetilde{S}_r(\mathbf{i}, s)$ can be in $I_r(\overline{\mathbf{x}})$ for each $s \in \{1, \ldots, \lfloor T^*/\Delta_r \rfloor\}$, given the previous elements in $I_r(\overline{\mathbf{x}})$: For $s = 1$, there are at most $(\ell/\Delta_r)^k$ possible choices. For $s \geqslant 2$, if $\overline{\mathbf{x}}$ intersects both $\widetilde{S}_r(\mathbf{i}, s-1)$ and $\widetilde{S}_r(\mathbf{j}, s)$, then $\mathbf{i}$ and $\mathbf{j}$ differ by at most 1 in each coordinate. Therefore, given the first elements in $I_r(\overline{\mathbf{x}})$ up to s, there are at most 3^k possible choices of the next element. Therefore,

$$|I_r| \leqslant (\ell/\Delta_r)^k \cdot (3^k)^{T^*/\Delta_r} \leqslant \exp\left(k\ln\ell + k\ln 3 \cdot \frac{T^*}{\Delta_r}\right) \leqslant \exp\left(2k \cdot \frac{T^*}{\Delta_r}\right).$$

Using a union bound we get that,

$$\mathbb{P}\left[\Psi_r(\mathbf{h}, v) \geqslant \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}\right] \leqslant \exp\left(2k \cdot \frac{T^*}{\Delta_r}\right) \cdot \exp\left(-\frac{T^*}{\Delta_r} \cdot C^{kr/8}\right) \leqslant e^{-T^*/\Delta_r}/m,$$

for $C \geqslant C_3'' = (2k + \ln m + 1)^8$. By another union bound over all m values of pair $\mathbf{h}, v$, we prove the desired result for $C_3 = \max\{C_2, C_3', C_3''\}$. □

Lemma 7.3.6. *For any constant $\kappa > 0$, there are constants c and C, such that*

$$\mathbb{P}\left[\Phi_1 \leqslant \frac{T^*}{4\Delta_1}\right] \geqslant 1 - 2n^{-\kappa}.$$

Proof. Let C be the smallest even integer that is at least C_3 (defined in Lemma 7.3.5). Let $\mathcal{E}_1$ be the event that $\Psi_{r+1} < e^{-(r+1)} \cdot \frac{T^*}{\Delta_{r+1}}$ for all $r \in \{1, \ldots, R-1\}$ at the same time, and $\mathcal{E}_2$ be the event that $\Phi_R = 0$. If $\mathcal{E}_1$ and $\mathcal{E}_2$ hold, then, by a recursive application of Lemma 7.3.3,

$$\begin{aligned}
\Phi_1 &\leqslant \Phi_R \cdot C^{4k(R-1)} + \sum_{r=1}^{R-1} \Psi_{r+1} \cdot C^{4kr} \\
&\leqslant T^* \cdot \sum_{r=1}^{R-1} \frac{e^{-(r+1)}}{\Delta_{r+1}} \cdot C^{4kr} \\
&= \frac{T^*}{C^{4k}} \cdot \sum_{r=1}^{R-1} e^{-(r+1)} \\
&\leqslant \frac{T^*}{e^2(1-1/e)\cdot\Delta_1} \\
&\leqslant \frac{T^*}{4\Delta_1}.
\end{aligned}$$

By a union bound and Lemma 7.3.5,

$$\mathbb{P}[\mathcal{E}_1] \geqslant 1 - \sum_{r=1}^{R-1} \mathbb{P}\left[\Psi_{r+1} \geqslant e^{-(r+1)} \cdot \frac{T^*}{\Delta_{r+1}}\right] \geqslant 1 - \sum_{r=1}^{R-1} e^{-T^*/\Delta_{r+1}} \geqslant 1 - R \cdot e^{-T^*/\Delta_R} \geqslant 1 - n^{-\kappa},$$

where the last inequality holds for $c > 0$ large enough and the corresponding value of R as determined by Lemma 7.3.4. Since $\mathbb{P}[\mathcal{E}_2] \geqslant 1 - n^{-\kappa}$, we have that $\mathbb{P}\left[\Phi_1 \leqslant \frac{T^*}{4\Delta_1}\right] \geqslant 1 - 2n^{-\kappa}$. □

7.3.3 Dissemination away from the boundary

Recall that $V' = \prod_{j=1}^{s}\{6\Delta_R, \ldots, \ell - 6\Delta_R\}$ is the "central" part of V. We first prove, using the bound on Φ_1, that if a vertex in V' is informed, then in at most $O(\ell)$ rounds, any other fixed vertex of V' becomes informed, w.h.p. Later, using standard random walk techniques, we prove that if the source vertex is an "edge" vertex in $V \setminus V'$, then some vertex in V' becomes informed in $O(\ell)$ rounds. We combine these facts to prove Theorem 7.1.1.

Lemma 7.3.7. *Let $\mathbf{x}, \mathbf{y} \in V'$. For any constant $\kappa > 0$, there is a large enough constant $c > 0$ such that if $\mathbf{x}$ is informed, then after at most $T^* = c\ell$ rounds, $\mathbf{y}$ also becomes informed, with probability at least $1 - 3n^{-\kappa}$.*

Proof. Fix any shortest path $\langle \mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_\lambda\rangle$ between $\mathbf{x}$ and $\mathbf{y}$, where $\lambda = \mathrm{dist}(\mathbf{x}, \mathbf{y})$. We consider a process $\hat{\mathbf{x}}(t)$, $t \geqslant 0$, on the vertices of that path, such that $\hat{\mathbf{x}}(0) = \mathbf{x}_0 = \mathbf{x}$, and for $t \geqslant 1$, $\hat{\mathbf{x}}(t)$ is defined as follows: if $\hat{\mathbf{x}}(t-1) = \mathbf{x}_i$ for some $i < \lambda$, and some agent moves from $\mathbf{x}_i$ to $\mathbf{x}_{i+1}$ in round t, then $\hat{\mathbf{x}}(t) = \mathbf{x}_{i+1}$; otherwise, $\hat{\mathbf{x}}(t) = \hat{\mathbf{x}}(t-1)$. Let $\tau = \min\{t \mid \hat{\mathbf{x}}(t) = \mathbf{y}\}$, which is an upper bound on the number of rounds until $\mathbf{y}$ is informed. It suffices to show that $\mathbb{P}[\tau \leqslant T^*] \geqslant 1 - 3n^{-\kappa}$.

Since $\mathbf{x}, \mathbf{y} \in V'$, clearly $\mathbf{x}_i \in V'$, for all $0 \leqslant i \leqslant \lambda$, and thus $\hat{\mathbf{x}}(t) \in V'$ for any $t \geqslant 0$. Then, for any integer $s \geqslant 0$, and for $t_s = s\Delta_1$, there is some index $\mathbf{i}_s$ such that $\hat{\mathbf{x}}(t_s) \in S_1(\mathbf{i}_s)$.

Suppose that the space-time set $\widetilde{S}_1(\mathbf{i}_s, s+1)$ is good. Then by definition of goodness, every set $Q_1(\mathbf{z})$ that is a subset of $B_1(\mathbf{i}_s)$ contains at least $\gamma_1 \cdot C^k$ agents at round $t_s = s\Delta_1$. Also, by construction $\hat{\mathbf{x}}(t_s) \in B_1(\mathbf{i}_s)$, and therefore, it is also the case that $\hat{\mathbf{x}}(t_s) \in Q_1(\mathbf{z}) \subset B_1(\mathbf{i}_s)$, for some vertex $\mathbf{z}$. Since $Q_1(\mathbf{z})$ is a k-dimensional finite grid with a side of length C, it follows that there is some agent g at a vertex $\mathbf{w}$ with $\mathrm{dist}(\hat{\mathbf{x}}(t_s), \mathbf{w}) \leqslant kC$, at time t_s. Suppose that $\hat{\mathbf{x}}(t_s) = \mathbf{x}_i$, for some $i < \lambda$. Then, with probability at least $\epsilon = (2k)^{-kC-1}$, agent g visits vertex $\mathbf{x}_i$ at some round $t \in \{t_s, \ldots, t_s + kC\}$, followed by a visit to $\mathbf{x}_{i+1}$ at round $t+1$ (note that, this probability bounds is very crude, but we only need a constant ϵ).

Let $m = \lfloor T^*/\Delta_1 \rfloor$. For $s \in \{0, \ldots, m-1\}$, let Z_s be the indicator random variable of the event that the space-time set $\widetilde{S}_1(\mathbf{i}_s, s+1)$ is good, and let $Z = \sum_{s=0}^{m-1} Z_s$. Then $Z \geqslant m - \Phi_1$, and choosing C and c as in Lemma 7.3.6, we have that

$$\mathbb{P}\left[Z \geqslant \frac{3}{4} \cdot m\right] \geqslant 1 - 2n^{-\kappa}. \tag{7.6}$$

For $s \in \{0, \ldots, m-1\}$, let Y_s be the indicator random variable of the event that $\hat{\mathbf{x}}(t_s) = \mathbf{y}$ or $\hat{\mathbf{x}}(t_{s+1}) \neq \hat{\mathbf{x}}(t_s)$, i.e., that the process makes progress towards $\mathbf{y}$ between rounds t_s and t_{s+1}, if it has not already reached $\mathbf{y}$ at time t_s. Recall that $\mathcal{K}_1(s)$ is the σ-algebra generated by the positions of all agents up to round t_s. Then, as we argued above,

$$\mathbb{P}[Y_s = 1 \mid \mathcal{K}_1(s); Z_s = 1] \geqslant \epsilon. \tag{7.7}$$

Note also that if $Y = \sum_{s=0}^{m-1} Y_s$, then $Y \geqslant k \cdot \ell$ implies $\tau \leqslant T^*$. Thus it suffices to show $\mathbb{P}[Y \geqslant k \cdot \ell] \geqslant 1 - n^{-\kappa}$.

Let $p_1, \ldots, p_\mu$ denote the sequence of all $s \in \{0, \ldots, m-1\}$, for which $Z_s = 1$. Define $X_j = Y_{p_j}$ for $1 \leqslant j \leqslant \mu$, and $X_j = 1$ for $j > \mu$. It follows from (7.7) that, for any $j \geqslant 1$,

$$\mathbb{P}[X_j = 1 \mid X_1, \ldots, X_{j-1}] \geqslant \epsilon.$$

Then for $X = \sum_{s=1}^{3m/4} X_s$, using Lemma 2.3.9, we can apply a standard Chernoff bound to obtain

$$\mathbb{P}\left[X \geqslant \frac{3\epsilon m}{8}\right] \geqslant 1 - e^{-3\epsilon m/32}.$$

Choosing c large enough that $\frac{3\epsilon m}{8} \geqslant k\ell$ and $e^{-3\epsilon m/32} \leqslant n^{-\kappa}$, we obtain

$$\mathbb{P}\left[X \geqslant k\ell\right] = 1 - n^{-\kappa}. \tag{7.8}$$

Note now that if $Z \geqslant \frac{3}{4} \cdot m$ then $Y \geqslant X$. It follows then from (7.6), (7.8), and union bound that $\mathbb{P}\left[Y \geqslant k\ell\right] \geqslant \mathbb{P}\left[\{Z \geqslant 3m/4\} \cap \{X \geqslant k\ell\}\right] = 1 - 3n^{-\kappa}$. □

7.4 Putting pieces together: Proof of Theorem 7.1.1

To complete the proof of Theorem 7.1.1, we have to consider the cases when G is a grid and a torus graph separately. If G is a grid graph, we have to deal with the edges of the graph as the analysis so far has only considered central vertices in V'. We show that if a non-central vertex is initially informed, then very quickly some central vertex becomes informed. The proof uses relatively standard cover-time arguments and avoids certain edge cases by projecting random walks to a lazy random walks in a lower-dimensional torus graph.

If G is a torus graph, then we do not need a special analysis for the vertices that are not in V', since the set V' can be chosen arbitrarily as there is no specific central part in a torus graph.

Lemma 7.4.1. *Consider the* VISIT-EXCHANGE *process on $G = G_{k,n}$ graph. For any constant $\kappa > 0$ and $\mathbf{x} \in V \setminus V'$, if $\mathbf{x}$ is informed, then after at most $O(\mathrm{poly}(\log \ell))$ rounds, some vertex in V' will become informed, with probability at least $1 - n^{-\kappa}$.*

Proof. We prove that in $O(\mathrm{poly}(\log \ell))$ rounds some agent will visit $\mathbf{x}$, and after that also visit a vertex in V'. Let $X(t)$ be a random walk starting from the stationary distribution of G and let $\tau_{\mathbf{x}}$ be the first round when it visits $\mathbf{x}$. If N is the number of times the walk $X(t)$ visits $\mathbf{x}$ in the first $\tau_1 = (\eta_1 \log \ell)^2$ rounds, where $\eta_1 > 0$ is a constant, then

$$\mathbb{E}\left[N\right] \geqslant \frac{\tau_1}{2n},$$

since the walk starts from stationarity and the vertex degrees are between k and $2k$. Also,

$$\begin{aligned}
\mathbb{E}\left[N \mid \tau_{\mathbf{x}}\right] &\leqslant \sum_{t=0}^{\tau_1 - \tau_{\mathbf{x}}} \mathbb{P}\left[X(\tau_{\mathbf{x}} + t) = \mathbf{x} \mid \tau_{\mathbf{x}}\right] \\
&= \sum_{t'=0}^{\lfloor(\tau_1 - \tau_x)/2\rfloor} \mathbb{P}\left[X(\tau_x + 2t') = \mathbf{x} \mid \tau_{\mathbf{x}}\right], \quad \text{since } G \text{ is bipartite,} \\
&\leqslant \sum_{t'=0}^{\lfloor(\tau_1 - \tau_x)/2\rfloor} \left(\frac{2k}{k \cdot n} + \frac{20 \cdot 2k}{k \cdot \sqrt{2t'+1}}\right), \quad \text{by Lemma 4.2.4 and } k \leqslant \deg(\mathbf{x}) \leqslant 2k, \\
&\leqslant \frac{2(\tau_1 + 1)}{n} + 40 \cdot \sum_{t=0}^{\tau_1} \frac{1}{\sqrt{t+1}}
\end{aligned}$$

$$
\begin{aligned}
&\leqslant \frac{2(\tau_1+1)}{n} + 80\cdot\sqrt{\tau_1+1} \\
&\leqslant 100\cdot\sqrt{\tau_1}, \quad \text{assuming a large enough constant } \eta_1.
\end{aligned}
$$

Thus,

$$
\mathbb{P}\left[N \geqslant 1\right] = \frac{\mathbb{E}\left[N\right]}{\mathbb{E}\left[N \mid \tau_x \leqslant \tau_1\right]} \geqslant \frac{\sqrt{\tau_1}}{200n}.
$$

Let $\mathcal{E}$ be the event that at least one agent visits $\mathbf{x}$ in the first τ_1 rounds. By the independence of the agents' walks, and by taking a sufficiently large η_1, we have

$$
\mathbb{P}\left[\mathcal{E}\right] \geqslant 1 - \left(1 - \frac{\sqrt{\tau_1}}{200n}\right)^{\alpha\cdot n} \geqslant 1 - e^{-\frac{\alpha}{200}\cdot\sqrt{\tau_1}} \geqslant 1 - n^{-\kappa}/2. \tag{7.9}
$$

Next, our goal is to bound the number of rounds until $X(t)$ reaches one of the vertices in V' after round $\tau_{\mathbf{x}}$. Consider a random walk $X'(t)$ on a k-dimensional *torus* with ℓ vertices on each side, for which too we use V' to denote the set of its "central" vertices. Since V' is equidistant from the sides of the grid G, $X(t)$ and $X'(t)$ can be coupled (in a natural way), so that they both hit a vertex in V' at the same round.

Let $h = 12\cdot\Delta_R$ and $\tau_2 = \eta_2 h^{k+1}\log\ell$ for some constant $\eta_2 > 0$. By symmetry of V', we can assume that one of the coordinates of $\mathbf{x}$ is less than $h/2$, w.l.o.g., $0 \leqslant x_1 < h/2$. Consider a walk $Y(t)$ on a smaller k-dimensional torus $\widehat{H}$ with h vertices on each side, where each coordinate of $Y(t)$ is the same as that coordinate of $X'(t)$, modulo h. By [LP17, Sec. 11.3.2], the expected cover time of Y is at most $O(kh^k\log h) = O(h^{k+1})$, therefore, for a sufficiently large η_2, the walk Y visits all vertices of $\widehat{H}$ in τ_2 rounds, with probability at least $1 - n^{-\kappa}/2$. Let $\mathbf{z}$ be the vertex of $\widehat{H}$ with all its coordinates equal to $h/2 = 6\Delta_R$. If $Y(t) = \mathbf{z}$ and $t < l - h$, then $X'(t) \in V'$, because no vertex in $V \setminus V'$ has its first coordinate equal to $h/2$ (modulo h) and is less than $l - h$ steps away from $\mathbf{x}$. It implies that in τ_2 rounds $X'(t)$ (and thus, $X(t)$) visits some vertex in V', with probability at least $1 - n^{-\kappa}/2$. Combining this with (7.9), we prove that in $\tau = \tau_1 + \tau_2$ rounds some vertex in V' becomes informed, with probability at least $1 - n^{-\kappa}$. Substituting the values of h and R, we see that $\tau = O(\mathrm{poly}(\log\ell))$, which completes the proof. □

We are ready to complete the proof of Theorem 7.1.1, where the grid and torus graphs are considered separately.

Proof of Theorem 7.1.1. First, suppose $G = G_{k,n}$ is the k-dimensional grid graph of n vertices. Fix any vertices $\mathbf{x} \in V'$ and $\mathbf{y} \in V$, and any constant $\kappa > 1$. First suppose that $\mathbf{y}$ is informed initially. If $\mathbf{y} \in V \setminus V'$, then by Lemma 7.4.1, some vertex $\mathbf{z} \in V'$ becomes informed in at most $O(\mathrm{poly}(\log\ell))$ rounds with probability at least $1 - n^{-\kappa}$. If $\mathbf{y} \in V'$, we just let $\mathbf{z} = \mathbf{y}$. Then, by Lemma 7.3.7 we conclude that $\mathbf{x}$ becomes informed (via $\mathbf{z}$) in at most $O(\ell)$ rounds with probability at least $1 - 3n^{-\kappa}$. Suppose now that $\mathbf{x}$ is informed initially, instead of $\mathbf{y}$. By Lemma 2.5.1, then $\mathbf{y}$ becomes informed in $O(\ell)$ rounds, with probability at least $1 - 3n^{-\kappa}$. Using a union bound over all vertices $\mathbf{y}$, we conclude that if $\mathbf{x} \in V'$ is informed then in at most $O(\ell)$ rounds, all vertices in V become informed, with probability at least $1 - 3n^{-\kappa+1}$. Finally, by Corollary 2.5.3, we obtain that for any source vertex in V, in at most $O(\ell)$ rounds, all vertices become informed, with probability at least $1 - 6n^{-\kappa+1}$.

Now, consider the case when $G = \widehat{G}_{k,n}$ is the torus graph. Suppose $\mathbf{x} \in V$ is the source vertex and $\kappa > 1$ is an arbitrary constant. Since G is cyclic, for any $\mathbf{y} \in V$ the maximum coordinate-wise difference between $\mathbf{x}$ and $\mathbf{y}$ is at most $\ell/2$. It implies that after a cyclic shift of the torus graph, we can make assume that $\mathbf{x}, \mathbf{y} \in V'$. Then by Lemma 7.3.7 we have that $\mathbf{y}$ becomes informed in at most $O(\ell)$ rounds with probability at least $1 - 3n^{-\kappa}$. By taking a union bound over all n vertices, we have that $T_{\text{visitx}} = O(\ell)$, with probability at least $1 - 3n^{-\kappa+1}$. $\square$

7.5 Auxiliary results for random walks on grids

Lemma 7.5.1. *Let $X(t)$ be a lazy simple random walk on $\mathbb{Z}$ starting from the origin, with a constant holding probability $\alpha > 0$. Then, there is a constant η, such that*

$$\mathbb{P}\left[X(t) = 0\right] \leqslant \frac{\eta}{\sqrt{t}}.$$

Proof. Suppose N is the number of times the walk moves (rather than staying put) in the first t rounds. N is a sum of independent Bernoulli trials with success probability $1 - \alpha$, thus, $\mathbb{E}\left[N\right] = (1 - \alpha)t$. Let $\mathcal{E}$ be the event that $N \geqslant (1 - \alpha)t/2$. By a Chernoff bound,

$$\mathbb{P}\left[\mathcal{E}\right] \geqslant 1 - e^{-(1-\alpha)t/8}.$$

Let t' be the smallest even integer that is at least $(1 - \alpha)t/2$ and let X' be a (non-lazy) simple random walk on $\mathbb{Z}$ starting from the origin. Then,

$$\begin{aligned}
\mathbb{P}\left[X(t) = 0\right] &\leqslant \mathbb{P}\left[X(t) = 0 \mid \mathcal{E}\right] + \mathbb{P}\left[\neg\mathcal{E}\right] \\
&\leqslant \mathbb{P}\left[X'(t') = 0\right] + e^{-(1-\alpha)t/8} \\
&= \binom{t'}{t'/2} \cdot 2^{-t'} + e^{-(1-\alpha)t/8} \\
&\leqslant \frac{1}{\sqrt{t'}} + e^{-(1-\alpha)t/8} \\
&\leqslant \sqrt{\frac{2}{(1-\alpha)t}} + e^{-(1-\alpha)t/8} \\
&\leqslant \frac{\eta}{\sqrt{t}},
\end{aligned}$$

for a sufficiently large constant η. $\square$

Lemma 7.5.2. *Let $X(t)$ be a simple random walk on $\mathbb{Z}^k$, starting from the origin. Consider vertices $\mathbf{x} = (x_1, \ldots, x_k)$ and $\mathbf{y} = (y_1, \ldots, y_k)$ such that for $1 \leqslant i \leqslant k$, $0 \leqslant |x_i| \leqslant |y_i|$. If $d(\mathbf{x}) = \sum_{i=1}^{k} |x_i|$ and $d(\mathbf{y}) = \sum_{i=1}^{k} |y_i|$ have the same parity, then for any $t \geqslant 0$,*

$$\mathbb{P}\left[X(t) = \mathbf{x}\right] \geqslant \mathbb{P}\left[X(t) = \mathbf{y}\right]. \tag{7.10}$$

Proof. We prove the result by an induction on t. For $t = 0$, the result is immediate, so suppose that (7.10) holds for $t \geqslant 0$. By symmetry we can assume that $0 \leqslant x_i \leqslant y_i$ for $1 \leqslant i \leqslant k$. We say that a bijection $h : \Gamma(\mathbf{x}) \to \Gamma(\mathbf{y})$ is *valid* if for $\mathbf{x}' = (x_1, \ldots, x_k) \in \Gamma(\mathbf{x})$ and its mapping $\mathbf{y}' = (y_1', \ldots, y_k') = h(\mathbf{x}')$ we have $|x_i'| \leqslant |y_i'|$ for $1 \leqslant i \leqslant k$.

If we are able to construct a valid bijection h, then

$$\mathbb{P}\left[X(t+1) = \mathbf{x}\right] = \frac{1}{2k} \cdot \sum_{\mathbf{x}' \in \Gamma(\mathbf{x})} \mathbb{P}\left[X(t) = \mathbf{x}'\right]$$

$$
\begin{aligned}
&\geqslant \frac{1}{2k} \cdot \sum_{\mathbf{x}' \in \Gamma(\mathbf{x})} \mathbb{P}\left[X(t) = h(\mathbf{x}')\right], \quad \text{by the inductive hypothesis,} \\
&= \frac{1}{2k} \cdot \sum_{\mathbf{y}' \in \Gamma(\mathbf{y})} \mathbb{P}\left[X(t) = y'\right] \\
&= \mathbb{P}\left[X(t+1) = \mathbf{y}\right],
\end{aligned}
$$

completing the proof. Thus, we describe such a bijection next.

Let $\mathbf{u}_i$ be the unit vector for which all coordinates are 0, except the ith, which is 1. Then, for any $\mathbf{x}$, $\Gamma(\mathbf{x}) = \{\mathbf{x} + \mathbf{u}_i \mid 1 \leqslant i \leqslant k\} \cup \{\mathbf{x} - \mathbf{u}_i \mid 1 \leqslant i \leqslant k\}$. We set $h(\mathbf{x} + \mathbf{u}_i) = \mathbf{y} + \mathbf{u}_i$ for which validity condition of the bijection holds trivially. This simple construction does not work for vertices $\mathbf{x}' = \mathbf{x} - \mathbf{u}_i$ because if $x_i = 0$ and $y_i = 1$ the bijection above is no longer valid.

Let $J = \{j_1, \ldots, j_l\} = \{j \mid x_j = 0 \text{ and } y_j = 1\}$. If l is odd, then by the fact that $d(\mathbf{x})$ and $d(\mathbf{y})$ have the same parity, there must be some $j' \notin J$ such that $1 \leqslant x_{j'} < y_{j'}$. Let $j_{l+1} = j'$. We set the values $h(\mathbf{x} + \mathbf{u}_{j_i})$ in pairs. For any odd $i \leqslant l$, we set $h(\mathbf{x} - \mathbf{u}_{j_i}) = \mathbf{y} - \mathbf{u}_{j_{i+1}}$ and $h(\mathbf{x} - \mathbf{u}_{j_{i+1}}) = \mathbf{y} - \mathbf{u}_{j_i}$. It is not hard to see that h stays valid and thus, we have set $h(\mathbf{x} - \mathbf{u}_j)$ for $j \in J$.

For the remaining vertices $\mathbf{x}' = \mathbf{x} - \mathbf{u}_i$ it holds that either $x_i = y_i = 0$, $x_i = 0$ and $y_i \geqslant 2$, or $1 \leqslant x_i \leqslant y_i$. In all three cases, we can set $h(\mathbf{x}') = \mathbf{y} - \mathbf{u}_i$. $\square$

Chapter 8

Conclusions

In this thesis we investigated a natural agent-based broadcasting protocol, called VISIT-EXCHANGE. Here we review the results of the thesis and give directions of possible future work.

8.1 Contributions

This thesis focused on the theoretical analysis of VISIT-EXCHANGE in the case when there is a linear number of agents in the network performing random walks. We first compared VISIT-EXCHANGE to traditional randomised rumour spreading algorithms, such as PUSH-PULL, and showed that in many cases, the two protocols have complementary properties. Namely, PUSH-PULL is slow when the network contains large hub nodes that are not very well connected directly, such as in high-degree balanced trees, while VISIT-EXCHANGE thrives in this setting. Conversely, in some graphs "node islands" may exist, which are visited rarely by agents in VISIT-EXCHANGE and thus the process has a large broadcast time. In such cases, it is possible that PUSH-PULL can be faster than VISIT-EXCHANGE. It therefore may make sense to combine the two protocols to reap the benefits from both protocols. Our preliminary experimental analysis, presented in Appendix A, suggests that such a combined protocol may indeed give the best of both worlds.

The first significant technical contribution of the thesis was the establishment of the equivalence of VISIT-EXCHANGE and randomised rumour spreading in sufficiently dense regular graphs. This motivated our study of VISIT-EXCHANGE for a variety of sparse graphs, and the specific analysis of the process for each type of network was the second main contribution of the thesis. These networks included regular graphs, expanders, grids and balanced trees, for which we proved tight or almost tight bounds. The summary of our main theoretical contributions can found in Table 1.1.

8.2 Future directions

There are three main directions in which this work can be extended. The first one is to tighten some of the bounds present here and to analyse a wider range of graphs. More concretely, our upper bound on T_{visitx} for a general regular expander is $O(\log n \cdot \log \log n)$ by Theorem 5.1.5. We conjecture that the $\log \log n$ factor is not necessary but the techniques used in this thesis have not been sufficient to show it. A similar improvement may be possible in Theorems 4.1.1 and 4.1.2 for the bound on general regular graphs and for the bound in terms of the average degree of the graph. A further analysis could also be done for graph classes not considered here, such as scale-free graphs as models for social and other real-world networks.

A second direction is the analysis of variations of VISIT-EXCHANGE, for example changing the number of agents in the process. A seminal paper by Alon et al. studied the speed-up of the cover time of a graph by multiple random walks [Alo+11]. A similar study on the dependence of the broadcast time of VISIT-EXCHANGE on the number of agents would be interesting. Besides varying the number of agents, it is also plausible to investigate various failure modes of VISIT-EXCHANGE, such as when agents disappear or forget the message with some probability. Similar work has been done on randomised rumour spreading [ES09].

Another possibility is to consider a quasirandom version of VISIT-EXCHANGE. In this case, the neighbours of each vertex are arranged in an arbitrary cyclic order with the initial vertex picked randomly. An agent arriving at a vertex, next goes to its neighbour according to that cyclic order (agents arriving simultaneously are ordered by their ID). This variant of VISIT-EXCHANGE is similar to the quasirandom rumour spreading [DFS14]. A significant advantage of quasirandom processes is that they require fewer random bits and only at the start of the process, making the quasirandomness attractive from a practical point of view. From a theoretical standpoint, quasirandomness adds additional dependencies in the process, which can complicate the analysis. [DFS14] shows that quasirandom rumour spreading is faster than standard rumour spreading on certain graph classes such as balanced trees, hypercubes and almost all random graphs. It is not clear whether quasirandom VISIT-EXCHANGE has similar properties.

Thirdly, it is possible to extend VISIT-EXCHANGE by making the agents more powerful. A simple extension is to assume that the agents have some limited amount of memory and perform non-backtracking random walks [LP16, Section 6.2], i.e., agents remember the vertex they arrived from in the previous round and do not visit it immediately after (unless the vertex has degree 1). It is not hard to see that this variant of VISIT-EXCHANGE is extremely fast for path graphs, which raises the question whether this modification can benefit VISIT-EXCHANGE in sparse graphs.

Bibliography

[AF02] David Aldous and James Allen Fill. *Reversible Markov Chains and Random Walks on Graphs*. Unfinished monograph, recompiled 2014. Available online. 2002.

[Alo+11] Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. "Many Random Walks Are Faster Than One". In: *Combinatorics, Probability & Computing* 20.4 (2011), pp. 481–502. DOI: 10.1017/S0963548311000125.

[AMP02] O. S. M. Alves, F. P. Machado, and S. Yu Popov. "The shape theorem for the frog model". In: *The Annals of Applied Probability* 12.2 (May 2002), pp. 533–546. DOI: 10.1214/aoap/1026915614.

[BA99] Albert-László Barabási and Réka Albert. "Emergence of Scaling in Random Networks". In: *Science* 286.5439 (1999), pp. 509–512. DOI: 10.1126/science.286.5439.509.

[BC19] Anna D. Broido and Aaron Clauset. "Scale-free networks are rare". In: *Nature Communications* 10.11 (Mar. 2019), p. 1017. DOI: 10.1038/s41467-019-08746-5.

[BEF16] Petra Berenbrink, Robert Elsässer, and Tom Friedetzky. "Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems". In: *Distributed Computing* 29.5 (Oct. 2016), pp. 317–339. DOI: 10.1007/s00446-016-0264-0.

[BES14] Petra Berenbrink, Robert Elsässer, and Thomas Sauerwald. "Randomised broadcasting: Memory vs. randomness". In: *Theor. Comput. Sci.* 520 (2014), pp. 27–42. DOI: 10.1016/j.tcs.2013.08.011.

[Bez+17] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. "Julia: A fresh approach to numerical computing". In: *SIAM Review* 59.1 (2017), pp. 65–98. DOI: 10.1137/141000671.

[BFc17] Seth Bromberger, James Fairbanks, and other contributors. *JuliaGraphs/LightGraphs.jl: an optimized graphs package for the Julia programming language.* 2017. DOI: 10.5281/zenodo.889971.

[Boy+06] Stephen P. Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. "Randomized gossip algorithms". In: *IEEE Trans. Information Theory* 52.6 (2006), pp. 2508–2530. DOI: 10.1109/TIT.2006.874516.

[Bro+94] Andrei Z. Broder, Anna R. Karlin, Prabhakar Raghavan, and Eli Upfal. "Trading Space for Time in Undirected s-t Connectivity". In: *SIAM J. Comput.* 23.2 (1994), pp. 324–334. DOI: 10.1137/S0097539790190144.

[Bro+98] Andrei Z. Broder, Alan M. Frieze, Stephen Suen, and Eli Upfal. "Optimal Construction of Edge-Disjoint Paths in Random Graphs". In: *SIAM Journal on Computing* 28.2 (Jan. 1998), pp. 541–573. DOI: 10.1137/S0097539795290805.

[Cen+17] Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. "Rumor Spreading with No Dependence on Conductance". In: *SIAM J. Comput.* 46.1 (2017), pp. 58–79. DOI: 10.1137/14099992X.

[CFR09] Colin Cooper, Alan M. Frieze, and Tomasz Radzik. "Multiple Random Walks in Random Regular Graphs". In: *SIAM J. Discrete Math.* 23.4 (2009), pp. 1738–1761. DOI: 10.1137/080729542.

[Chi+18] Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, and Alessandro Panconesi. "Rumor Spreading and Conductance". In: *Journal of the ACM* 65.4 (Apr. 2018), pp. 1–21. DOI: 10.1145/3173043.

[CL06] Fan R. K. Chung and Lincoln Lu. "Concentration Inequalities and Martingale Inequalities: A Survey". In: *Internet Mathematics* 3.1 (2006), pp. 79–127. DOI: 10.1080/15427951.2006.10129115.

[CLP10] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. "Rumour spreading and graph conductance". In: *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Moses Charikar. Society for Industrial and Applied Mathematics, Jan. 2010, pp. 1657–1663. DOI: 10.1137/1.9781611973075.135.

[Dat+20] George Datseris, Jonas Isensee, Sebastian Pech, and Tamás Gál. "DrWatson: the perfect sidekick for your scientific inquiries". In: *Journal of Open Source Software* 5.54 (2020), p. 2673. DOI: 10.21105/joss.02673.

[Dem+88] Alan J. Demers, Daniel H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard E. Sturgis, Daniel C. Swinehart, and Douglas B. Terry. "Epidemic Algorithms for Replicated Database Maintenance". In: *Operating Systems Review* 22.1 (1988), pp. 8–32. DOI: 10.1145/43921.43922.

[DFF11] Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. "Social Networks Spread Rumors in Sublogarithmic Time". In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. STOC '11. San Jose, California, USA: Association for Computing Machinery, 2011, pp. 21–30. DOI: 10.1145/1993636.1993640.

[DFS14] Benjamin Doerr, Tobias Friedrich, and Thomas Sauerwald. "Quasirandom Rumor Spreading". In: *ACM Trans. Algorithms* 11.2 (2014), 9:1–9:35. DOI: 10.1145/2650185.

[Die67] Klaus Dietz. "Epidemics and Rumours: A Survey". In: *Journal of the Royal Statistical Society. Series A (General)* 130.4 (1967), p. 505. DOI: 10.2307/2982521.

[DNS06] Tassos Dimitriou, Sotiris E. Nikoletseas, and Paul G. Spirakis. "The infection time of graphs". In: *Discrete Applied Mathematics* 154.18 (2006), pp. 2577–2589. DOI: 10.1016/j.dam.2006.04.026.

[Doe+08] Benjamin Doerr, Tobias Friedrich, Marvin Künnemann, and Thomas Sauerwald. "Quasirandom rumor spreading: An experimental analysis". In: *Journal of Experimental Algorithmics (JEA)* 16 (2008), pp. 3–1.

[DP09] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. DOI: 10.1017/CBO9780511581274.

[DPR21] Rami Daknama, Konstantinos Panagiotou, and Simon Reisser. "Robustness of randomized rumour spreading". In: *Combinatorics, Probability and Computing* 30.1 (Jan. 2021), pp. 37–78. DOI: 10.1017/S0963548320000310.

[DR96] Devdatt P. Dubhashi and Desh Ranjan. "Balls and Bins: A Study in Negative Dependence". In: *BRICS Report Series* 3.25 (Jan. 1996). DOI: 10.7146/brics.v3i25.20006.

[ELS04] Robert Elsässer, Ulf Lorenz, and Thomas Sauerwald. "Agent-Based Information Handling in Large Networks". In: *Mathematical Foundations of Computer Science 2004*. Ed. by Jiří Fiala, Václav Koubek, and Jan Kratochvíl. Vol. 3153. Springer Berlin Heidelberg, 2004, pp. 586–598. DOI: 10.1007/978-3-540-28629-5_45.

[ES08] Robert Elsässer and Thomas Sauerwald. "The Power of Memory in Randomized Broadcasting". In: *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '08. San Francisco, California: Society for Industrial and Applied Mathematics, 2008, pp. 218–227.

[ES09] R. Elsässer and T. Sauerwald. "On the runtime and robustness of randomized broadcasting". In: *Theoretical Computer Science*. Graphs, Games and Computation: Dedicated to Professor Burkhard Monien on the Occasion of his 65th Birthday 410.36 (Aug. 2009), pp. 3414–3427. DOI: 10.1016/j.tcs.2008.04.017.

[ES11] Robert Elsässer and Thomas Sauerwald. "Tight bounds for the cover time of multiple random walks". In: *Theoretical Computer Science* 412.24 (May 2011), pp. 2623–2641. DOI: 10.1016/j.tcs.2010.08.010.

[Eug+04] P.T. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulie. "Epidemic information dissemination in distributed systems". In: *Computer* 37.5 (May 2004), pp. 60–67. DOI: 10.1109/MC.2004.1297243.

[Fei+90] Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. "Randomized Broadcast in Networks". In: *Random Struct. Algorithms* 1.4 (1990), pp. 447–460. DOI: 10.1002/rsa.3240010406.

[Fel68] William Feller. *An Introduction to Probability Theory and its Applications. Vol. 1*. eng. Wiley series in probability and mathematical statistics. Wiley, 1968.

[FG85] A.M. Frieze and G.R. Grimmett. "The shortest-path problem for graphs with random arc-lengths". In: *Discrete Applied Mathematics* 10.1 (Jan. 1985), pp. 57–77. DOI: 10.1016/0166-218X(85)90059-9.

[FP10] Nikolaos Fountoulakis and Konstantinos Panagiotou. "Rumor Spreading on Random Regular Graphs and Expanders". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Ed. by Maria Serna, Ronen Shaltiel, Klaus Jansen, and José Rolim. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 560–573. DOI: 10.1002/rsa.20432.

[FPS12] Nikolaos Fountoulakis, Konstantinos Panagiotou, and Thomas Sauerwald. "Ultra-Fast Rumor Spreading in Social Networks". In: *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '12. Kyoto, Japan: Society for Industrial and Applied Mathematics, 2012, pp. 1642–1660. DOI: 10.5555/2095116.2095246.

[Gia11] George Giakkoupis. "Tight bounds for rumor spreading in graphs of a given conductance". In: *Symposium on Theoretical Aspects of Computer Science (STACS2011)* 9 (Mar. 2011). DOI: 10.4230/LIPIcs.STACS.2011.57.

[Gia14] George Giakkoupis. "Tight Bounds for Rumor Spreading with Vertex Expansion". In: *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms, SODA*. 2014, pp. 801–815. DOI: 10.1137/1.9781611973402.59.

[GK18] Mohsen Ghaffari and Fabian Kuhn. "Distributed MST and Broadcast with Fewer Messages, and Faster Gossiping". In: *32nd International Symposium on Distributed Computing, DISC 2018*. Vol. 121. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 30:1–30:12. DOI: 10.4230/LIPIcs.DISC.2018.30.

[GMS19] George Giakkoupis, Frederik Mallmann-Trenn, and Hayk Saribekyan. "How to Spread a Rumor: Call Your Neighbors or Take a Walk?" In: *Proc. 38th ACM Symposium on Principles of Distributed Computing, PODC*. 2019, pp. 24–33. DOI: 10.1145/3293611.3331622.

[GNW16] George Giakkoupis, Yasamin Nazari, and Philipp Woelfel. "How Asynchrony Affects Rumor Spreading Time". In: *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*. PODC '16. Chicago, Illinois, USA: Association for Computing Machinery, 2016, pp. 185–194. DOI: 10.1145/2933057.2933117.

[GS01] Geoffrey Grimmett and David Stirzaker. *Probability and random processes*. Oxford university press, 2001.

[GS18] Peter Gracar and Alexandre Stauffer. "Percolation of Lipschitz Surface and Tight Bounds on the Spread of Information Among Mobile Agents". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Vol. 116. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 39:1–39:17. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2018.39.

[GSS20] George Giakkoupis, Hayk Saribekyan, and Thomas Sauerwald. "Spread of Information and Diseases via Random Walks in Sparse Graphs". In: *DISC 2020 - 34rd International Symposium on Distributed Computing*. full version. Freiburg, Germany, Oct. 2020, pp. 1–42.

[Hae15] Bernhard Haeupler. "Simple, Fast and Deterministic Gossip and Rumor Spreading". In: *J. ACM* 62.6 (2015), 47:1–47:18. DOI: 10.1145/2767126.

[Her18] Jonathan Hermon. "Frogs on trees?" In: *Electron. J. Probab.* 23 (2018), 40 pp. DOI: 10.1214/18-EJP144.

[HHL88] Sandra M. Hedetniemi, Stephen T. Hedetniemi, and Arthur L. Liestman. “A survey of gossiping and broadcasting in communication networks”. In: *Networks* 18.4 (1988), pp. 319–349. DOI: https://doi.org/10.1002/net.3230180406.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. “Expander graphs and their applications”. In: *Bulletin of the American Mathematical Society* 43.04 (Aug. 2006), pp. 439–562. DOI: 10.1090/S0273-0979-06-01126-8.

[HP20] Reaz Huq and Pawel Pralat. “Broadcasting on Paths and Cycles”. In: (Nov. 2020). URL: http://arxiv.org/abs/2011.12733.

[Jan17] Svante Janson. *Tail bounds for sums of geometric and exponential variables*. 2017.

[Kar+00] Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. “Randomized Rumor Spreading”. In: *Proc. 41st IEEE Symposium on Foundations of Computer Science, FOCS*. 2000, pp. 565–574. DOI: 10.1109/SFCS.2000.892324.

[KMW27] William Ogilvy Kermack, A. G. McKendrick, and Gilbert Thomas Walker. “A contribution to the mathematical theory of epidemics”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115.772 (Aug. 1927), pp. 700–721. DOI: 10.1098/rspa.1927.0118.

[KS03] Harry Kesten and Vladas Sidoravicius. “Branching Random Walk with Catalysts”. In: *Electronic Journal of Probability* 8.0 (2003). DOI: 10.1214/EJP.v8-127.

[KS05] Harry Kesten and Vladas Sidoravicius. “The spread of a rumor or infection in a moving population”. In: *The Annals of Probability* 33.6 (Nov. 2005). Zbl: 1111.60074, pp. 2402–2462. DOI: 10.1214/009117905000000413.

[KS08] Harry Kesten and Vladas Sidoravicius. “A shape theorem for the spread of an infection”. In: *Annals of Mathematics* 167.3 (May 2008), pp. 701–766. DOI: 10.4007/annals.2008.167.701.

[Lam+12] Henry Lam, Zhenming Liu, Michael Mitzenmacher, Xiaorui Sun, and Yajun Wang. “Information dissemination via random walks in d-dimensional space”. In: *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms, SODA*. 2012, pp. 1612–1622. DOI: 10.1137/1.9781611973099.128.

[Lia+18] Xiangru Lian, Wei Zhang, Ce Zhang, and Ji Liu. “Asynchronous decentralized parallel stochastic gradient descent”. In: *International Conference on Machine Learning*. PMLR. 2018, pp. 3043–3052.

[LK14] Jure Leskovec and Andrej Krevl. *SNAP Datasets: Stanford Large Network Dataset Collection*. http://snap.stanford.edu/data. June 2014.

[Lov93] László Lovász. “Random Walks on Graphs: A Survey, Combinatorics, Paul Erdos is Eighty”. In: *Bolyai Soc. Math. Stud.* 2 (Jan. 1993).

[LP16] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Vol. 42. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, 2016, pp. xv+699. DOI: 10.1017/9781316672815.

[LP17] David A Levin and Yuval Peres. *Markov chains and mixing times*. Vol. 107. American Mathematical Soc., 2017.

[MN14] Charles Macal and Michael North. “Introductory tutorial: Agent-based modeling and simulation”. In: *Proceedings of the Winter Simulation Conference 2014*. 2014, pp. 6–20. DOI: 10.1109/WSC.2014.7019874.

[MS08] Damon Mosk-Aoyama and Devavrat Shah. “Fast Distributed Algorithms for Computing Separable Functions”. In: *IEEE Transactions on Information Theory* 54.7 (July 2008), pp. 2997–3007. DOI: 10.1109/TIT.2008.924648.

[MU17] Michael Mitzenmacher and Eli Upfal. *Probability and Computing*. Second. Cambridge University Press, 2017.

[OP19] Roberto I. Oliveira and Yuval Peres. “Random walks on graphs: New bounds on hitting, meeting, coalescing and returning”. In: *Proc. 16th Workshop on Analytic Algorithmics and Combinatorics, ANALCO*. 2019, pp. 119–126. DOI: 10.1137/1.9781611975505.13.

[Ora01] Andy Oram. *Peer-to-Peer: Harnessing the power of disruptive technologies*. "O’Reilly Media, Inc.", 2001.

[Pet+10] Alberto Pettarin, Andrea Pietracaprina, Geppino Pucci, and Eli Upfal. “Infectious Random Walks”. In: *CoRR* abs/1007.1604 (2010), 21 pp. URL: http://arxiv.org/abs/1007.1604.

[Pit87] Boris Pittel. “On Spreading a Rumor”. In: *SIAM Journal on Applied Mathematics* 47.1 (1987), pp. 213–223.

[Pop03] Serguei Yu. Popov. “Frogs and some other interacting random walks models”. In: *Proc. Discrete Random Walks, DRW*. 2003, pp. 277–288. URL: http://dmtcs.episciences.org/3328.

[Sau10] Thomas Sauerwald. “On Mixing and Edge Expansion Properties in Randomized Broadcasting”. In: *Algorithmica* 56.1 (Jan. 2010), pp. 51–88. DOI: 10.1007/s00453-008-9245-4.

[Sha07] Devavrat Shah. “Gossip Algorithms”. In: *Foundations and Trends® in Networking* 3.1 (2007), pp. 1–125. DOI: 10.1561/1300000014.

[SS11] Thomas Sauerwald and Alexandre Stauffer. “Rumor Spreading and Vertex Expansion on Regular Graphs”. In: *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Dana Randall. SIAM, 2011, pp. 462–475. DOI: 10.1137/1.9781611973082.37.

[ST11] Daniel A. Spielman and Shang-Hua Teng. “Spectral Sparsification of Graphs”. In: *SIAM J. Comput.* 40.4 (2011), pp. 981–1025. DOI: 10.1137/08074489X.

[VRB03] Werner Vogels, Robbert van Renesse, and Ken Birman. “The Power of Epidemics: Robust Communication for Large-Scale Distributed Systems”. In: *SIGCOMM Comput. Commun. Rev.* 33.1 (Jan. 2003), pp. 131–135. DOI: 10.1145/774763.774784.

[Waj17] David Wajc. “Negative association: definition, properties, and applications”. In: (2017), p. 10. URL: https://web.stanford.edu/~wajc/notes/Negative%20Association.pdf (visited on 06/05/2021).

Appendix A

Experimental evaluation

The theoretical results presented in this thesis give asymptotic bounds on the broadcast time of VISIT-EXCHANGE with respect to graph parameters, and compare it to standard rumour spreading protocols. In this chapter we evaluate agent-based information dissemination experimentally and present results from two sets of simulations. The first one concerns some of our theoretical results. These simulations indicate the magnitude of the hidden constants in our bounds as well as show how the processes evolve over time compared to one another in different types of networks. The second set of experiments analyses the sensitivity of the broadcast time of VISIT-EXCHANGE depending on the number of agents in the system. These results are presented on real-world networks.

It appeared from our theoretical study that VISIT-EXCHANGE is able to disseminate information relatively quickly when the network contains large hubs, and in networks with vertices of small degree in their "periphery" PUSH-PULL can be fast. In light of this fact, we design a new protocol VX-PUSH-PULL, which combines the two. When comparing the combined process to the others, we use half as many agents as the original VISIT-EXCHANGE, and execute its PUSH-PULL component with failures, where each message transmission succeeds with probability $1/2$, independently of all others. Thus, in expectation, the amount of communication per round is the same in all cases.

The main conclusions from the experimental analysis are the following.

- In the theoretical bounds we prove, the constants hidden in the big-O notation appear to be small, and, as a result, the asymptotic trends are visible in our experiments.
- The combined protocol VX-PUSH-PULL seems to be as fast or faster than both PUSH-PULL and VISIT-EXCHANGE, adding weight to our hypothesis that this design can give a practical dissemination algorithm that is more efficient on a wider range of networks than both traditional rumour spreading and agent-based dissemination.

A.1 Notation

The experiments in this chapter involve processes PUSH, PUSH-PULL, VISIT-EXCHANGE and VX-PUSH-PULL (which is defined later). In order to make a fair comparison of these processes, we only consider settings where they have the same amount of communication per round, in expectation. We measure the communication by the number of edges of the graph used in each round (if an edge is used more than once in one round, we count all uses). For VISIT-EXCHANGE with k agents, in one round k edges are used. In standard

PUSH and PUSH-PULL, in each round exactly n edges are used.[1] For a fair comparison, we change these processes based on a parameter $\alpha > 0$, as detailed next.

Let $q = \lfloor \alpha \rfloor$ and $\beta = \alpha - r \in [0, 1)$ be the integer and fractional parts of α. We denote by VISIT-EXCHANGE$_\alpha$ the VISIT-EXCHANGE process that uses $k = \alpha n$ agents. PUSH$_\alpha$ is a variant of PUSH, where in every round each vertex u contacts q randomly selected neighbours and informs them if it is informed at the start of the round. After these q contacts with neighbours, u finishes the round with probability $1 - \beta$, but also contacts and informs another randomly selected neighbour with probability β. PUSH-PULL$_\alpha$ is defined similarly. The VX-PUSH-PULL$_\alpha$ process is a composition of VISIT-EXCHANGE$_{\alpha/2}$ and PUSH-PULL$_{\alpha/2}$ protocols: In every round, one step of each of these protocols is taken by all nodes and agents. For any $\alpha > 0$, all four protocols with parameter α use exactly αn edges per round, in expectation. This allows for a fair comparison.

For the experiments related to our theoretical results we use $\alpha = 1$. This is equivalent to using exactly n agents in VISIT-EXCHANGE and using unmodified PUSH and PUSH-PULL processes. In this setting we also omit the subscript α.

A.2 Methods and data

For each process P $\in$ {PUSH, PUSH-PULL, VISIT-EXCHANGE, VX-PUSH-PULL}, we mainly study its expected broadcast time, denoted by T_{P}. In each experiment, we approximate T_{P} by averaging the broadcast times from 20 independent executions. (This number of repetitions ensures that the generation of results completes in a reasonable amount of time and that the overall trends are visible). The lower bounds in our theoretical analysis relied on the existence of small "node islands," which are rarely visited by agents (see Sections 3.2.4 and 6.5). Hence, we record the number of rounds $\widetilde{T}_{\text{P}}$ until 90% of the vertices become informed.

We conducted experiments on the following classes of graphs:

- Path graphs of n vertices, denoted by P_n, for $n \in [100, 1000]$ range.
- Rooted balanced trees $R_{b,h}$, where b is the number of children of each internal vertex and h is the height of the tree. We selected a set of (b, h) pairs such that the number of vertices in the corresponding trees are almost equal, so the comparison is meaningful. All the balanced trees considered have about $2,200,000$ vertices.
- Dense regular, non-expanding graphs D_n with n vertices, where n is a square number. D_n is obtained by connecting all vertices in the same row of a n-vertex 2-dimensional torus graph. Thus, every vertex of D_n has degree $\sqrt{n} + 1$, and $\text{diam}(D_n) = \sqrt{n}$. We considered values of n in the range $[1000, 10000]$.
- Real-world networks from Stanford Large Network Database (SNAP) [LK14]. The database contains a many networks but we used the largest ones for each type (social, citation, peer-to-peer, web) for which the simulations would finish in a reasonable amount of time. Some of these graphs are not connected, in which case we took the largest connected component as an input. The details of the networks can be found in Table A.1 along with the results.

[1] We assume that in PUSH, even uninformed vertices initiate a connection, since in practice many messages will be circulated simultaneously so each vertex is likely to be informed by some message.

To conduct the experiments we developed a package, written in Julia language [Bez+17].[2] The main packages we used were LightGraphs, for storing graphs and working with them [BFc17], and DrWatson for organising the codebase and plots [Dat+20].

A.3 Validating theoretical results

First, we look at path graphs P_n with n vertices. By Theorem 7.1.1 and [Fei+90], the broadcast time of all the processes we consider is linear in n. Fig. A.1(a) confirms that theoretical result and we can additionally observe that the broadcast time of PUSH and VISIT-EXCHANGE is very close to $2n$ for the graph with n vertices, while PUSH-PULL is faster by about 30%. Unsurprisingly, VX-PUSH-PULL's performance is between the latter two. It is not very hard to obtain a tight analysis of this precise bound for PUSH, while it is not obvious that $T_{\text{visitx}} \approx 2n$ as well. Fig. A.1(b) shows that the processes make progress through the path graph at a constant rate, on average, which is not surprising as otherwise the broadcast time for P_n would not be linear in n. Notice that the curves become flatter at the very end of the process. This is because we take an average over a number of executions. The fact that for VISIT-EXCHANGE the flatness is most pronounced indicates that T_{visitx} has a larger upper tail, compared to the other processes.

The next experiment compares the processes on dense regular graphs D_n. By Theorem 3.1.1, all processes must have the same asymptotic broadcast time, which is $O(\sqrt{n})$ in expectation. These facts are reflected in Fig. A.2(a), where the gap between the slowest and the fastest process does not exceed the factor of 2.

Next we present experiments on balanced trees. Recall that $R_{b,h}$ is a balanced tree of branching b and height h. By our result for VISIT-EXCHANGE in Chapter 6 we have that $T_{\text{visitx}} = O(h \log h + \log n)$, w.h.p., and by a standard result, $T_{\text{push}} = O(b \log n)$, w.h.p. [Fei+90]. Thus, we expect that for small values of b both PUSH and PUSH-PULL are faster than VISIT-EXCHANGE, and, as b increases, the order reverses. Indeed, the asymptotic bound can be observed experimentally. In Fig. A.3(a), we use a collection of balanced trees of roughly the same size. While VISIT-EXCHANGE is slightly slower than the other processes for small values of b, it never has a very large broadcast time. On the other hand, PUSH-PULL becomes very slow for large values of b.[3] VX-PUSH-PULL is perhaps an excellent compromise as its broadcast time is very close to the minimum of the broadcast times of VISIT-EXCHANGE and PUSH-PULL in all cases. Fig. A.3(b) shows how the processes evolve on a fixed balanced tree with large branching. It should be noted that in all processes, but particularly in VISIT-EXCHANGE and VX-PUSH-PULL, the majority of the time is spent in informing the last small fraction of vertices. Hence, if the aim is to only inform almost all vertices the agent-based protocols still outperform randomised rumour spreading for balanced trees with large branching.

We also simulated the processes on binary trees. Fig. A.3(c) illustrates the lower bound on the broadcast time of VISIT-EXCHANGE in binary trees (Theorem 6.1.1). Recall that for the lower bound, we showed that, w.h.p., there are subtrees of height $\Theta(\log \log n)$ that are not visited in $O(\log n \cdot \log \log n)$ rounds. This explains the larger variations from the expected curve in VISIT-EXCHANGE in these executions. We can see the effect of these

[2] The code can be found at `https://github.com/saribekyan/RumourSpreading`.

[3] For PUSH-PULL the case of $h = 1$, i.e., the star graph, is a special case, where $T_{\text{ppull}} \leqslant 2$. For $h = 2$, for example, $\mathbb{E}[T_{\text{ppull}}] = \Theta(\sqrt{n})$.

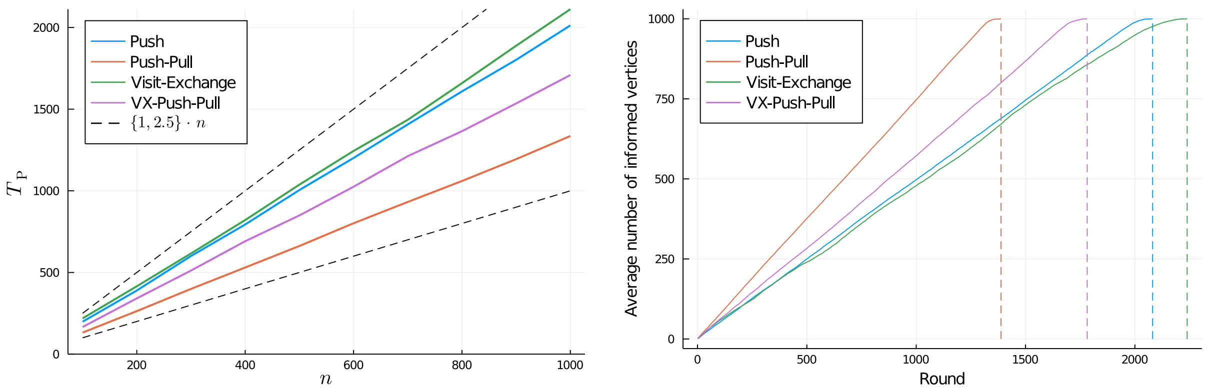

(a) (b)

Figure A.1 Broadcasting on rooted balanced trees P_n, where the source is one of the endpoints. **(a)** For each process P, the broadcast time T_{P} on P_n with respect to n. The dashed lines are linear functions with slopes 1.0 and 2.5. **(b)** The average number of informed vertices at each round of the process on a path graph with 1000 vertices.

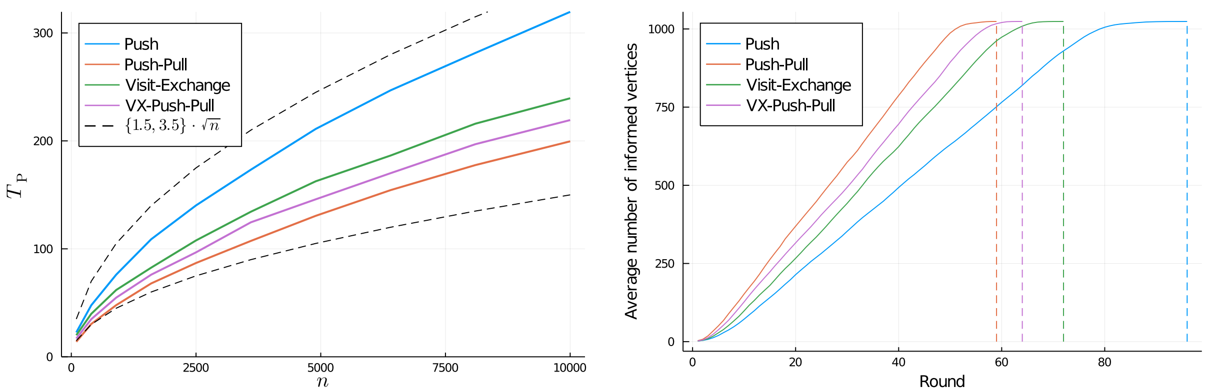

(a) (b)

Figure A.2 Broadcasting on dense regular graphs D_n. The source vertex does not matter, due to symmetry. **(a)** The broadcast time T_{P} for each process P with respect to the number of vertices n. **(b)** The average number of informed vertices in each process at each round.

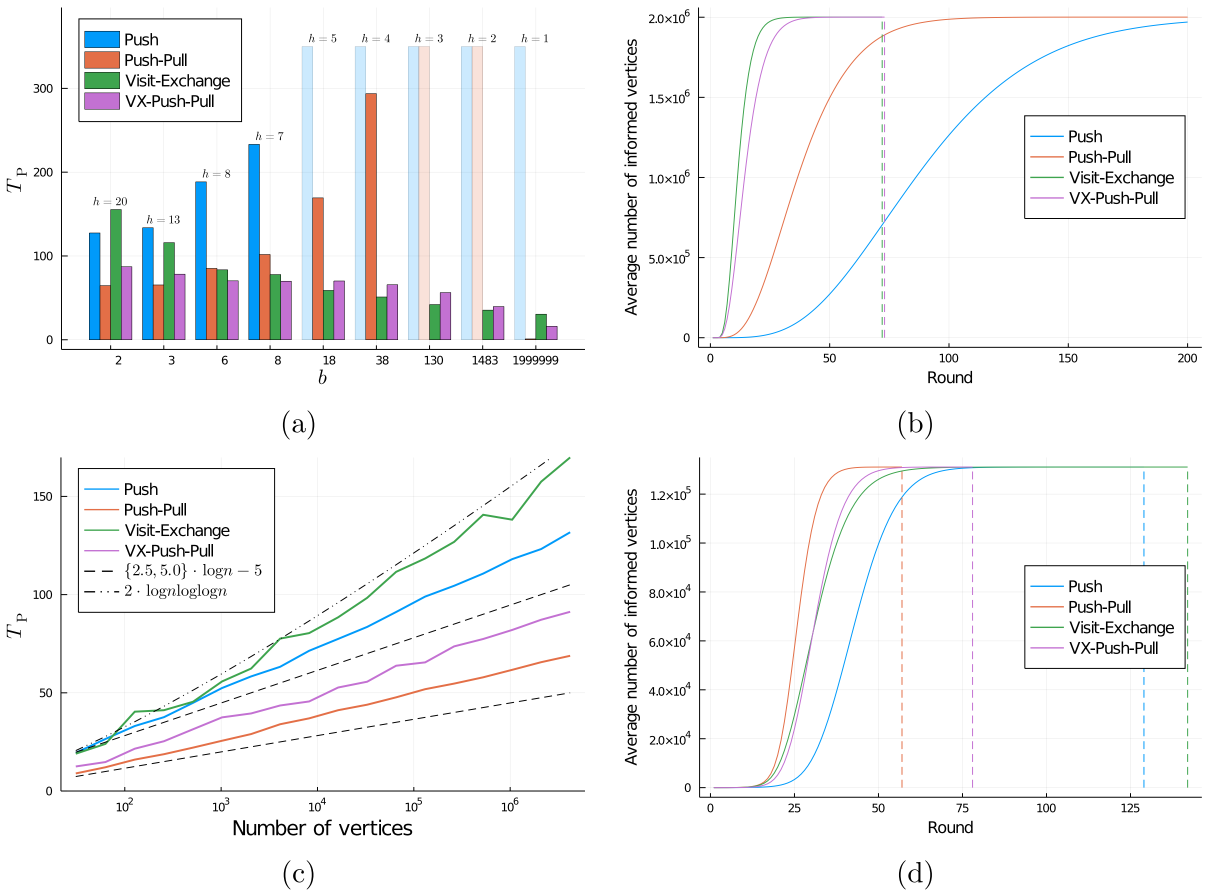

Figure A.3 Broadcasting on rooted balanced trees $R_{b,h}$, where the source is the root vertex. **(a)** The broadcast times for all processes in a group of balanced trees chosen such that $n \in [2 \cdot 10^6, 2.4 \cdot 10^6]$. The transparent bars indicate that the process' execution was cut short since it took more time than would reasonably fit in the graph. The last example corresponds to a star graph with $2 \cdot 10^6$ vertices. **(b)** The average number of informed vertices in $R_{18,5}$ in a given round. The PUSH and PUSH-PULL processes were cut short early to fit the data in the graph. **(c)** The broadcast times for binary trees. **(d)** Average number of informed vertices at any round in a binary tree of height 16.

SNAP Graph ID	Type	n	diam	$d_{\max}$	d_{avg}	T_{ppull}	T_{visitx}	T_{vxpull}
ego-Facebook	Social	4039	8	1045	43.7	**18.7**	238.8	27.1
CA-AstroPh	Citation	17903	14	504	22.0	**28.5**	256.6	41.1
p2p-Gnutella31	P2P sharing	62561	11	95	4.7	**22.5**	74.6	26.4
web-Google	Web	855802	21	6332	10.0	105.0	326.1	**88.3**

Table A.1 Comparison of PUSH-PULL, VISIT-EXCHANGE and VX-PULL on some real graph instances. Average broadcast times is taken over 20 executions.

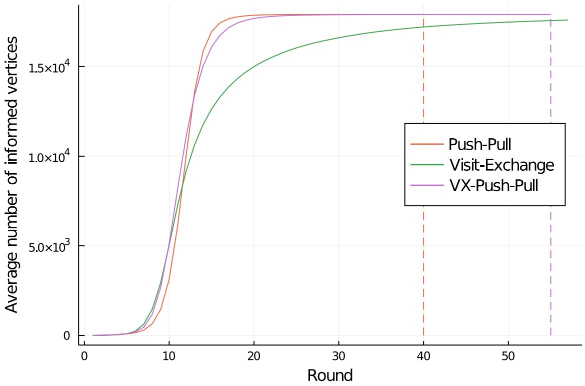

(a) CA-AstroPh

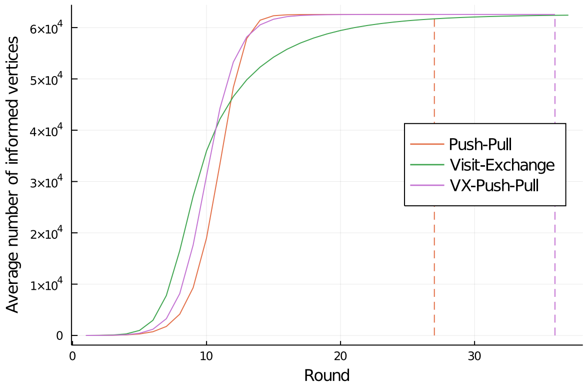

(b) p2p-Gnutella31

Figure A.4 Experiments on real-world networks where processes have parameter $\alpha = 1$. SNAP graph IDs are given in the captions.

"node islands" that become informed very late from the execution on the binary tree of height 16, in Fig. A.3(d). Note that VISIT-EXCHANGE spends the majority of the time to inform the last small fraction (e.g., 1%) of the vertices, and up to that point it is actually faster than PUSH.

A.4 Results on real-world networks

In this section we evaluate PUSH-PULL, VISIT-EXCHANGE and VX-PUSH-PULL on a few real networks. We omit the PUSH protocol since it is strictly dominated by PUSH-PULL and for many instances it takes very long to complete. We set different values of α in the range $[0.05, 4.0]$ to investigate the effect of the number of agents on the broadcast time. We also consider the 90% broadcast time $\widetilde{T}_{\text{P}}$, that is, the number of rounds until 90% of the vertices become informed.

First, in Table A.1 we present the broadcast times of the processes in the setting when $\alpha = 1$, together with some key graph parameters. The source node is fixed to an arbitrary vertex. It is immediately obvious that VISIT-EXCHANGE does not perform a complete broadcast very well, and while VX-PUSH-PULL is not much slower than PUSH-PULL on any of the cases, it is only faster in one instance. By a closer inspection of how the processes progress in Fig. A.4, we see that agent-based broadcasting is fast initially but takes a long time to broadcast to all vertices. Even though in the figure we present only two examples, the trend exists in other graphs as well.

For this reason we next focus on the 90% broadcast time, denoted by $\widetilde{T}_{\text{P}}$ for process

SNAP Graph ID	$\widetilde{T}_{\mathrm{P}}$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1.0$	$\alpha = 2.0$	$\alpha = 4.0$
ego-Facebook	$\widetilde{T}_{\mathrm{ppull}}$	179.8	**90.8**	**22.0**	**14.4**	**10.3**	**7.8**
	$\widetilde{T}_{\mathrm{visitx}}$	327.4	183.9	49.2	31.0	18.8	13.0
	$\widetilde{T}_{\mathrm{vxpull}}$	**174.0**	93.2	28.8	17.6	12.2	8.9
CA-AstroPh	$\widetilde{T}_{\mathrm{ppull}}$	177.6	92.7	24.0	**14.5**	**10.5**	**8.4**
	$\widetilde{T}_{\mathrm{visitx}}$	322.5	176.3	44.8	25.8	16.4	11.0
	$\widetilde{T}_{\mathrm{vxpull}}$	**155.8**	**80.6**	**23.5**	15.8	11.0	8.4
p2p-Gnutella31	$\widetilde{T}_{\mathrm{ppull}}$	178.4	91.8	22.0	13.2	9.3	7.8
	$\widetilde{T}_{\mathrm{visitx}}$	210.1	112.6	28.2	17.2	11.0	8.0
	$\widetilde{T}_{\mathrm{vxpull}}$	**122.0**	**68.0**	**19.8**	**13.0**	**9.0**	**7.0**
web-Google	$\widetilde{T}_{\mathrm{ppull}}$	221.4	114.4	29.3	**18.0**	13.2	11.0
	$\widetilde{T}_{\mathrm{visitx}}$	338.6	179.5	44.6	26.0	16.8	12.0
	$\widetilde{T}_{\mathrm{vxpull}}$	**184.0**	**100.0**	**28.4**	18.7	**12.5**	**10.0**

Table A.2 Comparison of PUSH-PULL$_\alpha$, VISIT-EXCHANGE$_\alpha$ and VX-PULL$_\alpha$ on real graph instances for a range of values for α. They are executed until 90% of the vertices become informed. Additionally, we vary the parameter α which determines the number of agents present in the system.

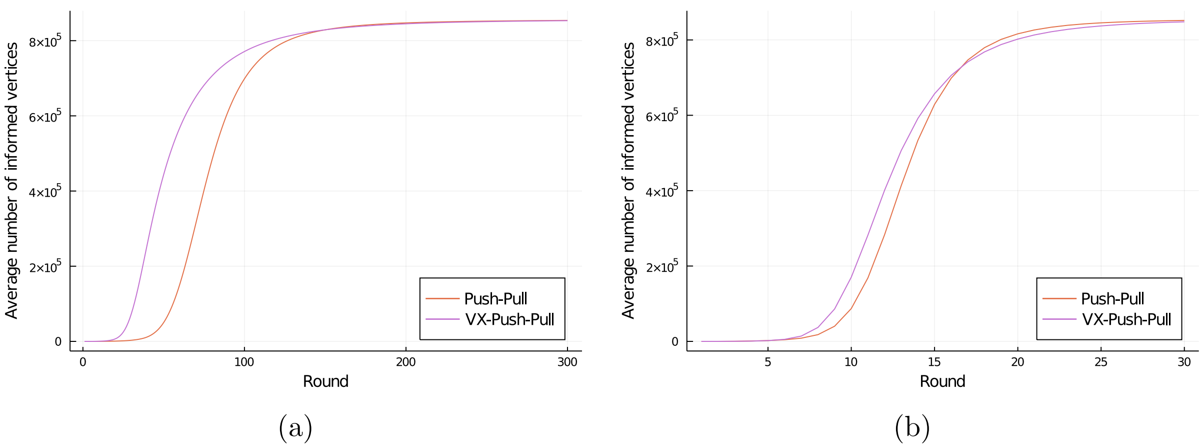

Figure A.5 Two executions of PUSH-PULL and VX-PUSH-PULL on web-Google graph for two values of α. **(a)** $\alpha = 0.1$ **(b)** $\alpha = 1.0$.

P, and investigate the effect of α on $\widetilde{T}_{\text{P}}$. Table A.2 summarises our results. We can see that VX-PUSH-PULL is almost always faster than both PUSH-PULL and VISIT-EXCHANGE. For smaller values of α this fact is even more pronounced. For example, when $\alpha = 0.05$, VX-PUSH-PULL is the fastest to achieve partial broadcast in all graphs. Furthermore, consider the setting when $\alpha < 1$. In this case, the expected amount of communication per round compared to the default process is multiplied by α. Observe that at the same time, especially for VX-PUSH-PULL, the partial broadcast time increases by a factor significantly less than $1/\alpha$. For instance, when the expected communication drops by a factor of 20 ($\alpha = 0.05$), the partial broadcast time increases by a factor of about 10, in all presented graphs. Fig. A.5 shows a particular example of PUSH-PULL and VX-PUSH-PULL on a graph of about 900,000 vertices with $\alpha = 0.1$ and $\alpha = 1.0$, where this phenomenon can also be seen.

A.5 Summary

The main contribution of this thesis is the theoretical analysis of information dissemination algorithms. In this experimental chapter our aim was twofold: To validate our theoretical findings, and, perhaps more interestingly, to investigate the practicality of agent-based information dissemination. The observed experimental findings matched our expectations based on the theoretical results. Moreover, by the example of VX-PUSH-PULL, we saw that agent-based information dissemination can be a useful routine in practice if used in combination with other, well-known dissemination protocols. The advantage of VX-PUSH-PULL$_\alpha$ was particularly visible in contrived settings with small values of α, where the amount of communication per unit of time was reduced.

Our experiments are by no means comprehensive. Further work is needed to evaluate these broadcasting protocols as well as other potential variations on more diverse range of networks, both real-world and generated. For example, we have used the two components of VX-PUSH-PULL equally, but have not determined how the "division of labour" between VISIT-EXCHANGE and PUSH-PULL affects the broadcast time. Furthermore, experiments on larger real-world networks from different domains can shed more light on agent-based dissemination. Lastly, more complex protocols, such as load balancing, can also be evaluated experimentally in an agent-based setting. As is the case for simple broadcasting, introducing agents in these settings may result in performance benefits.