

ARTICLE

Robustness of randomized rumour spreading¹

Rami Daknama, Konstantinos Panagiotou* and Simon Reisser

Institute for Mathematics, Ludwig-Maximilians-Universität München, 80333 Munich, Germany

*Corresponding author. Email: kpanagio@math.lmu.de

(Received 27 September 2019; revised 8 June 2020; accepted 15 May 2020)

Abstract

In this work we consider three well-studied broadcast protocols: *push*, *pull* and *push&pull*. A key property of all these models, which is also an important reason for their popularity, is that they are presumed to be very robust, since they are simple, randomized and, crucially, do not utilize explicitly the global structure of the underlying graph. While sporadic results exist, there has been no systematic theoretical treatment quantifying the robustness of these models. Here we investigate this question with respect to two orthogonal aspects: (adversarial) modifications of the underlying graph and message transmission failures.

We explore in particular the following notion of *local resilience*: beginning with a graph, we investigate up to which fraction of the edges an adversary may delete at each vertex, so that the protocols need significantly more rounds to broadcast the information. Our main findings establish a separation among the three models. On one hand, *pull* is robust with respect to all parameters that we consider. On the other hand, *push* may slow down significantly, even if the adversary may modify the degrees of the vertices by an arbitrarily small positive fraction only. Finally, *push&pull* is robust when no message transmission failures are considered, otherwise it may be slowed down.

On the technical side, we develop two novel methods for the analysis of randomized rumour-spreading protocols. First, we exploit the notion of self-bounding functions to facilitate significantly the round-based analysis: we show that for any graph the variance of the growth of informed vertices is bounded by its expectation, so that concentration results follow immediately. Second, in order to control adversarial modifications of the graph we make use of a powerful tool from extremal graph theory, namely Szemerédi's Regularity Lemma.

2020 MSC Codes: Primary 05C85; Secondary 68R10

1. Introduction

Randomized broadcast protocols are highly relevant for data distribution in large networks of various kinds, including technological, social and biological networks. Among many others there are three basic models in the literature, introduced in [9], [19] and [27], namely *push*, *pull* and *push&pull* (or *pp* for short). Consider a connected graph in which some vertex holds a piece of information; we call this vertex (initially) informed. All three models have the common characteristic that they proceed in rounds. In the *push* model, in every round every informed vertex chooses a neighbour independently and uniformly at random and informs it; this of course only

¹An extended abstract of this paper was published in the *Proceedings of the European Symposium on Algorithms 2019 (ESA '19)*.

has an effect if the target vertex was previously uninformed. Conversely, in the *pull* model every round every *uninformed* vertex chooses a neighbour independently and uniformly at random and asks for the information. If the asked vertex has the information, then the asking vertex becomes informed as well. The third model, *push&pull*, combines both worlds: in each round, each vertex chooses a neighbour independently and uniformly at random, and if one of two vertices is informed, then afterwards both become so. We additionally assume that each message transmission succeeds independently with probability $q \in (0, 1]$. For these algorithms, the main parameter that we consider is the random variable that counts how many rounds are needed until all vertices are informed, and we call these quantities the *runtimes* of the respective algorithms.

In the remainder we will denote the runtime of *push* by $T_{push}(G, v, q)$, where G is the underlying graph, initially the vertex v is informed, and we have a transmission success probability of $q \in (0, 1]$. Analogously we denote the runtimes of *pull* and *push&pull* by $T_{pull}(G, v, q)$ and $T_{pp}(G, v, q)$ respectively. If the choice of v does not matter we will omit it in our notation. The most basic case is when G is the complete graph K_n with n vertices. Then (see e.g. Doerr and Kostrygin [11]) it is known that, for $\mathcal{P} \in \{push, pull, pp\}$ and $q \in (0, 1]$ in expectation and with probability tending to 1 as $n \rightarrow \infty$,

$$T_{\mathcal{P}}(K_n, q) = c_{\mathcal{P}}(q) \log n + o(\log n),$$

where, for $q \in (0, 1)$,

$$c_{push}(q) := \frac{1}{\log(1+q)} + \frac{1}{q},$$

$$c_{pull}(q) := \frac{1}{\log(1+q)} - \frac{1}{\log(1-q)},$$

$$c_{pp}(q) := \frac{1}{\log(1+2q)} + \frac{1}{q - \log(1-q)},$$

and where we set $c_{\mathcal{P}}(1) := \lim_{q \rightarrow 1} c_{\mathcal{P}}(q)$. If q is clear from the context, we write $c_{\mathcal{P}}$ instead of $c_{\mathcal{P}}(q)$. In fact, the results in [11] and also [12] are much more precise, but the stated forms will be sufficient for what follows.

Contribution and related work. In this article our focus is on quantifying the *robustness* of all three models. Indeed, robustness is a key property that is often attributed to them, since they are simple, randomized and, crucially, do not exploit explicitly the structure of the underlying graph (apart from considering the neighbourhoods of the vertices, of course). Clearly the runtime can vary tremendously between different graphs with the same number of vertices. Hence it is essential to understand the impact of structural graph characteristics on the runtime of rumour-spreading algorithms.

One result in this spirit for the *push* model was shown in [28]. Roughly speaking, in that paper it is shown that even on graphs with low density, if the edges are distributed rather uniformly, then *push* is as fast as on the complete graph. This can be interpreted as a robustness result: starting with a complete graph, one can delete a vast amount of edges, and as long as this is done rather uniformly, the runtime of *push* is affected insignificantly. To state the result more precisely, we need the following notion.

Definition 1.1 ($(n, \delta, \Delta, \lambda)$ -graph). Let G be a connected graph with n vertices that has minimum degree δ and maximum degree Δ . Let $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$ be the eigenvalues of the adjacency matrix of G , and set

$$\lambda = \max_{2 \leq i \leq n} |\mu_i| = \max\{|\mu_2|, |\mu_n|\}.$$

We will call G an $(n, \delta, \Delta, \lambda)$ -graph.

In this paper we are interested in the case where G gets large, that is, when $n \rightarrow \infty$. Hence all asymptotic notation in this paper is with respect to n ; in particular, ‘with high probability’, or w.h.p. for short, means with probability $1 - o(1)$ when $n \rightarrow \infty$.

Definition 1.2 (expander sequence). Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be a sequence of graphs, where G_n is a $(n, \delta_n, \Delta_n, \lambda_n)$ -graph for each $n \in \mathbb{N}$. We say that \mathcal{G} is an *expander sequence* if $\Delta_n/\delta_n = 1 + o(1)$ and $\lambda_n = o(\Delta_n)$.

Note that if we consider any sequence $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ of graphs this always implicitly defines δ_n, Δ_n and λ_n as in Definition 1.2. Expander graphs have found numerous applications in computer science and mathematics; see *e.g.* the survey [25]. If \mathcal{G} is an expander sequence, then intuitively this means that for n large enough, the edges of G_n are rather uniformly distributed. For a more formal statement see Lemma 2.7. Moreover, note that our definition of expander sequences excludes the case when Δ_n is bounded. This is actually a necessary condition for our robustness results to hold; see [13]. With all these definitions at hand we can state the result from [28] that quantifies the robustness of *push* with respect to the network topology, that is, the runtime is asymptotically the same as on the complete graph K_n .

Theorem 1.1. *Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence. Then w.h.p.*

$$T_{\text{push}}(G_n) = c_{\text{push}}(1) \log n + o(\log n).$$

Apart from expander sequences, results in the form of Theorem 1.1 (where the asymptotic runtimes of one or more of these algorithms are determined) were also shown for sufficiently dense Erdős–Rényi random graphs [16], random regular graphs [15] as well as hypercubes [28]. Moreover, the order of the runtime on various models that describe social networks was investigated. The Chung–Lu model was studied in [17], preferential attachment graphs were explored in [10], and geometric graphs were examined in [18]. A somewhat different approach is to derive general runtime bounds that hold for all graphs and depend only on some graph parameter, *e.g.* conductance [6, 20], vertex expansion [21] or diameter [5, 14, 23]. Furthermore, several variants of *push*, *pull* and *push&pull* were studied. These include vertices being restricted to answer only one *pull* request per round [7], vertices being allowed to contact multiple neighbours per round [11, 28], vertices not calling the same neighbour twice [10] and asynchronous versions [1, 2, 4, 29]. Finally, besides [11], robustness of these rumour-spreading algorithms with respect to message transmission failures was also studied by Elsässer and Sauerwald in [13]. It was shown for any graph that if a message fails with probability $1 - p$, then the runtime of *push* increases at most by a factor of $6/p$.

In this work our focus is on three subjects concerning the robustness of rumour spreading. Our first (and not unexpected) result extends Theorem 1.1 to the runtimes of *pull* and *push&pull*. In particular, we show that none of the three protocols slows down or speeds up on graphs with good expansion properties compared to its runtime on the complete graph. This motivates us to investigate how severely a graph with good expansion properties has to be modified to increase the respective runtimes.

In our second contribution, which is also the main result and which differs from what was treated in previous works, we propose and study a novel approach to quantifying robustness. In particular, we investigate the impact of adversarial edge deletions, where we use the well-known concept of *local resilience*; see *e.g.* [8, 31]. To be specific, we explore up to which fraction of edges an adversary needs to be allowed to delete at each vertex to slow down the process by a significant amount of time, *i.e.* by $\Omega(\log n)$ rounds. Here we discover a surprising dichotomy in the following sense. On the one hand, we show that neither *pull* and *push&pull* can be slowed down by such

adversarial edge deletions – in essentially all but trivial cases, where the fraction is so large that the graph may become (almost) disconnected. On the other hand, we demonstrate that even a small number of edge deletions is sufficient to slow down *push* by $\Omega(\log n)$ rounds. In other words, we find that in contrast to *pull* and *push&pull*, the *push* protocol is not resilient to adversarial deletions and lacks (in this specific sense) the robustness of the other two protocols.

As our third subject, we generalize the previous results by additionally considering message transmission failures that occur independently with probability $1 - q \in [0, 1)$. On the positive side, we show that for arbitrary $q \in (0, 1]$, all three algorithms inform *almost* all vertices at least as fast as in an expander sequence in spite of adversarial edge deletions. However, if we want to inform all vertices, only *pull* is not slowed down by adversarial edge deletions for all values of q ; *push* can be slowed down as before, and *push&pull* is a mixed bag, in that for $q = 1$ it cannot be slowed down whereas for $q < 1$ it can. Furthermore, in general it is also possible to speed up *push&pull* by deleting edges, which is however not surprising as the star-graph deterministically finishes in at most two rounds.

Summarizing, this work enhances previous (robustness) results, particularly the ones concerning precise asymptotic runtimes and random transmission failures. Crucially, we introduce and study the concept of local resilience as a method to investigate robustness. However, apart from that, in this paper we develop two new general methods for the analysis of rumour-spreading algorithms.

- The most common approach in the current literature for the study of the runtime is to determine the expected number of newly informed vertices in one or more rounds and to show concentration, for example by bounding the variance. Achieving this, however, is often quite complex and makes laborious and lengthy technical arguments necessary. Here we use the theory of *self-bounding* functions (see Section 2), which allows us to cleanly upper-bound the variance by the *expected value*. The argument works for all three investigated algorithms and the bound is valid for all graphs. We are certain that this method will also facilitate future work on the analysis of rumour-spreading algorithms.
- Studying the robustness of the protocols is a challenging task, as the adversary (as described previously) has various options to modify the graph, for example by introducing a high variance in the degrees of the vertices; this turns out to be particularly problematic in the case of *push&pull*. Here we demonstrate that such types of irregularities can be handled universally by applying a powerful tool from a completely different area, namely extremal graph theory. In particular, we use Szemerédi's Regularity Lemma (see e.g. [30]), which allows us to partition the vertex set of a graph such that nearly all pairs of sets in the partition behave nearly like perfect regular bipartite graphs. This allows us to apply our methods on these regular pairs; eventually we obtain a linear recursion that can be solved by analysing the maximal eigenvalue of the underlying matrix.

1.1 Results

Our first result addresses the question about how fast rumours spread on expander graphs; in order to obtain a concise statement, the occurrence of independent message transmission failures is also considered.

Theorem 1.2. *Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence and let $q \in (0, 1]$. Then w.h.p.*

- $T_{push}(G_n, q) = c_{push}(q) \log n + o(\log(n)),$
- $T_{pull}(G_n, q) = c_{pull}(q) \log n + o(\log(n)),$
- $T_{pp}(G_n, q) = c_{pp}(q) \log n + o(\log(n)).$

The first statement is an extension of Theorem 1.1 and its proof is a straightforward adaptation of the proof in [28]. We omit it. The contribution here is the proof of (b) and (c). Next we consider the case with edge deletions in addition to the message transmission failures.

Theorem 1.3. *Let $0 < \varepsilon < 1/2, q \in (0, 1]$ and $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence. Let $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$ be such that each \tilde{G}_n is obtained by deleting edges of G_n such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of its edges. Then w.h.p.*

- (a) $T_{pull}(\tilde{G}_n, q) = c_{pull}(q) \log n + o(\log n)$,
- (b) $T_{pp}(\tilde{G}_n, 1) \leq c_{pp}(1) \log n + o(\log n)$, when additionally assuming that $\delta(G_n) \geq \alpha n$ for some constant $0 < \alpha \leq 1$.

This result demonstrates unconditionally the robustness of *pull*, and conditionally on $q = 1$ the robustness of *push&pull* on dense graphs, in the case of edge deletions, that is, the runtime is asymptotically the same as in the complete graph. Moreover, we even show that *push&pull* may profit from edge deletions in contrast to being slowed down; see Subsection 3.6 for an example. The proof of this result, especially the statement about *push&pull*, is rather involved, since the original graph may become quite irregular after the edge deletions. Here we use, among many other ingredients, the aforementioned decomposition of the graph given by Szemerédi’s Regularity Lemma.

Note that Theorem 1.3 does not consider *push* and *push&pull* (when $q \neq 1$) at all. Indeed, our next result states that in these cases the behaviour is rather different and that the algorithms may be slowed down.

Theorem 1.4. *Let $\varepsilon > 0$ and $q \in (0, 1]$. Then there is an expander sequence $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ and a sequence of graphs $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$ with the following properties. Each \tilde{G}_n is obtained by deleting edges of G_n such that each vertex keeps at least a $(1 - \varepsilon)$ fraction of its edges. Moreover, w.h.p.*

- (a) $T_{push}(\tilde{G}_n, q) \geq c_{push}(q) \log n + \varepsilon/(2q) \log n + o(\log n)$,
- (b) $T_{pp}(\tilde{G}_n, q) \geq c_{pp}(q) \log n + (\varepsilon/(8q) - \varepsilon q^3/5) \log n + o(\log n)$.

Nevertheless, not all hope is lost. On the positive side, the next result states that *push* and *push&pull* are able to inform almost all vertices as fast as on the complete graph in spite of adversarial edge deletions. In this sense, we obtain an almost-robustness result for these cases.

Theorem 1.5. *Let $0 < \varepsilon < 1/2, q \in (0, 1]$ and $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence. Let $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$ be such that each \tilde{G}_n is obtained by deleting edges of G_n such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of its edges. For $\mathcal{P} \in \{push, pp\}$, let $\tilde{T}_{\mathcal{P}}$ denote the number of rounds needed to inform at least $n - n/\log n$ vertices. Then w.h.p.*

- (a) $\tilde{T}_{push}(\tilde{G}_n, q) = \log_{1+q}(n) + o(\log n)$,
- (b) $\tilde{T}_{pp}(\tilde{G}_n, q) \leq \log_{1+2q}(n) + o(\log n)$, when additionally assuming that $\delta(G_n) \geq \alpha n$ for some $0 < \alpha \leq 1$.

We conjecture that there is also a version of Theorem 1.5(b) that is true for *push&pull* on sparse graphs; to be precise we conjecture that in the setting of Theorem 1.5(b), $\tilde{T}_{pp}(\tilde{G}_n) \leq \log_{1+2q}(n) + o(\log n)$, without further restrictions on G_n , that is, *push&pull* cannot be slowed down, informing almost all vertices.

As a final remark, note that Theorems 1.3 and 1.5 are tight in the sense that if an adversary may delete up to half of the edges at each vertex, then there are expander graphs that become

disconnected such that their components have linear size. On those graphs a linear fraction of the vertices will remain uninformed forever.

Outline. The rest of this paper is structured as follows. In Section 2 we collect and prove several important facts; this part of the paper also contains our technical contribution concerning the analysis through self-bounding functions. In Section 3.1 we show that *pull* is as fast on expanders with (or without) deleted edges as it is on the complete graph. Section 3.2 treats *push&pull* on expanders without deleted edges. In the remaining subsections we focus on the cases that may be slowed down by edge deletions. In Section 3.3 we show that adversarial edge deletions cannot slow down the time until *push* has informed almost all vertices, by giving a coupling to the case without edge deletions. Conversely, in Section 3.4 we show that the time until *push* has informed all vertices can be slowed down by edge deletions, even if only a few edges are deleted. Then in Section 3.5 we show that *push&pull* informs almost all vertices of dense graphs fast in spite of adversarial edge deletions. We utilize a version of Szemerédi’s Regularity Lemma to get a well-behaved partition of the vertex set that is suitable for performing a round-based analysis. However, if $q < 1$, adversarial edge deletions can slow down or speed up the time until *push&pull* has informed all vertices for nearly all values of q ; we show this in Section 3.6.

Further notation. Let $G = (V, E)$ denote a graph with vertex set V and edge set $E \subseteq \binom{V}{2}$. We will denote the set of neighbours of any vertex $v \in V$ by $N_G(v)$ or by $N(v)$, and we will denote its degree by $d_G(v) := |N_G(v)|$ or by $d(v)$; δ_G or δ and Δ_G or Δ denote the minimum and maximum degree of G . Similarly the neighbourhood of any set of vertices $S \subseteq V$ is defined by $N_G(S) := \cup_{v \in S} N_G(v)$. Furthermore, let $U, W \subseteq V$ with $U \cap W = \emptyset$ be two disjoint vertex sets; then $E(U, W) = E_G(U, W)$ denotes the set of edges with one vertex in U and one vertex in W and let $e(U, W) = e_G(U, W) := |E_G(U, W)|$. With $E_G(U)$ we denote the set of edges with both vertices in U ; $e_G(U) := |E_G(U)|$. For any round $t \in \mathbb{N}$ and $\mathcal{P} \in \{\text{push}, \text{pull}, \text{pp}\}$, we let $I_t^{(\mathcal{P})}(G)$ denote the set of vertices of G informed by *push*, *pull* and *push&pull* respectively at the beginning of round t and $|I_1^{(\mathcal{P})}| = 1$; if the underlying graph is clear from the context we will omit it; if we consider a sequence of graphs $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ and a sequence of times $t = (t(n))_{n \in \mathbb{N}}$, then

$$I_t^{(\mathcal{P})}(\mathcal{G}) = (I_{t(n)}^{(\mathcal{P})}(G_n))_{n \in \mathbb{N}}$$

is also a sequence. Similarly, $U_t^{(\mathcal{P})} := V \setminus I_t^{(\mathcal{P})}$ denotes the set of uninformed vertices. By \log we refer to the natural logarithm. For any event A we will write $\mathbb{E}_t[A]$ instead of $\mathbb{E}[A|I_t]$ for the conditional expectation and $\mathbb{P}_t[A]$ instead of $\mathbb{P}[A|I_t]$ for the conditional probability. Finally we want to clarify our use of Landau symbols. Let $a, b \in \mathbb{R}$ and f be a function. The terms $a \leq b + o(f)$ and $a \geq b - o(f)$ mean that there exist positive functions $g, h \in o(f)$ such that $a \leq b + g$ and $a \geq b - h$. Consequently $a = b + o(f)$ means that there exists a positive function $g \in o(f)$ such that $a \in [b - g, b + g]$.

2. Tools and techniques

In this section we collect and prove statements about our protocols and properties of expander sequences. We begin by applying the previously mentioned notion of self-bounding functions to derive universal and simple-to-apply concentration results for our random variables, *i.e.* the number of informed vertices after a particular round. Then we extend the concentration results to more than one round. In the last part we recall the well-known Expander Mixing Lemma and utilize it to derive properties (weak expansion, path enumeration) for the case where we delete edges from our graphs.

Self-bounding functions. Our main technical new result in this section is the following bound on the variance for the number of informed vertices in any given round; it is true for any graph and any set of informed vertices.

Lemma 2.1. *Let G be a graph, $t \in \mathbb{N}$ and $I_t = I_t^{(\mathcal{P})}(G)$ for $\mathcal{P} \in \{\text{push}, \text{pull}, \text{pp}\}$. Then*

$$\text{Var}[|I_{t+1}| \mid I_t] \leq \mathbb{E}[|I_{t+1}| \mid I_t].$$

Lemma 2.1 follows directly from Lemmas 2.2 and 2.3. Before stating them we introduce the notion of self-bounding functions.

Definition 2.1 (self-bounding function). Let X be a set and $m \in \mathbb{N}$. A non-negative function $f: X^m \rightarrow \mathbb{R}$ is self-bounding if there exist functions $f_i: X^{m-1} \rightarrow \mathbb{R}$ such that, for all $x_1, \dots, x_m \in X$ and all $i = 1, \dots, m$,

$$0 \leq f(x_1, \dots, x_m) - f_i(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_m) \leq 1$$

and

$$\sum_{1 \leq i \leq m} (f(x_1, \dots, x_m) - f_i(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_m)) \leq f(x_1, \dots, x_m).$$

A striking property of self-bounding functions is the following bound on the variance.

Lemma 2.2 ([3]). *For a self-bounding function f and independent random variables X_1, \dots, X_m , $m \in \mathbb{N}$,*

$$\text{Var}[f(X_1, \dots, X_m)] \leq \mathbb{E}[f(X_1, \dots, X_m)].$$

Lemma 2.3. *Let G be a graph, $t \in \mathbb{N}$, and let $I_t = I_t^{(\mathcal{P})}(G)$ for $\mathcal{P} \in \{\text{push}, \text{pull}, \text{pp}\}$. Then, conditional on I_t , there exists $m \in \mathbb{N}$, independent random variables X_1, \dots, X_m and a self-bounding function $f = f^{(\mathcal{P})}$ such that $|I_{t+1}| = f(X_1, \dots, X_m)$.*

Proof. We will prove in detail the result for *push*, and then we show what needs to be modified in order to obtain the statement in the case of *pull* and *push&pull*. Let $I_t = I_t^{(\text{push})}$, $n \in \mathbb{N}$ be the number of vertices of G , i.e. $V = [n]$, and $f: [n]^{|I_t|} \rightarrow \mathbb{R}$ with

$$(x_1, \dots, x_{|I_t|}) \mapsto |I_t| + \sum_{1 \leq k \leq |I_t|} \mathbb{1}[x_k \in U_t] \mathbb{1}[\forall \ell < k: x_k \neq x_\ell].$$

Moreover, let $(X_i)_{1 \leq i \leq |I_t|}$ be independent random variables, where X_i is a uniformly random neighbour of the i th vertex – according to an arbitrary ordering – in I_t . We argue that $f(X_1, \dots, X_{|I_t|}) = |I_{t+1}|$. Consider $v \in I_t$; then v is counted by the $|I_t|$ term in f . For $v \in I_{t+1} \setminus I_t$, let $v_1, \dots, v_s \in I_t$, $s \in \mathbb{N}$ be the informed vertices with random neighbour v in round t , i.e. $X_{v_1} = \dots = X_{v_s} = v$ and $X_u \neq v$ for all other $u \in I_t$. Assume further that $v_1 < v_2 < \dots < v_s$. For $k = v_1$ the term $\mathbb{1}[X_k \in U_t] \mathbb{1}[\forall \ell < k: x_k \neq x_\ell] = 1$ as $X_{v_1} = v \in U_t$ and for all $i \leq v_1$ it holds that $X_i \neq X_{v_1}$. For $k = v_r$, $2 \leq r \leq s$ the term $\mathbb{1}[\forall \ell < k: x_k \neq x_\ell] = 0$ as $v_1 < v_r$ and $X_{v_1} = X_{v_r} = v$. Thus every vertex $v \in I_{t+1} \setminus I_t$ is counted exactly once by f . Further, set

$$f_i(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{|I_t|}) = |I_t| + \sum_{k=1, k \neq i}^{|I_t|} \mathbb{1}[x_k \in U_t] \mathbb{1}[\forall j < k, j \neq i: x_j \neq x_k], \quad 1 \leq i \leq |I_t|.$$

The function f_i arises from f by leaving the i th variable out of consideration, that is, the push of the i th vertex has no effect. Then by definition $f - f_i \in \{0, 1\}$ for all $1 \leq i \leq |I_t|$, and in fact we have

$$f - f_i = \mathbb{1}[x_i \in U_t] \mathbb{1}[\forall j \neq i: x_i \neq x_j].$$

This quantity is precisely the difference in informed vertices after round t , assuming the i th vertex did not push. Furthermore

$$\sum_{1 \leq i \leq |I_t|} (f - f_i) \leq \sum_{1 \leq i \leq |I_t|} \mathbb{1}[x_i \in U_t] \mathbb{1}[\forall j \neq i: x_i \neq x_j] \leq f.$$

Thus f has the self-bounding property, which establishes the claim in the case of *push*. The proof for *pull* is completely analogous, where we use

$$f^{(pull)} : [n]^{|U_t|} \rightarrow \mathbb{R}, (x_1, \dots, x_{|U_t|}) \mapsto |I_t| + \sum_{k \in U_t} \mathbb{1}[x_k \in I_t]$$

and, similarly, for *push&pull* we use $f^{(pp)} : [n]^n \rightarrow \mathbb{R}$ with

$$\begin{aligned} (x_1, \dots, x_n) \mapsto & |I_t| + \sum_{1 \leq k \leq n} \mathbb{1}[k \in I_t] \mathbb{1}[x_k \in U_t] \mathbb{1}[\forall j \in \{1, \dots, k\} \cap I_t: x_k \neq x_j] \\ & + \sum_{1 \leq k \leq n} \mathbb{1}[k \in U_t] \mathbb{1}[x_k \in I_t] \mathbb{1}[\forall w \in I_t: x_w \neq k]. \end{aligned}$$

Here it is useful to see that the two sums in $f^{(pp)}$ are complementary, that is, only one of the summands for index k can be 1. Thus the functions $f_i^{(pull)}$ and $f_i^{(pp)}$ are obtained analogously to the push case. □

Remark 2.1. Let $G = (V, E)$ be a graph. Lemma 2.3 also applies to subsets of I_{t+1} , that is, for any $U \subset V$ and conditioned on I_t we have that $|I_{t+1} \cap U|$ and $|(I_{t+1} \cap U) \setminus I_t|$ are self-bounding.

The following proposition gives a tool that we will use in order to extend our round-wise analysis to longer phases.

Proposition 2.4. *Let $(\mathcal{A}_i)_{i \in \mathbb{N}_0}$ be a sequence of events, $0 < c < 1$, $\delta > 0$ and $t_1 \geq t_0 \geq 1$, such that*

$$\mathbb{P}[\mathcal{A}_t \mid \mathcal{A}_{t_0}, \dots, \mathcal{A}_{t-1}, \mathcal{A}_0] \geq 1 - c^{t-t_0} \delta \quad \text{for all } t_0 \leq t \leq t_1.$$

Then

$$\mathbb{P}\left[\bigcap_{t=t_0}^{t_1} \mathcal{A}_t \mid \mathcal{A}_0\right] \geq 1 - \delta/(1 - c).$$

Proof. Using the definition of conditional probability we obtain, as $c < 1$,

$$\begin{aligned} \mathbb{P}\left[\bigcap_{t=t_0}^{t_1} \mathcal{A}_t \mid \mathcal{A}_0\right] &= \prod_{t=t_0}^{t_1} \mathbb{P}[\mathcal{A}_t \mid \mathcal{A}_{t_0}, \dots, \mathcal{A}_{t-1}, \mathcal{A}_0] \\ &\geq \prod_{t=t_0}^{t_1} (1 - c^{t-t_0} \delta) \\ &\geq 1 - \sum_{t=t_0}^{t_1} (c^{t-t_0} \delta) \end{aligned}$$

$$\begin{aligned}
 &= 1 - \delta \sum_{t=0}^{t_1-t_0} c^t \\
 &\geq 1 - \delta/(1 - c). \quad \square
 \end{aligned}$$

We give two typical applications of the previous lemmas, similar to what we will encounter several times later in the paper. The first lemma addresses the case where we have a lower bound for the expected number of informed vertices after one round.

Lemma 2.5. *Let $\mathcal{P} \in \{\text{push, pull, pp}\}$ and $I_t = I_t^{(\mathcal{P})}$. Assume that there is $c > 1$ such that $\mathbb{E}_t[|I_{t+1}|] \geq c|I_t|$ for all t as long as $n/f(n) \leq |I_t| \leq n/g(n)$ for some functions $1 \leq f(n) \leq g(n) \leq n, f = o(n)$. Assuming $|I_{t_0}| \geq n/f(n)$, then there is $\tau = \log_c(f(n)/g(n)) + o(\log n)$ such that w.h.p.*

$$|I_{t_0+\tau}| \geq n/g(n).$$

Proof. Let $t \geq t_0$ and $n/f(n) \leq |I_t| \leq n/g(n)$. Lemma 2.1 guarantees that $\text{Var}_t[|I_{t+1}|] \leq \mathbb{E}_t[|I_{t+1}|]$, and applying Chebyshev’s inequality gives

$$\mathbb{P}_t[| |I_{t+1}| - \mathbb{E}_t[|I_{t+1}|] | \leq \mathbb{E}_t[|I_{t+1}|]^{2/3}] \geq 1 - \mathbb{E}_t[|I_{t+1}|]^{-1/3} \geq 1 - |I_t|^{-1/3}. \quad (2.1)$$

Consider the events

$$\mathcal{A}_t = |I_t| \geq \mathbb{E}_{t-1}[|I_t|] - \mathbb{E}_{t-1}[|I_t|]^{2/3} \quad \text{or} \quad |I_t| \geq n/g(n).$$

The intersection of $\mathcal{A}_{t_0+1}, \dots, \mathcal{A}_t$ implies inductively that either $|I_t| \geq n/g(n)$ or

$$\begin{aligned}
 |I_t| &\geq (1 - \mathbb{E}_{t-1}[|I_t|]^{-1/3})\mathbb{E}_{t-1}[|I_t|] \\
 &\geq (1 - (c|I_{t-1}|)^{-1/3})c|I_{t-1}| \\
 &\geq ((1 - (c|I_{t_0}|)^{-1/3})c)^{t-t_0}|I_{t_0}|. \quad (2.2)
 \end{aligned}$$

We obtain with (2.1)

$$\mathbb{P}_{t_0}[\mathcal{A}_{t+1} \mid \mathcal{A}_{t_0+1}, \dots, \mathcal{A}_t, |I_t| < n/g(n)] \geq 1 - ((1 - (c|I_{t_0}|)^{-1/3})c)^{-(t-t_0)/3}|I_{t_0}|^{-1/3},$$

and otherwise

$$\mathbb{P}_{t_0}[\mathcal{A}_{t+1} \mid \mathcal{A}_{t_0+1}, \dots, \mathcal{A}_t, |I_t| \geq n/g(n)] = 1.$$

Choose

$$\tau := t - t_0 = \log_c(f(n)/g(n)) + o(\log n)$$

as small as possible such that the lower bound for $|I_{t+1}|$ in (2.2) is $\geq n/g(n)$, that is, the lower bound in (2.2) is $< n/g(n)$ for $t = t_0 + \tau$. Combining the two conditional probabilities we obtain for all $t_0 \leq t \leq t_0 + \tau$

$$\mathbb{P}_{I_0}[\mathcal{A}_{t+1} \mid \mathcal{A}_{t_0+1}, \dots, \mathcal{A}_t] \geq 1 - ((1 - (c|I_{t_0}|)^{-1/3})c)^{-(t-t_0)/3} |I_{t_0}|^{-1/3}.$$

Applying Proposition 2.4 then yields the claim. □

In the second lemma we make the stronger assumption that we can determine asymptotically the expected number of informed vertices after one round. Here we assume that we begin with a ‘small’ set of informed vertices, say of size $\sqrt{\log n}$, and want to reach a set of size nearly linear in n .

Lemma 2.6. *Assume that there is some $c > 1$ such that $\mathbb{E}_t[|I_{t+1}|] = (1 + o(1))c|I_t|$ for all t as long as $\sqrt{\log n} \leq |I_t| \leq n/\log n$. Assume furthermore that $|I_{t_0}| \geq \sqrt{\log n}$. Then there are $\tau_1, \tau_2 = \log_c(n/|I_{t_0}|) + o(\log n)$ such that w.h.p.*

$$|I_{t_0+\tau_1}| \leq \frac{n}{\log n} \leq |I_{t_0+\tau_2}|.$$

Proof. Lemma 2.5, setting $f = n/\sqrt{\log n}$ and $g = \log n$ directly implies the existence of τ_1 . To find τ_2 , let \mathcal{A}_t be the event $||I_t| - \mathbb{E}_{t-1}[|I_t|]| \leq \mathbb{E}_{t-1}[|I_t|]^{2/3}$. There is $h(n) \in o(1)$ such that, for $c^- := (1 - h(n))c$ and $c^+ := (1 + h(n))c$, we have that $\mathbb{E}_t[|I_{t+1}|] \leq c^+|I_t|$ and $\mathbb{E}_t[|I_{t+1}|] \geq c^-|I_t|$. Using this notation, the events $\mathcal{A}_{t_0+1}, \dots, \mathcal{A}_{t+1}$ together imply inductively that

$$\begin{aligned} |I_{t+1}| &\leq (1 + \mathbb{E}_t[|I_{t+1}|]^{-1/3})\mathbb{E}_t[|I_{t+1}|] \\ &\leq (1 + (c^-|I_t|)^{-1/3})c^+|I_t| \\ &\leq ((1 + (c^-|I_{t_0}|)^{-1/3})c^+)^{t-t_0}|I_{t_0}| \end{aligned}$$

for all t such that the right-hand side is bounded by $n/\log n$. Moreover, for all such t ,

$$\begin{aligned} |I_{t+1}| &\geq (1 - \mathbb{E}_t[|I_{t+1}|]^{-1/3})\mathbb{E}_t[|I_{t+1}|] \\ &\geq (1 - (c^-|I_t|)^{-1/3})c^-|I_t| \\ &\geq ((1 - (c^-|I_{t_0}|)^{-1/3})c^-)^{t-t_0}|I_{t_0}|. \end{aligned}$$

Thus, as \mathcal{A}_t only depends on I_t , it follows with (2.1) that

$$\mathbb{P}_{I_0}[\mathcal{A}_{t+1} \mid \mathcal{A}_{t_0+1}, \dots, \mathcal{A}_t] \geq 1 - ((1 - (c^-|I_{t_0}|)^{-1/3})c^-)^{-(t-t_0)/3} |I_{t_0}|^{-1/3}.$$

Applying Proposition 2.4 yields the existence of τ_2 . □

Expander sequences. In this section we collect some important properties of expander sequences that we are going to use later. We start by stating a version of the well-known Expander Mixing Lemma applied to our setting of expander sequences.

Lemma 2.7 ([28, Corollary 2.4]). *Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}} = ((V_n, E_n))_{n \in \mathbb{N}}$ be an expander sequence. Then, for $S_n \subseteq V_n$ such that $1 \leq |S_n| \leq n/2$, it is*

$$\left| e(S_n, V_n \setminus S_n) - \frac{\Delta_n |S_n| (n - |S_n|)}{n} \right| = o(\Delta_n |S_n|).$$

The following result is a consequence of the Expander Mixing Lemma that applies to graphs in which some edges were removed. It seems very simple but it turns out to be surprisingly useful.

Lemma 2.8. *Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}} = ((V_n, E_n))_{n \in \mathbb{N}}$ be an expander sequence. Let $\varepsilon > 0$ and set $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$, where each \tilde{G}_n is obtained from G_n by deleting edges such that each vertex keeps at least*

a $(1/2 + \varepsilon)$ fraction of its edges. For each $n \in \mathbb{N}$ let $S_n \subseteq V_n$. Then there is $n_0 \in \mathbb{N}$ such that, for all $n \geq n_0$,

$$e_{\tilde{G}_n}(S_n, V_n \setminus S_n) \geq \varepsilon e_{G_n}(S_n, V_n \setminus S_n).$$

Proof. Without loss of generality we assume that $|S_n| \leq n/2$. Since at most $(1/2 - \varepsilon)\Delta_n$ edges are deleted at each vertex, we immediately obtain that

$$e_{\tilde{G}_n}(S_n, V_n \setminus S_n) \geq e_{G_n}(S_n, V_n \setminus S_n) - \Delta_n(1/2 - \varepsilon)|S_n|.$$

Using Lemma 2.7 and choosing n_0 large enough such that

$$\frac{o(\Delta_n)}{\Delta_n} \frac{n}{n - |S_n|} < \varepsilon \quad \text{for all } n \geq n_0,$$

we obtain that

$$\begin{aligned} & (1 - \varepsilon)e_{G_n}(S_n, V_n \setminus S_n) - \Delta_n(1/2 - \varepsilon)|S_n| \\ & \geq (1 - \varepsilon) \frac{\Delta_n |S_n| (n - |S_n|)}{n} - o(\Delta_n)|S_n| - \Delta_n(1/2 - \varepsilon)|S_n| \\ & = \frac{\Delta_n |S_n| (n - |S_n|)}{n} \left(1 - \varepsilon - \frac{o(\Delta_n)}{\Delta_n} \frac{n}{n - |S_n|} - \frac{n(1/2 - \varepsilon)}{n - |S_n|} \right). \end{aligned}$$

As $n - |S_n| \geq n/2$, the last expression is > 0 . Hence

$$\begin{aligned} e_{\tilde{G}_n}(S_n, V_n \setminus S_n) & \geq \varepsilon e_{G_n}(S_n, V_n \setminus S_n) + (1 - \varepsilon)e_{G_n}(S_n, V_n \setminus S_n) - \Delta_n(1/2 - \varepsilon)|S_n| \\ & \geq \varepsilon e_{G_n}(S_n, V_n \setminus S_n). \quad \square \end{aligned}$$

Next we give a lemma that counts the number of paths between two arbitrary vertices of a dense graph satisfying a weak expander property (as for example guaranteed by Lemma 2.8). This will later be used to give a lower bound on the probability of any vertex being informed after a given constant number of rounds.

Lemma 2.9. *Let $G = (V, E)$, $|V| = n$. Assume that there is $\alpha > 0$ such that $d(v) \geq \alpha n$ for all $v \in V$ and $e(W, V \setminus W) \geq \alpha |W| |V \setminus W|$ for all $W \subseteq V$. Then, for all $u, w \in V$, there is $1 \leq d \leq 8/\alpha^2 + 2$ such that there are at least $(\alpha^4/64)^{d+1} n^{d-1}$ paths of length d from u to w .*

Proof. Assume $\alpha \leq 1/2$, as otherwise the claim is trivial (with $d \in \{1, 2\}$). We define sequences $(S_i)_{i \in \mathbb{N}}$ and $(H_i)_{i \in \mathbb{N}} \subseteq V$ as follows. Set $S_1 = \{u\} \cup N(u)$, $W = \{w\} \cup N(w)$ and $H_1 = V \setminus (S_1 \cup W)$ and proceed for $i \geq 1$ as follows. Let $\tilde{S}_{i+1} \subseteq H_i$ be the set of vertices $v \in H_i$ with $|N(v) \cap S_i| \geq \alpha^2 n/8$. Set $S_{i+1} = S_i \cup \tilde{S}_{i+1}$ and $H_{i+1} = H_i \setminus \tilde{S}_{i+1}$. Then we claim that, for all $i \geq 1$,

$$e(S_i, W) \geq \alpha^3 n^2/2 \quad \text{or} \quad |S_{i+1}| \geq |S_i| + \alpha^2 n/8. \tag{2.3}$$

To see this, assume that $e(S_i, W) \leq \alpha^3 n^2/2$. Since $|S_i|, |W| \geq \alpha n$, the weak expansion property guarantees that

$$e(S_i, H_i) = e(S_i, H_i \cup W) - e(S_i, W) \geq \alpha |S_i| |H_i \cup W| - \alpha^3 n^2/2 \geq \alpha^2 (1 - \alpha) n^2 - \alpha^3 n^2/2,$$

and using $\alpha \leq 1/2$ we obtain that $e(S_i, H_i) \geq \alpha^2 n^2/4$. To complete the proof of (2.3) we compute the size of \tilde{S}_{i+1} . As $|N(v) \cap S_i| \leq \alpha^2 n/8$ for all $v \in H_i \setminus \tilde{S}_{i+1}$ and $|N(v) \cap S_i| \leq n$, we get

$$\frac{\alpha^2 n^2}{4} \leq e(S_i, H_i) \leq |\tilde{S}_{i+1}| n + |H_i| \frac{\alpha^2 n}{8}.$$

Since $|H_i| \leq n$ we immediately get that $|\tilde{S}_{i+1}| \geq \alpha^2 n/8$, which shows (2.3). We next show that there are (sufficiently) many paths for each vertex in S_i to u . More precisely, let $1 \leq j \leq 8/\alpha^2$ be such that $e(S_i, W) < \alpha^3 n^2/2$ for all $1 \leq i \leq j$. For those i we have by (2.3) that $|S_i| \geq i \cdot \alpha^2 n/8$. We claim that for all $v \in S_i \setminus \{u\}$ there is $d \leq i$ such that v has at least $(\alpha^4/64)^d \cdot n^{d-1}$ paths of length d with endpoint u . We show the claim by induction on i . The base case $v \in S_1 \setminus \{u\}$ is clear, as $1 \geq \alpha^4/64$. For the induction step assume that $v \in S_{i+1} \setminus S_i$, $v \neq u$. Then by construction $|N(v) \cap S_i| \geq \alpha^2 n/8$. Thus, by induction hypothesis, there is $d \leq i$ such that v has at least $\alpha^2 n/(8i)$ neighbours with at least $(\alpha^4/64)^d n^{d-1}$ paths with endpoint u . As $i \leq 8/\alpha^2$ this gives that v has at least $\alpha^2 n/(8i) \cdot (\alpha^4/64)^d n^{d-1} \geq (\alpha^4/64)^{d+1} n^d$ paths of length $d+1 \leq i+1$ with endpoint u , and this accomplishes the induction step. With all these facts at hand we finally show the claim of the lemma. Let $j \leq 8/\alpha^2$ be the first index such that $e(S_j, W) \geq \alpha^3 n^2/2$, and let $W' \subseteq W$ be such that $|N(v) \cap S_j| \geq \alpha^3 n/4$ for all $v \in W'$. Thus

$$\frac{\alpha^3 n^2}{2} \leq e(S_j, W) \leq |W'|n + |W| \frac{\alpha^3 n}{4},$$

and thus $|W'| \geq \alpha^3 n/4$. Then there is $d \leq j$ and $W'' \subseteq W'$ such that $|W''| \geq |W'|/j$ and every v in W'' has at least $\alpha^3 n/(4j)$ neighbours with at least $(\alpha^4/64)^d n^{d-1}$ paths of length d with endpoint u . Therefore every $v \in W''$ has at least

$$(\alpha^4/64)^d n^{d-1} \cdot \alpha^3 n/(4j) \geq (\alpha^4/64)^{d+1} n^d/j$$

paths of length $d+1$ with endpoint u . This in turn gives that there are at least

$$|W''|/j \cdot (\alpha^4/64)^{d+1} n^d/j \geq \alpha^3/4 \cdot (\alpha^4/64)^{d+2} n^{d+1}$$

paths of length $d+2$ from w to u , and the proof is completed. □

Next comes a technical lemma that, given a small set, quantifies the number of vertices for which only a small fraction of their neighbourhood intersects that given set.

Lemma 2.10. *Let $\mathcal{G} = (G_n)_{n \in \mathbb{N}} = ((V_n, E_n))_{n \in \mathbb{N}}$ be an expander sequence. Let $\varepsilon > 0$ and let $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$, where each \tilde{G}_n it is obtained from G_n by deleting edges such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of its edges. Let $A_n \subseteq V_n$ with $|A_n| = o(n)$.*

(a) *There is $B_n \subseteq A_n$ with $|B_n| = (1 - o(1))|A_n|$ such that, for all $u \in B_n$,*

$$\frac{|N_{\tilde{G}_n}(u) \cap A_n|}{|N_{\tilde{G}_n}(u)|} = o(1).$$

(b) *There is $B_n \subseteq V_n \setminus A_n$ with $|V_n \setminus (A_n \cup B_n)| = o(|A_n|)$ such that, for all $v \in B_n$,*

$$\frac{|N_{\tilde{G}_n}(v) \cap A_n|}{|N_{\tilde{G}_n}(v)|} = o(1).$$

Proof. Let δ_n, Δ_n denote the minimum and maximum degree of G_n . Lemma 2.7 yields that

$$e_{G_n}(A_n, V_n \setminus A_n) = \frac{\Delta_n |A_n| |V_n \setminus A_n|}{n} + o(\Delta_n) |A_n| = (1 + o(1)) \Delta_n |A_n|.$$

As there are a maximum of $\Delta_n|A_n|$ edges with at least one point in A_n , we get that $e_{G_n}(A_n) = o(\Delta_n)|A_n|$. Since we obtain \tilde{G}_n from G_n by deleting edges,

$$e_{\tilde{G}_n}(A_n) = o(\Delta_n)|A_n|. \tag{2.4}$$

With this fact at hand we show (a). Let $\eta > 0$ and call a vertex $u \in A_n$ bad if $|N_{\tilde{G}_n}(u) \cap A_n| \geq \eta|N_{\tilde{G}_n}(u)|$. Since $N_{\tilde{G}_n}(u) \geq \delta_n/2$, we obtain for any bad u that $|N_{\tilde{G}_n}(u) \cap A_n| \geq \eta\delta_n/2$. As $\delta_n = (1 - o(1))\Delta_n$, we infer from (2.4) that the number of bad vertices is $o(|A_n|)$.

To see (b), again let $\eta > 0$ and this time call a vertex $v \in V_n \setminus A_n$ bad if $|N_{\tilde{G}_n}(v) \cap A_n| \geq \eta|N_{\tilde{G}_n}(v)|$. Then, for any such bad v , we know that $|N_{\tilde{G}_n}(v) \cap A_n| \geq \eta\delta_n/2$. As before, using (2.4) we readily get that the number of bad v is $o(|A_n|)$. \square

We conclude our preliminary section by giving a lemma that crudely bounds the time needed until at least $\omega(1)$ vertices are informed.

Lemma 2.11. *Let $0 < \varepsilon \leq 1/2, q \in (0, 1]$ and $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence. Let $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$ be such that each \tilde{G}_n is obtained by deleting edges of G_n such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of its edges. Let $\mathcal{P} \in \{\text{push, pull, pp}\}$ and suppose that $|I_t^{(\mathcal{P})}| < \sqrt{\log n}$. Then there is $\tau = o(\log n)$ such that w.h.p. $|I_{t+\tau}^{(\mathcal{P})}| \geq \sqrt{\log n}$.*

Proof. Recall that the probability that $v \in U_t$ gets informed by pull is $q|N(v) \cap I_t|/|N(v)|$. Thus

$$\mathbb{P}_t[|I_{t+1}^{(\text{pull})} \setminus I_t| = 0] = \prod_{u \in N(I_t) \cap U_t} \left(1 - \frac{q|N(u) \cap I_t|}{|N(u)|}\right) \leq e^{-qe(U_t, I_t)/\Delta_n}.$$

Similarly we obtain for push

$$\mathbb{P}_t[|I_{t+1}^{(\text{push})} \setminus I_t| = 0] = \prod_{v \in I_t} \frac{|N(v) \cap I_t|}{|N(v)|} = \prod_{v \in I_t} \left(1 - \frac{|N(v) \cap U_t|}{|N(v)|}\right) \leq e^{-qe(I_t, U_t)/\Delta_n}.$$

The same bound is obviously also true for push&pull. Thus, for all $\mathcal{P} \in \{\text{push, pull, pp}\}$,

$$\mathbb{P}_t[|I_{t+1}^{(\mathcal{P})} \setminus I_t| \geq 1] \geq 1 - e^{-qe(U_t, I_t)/\Delta_n}.$$

As Lemmas 2.7 and 2.8 imply that $e(U_t, I_t) \geq (1 + o(1))\varepsilon\Delta_n|I_t|$, there is $c \in (0, 1)$ such that $\mathbb{P}[|I_{t+1}^{(\mathcal{P})} \setminus I_t| \geq 1] > c$. Define $\tau := \lceil (2/c)\sqrt{\log n} \rceil$ and $X = \text{Bin}(\tau, c)$ with $\mathbb{E}[X] = c\tau$ and $\text{Var}[X] = \tau(1 - c)c$. Then, using Chebyshev,

$$\mathbb{P}_t[|I_{t+\tau}^{(\mathcal{P})}| \leq \sqrt{\log n}] \leq \mathbb{P}_t[X \leq \sqrt{\log n}] \leq \mathbb{P}_t[|X - \mathbb{E}[X]| \leq \mathbb{E}[X]/2] \leq 4\text{Var}[X]/\mathbb{E}[X]^2 = o(1). \quad \square$$

3. Proofs

3.1 Proof of Theorems 1.2(b) and 1.3(a) – edge deletions do not slow down pull

Let $0 < \varepsilon \leq 1/2$. In this section we study the runtime of pull in the case in which the input graph is an expander, and where at each vertex at most a $(1/2 - \varepsilon)$ fraction of the edges is deleted. The runtime on expander sequences without edge deletions, i.e. the setting in Theorem 1.2(b), is included as the special case where we set $\varepsilon = 1/2$. In contrast to previous proofs, in the analysis of pull the ‘standard’ approach that consists of showing, for example, that $\mathbb{E}_t[|I_{t+1} \setminus I_t|] \approx |I_t|$ fails. The main reason is that the graph between I_t and U_t might be quite irregular, so that, depending on the actual state, $\mathbb{E}_t[|I_{t+1} \setminus I_t|] \approx c|I_t|$ for some $c < 1$. However, we discover a different invariant that is preserved, namely that the number of edges between I_t and U_t behaves in an exponential way. With Lemmas 2.7 and 2.8 we can then relate this to the number of informed vertices.

Lemma 3.1. Consider the setting of Theorem 1.3(a) and let $I_t = I_t^{(pull)}$.

(a) Let $\sqrt{\log n} \leq |I_t| \leq n/\log n$. Then

$$|e(U_{t+1}, I_{t+1}) - (1 + q)e(U_t, I_t)| \leq |I_t|^{-1/3} e(U_t, I_t)$$

with probability at least $1 - O(|I_t|^{-1/3})$.

(b) Let $|U_t| \leq n/\log n$. Then $\mathbb{E}_t[|U_{t+1}|] = (1 - q + o(1))|U_t|$.

Proof. We start with (a). Let $D_t = e(U_{t+1}, I_{t+1}) - e(U_t, I_t)$ and for $u \in U_t$ let X_u be the random variable that indicates whether u gets informed in round $t + 1$. Then

$$\begin{aligned} \mathbb{E}_t[D_t] &= \sum_{u \in U_t} \sum_{v \in N(u) \cap U_t} \mathbb{E}_t[X_u(1 - X_v)] - \sum_{u \in U_t} \mathbb{E}_t[X_u] \cdot |N(u) \cap I_t| \\ &= \sum_{u \in U_t} q \frac{|N(u) \cap I_t|}{|N(u)|} \left(\left(\sum_{v \in N(u) \cap U_t} 1 - q \frac{|N(v) \cap I_t|}{|N(v)|} \right) - |N(u) \cap I_t| \right). \end{aligned}$$

The second sum is at most $|N(u)|$, so obviously $\mathbb{E}_t[D_t] \leq qe(U_t, I_t)$. To get a lower bound consider a largest set $\tilde{U} \subseteq U_t$ such that $|N(u) \cap I_t|/|N(u)| = o(1)$ for all $u \in \tilde{U}$. From Lemma 2.10(b) we obtain that $|U_t \setminus \tilde{U}| = o(|I_t|)$, and so

$$\mathbb{E}_t[D_t] \geq \sum_{u \in U_t} q|N(u) \cap I_t| \left(\left(\sum_{v \in N(u) \cap \tilde{U}} \frac{1}{|N(u)|} - o\left(\frac{1}{|N(u)|}\right) \right) - \frac{|N(u) \cap I_t|}{|N(u)|} \right).$$

Consider furthermore $\hat{U} \subseteq \tilde{U}$ such that $|N(u) \cap \tilde{U}|/|N(u)| = 1 - o(1)$ and thus also $|N(u) \cap I_t|/|N(u)| = o(1)$ for all $u \in \hat{U}$. Lemma 2.10(b) again yields that we can choose \hat{U} such that $|U_t \setminus \hat{U}| = o(|I_t|)$, and thus

$$\begin{aligned} \mathbb{E}_t[D_t] &\geq (1 - o(1)) \sum_{u \in \hat{U}} q|N(u) \cap I_t| \left(\frac{|N(u) \cap \tilde{U}|}{|N(u)|} - \frac{|N(u) \cap I_t|}{|N(u)|} \right) - \sum_{u \in U_t \setminus \hat{U}} |N(u) \cap I_t| \\ &\geq (q - o(1))e(U_t, I_t) - 2e(U_t \setminus \hat{U}, I_t). \end{aligned}$$

According to Lemmas 2.7 and 2.8 we have that $e(U_t, I_t) = \Theta(|I_t| \Delta_n)$. But

$$e(U_t \setminus \tilde{U}, I_t) \leq |U_t \setminus \tilde{U}| \Delta_n = o(|I_t| \Delta_n).$$

Thus $\mathbb{E}_t[e(U_{t+1}, I_{t+1})] = (1 + q - o(1))e(U_t, I_t)$. In the next step we bound the variance. For each edge e let X_e be the indicator random variable that denotes the events that $e \in E(U_{t+1}, I_{t+1})$. Thus

$$e(U_{t+1}, I_{t+1}) = \sum_{e \in E} X_e = \frac{1}{2} \sum_{u \in V} \sum_{v \in N(u)} X_{\{u,v\}}.$$

Using the fact that X_e and $X_{e'}$ are independent for all $e, e' \in E$ with $e \cap e' = \emptyset$,

$$\begin{aligned} \text{Var}[e(U_{t+1}, I_{t+1})] &= \text{Var} \left[\sum_{e \in E} X_e \right] \\ &= \sum_{e, e' \in E} \mathbb{E}[X_e X_{e'}] - \mathbb{E}[X_e] \mathbb{E}[X_{e'}] \\ &\leq \sum_{u \in V} \sum_{v, v' \in N(u)} \mathbb{E}[X_{\{u,v\}} X_{\{u,v'\}}] \end{aligned}$$

$$\begin{aligned} &\leq \Delta_n \sum_{u \in V} \sum_{v \in N(u)} \mathbb{E}[X_{\{u,v\}}] \\ &= 2\Delta_n \mathbb{E}[e(U_{t+1}, I_{t+1})]. \end{aligned}$$

Since

$$\mathbb{E}_t[e(U_{t+1}, I_{t+1})] = (1 + q - o(1))e(U_t, I_t) = \Theta(\Delta_n |I_t|),$$

by Lemmas 2.7 and 2.8 and

$$\text{Var}[e(U_{t+1}, I_{t+1})] \leq 2\Delta_n \mathbb{E}_t[e(U_{t+1}, I_{t+1})]$$

we immediately obtain for $|I_t| \geq \sqrt{\log n}$ with Chebyshev's inequality that

$$\mathbb{P}[|e(U_{t+1}, I_{t+1}) - \mathbb{E}_t[e(U_{t+1}, I_{t+1})]| \geq e(U_t, I_t)|I_t|^{-1/3}] \leq O(|I_t|^{-1/3}).$$

Next we show (b). We bound the expected number of uninformed vertices after one additional round. Lemma 2.10(a) asserts that there is a set $\tilde{U} \subseteq U_t$ such that $|\tilde{U}| = (1 - o(1))|U_t|$ and $|N(u) \cap I_t|/|N(u)| = 1 - o(1)$ for all $u \in \tilde{U}$. Thus

$$\begin{aligned} \mathbb{E}_t[|U_{t+1}|] &= \sum_{u \in U_t} 1 - q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &\leq |U_t| - q \sum_{u \in \tilde{U}} \frac{|N(u) \cap I_t|}{|N(u)|} \\ &= |U_t| - q(1 - o(1))|\tilde{U}| \\ &= (1 - q - o(1))|U_t|. \end{aligned}$$

As $|N(u) \cap I_t| \leq |N(u)|$ we also have

$$\mathbb{E}_t[|U_{t+1}|] = \sum_{u \in U_t} 1 - q \frac{|N(u) \cap I_t|}{|N(u)|} \geq \sum_{u \in U_t} (1 - q) = (1 - q)|U_t|. \quad \square$$

Lemmas 3.2 and 2.11 give lower bounds which, together with an upper bound provided by Lemma 3.3, imply Theorems 1.2(b) and 1.3(a).

Lemma 3.2. (upper bound in Theorem 1.3(a)). *Consider the setting of Theorem 1.3(a) and let $I_t = I_t^{(pull)}$. Then the following statements hold w.h.p.*

- (a) Let $\sqrt{\log n} \leq |I_t| \leq n/\log n$. Then there are $\tau_1, \tau_2 = \log_{1+q}(n/|I_t|) + o(\log n)$ such that $|I_{t+\tau_2}| < n/\log n < |I_{t+\tau_1}|$.
- (b) Let $n/\log n \leq |I_t| \leq n - n/\log n$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| > n - n/\log n$.
- (c) Let $|I_t| \geq n - n/\log n$.
 - (i) Case $q = 1$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| = n$.
 - (ii) Case $q \neq 1$. Then there is $\tau \leq -\log n/\log(1 - q) + o(\log n)$ such that $|I_{t+\tau}| = n$.

Proof. We start with (a). Let $|I_t| \in [\sqrt{\log n}, n/\log n]$. First note that any bound on $e(U_t, I_t)$ translates to a bound for $|I_t|$, as with Lemmas 2.7 and 2.8 we obtain

$$(1 - o(1))\varepsilon \Delta_n |I_t| \leq e(U_t, I_t) \leq \Delta_n |I_t|. \tag{3.1}$$

In particular, up to constant factors, $|I_t|$ is $e(U_t, I_t)/\Delta_n$ and vice versa. From Lemma 3.1(a) we obtain that $e(U_{t+1}, I_{t+1}) = (1 + q \pm |I_t|^{-1/3})e(U_t, I_t)$ with probability $1 - O(|I_t|^{-1/3})$. Proceeding as in Lemmas 2.5 and 2.6 and their proofs, where we replace the events

$$|I_t| \geq \mathbb{E}_{t-1}[|I_t|] - \mathbb{E}_{t-1}[|I_t|]^{2/3} \quad \text{or} \quad |I_t| \geq n/g(n)$$

and

$$||I_t| - \mathbb{E}_{t-1}[|I_t|]| \leq \mathbb{E}_{t-1}[|I_t|]^{2/3}$$

with

$$e(U_t, I_t) \geq (1 + q - |I_{t-1}|^{-1/3})e(U_{t-1}, I_{t-1}) \quad \text{or} \quad |I_t| \geq n/\log n$$

and

$$e(U_{t+1}, I_{t+1}) = (1 + q \pm |I_t|^{-1/3})e(U_t, I_t),$$

we obtain the statement.

We continue with (b). Consider first the case $|I_t| \in [n/\log n, n/2]$. Using Lemmas 2.7 and 2.8, i.e. $e(U_t, I_t) \geq \varepsilon|U_t| |I_t|\Delta_n/n + o(\Delta_n)|I_t|$, together with $|U_t| \geq n/2$ implies

$$\begin{aligned} \mathbb{E}_t[|I_{t+1} \setminus I_t|] &= \sum_{u \in U_t} q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &\geq \frac{q \cdot e(U_t, I_t)}{\Delta_n} \\ &\geq \frac{q\varepsilon|U_t| |I_t|\Delta_n/n + o(\Delta_n)|I_t|}{\Delta_n(1 + o(1))} \\ &\geq \left(\frac{q\varepsilon}{2} + o(1)\right)|I_t|. \end{aligned}$$

Applying Lemma 2.5, where we set $g = 2, f = \log n$ and $c = q\varepsilon/2 + o(1)$, we are finished with this part as well. Now let $|I_t| \in [n/2, n - n/\log n]$. We switch our focus to the set of uninformed vertices. Using again the fact that $e(U_t, I_t) \geq \varepsilon|U_t| |I_t|\Delta_n/n + o(\Delta_n)|U_t|$, we have

$$\begin{aligned} \mathbb{E}_t[|U_{t+1}|] &= \sum_{u \in U_t} 1 - q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &= \sum_{u \in U_t} 1 - q \frac{|N(u) \cap I_t|}{\Delta_n(1 + o(1))} \\ &= |U_t| - \frac{q \cdot e(U_t, I_t)}{\Delta(1 + o(1))} \\ &= |U_t| - \frac{q\varepsilon|U_t| |I_t|\Delta_n/n + o(\Delta_n)|U_t|}{\Delta_n(1 + o(1))} \\ &\leq \left(1 - \frac{q\varepsilon}{2} + o(1)\right)|U_t|. \end{aligned}$$

Inductively we obtain for any integer $\tau \geq 1$ the bound $\mathbb{E}_t[|U_{t+\tau}|] \leq (1 - q\varepsilon/2 + o(1))^\tau |U_t|$, and so for some $\tau := 2 \log \log n / \log(1/(1 - q\varepsilon/2 + o(1))) = o(\log n)$ we have

$$\mathbb{E}_t[|U_{t+\tau}|] \leq |U_t|/\log^2 n = o(n/\log n).$$

Hence, by Markov's inequality, $\mathbb{P}_t[|U_{t+\tau}| \geq n/\log n] = o(1)$.

In order to show (c), let $|I_t| \in [n - n/\log n, n]$. As for $q = 1$ the term $1 - q$ in Lemma 3.1(b) vanishes, we distinguish the cases $q = 1$ and $q \neq 1$. We start with $q = 1$. By induction, it follows that for any round $\tau > 0$ and suitable $f = o(1)$,

$$\mathbb{E}_t[|U_{t+\tau}|] \leq (f(n))^\tau |U_t|.$$

We choose $\tau = \log_{1/f(n)}(n) = o(\log n)$ as $1/f = \omega(1)$. Hence we obtain $\mathbb{E}_t[|U_{t+\tau}|] \leq |U_t|/n \leq 1/\log n$. Therefore we have $\mathbb{P}_t[|U_{t+\tau}| \geq 1] \leq o(1)$ by Markov's inequality. For $q \neq 1$ we have by induction, for any number of rounds $\tau \geq 1$,

$$\mathbb{E}_t[|U_{t+\tau}|] \leq (1 - q + o(1))^\tau |U_t|.$$

We choose

$$\tau = \log_{1/(1-q+o(1))}(n) = -\log n / \log(1 - q) + o(\log n).$$

Thus, using Markov's inequality, analogously to the case $q = 1$, we obtain the desired upper bound. \square

Note that for $q = 1$ this already implies Theorems 1.2(b) and 1.3(a). This leaves the case for $q \neq 1$.

Lemma 3.3. *Let $0 < \varepsilon \leq 1/2, q \in (0, 1]$ and $\mathcal{G} = (G_n)_{n \in \mathbb{N}}$ be an expander sequence. Let $\tilde{\mathcal{G}} = (\tilde{G}_n)_{n \in \mathbb{N}}$ be such that each \tilde{G}_n is obtained by deleting edges of G_n , such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of its edges and abbreviate $I_t = I_t^{(pull)}$. Let $q \in (0, 1)$ and $|I_t| \leq n/2$. Then, for $\tau = -\log n / \log(1 - q)$ and all $c < 1$, w.h.p. $|I_{t+c\tau}| < n$.*

Proof. We consider a modified process in which vertices have a higher chance of getting informed. In particular, note that the probability that $u \in U_t$ gets informed is at most $q|N(u) \cap I_t|/|N(u)| \leq q$ and that all these events are independent; now we assume that each such u gets independently informed with probability exactly q . Then the runtime in this modified model constitutes a lower bound for the runtime in the original model.

Let $c < 1, u \in U_t$ and E_u be the event that u does not get informed in $c\tau$ rounds in this model. Thus

$$\mathbb{P}[E_u] = (1 - q)^{c\tau} = (1 - q)^{-c \log n / \log(1 - q)} = n^{-c} = \omega(1/n),$$

and as the events E_u are independent and $|U_t| = \Theta(n)$,

$$\mathbb{P}\left[\bigwedge_{u \in U_t} \overline{E_u}\right] \leq \prod_{u \in U_t} \mathbb{P}[\overline{E_u}] \leq \exp\left(-\sum_{u \in U_t} \mathbb{P}[E_u]\right) = o(1). \quad \square$$

3.2 Proof of Theorem 1.2(c) – push&pull is fast on expanders

As we are now in the case without edge deletions, we begin with a lemma that determines the expected number of informed vertices in one round. Intuitively we will show that *push* and *pull* do not interact badly, and therefore *push&pull* is given as a straightforward combination of *push* and *pull*.

Lemma 3.4. *Let \mathcal{G} be an expander sequence and abbreviate $I_t = I_t^{(pp)}$.*

- (a) *Let $|I_t| \leq n / \log n$. Then $\mathbb{E}_t[|I_{t+1} \setminus I_t|] = (2q + o(1))|I_t|$.*
- (b) *Let $|U_t| \leq n / \log n$. Then $\mathbb{E}_t[|U_{t+1}|] = (1 + o(1))e^{-q}(1 - q)|U_t|$.*

Proof. We begin with (a). The probability that $v \in U_t$ gets informed by *pull* is $q|N(v) \cap I_t|/|N(v)|$. Thus, using Lemma 2.7,

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(pull)} \setminus I_t|] &= \sum_{u \in U_t} q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &= q \sum_{u \in U_t} \frac{|N(u) \cap I_t|}{\Delta_n(1 + o(1))} \\ &= (q + o(1)) \frac{e(U_t, I_t)}{\Delta_n} \\ &= (q + o(1)) \frac{|U_t| |I_t| \Delta_n/n + o(\Delta_n) |I_t|}{\Delta_n}. \end{aligned} \tag{3.2}$$

Since $|I_t| = o(n)$ we obtain that $|U_t| = (1 - o(1))n$, and this expression simplifies to $(q + o(1))|I_t|$.

Before we switch our attention to *push* we make a simple observation. Let $a_1, \dots, a_k, k \in \mathbb{N}$ be real numbers. Then, using the fact that for any $a = o(1)$ it is $e^{-a+o(a)} = 1 - a$ and $e^{-a} = 1 - a + o(a)$, we have

$$\prod_{1 \leq i \leq k} (1 - a_i) = \exp\left(- (1 + o(1)) \sum_{1 \leq i \leq k} a_i\right) = 1 - (1 + o(1)) \sum_{1 \leq i \leq k} a_i \quad \text{if } \sum_{1 \leq i \leq k} a_i = o(1). \tag{3.3}$$

The probability that $v \in U_t$ gets informed by *push* is

$$1 - \prod_{i \in N(v) \cap I_t} (1 - q/|N(v)|).$$

According to Lemma 2.10(b) there is $B_t \subseteq U_t$ such that $|N(u) \cap I_t| = o(|N(u)|)$ for all $u \in B_t$ and $|U_t \setminus B_t| = o(|I_t|)$. Thus (3.3) is applicable, and in a similar fashion to (3.2) we get

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(push)} \setminus I_t|] &= \sum_{u \in U_t} 1 - \prod_{i \in N(u) \cap I_t} \left(1 - \frac{q}{|N(i)|}\right) \\ &= q \sum_{u \in B_t} \frac{|N(u) \cap I_t|}{\Delta_n(1 + o(1))} + o(|I_t|) \\ &= (q + o(1))|I_t|. \end{aligned} \tag{3.4}$$

We express the expected number of vertices informed by *push* & *pull* after one additional round in terms of the expected values we just calculated ((3.2) and (3.4)):

$$\begin{aligned} \mathbb{E}_t[|I_{t+1} \setminus I_t|] &= \mathbb{E}_t[|I_{t+1}^{(pull)} \setminus I_t| + |I_{t+1}^{(push)} \setminus I_t| - |(I_{t+1}^{(push)} \setminus I_t) \cap (I_{t+1}^{(pull)} \setminus I_t)|] \\ &= (2q - o(1))|I_t| - \mathbb{E}_t[|(I_{t+1}^{(push)} \setminus I_t) \cap (I_{t+1}^{(pull)} \setminus I_t)|]. \end{aligned} \tag{3.5}$$

Lemma 2.10(a) gives a set

$$A_t \subseteq I_{t+1}^{(push)}, |A_t| = (1 - o(1))|I_{t+1}^{(push)}|,$$

such that

$$|N(u) \cap I_{t+1}^{(push)}| = o(1)|N(u)| \quad \text{for all } u \in A_t.$$

Since push and pull happen independently,

$$\begin{aligned} \mathbb{E}_t[|(I_{t+1}^{(pull)} \setminus I_t) \cap (I_{t+1}^{(push)} \setminus I_t)| | I_{t+1}^{(push)}] &= \sum_{u \in I_{t+1}^{(push)} \setminus I_t} \mathbb{P}_t[u \in I_{t+1}^{(pull)} \setminus I_t] \\ &= \sum_{u \in I_{t+1}^{(push)} \setminus I_t} q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &\leq \sum_{u \in A_t} q \frac{|N(u) \cap I_t|}{|N(u)|} + \sum_{u \in I_{t+1}^{(push)} \setminus A_t} q \frac{|N(u) \cap I_t|}{|N(u)|}. \end{aligned}$$

Using the fact that $|N(u) \cap I_t| = o(|N(u)|)$ for all $u \in A_t$, we obtain

$$\mathbb{E}_t[|(I_{t+1}^{(pull)} \setminus I_t) \cap (I_{t+1}^{(push)} \setminus I_t)|] \leq \mathbb{E}_t[o(|A_t|) + |I_{t+1}^{(push)} \setminus A_t|] = o(|I_t|),$$

as

$$|A_t| \leq |I_{t+1}^{(push)}| \leq 2|I_t| \quad \text{and} \quad |I_{t+1}^{(push)} \setminus A_t| = o(|I_{t+1}^{(push)}|) = o(|I_t|).$$

Combining this with (3.5) we get $\mathbb{E}_t[|I_{t+1} \setminus I_t|] = (2q + o(1))|I_t|$, as claimed.

Next we show (b). Let A_u be the event that an uninformed vertex u does not get informed by the *push* algorithm, let B_u be the corresponding event for *pull*. Then A_u and B_u are independent and $A_u \cap B_u$ is the event that u does not get informed in the current round. We obtain

$$\begin{aligned} \mathbb{P}_t[A_u] &= \prod_{i \in N(u) \cap I_t} \left(1 - \frac{q}{|N(i)|}\right) \\ &\leq \left(1 - \frac{q}{\Delta_n}\right)^{|N(u) \cap I_t|} \\ &\leq \exp\left(-q \frac{|N(u) \cap I_t|}{\Delta_n}\right) \\ &= \exp\left(\frac{-q|N(u) \cap I_t|}{(1 + o(1))|N(u)|}\right) \end{aligned}$$

and

$$\mathbb{P}_t[B_u] = 1 - \frac{q|N(u) \cap I_t|}{|N(u)|}.$$

According to Lemma 2.10(a) there is a set $C_t \subseteq U_t$, $|C_t| = (1 - o(1))|U_t|$ such that $|N(u) \cap I_t| = (1 - o(1))|N(u)|$ for all $u \in C_t$. As $\mathbb{P}_t[A_u \cap B_u] \leq 1$, we therefore get

$$\mathbb{E}_t[|U_{t+1}|] = \sum_{u \in U_t} \mathbb{P}_t[A_u \cap B_u] \leq \sum_{u \in C_t} \mathbb{P}_t[A_u] \cdot \mathbb{P}_t[B_u] + |U_t \setminus C_t| \leq (1 + o(1))e^{-q}(1 - q)|U_t|.$$

For the lower bound we need to find a lower bound on the probability of a single uninformed vertex not getting informed in one round by *push*. Indeed, for any $u \in U_t$ and sufficiently large n ,

$$\mathbb{P}_t[A_u] = \prod_{v \in N(u) \cap I_t} \left(1 - \frac{q}{|N(v)|}\right) \geq \left(1 - \frac{q}{\delta_n}\right)^{|N(u) \cap I_t|} \geq e^{-q\Delta_n/\delta_n}. \tag{3.6}$$

Combining this inequality with the trivial bound $\mathbb{P}[B_u] \geq 1 - q$, we get a lower bound on the expected number of uninformed vertices after one round using *push&pull*:

$$\begin{aligned} \mathbb{E}_t[|U_{t+1}|] &= \sum_{u \in U_t} \mathbb{P}_t[A_u \cap B_u] \\ &= \sum_{u \in U_t} \mathbb{P}_t[A_u] \cdot \mathbb{P}_t[B_u] \\ &\geq e^{-q\Delta_n/\delta_n}(1 - q)|U_t| \\ &= (1 + o(1))e^{-q}(1 - q)|U_t|. \end{aligned} \quad \square$$

Next we show upper and lower bounds that together with Lemma 2.11 imply Theorem 1.2(c).

Lemma 3.5. *Let \mathcal{G} be an expander sequence and abbreviate $I_t = I_t^{(pp)}$. Let $q \in (0, 1]$. Then the following statements hold w.h.p.*

- (a) *Let $\sqrt{\log n} \leq |I_t| \leq n/\log n$. Then there are $\tau_1, \tau_2 = \log_{1+2q}(n/|I_t|) + o(\log n)$ such that $|I_{t+\tau_2}| < n/\log n < |I_{t+\tau_1}|$.*
- (b) *Let $n/\log n \leq |I_t| \leq n - n/\log n$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| > n - n/\log n$.*
- (c) *Let $|I_t| \geq n - n/\log n$.*
 - (i) *Case $q = 1$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| = n$.*
 - (ii) *Case $q \neq 1$. Then there is $\tau \leq \log n/(q - \log(1 - q)) + o(\log n)$ such that $|I_{t+\tau}| = n$.*

Proof. Since $|I_t| \geq |I_t^{(pull)}|$, statements (b) and (c) for $q = 1$ follow immediately from Lemma 3.2. To see (a), note that by using Lemma 3.4 we get $\mathbb{E}_t[|I_{t+1} \setminus I_t|] = (2q + o(1))|I_t|$, and applying Lemma 2.6 implies the claim.

Finally we show (c) for $q \neq 1$. Let $|I_t| \geq n - n/\log n$. By Lemma 3.4, we obtain that, for any $\tau \in \mathbb{N}$,

$$\mathbb{E}_t[|U_{t+\tau}|] = ((1 + o(1))e^{-q}(1 - q))^\tau |U_t|.$$

Thus we may choose $\tau = \log n/(q - \log(1 - q)) + o(\log n)$ such that, say, $\mathbb{E}_t[|U_{t+\tau}|] \leq |U_t|/n \leq 1/\log n$. Thus $\mathbb{P}_t[|U_{t+\tau}| \geq 1] \leq o(1)$ by Markov’s inequality. \square

Note that for $q = 1$ this already implies Theorem 1.2(c). This leaves the case for $q \neq 1$.

Lemma 3.6. *Let \mathcal{G} be an expander sequence and abbreviate $I_t = I_t^{(pp)}$, let $q \in (0, 1)$ and $|I_t| \leq n/2$. Then for $\tau = \log n/(q - \log(1 - q))$ and all $c < 1$ w.h.p. $|I_{t+c\tau}| < n$.*

Proof. We consider a modified process in which vertices have a higher chance of getting informed. In particular, note that the probability that $u \in U_t$ gets informed by *pull* is at most $q|N(u) \cap I_t|/|N(u)| \leq q$ and that all these events are independent; according to (3.6) the probability that $u \in U_t$ gets informed by *push* is at most $1 - e^{-q\Delta_n/\delta_n}$. Now we assume that each such u gets independently informed with probability exactly $1 - e^{-q\Delta_n/\delta_n}(1 - q)$. Then the runtime in this modified model constitutes a lower bound for the runtime in the original model. Let $u \in U_t$ and E_u be the event that u does not get informed in this modified model in $c\tau$ rounds. Thus, for $c < 1$,

$$\mathbb{P}[E_u] \geq ((1 - q)e^{-q\Delta_n/\delta_n})^{c\tau} = \omega(n^{-1}),$$

and as the events E_u are independent and $|U_t| = \Theta(n)$,

$$\mathbb{P}\left[\bigwedge_{u \in U_t} \overline{E_u}\right] \leq \prod_{u \in U_t} \mathbb{P}[\overline{E_u}] \leq \exp\left(-\sum_{u \in U_t} \mathbb{P}[E_u]\right) = o(1). \quad \square$$

3.3 Proof of Theorem 1.5(a) – push informs almost all vertices fast in spite of edge deletions

To shorten the notation, let us call the setting with deleted edges the ‘new model’ and the setting without deleted edges the ‘old model’, that is, the term ‘new model’ corresponds to the graphs in $\tilde{\mathcal{G}}$ while ‘old model’ refers to the (original) graphs in \mathcal{G} . We prove Lemma 3.7, which directly implies Theorem 1.5(a). We write $I_t = I_t^{(push)}$ throughout.

Lemma 3.7. *Under the assumptions of Theorem 1.5(a), the following holds for the new model:*

- (a) *There are $\tau, \tilde{\tau} = \log_{1+q}(n) + o(\log n)$ such that w.h.p. $|I_{\tilde{\tau}}| < n/\log n < |I_{\tau}|$.*
- (b) *Assume $|I_t| \geq n/\log n$. Then there is a $\tau = o(\log n)$ such that w.h.p. $|I_{t+\tau}| \geq n - n/\log n$.*

For the proof of Lemma 3.7 we will need the following statements, the first one taken from [28].

Lemma 3.8 (proof of Lemma 2.5 in [28]). *Consider the old model. Assume $|I_t| < n/\log n$ and $q = 1$. Then*

$$\mathbb{P}_t[|I_{t+1}| = |I_t| + (1 - o(1))|I_t|] = 1 - o(1). \quad (3.7)$$

Lemma 3.9. *Consider push on a sequence of graphs $(G_n)_{n \in \mathbb{N}}$, where G_n has n vertices. Assume that $|I_t| = \omega(1)$ and that (3.7) holds for $q = 1$, that is, assume that*

$$\mathbb{P}_t[|I_{t+1}| = |I_t| + (1 - o(1))|I_t|] = 1 - o(1) \quad \text{for } q = 1.$$

Then, for $q \in (0, 1]$,

$$\mathbb{P}_t[|I_{t+1}| = |I_t| + (q - o(1))|I_t|] = 1 - o(1). \quad (3.8)$$

Moreover, assume that whenever $|I_t| < n/\log n$, for $q = 1$, (3.7) holds. Then there are $\tau, \tilde{\tau} = \log_{1+q}(n) + o(\log n)$ such that w.h.p.

$$|I_{\tilde{\tau}}| < n/\log n < |I_{\tau}|. \quad (3.9)$$

Proof. For a graph G and for $v \in I_t$, let $X_v(G)$ denote the vertex to which v pushes in round t . Let

$$N_{t+1} := \{X_v(G_n) \mid v \in I_t\} \cap U_t.$$

Note that whenever $|I_t| < n/\log n$, w.h.p. $|N_{t+1}| = (1 - o(1))|I_t|$ from (3.7). For $q \in (0, 1]$ each vertex in N_{t+1} has a probability at least q of being informed and all these events are independent; thus (3.8) follows directly by applying the Chernoff bounds whenever $|I_t| = \omega(1)$.

In order to prove the second statement we call a round t that does not satisfy (3.8) a *failed* round. Note that we just argued that the probability that a round fails is $o(1)$ whenever $|I_t| = \omega(1)$ and $|I_t| < n/\log n$, and the events that distinct rounds fail are independent. In particular, the number of failed rounds among the next R rounds, assuming that $|I_t|$ stays below $n/\log n$, is w.h.p. $o(R)$. Moreover, if a round does not fail, the number of informed vertices increases by a factor of $(1 + q + o(1))$ and otherwise it may increase by an arbitrary factor in the interval $[1, 2]$.

Finally, Lemma 2.11 yields that there is $t^* = o(\log n)$ such that w.h.p. $|I_{t^*}| = \omega(1)$, which implies that after $R + t^*$ rounds, the number of informed vertices is w.h.p. in the interval

$$[(1 + q + o(1))^{R-o(R)}, (1 + q + o(1))^{R-o(R)} \cdot 2^{o(R)}],$$

and choosing $R = \log_{1+q}(n) + o(\log n)$ in two ways establishes (3.9). □

In the subsequent proof of Lemma 3.7 we will use the simple observations that, for any $n \in \mathbb{N}_0$,

$$\mathbb{P}[\text{Bin}(n, 1/2) \geq n/2] \geq 1/2 \quad \text{and} \quad \mathbb{P}[\text{Bin}(n, 1/4) \geq n/4] \geq 1/4 \tag{3.10}$$

(see e.g. [22] when $n > 4$), and the other cases are checked easily.

Proof of Lemma 3.7. We first show (a). We assume $q = 1$ and prove that, for $|I_t| < n/\log n$, (3.7) also holds in the new model; then claim (a) follows directly from Lemma 3.9. Let $G = (V, E)$ be a graph. For $v \in I_t$ let $X_v(G)$ denote the vertex to which v pushes in round t . For $u \in V$ let $c_u(G) := |\{v \in I_t \mid X_v(G) = u\}|$ denote the number of times u is pushed in round t . Let

$$\mathcal{Y}_t(G) := \{v \in I_t \mid c_v(G) = 1\} \quad \text{and} \quad \mathcal{H}_t(G) := \{v \in I_t \mid c_v(G) \geq 1\}$$

denote the set of informed vertices that are being pushed exactly once in round t and the set of informed vertices that are being pushed at least once in round t respectively. Let

$$\mathcal{Z}_t(G) := \{v \in V \mid c_v(G) \geq 2\}$$

denote the set of vertices that are being pushed more than once in round t . Let $Y_t(G) := |\mathcal{Y}_t(G)|$ and $H_t(G) := |\mathcal{H}_t(G)|$ and, in a slight abuse of notation, let

$$Z_t(G) := \sum_{k \geq 2} (k - 1) \cdot |\{v \in V \mid c_v(G) = k\}|$$

denote the number of vertices that are being pushed multiple times in round t counted with multiplicity. Note that the quantity $Y + Z$ denotes the number of pushes that have no effect in the respective round, that is, there are $Y + Z$ pushes that are useless in the sense that even without them, the same number of vertices would become informed in the respective round. In the following paragraphs we condition on I_t implicitly, that is, we write $\mathbb{P}[\dots]$ instead of $\mathbb{P}_t[\dots]$, etc., to lighten the notation. We want to show that (3.7) does hold in the new model; for contradiction we assume that this is not the case. Hence we can infer that there is a constant $c > 0$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{P}[Y_t(\tilde{G}_n) \geq c|I_t|] > 0 \quad \text{or} \quad \limsup_{n \rightarrow \infty} \mathbb{P}[Z_t(\tilde{G}_n) \geq c|I_t|] > 0.$$

Thus, without loss of generality, we can assume that there is $f^* > 0$ and $n_0 \in \mathbb{N}$ such that

$$\mathbb{P}[Y_t(\tilde{G}_n) \geq c|I_t|] > f^* \text{ for all } n \geq n_0 \quad \text{or} \quad \mathbb{P}[Z_t(\tilde{G}_n) \geq c|I_t|] > f^* \text{ for all } n \geq n_0.$$

If this is not the case we can restrict ourselves to a suitable subsequence of $(n)_{n \in \mathbb{N}}$ on which it is true. Next, we describe an explicit coupling between the new and the old model. For any vertex v consider $X_v(G_n)$. If $X_v(G_n) \in N_{\tilde{G}_n}^-(v)$, then set $X_v(\tilde{G}_n) := X_v(G_n)$ and otherwise choose $X_v(\tilde{G}_n)$ uniformly at random from $N_{\tilde{G}_n}^-(v)$. Note that $X_v(G_n), X_v(\tilde{G}_n)$ have by construction the correct marginal distribution. Moreover, note that by construction, the family

$$(X_v(G_n) \mid (X_u(\tilde{G}_n))_{u \in V_n})_{v \in V_n} \tag{3.11}$$

of random variables is independent, since $X_v(G_n)$ depends only on $X_v(\tilde{G}_n)$ for all $v \in V_n$.

We begin with the case that $\mathbb{P}[Y_t(\tilde{G}_n) \geq c|I_t|] > f^*$. We will show

$$\mathbb{P}[H_t(G_n) \geq Y_t(\tilde{G}_n)/2 \mid \mathcal{Y}_t(\tilde{G}_n)] \geq 1/2$$

and then, since by assumption $\mathbb{P}[Y_t(\tilde{G}_n) \geq c|I_t|] > f^*$, we can infer $\mathbb{P}[H_t(G_n) \geq c|I_t|/2] \geq f^*/2$, which contradicts Lemma 3.8. Let $\mathcal{Y}_t(\tilde{G}_n) = \{y_1, \dots, y_{Y_t(\tilde{G}_n)}\}$. Then there are distinct vertices $v_1, \dots, v_{Y_t(\tilde{G}_n)} \in I_t$ such that $X_{v_i}(\tilde{G}_n) = y_i$ for all $i \in \{1, \dots, Y_t(\tilde{G}_n)\}$. Due to (3.11) the events $\{X_{v_i}(G_n) = X_{v_i}(\tilde{G}_n)\}_{1 \leq i \leq Y_t}$ are independent. Moreover, for all $i \in \{1, \dots, Y_t(\tilde{G}_n)\}$,

$$\mathbb{P}[X_{v_i}(G_n) = X_{v_i}(\tilde{G}_n) \mid \mathcal{Y}_t(\tilde{G}_n)] = \frac{d_{\tilde{G}_n}(v_i)}{d_{G_n}(v_i)} \geq 1/2 + \varepsilon$$

and therefore, given $\mathcal{Y}_t(\tilde{G}_n)$, $H_t(G_n)$ dominates a binomially distributed random variable $\text{Bin}(Y_t(\tilde{G}_n), 1/2)$. In particular, this implies with (3.10) that

$$\mathbb{P}[H_t(G_n) \geq Y_t(\tilde{G}_n)/2 \mid \mathcal{Y}_t(\tilde{G}_n)] \geq 1/2,$$

as claimed.

We continue with the case $\mathbb{P}[Z_t(\tilde{G}_n) \geq c|I_t|] > f^*$. Let $\mathcal{Z}_t(\tilde{G}_n) = \{z_1, \dots, z_{|\mathcal{Z}_t(\tilde{G}_n)|}\}$. Then, for any $i \in \{1, \dots, |\mathcal{Z}_t(\tilde{G}_n)|\}$ let $n_i := c_{z_i}(\tilde{G}_n) \geq 2$, that is, there are distinct vertices $v_{i,1}, \dots, v_{i,n_i}$ such that $X_{v_{i,j}}(\tilde{G}_n) = z_i$ for all $v \in \{v_{i,1}, \dots, v_{i,n_i}\}$. We will show that

$$\mathbb{P}[Z_t(G_n) \geq Z_t(\tilde{G}_n)/8 \mid \mathcal{Z}_t(\tilde{G}_n), n_1, \dots, n_{|\mathcal{Z}_t(\tilde{G}_n)|}] \geq 1/8 \tag{3.12}$$

and then, since by assumption $\mathbb{P}[Z_t(\tilde{G}_n) \geq c|I_t|] > f^*$, we obtain $\mathbb{P}[Z_t(G_n) \geq c/8|I_t|] \geq f^*/8$, which contradicts Lemma 3.8. Due to (3.11) the events

$$\{X_{v_{i,j}}(G_n) = X_{v_{i,j}}(\tilde{G}_n)\}_{1 \leq i \leq |\mathcal{Z}_t(\tilde{G}_n)|, 1 \leq j \leq n_i} \tag{3.13}$$

are independent. Moreover, for all $1 \leq i \leq |\mathcal{Z}_t(\tilde{G}_n)|, 1 \leq j \leq n_i$,

$$\mathbb{P}[X_{v_{i,j}}(G_n) = X_{v_{i,j}}(\tilde{G}_n) \mid \mathcal{Z}_t(\tilde{G}_n), n_1, \dots, n_{|\mathcal{Z}_t(\tilde{G}_n)|}] = \frac{d_{\tilde{G}_n}(v_{i,j})}{d_{G_n}(v_{i,j})} \geq 1/2 + \varepsilon. \tag{3.14}$$

For $1 \leq i \leq |\mathcal{Z}_t(\tilde{G}_n)|$ let $B_i \sim \text{Bin}(n_i, 1/2)$ be independent random variables. Moreover, let $M_1 := \{i \mid 1 \leq i \leq |\mathcal{Z}_t(\tilde{G}_n)|, n_i = 2\}$ and $M_2 := \{i \mid 1 \leq i \leq |\mathcal{Z}_t(\tilde{G}_n)|, n_i > 2\}$. Using (3.13) and (3.14), given $\mathcal{Z}_t(\tilde{G}_n), n_1, \dots, n_{|\mathcal{Z}_t(\tilde{G}_n)|}$, we infer that $Z_t(G_n)$ dominates

$$\sum_{i=1}^{|\mathcal{Z}_t(\tilde{G}_n)|} \max\{B_i - 1, 0\} \geq \sum_{i \in M_1} \max\{B_i - 1, 0\} + \sum_{i \in M_2} B_i - |M_2|.$$

We treat the two sums individually. Note that $\sum_{i \in M_1} \max\{B_i - 1, 0\} \sim \text{Bin}(|M_1|, 1/4)$; in particular,

$$\mathbb{P}\left[\sum_{i \in M_1} \max\{B_i - 1, 0\} \geq |M_1|/4\right] \geq 1/4$$

by (3.10). Regarding the second sum, since $\sum_{i \in M_2} B_i \sim \text{Bin}(\sum_{i \in M_2} n_i, 1/2)$, we obtain

$$\mathbb{P}\left[\sum_{i \in M_2} B_i \geq 1/2 \sum_{i \in M_2} n_i\right] \geq 1/2.$$

Thus, given $\mathcal{Z}_t(\tilde{G}_n), n_1, \dots, n_{|\mathcal{Z}_t(\tilde{G}_n)|}$ and using $2|M_1| = \sum_{i \in M_1} n_i$ and $\sum_{i \in M_2} n_i \geq 3|M_2|$, we infer that with probability at least $1/4 \cdot 1/2 = 1/8$

$$\begin{aligned} Z_t(G_n) &\geq \frac{1}{4}|M_1| + \frac{1}{2} \sum_{i \in M_2} n_i - |M_2| \\ &= \frac{1}{8} \sum_{i \in M_1} n_i + \frac{1}{2} \sum_{i \in M_2} n_i - |M_2| \\ &\geq \frac{1}{8} \sum_{i \in M_1} n_i + \frac{1}{6} \sum_{i \in M_2} n_i \\ &\geq \frac{1}{8} \sum_{i=1}^{|\mathcal{Z}_t(\tilde{G}_n)|} n_i \\ &= \frac{1}{8}(Z_t(\tilde{G}_n) + |\mathcal{Z}_t(\tilde{G}_n)|) \\ &\geq \frac{1}{8}Z_t(\tilde{G}_n). \end{aligned}$$

This establishes (3.12). All in all, for $q = 1$ we have shown that (3.7) also holds in the new model. Hence claim (a) follows directly from Lemma 3.9.

Next we prove claim (b). We write $\Delta_n := \Delta(G_n), \tilde{\Delta}_n := \Delta(\tilde{G}_n), \delta_n := \delta(G_n)$ and $\tilde{\delta}_n := \delta(\tilde{G}_n)$; moreover, we write $\tilde{N}(\cdot)$ instead of $N_{\tilde{G}_n}(\cdot)$. We assume that $|I_t| \in [n/\log n, n - n/\log n]$. We further distinguish two cases, namely $|I_t| \in [n/\log n, n/2]$ and $|I_t| \in [n/2, n - n/\log n]$. We start with the case $|I_t| \in [n/\log n, n/2]$. Using Lemmas 2.7 and 2.8 and the assumption that $\Delta_n/\delta_n = 1 + o(1)$ we obtain, for any $0 < \bar{\varepsilon} < \varepsilon/2$, for n sufficiently large,

$$e(I_t, U_t) > \bar{\varepsilon}\delta_n|I_t|. \tag{3.15}$$

Using the fact that $e^x \geq (1 + x/n)^n$ for $n \in \mathbb{N}$ and $|x| \leq n$, we obtain

$$\mathbb{E}_t[|I_{t+1} \setminus I_t|] \geq \sum_{u \in \tilde{N}(I_t) \setminus I_t} \left[1 - \prod_{v \in \tilde{N}(u) \cap I_t} \left(1 - \frac{q}{\tilde{\Delta}_n} \right) \right] \geq \sum_{u \in \tilde{N}(I_t) \setminus I_t} 1 - e^{-|\tilde{N}(u) \cap I_t|q/\tilde{\Delta}_n}.$$

Further, using the fact that $e^{-x} \leq 1 - x/2$ for any $x \in (0, 1)$ and (3.15) yields the bound

$$\mathbb{E}_t[|I_{t+1} \setminus I_t|] \geq \sum_{u \in \tilde{N}(I_t) \setminus I_t} \frac{q|\tilde{N}(u) \cap I_t|}{2\tilde{\Delta}_n} = \frac{qe(I_t, U_t)}{2\tilde{\Delta}_n} \geq \frac{\bar{\varepsilon}q\delta_n}{2\Delta_n}|I_t|.$$

For this case the claim follows by Lemma 2.5, when setting $f = n/\log n, g = \log n$ and $c = \bar{\varepsilon}q\delta_n/(2\Delta_n)$.

Finally we consider the case $|I_t| \in [n/2, n - n/\log n]$; here we examine the shrinking of U_t . Using Lemmas 2.7 and 2.8 we obtain, for any $0 < \bar{\varepsilon} < \varepsilon/2$, for n sufficiently large, $e(I_t, U_t) > \bar{\varepsilon}\delta_n|U_t|$. Hence, again using the fact that for any $x \in (0, 1)$ it holds that $e^{-x} \leq 1 - x/2$, and that for $n \in \mathbb{N}$ and $|x| \leq n$ it is $e^x \geq (1 + x/n)^n$, we obtain

$$\begin{aligned} \mathbb{E}_t[|U_{t+1}|] &= \sum_{u \in U_t} \prod_{v \in \tilde{N}(u) \cap I_t} \left(1 - \frac{q}{d_{\tilde{G}_n}(v)} \right) \\ &\leq \sum_{u \in U_t} e^{-|\tilde{N}(u) \cap I_t|q/\tilde{\Delta}_n} \end{aligned}$$

$$\begin{aligned} &\leq \sum_{u \in U_t} 1 - \frac{q|\tilde{N}(u) \cap I_t|}{2\tilde{\Delta}_n} \\ &\leq |U_t| - \frac{\bar{\varepsilon}q\delta_n}{2\tilde{\Delta}_n}|U_t| \\ &\leq \left(1 - \frac{\bar{\varepsilon}q\delta_n}{2\Delta_n}\right)|U_t|. \end{aligned}$$

Using the tower property of conditional expectation, we immediately get

$$\mathbb{E}_t[|U_{t+\tau}|] \leq \left(1 - \frac{\bar{\varepsilon}q\delta_n}{2\Delta_n}\right)^\tau |U_t|, \quad \tau \in \mathbb{N}.$$

Thus, for

$$\tau := -2 \log \log(n) / \log(1 - \bar{\varepsilon}q\delta_n / (2\Delta_n)) = o(\log n),$$

we have $\mathbb{E}_t[|U_{t+\tau}|] = o(n / \log n)$. Hence by Markov's inequality, $\mathbb{P}[|U_{t+\tau}| \geq n / \log n] = o(1)$. \square

3.4 Proof of Theorem 1.4(a) – edge deletions slow down push

Let $I_t^{(push)} := I_t$. In order to show the claim we construct an explicit sequence of graphs that has the desired property. More precisely, for any $\varepsilon > 0$, each $q \in (0, 1]$ and $n \in \mathbb{N}$ we will define a graph $G_n(\varepsilon)$ that is obtained by deleting edges from the complete graph on n vertices such that each vertex keeps at least a $(1 - \varepsilon)$ fraction of its edges and such that *push* slows down significantly.

We define $G_n(\varepsilon) = (V_1 \cup V_2, E)$ with vertex set $V = V_1 \cup V_2$, where $V_1 := \{1, \dots, \lfloor n/2 \rfloor\}$ and $V_2 := \{\lfloor n/2 \rfloor + 1, \dots, n\}$, as follows. We include in E all pairs of vertices that intersect V_1 and, moreover, we add edges (that now have endpoints only in V_2) such that all vertices in V_2 have degree $\lceil (1 - \varepsilon)n \rceil + 1 \pm 1$. According to Lemma 3.7(a) there is a $t = \log_{1+q}(n) + o(\log n)$ such that w.h.p. $|I_t| < n / \log n$. It thus suffices to show that it takes w.h.p. at least $(1 + \varepsilon/2)q^{-1} \log n$ more rounds to inform all remaining vertices.

Let $U'_t := U_t^{(push)} \cap V_2$. As $|I_t| < n / \log n$ we have $|U'_t| \geq n/4$ with plenty of room to spare. In the remainder of this proof we will consider a modified process in which vertices have a higher chance of getting informed; in particular we assume that in each round, *all* vertices choose a neighbour independently and uniformly at random and after this round the chosen vertices are informed. Let E_u denote the event that $u \in U'_t$ does not get informed within the next $\tau := (1 + \varepsilon/2)q^{-1} \log n$ rounds in this modified model. Each vertex $u \in U'_t$ has $\lfloor n/2 \rfloor$ neighbours that have degree $n - 1$, at most $\lceil (1 - \varepsilon)n \rceil + 1 \pm 1 - \lfloor n/2 \rfloor \leq (1/2 - \varepsilon)n + 4$ neighbours that have at least degree $(1 - \varepsilon)n$ and no further neighbours. Therefore, using the fact that for any $a \in \mathbb{R}$ we have $(1 + a/n)^n = e^a + O(1/n)$, we obtain for each $u \in U'_t$

$$\begin{aligned} \mathbb{P}_t[E_u] &\geq \left(\left(1 - \frac{q}{n-1}\right)^{n/2} \left(1 - \frac{q}{(1-\varepsilon)n}\right)^{(1/2-\varepsilon)n+4} \right)^\tau \\ &= (1 + o(1))(e^{-q(1/2+(1/2-\varepsilon)/(1-\varepsilon))})^\tau \\ &= (1 + o(1)) \exp\left(-\frac{4 - 4\varepsilon - 3\varepsilon^2}{4 - 4\varepsilon} \log n\right) \\ &= \omega(n^{-1}). \end{aligned}$$

In this modified model the events $\{\overline{E_u} \mid u \in U'_t\}$ also satisfy $\mathbb{P}_t[\overline{E_u} \mid \{E_v : v \in U\}] \leq 1 - p$ for all $u \in V_2$ and $U \subseteq V \setminus \{u\}$ and for some $p = \omega(n^{-1})$. This follows immediately from the previous

calculation, as conditioning on an event like $\{\overline{E}_v: v \in U\}$ only decreases the number of vertices that can push to u . Thus, as $|U'_t| = \Theta(n)$,

$$\mathbb{P}_t \left[\bigwedge_{u \in U'_t} \overline{E}_u \right] \leq \prod_{u \in U'_t} (1 - p) \leq \exp \left(- \sum_{u \in U'_t} p \right) = o(1).$$

3.5 Proof of Theorems 1.3(b) and 1.5(a) – push&pull informs almost all vertices fast in spite of edge deletions

Before we show the actual proof we will first present an informal argument that contains all relevant ideas and important observations. Let $\sqrt{\log n} \leq |I_t| \leq n/\log n$ and assume $q = 1$. In Section 3.3 we proved that for *push* the informed vertices nearly double in every round for an arbitrary expander sequence with edge deletions and an otherwise arbitrary set I_t . For *pull* this is not true; however, we proved in Section 3.1 that the number of edges between the informed and the uninformed vertices nearly doubles in every round. The first attempt towards the proof of Theorems 1.3(b) and 1.5(b) then seems obvious: one would try to show that either the vertices triple every round, or the edges do so, or for example that the product of the two quantities increases by a factor of 9. As it turns out, this is in general not the case; indeed, it is possible to choose an expander sequence, to delete edges such that each vertex keeps at least a $(1/2 + \varepsilon)$ -fraction of its neighbours, and to choose a (large) set of informed vertices I_t such that after one round w.h.p. either $|I_{t+1}| < c|I_t|$ or $e(I_{t+1}, U_{t+1}) < ce(I_t, U_t)$ or $|I_{t+1}|e(I_{t+1}, U_{t+1}) < c^2|I_t|e(I_t, U_t)$ for some $c < 3$. On the other hand, and although we have no explicit description of these ‘malicious’ sets, it seems rather unlikely that such sets will occur several times during the execution of *push&pull*.

In order to show the claimed running time of *push&pull* we will impose some additional structure. Let $\varepsilon > 0$. In the subsequent exposition we assume that our graph G – obtained from an expander by deleting edges such that each vertex keeps at least a $(1/2 + \varepsilon)$ fraction of the edges – has a *very special* structure. In particular, we assume that there is a partition $\Pi = (V_i)_{i \in [k]}$ of the vertex set of G into a bounded number k of equal parts such that $E_G(V_i) = \emptyset$ for all $1 \leq i \leq k$ and such that the induced subgraph (V_i, V_j) looks like a random regular bipartite graph for all $1 \leq i < j \leq k$. Of course, not every relevant G admits such a partition; however, Szemerédi’s Regularity Lemma guarantees that every sufficiently large graph has a partition that is in a well-defined sense *almost* like the one described previously, and a substantial part of our proof is concerned with showing that being ‘almost special’ does not hurt significantly.

Assuming that G is very special, let us collect some easy facts. Denote the degree of $u \in V_i$ in the induced subgraph (V_i, V_j) by d_{ij} ; this immediately gives $d_G(u) = \sum_{\ell=1}^k d_{i\ell}$, and note that $d_{ii} = 0$ as there are no edges in V_i . Moreover, regular bipartite random graphs satisfy an expander property, that is,

$$\begin{aligned} e(W_i, W_j) &= d_{ij}|W_i| |W_j|/|V_j| + o(d_{ij})|W_i| \\ &\approx |W_i| |W_j|d_{ij}k/n \quad \text{for all } W_i \subseteq V_i, W_j \subseteq V_j, 1 \leq i < j \leq k, \end{aligned}$$

where we used the fact that all $|V_i|$ are of equal size. This is quite similar to the property that we used in our preceding analysis on expander sequences; see Lemma 2.7. As a pair in Π behaves like a bipartite expander sequence, we can easily compute the expected number of informed vertices like we did in Section 3.2. We do so now for *pull*. Let $|I_{t+1}^{ij}|$ be the number of vertices in V_i informed after round $t + 1$ by *pull* from vertices only in V_j and set $I_t^i := I_t \cap V_i, U_t^i := U_t \cap V_i$ for all $1 \leq i \leq k$. Thus, as long as I_t^i is much smaller than V_i (and thus also $U_t^i \approx |V_i| = n/k$), we get

$$\mathbb{E}_t[|I_{t+1}^{(pull),ij} \setminus I_t|] = \sum_{u \in U_t^i} \frac{|N(u) \cap I_t^j|}{d(u)} = \frac{e(U_t^i, I_t^j)}{\sum_{1 \leq \ell \leq k} d_{i\ell}} \approx \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{i\ell}} |I_t^j|.$$

A similar calculation, which we do not perform in detail, yields for *push*

$$\mathbb{E}_t[|I_{t+1}^{(push),ij} \setminus I_t|] \approx \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{\ell j}} |I_t^j|.$$

Moreover, as in previous proofs it turns out that the number of vertices informed simultaneously by *push* as well as *pull* is negligible; compare with the proof of Lemma 3.4. Thus we obtain that more or less

$$\mathbb{E}_t[|I_{t+1}^{(pp),ij}|] \approx |I_t^i| + \left(\frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right) |I_t^j|,$$

and by linearity of expectation

$$\mathbb{E}_t[|I_{t+1}^{(pp),i}|] \approx |I_t^i| + \sum_{1 \leq j \leq k} \left(\frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right) |I_t^j|.$$

Set $X_t = (|I_t^i|)_{i \in [k]}$ and $A = (A_{ij})_{1 \leq i, j \leq k}$, the matrix with entries

$$A_{ij} = \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij}}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \quad \text{for } 1 \leq i \neq j \leq k$$

and $A_{ii} = 1$ for $1 \leq i \leq k$. With this notation we obtain the recursive relation

$$\mathbb{E}_t[X_{t+1}] \approx A \cdot X_t, \tag{3.16}$$

that is, we may expect that $X_t \approx \mathbb{E}_t[X_t] \approx A^t X_0$. If we then let λ_{\max} denote the greatest eigenvalue of A , then we obtain that to leading order

$$|I_t| \approx \lambda_{\max}^t.$$

Our aim is to show that *push&pull* is (at least) as fast as on the complete graph, that is, $|I_t| \gtrsim 3^t$, so we take a closer look at the eigenvalues of A . By construction A is symmetric, so the largest eigenvalue equals $\sup_{\|x\|=1} \|x^T A x\|$, and the simple choice $x = k^{-1/2} \mathbf{1}$ yields

$$\lambda_{\max} \geq \frac{\sum_{(i,j)} A_{ij}}{k} = \frac{\sum_{j=1}^k 1 + \sum_{i=1}^k \sum_{j=1}^k d_{ij} / (\sum_{\ell=1}^k d_{i\ell}) + \sum_{j=1}^k \sum_{i=1}^k d_{ij} / (\sum_{\ell=1}^k d_{\ell j})}{k} = 3.$$

This neat property leads us to the expected result

$$T_{pp}(G) = (1 + o(1)) \log_{\lambda_{\max}} n \leq (1 + o(1)) \log_3 n,$$

and it also completes the informal argument that justifies the claim made in Theorems 1.3(b) and 1.5(b). In the rest of this section we will turn this argument step by step into a formal proof by filling in all missing pieces.

Obtaining an appropriate regular partition. An important ingredient in the previous sketch was the assumption that the given graph has a partition into a bounded number of equal parts, such that the bipartite graph induced by any two different parts looks like a random regular graph. This assumption is quite strong and very much not true in general. However, restricting ourselves to dense graphs we can actually come quite close to that. Let us begin with some definitions; the statements are taken from [30].

Definition 3.1 (density). Given a graph $G = (V, E)$ and two disjoint non-empty sets of vertices $X, Y \subseteq V$, we define the *density* of the pair (X, Y) as

$$d_G(X, Y) = \frac{e_G(X, Y)}{|X| |Y|}.$$

As usual, if the graph is clear from the context the index will be omitted. The next definition gives a partition that is close to the previously described properties; all sets in the partition have nearly the same size and nearly all pairs behave in a well-defined sense like regular bipartite random graphs.

Definition 3.2 ((ε, k_0, K_0) -Szemerédi partition). Let $G = (V, E)$ and $k \in \mathbb{N}$. We call $\Pi = \{V_i\}_{i \in [k]}$ an (ε, k_0, K_0) -Szemerédi partition of G if the following conditions are fulfilled.

- (a) $V_1 \dot{\cup} \dots \dot{\cup} V_k = V$.
- (b) $k_0 \leq k \leq K_0$.
- (c) $|V_1| \leq \dots \leq |V_k| \leq |V_1| + 1$.
- (d) For all but at most εk^2 pairs (V_i, V_j) of Π with $i < j$, we have that for all subsets $U_i \subseteq V_i$ and $U_j \subseteq V_j$ with $|U_i| \geq \varepsilon |V_i|$ and $|U_j| \geq \varepsilon |V_j|$,

$$|d(U_i, U_j) - d(V_i, V_j)| \leq \varepsilon.$$

A pair (V_i, V_j) satisfying the last condition is called ε -regular. For pairs (V_i, V_j) in Π we will abbreviate $d(V_i, V_j)$ to d_{ij} .

Next we state Szemerédi’s Regularity Lemma. It guarantees that we will have a Szemerédi partition if the underlying graph is large enough.

Lemma 3.10 ([30], The Regularity Lemma). *For every $\varepsilon > 0$ and every $k_0 \in \mathbb{N}$ there exist $K_0 = K_0(\varepsilon, k_0)$ and n_0 such that every graph $G = (V, E)$ with at least $|V| = n \geq n_0$ vertices admits an (ε, k_0, K_0) -Szemerédi partition.*

The next lemma gives a useful property of regular pairs. In particular, with the exception of a small set only, all other vertices have a degree that is close to dN , where d is the density of the pair and N is the number of vertices in each part. In fact the statement also is true for arbitrary but not too small subsets of the parts.

Lemma 3.11. *Let $G = (V, E)$ be a graph, $\varepsilon > 0$ and $U, U' \subseteq V$. Suppose that (U, U') is an ε -regular pair, and let $W \subseteq U'$, $|W| \geq \varepsilon |U'|$. Furthermore, let $\mathcal{E}(U, W) \subseteq U$ be the largest set such that $|d(u, W) - d(U, U')| \geq \varepsilon$ for all $u \in \mathcal{E}(U, W)$. Then $|\mathcal{E}(U, W)| \leq 2\varepsilon |U|$.*

Proof. We will prove this by contradiction. Assume that $|\mathcal{E}(U, W)| \geq 2\varepsilon |U|$. Let us write $\mathcal{E}(U, W) = S \cup L$, where

$$S = \{u \in \mathcal{E}(U, W) : d(u, W) < d(U, U') - \varepsilon\}, \quad L = \{u \in \mathcal{E}(U, W) : d(u, W) > d(U, U') + \varepsilon\}.$$

Then $|S| \geq \varepsilon |U|$ or $|L| \geq \varepsilon |U|$. In the former case

$$d(S, W) = \frac{\sum_{u \in S} e(u, W)}{|S| |W|} = \frac{\sum_{u \in S} d(u, W)}{|S|} < d(U, U') - \varepsilon.$$

As $|S| \geq \varepsilon |U|$, $|W| \geq \varepsilon |U'|$, this contradicts the assumption that (U, U') is an ε -regular pair. The case $|L| \geq \varepsilon |U|$ follows analogously by showing that $d(L, W) > d(U, U') + \varepsilon$. □

We call the set $\mathcal{E}(U, W)$ in Lemma 3.11 the exceptional set of U with respect to W . In particular, Lemma 3.11 implies that for every ε -regular pair (U, U') and all $W \subseteq U', |W| \geq (1 - c\varepsilon)|U'|, c > 0$ we have

$$\begin{aligned} |d(u, W) - d(U, U')| &\leq |d(u, W) - d(u, U')| + |d(u, U') - d(U, U')| \\ &\leq (c + 1)\varepsilon \quad \text{for all } u \in U \setminus \mathcal{E}(U, U'). \end{aligned} \tag{3.17}$$

Having done these preparations we can now determine a partition that comes close to the initially described properties.

Lemma 3.12. *Let $G_n = (V, E)$ be a graph on n vertices such that $\delta_{G_n} \geq \alpha n$ for some $\alpha > 0$. Then for all $\eta > 0$ and $k_0 > 1/\sqrt{\eta}$ there exists $n_0, K_0 \in \mathbb{N}$ such that for all G_n with $n \geq n_0$ there is an (η, k_0, K_0) -Szemerédi partition $\Pi = \{V_i\}_{i \in [k]}$ of G_n with the following property. There is $F \subseteq \Pi$ with $|F| \leq \eta k$ such that, for all $V_i \in \Pi \setminus F$,*

- there are at most ηk non- η -regular pairs $(V_i, V_j), j \in [k]$, and
- there exists an exceptional set $N_i, |N_i| \leq \eta|V_i|$ such that

$$d(u) \leq (1 + \eta) \frac{n}{k} \sum_{1 \leq j \leq k} d(V_i, V_j) \quad \text{for all } u \in V_i \setminus N_i.$$

Proof. According to Lemma 3.10, for all $\xi > 0$ and $k_0 > 1/\sqrt{\xi}$, there are $n_0, K_0 \in \mathbb{N}$ such that for all G_n with $n \geq n_0$ there is a $k \in \mathbb{N}$ and a (ξ, k_0, K_0) -Szemerédi partition $\Pi = \{V_i\}_{i \in [k]}$ of G_n . Let $F \subseteq \Pi$ contain the parts $V_i \in \Pi$ such that there are at least $\sqrt{\xi}k$ other parts $V_j \in \Pi$ such that the pair (V_i, V_j) is not ξ -regular. As there are at most ξk^2 non- ξ -regular pairs, we infer that $|F| \leq \sqrt{\xi}k$. Let $V_i \in \Pi \setminus F$. Further, let $A_i \subseteq \Pi$ be such that (V_i, V_j) is a ξ -regular pair for all $V_j \in \Pi \setminus A_i$ and (V_i, V_j) is not ξ -regular for all $V_j \in A_i$. The definition of F implies that $|A_i| \leq \sqrt{\xi}k$. For these $V_j \in \Pi \setminus A_i$ let $\mathcal{E}_i(V_j) = \mathcal{E}(V_i, V_j)$ be the exceptional set of V_i with respect to V_j . On top of that let $N_i \subseteq V_i$ be the set of points in V_i that are in at least $\sqrt{\xi}k$ exceptional sets with respect to parts in $\Pi \setminus A_i$. As there are at most k exceptional sets and by Lemma 3.11 each exceptional set has at most $2\xi|V_i|$ vertices, we get that $|N_i| \leq 2\sqrt{\xi}|V_i|$. Let $V_i \in \Pi \setminus F, u \in V_i \setminus N_i$ and let $B(u) \subseteq \Pi \setminus A_i$ be the set of parts such that $u \in \mathcal{E}_i(V_j)$ for all $V_j \in B$. Then $|B| \leq \sqrt{\xi}k$ and

$$\begin{aligned} d(u) &= \sum_{1 \leq j \leq k} |V_j|d(u, V_j) \\ &= \left(\sum_{V_j \in A_i \cup B} |V_j|d(u, V_j) + \sum_{V_j \in \Pi \setminus (A_i \cup B)} |V_j|d(u, V_j) \right) \\ &\leq \left| N(u) \cap \left(\bigcup_{V_j \in A_i \cup B \cup \{V_i\}} V_j \right) \right| + \sum_{1 \leq j \leq k} |V_j|(d(V_i, V_j) + \xi). \end{aligned}$$

By the definition of F and as $u \in V_i \setminus N_i$, we get that

$$\left| \bigcup_{V_j \in A_i \cup B \cup \{V_i\}} V_j \right| \leq (\sqrt{\xi}k + \sqrt{\xi}k + 1)(n/k + 1) \leq 3\sqrt{\xi}n.$$

With that at hand and by using $d(u) \geq \alpha n$ and the fact that the sizes of the parts in Π differ by at most one, we obtain

$$d(u) \leq 3\sqrt{\xi}n + \frac{n}{k} \sum_{1 \leq j \leq k} d(V_i, V_j) + 2\xi n \leq \frac{n}{k} \sum_{1 \leq j \leq k} d(V_i, V_j) + 5\sqrt{\xi}d(u)/\alpha.$$

Let $\eta > 0$. Choosing ξ small enough such that

$$\max\{\xi, 2\sqrt{\xi}, 1/(1 - 5\sqrt{\xi}/\alpha) - 1\} \leq \eta$$

implies the claim. □

The recursion relation. In this section we exploit the properties of the partition to study the expected number of informed vertices after one additional round; our aim is to establish a precise version of (3.16). In the remainder let $\|A\|_F = (\sum_{1 \leq i \leq n} \sum_{1 \leq j \leq n} |a_{i,j}|^2)^{1/2}$ denote the Frobenius norm of a matrix $A \in \mathbb{R}^{n \times n}$.

For the next lemma consider the setting of Theorems 1.3(b) and 1.5(b), that is, we are given an expander sequence $(G_n)_{n \in \mathbb{N}}$ with minimal degree $\delta_n \geq \alpha n$ for some $\alpha > 0$ and an $\varepsilon > 0$. We obtain a sequence of graphs $(\tilde{G}_n)_{n \in \mathbb{N}}$ by deleting up to a $1/2 - \varepsilon$ fraction of the edges at each vertex in G_n . Further, let $\eta > 0, k_0 \in \mathbb{N}$ and $\Pi = \{V_i\}_{i \in [k]}$ be the (η, k_0, K_0) -Szemerédi partition of \tilde{G}_n as given by Lemma 3.12. For that partition define $\mathcal{E}_{i,j} := \mathcal{E}(V_i, V_j)$ as the exceptional set of V_i with respect to V_j given by Lemma 3.11, $i \neq j \in [k]$, F and N_i as the exceptional sets from Lemma 3.12, $i \in \Pi \setminus F$. Moreover, let $\Pi_i = \{V_j \in \Pi \setminus F : (V_i, V_j) \text{ is } \eta\text{-regular}\}$ and note that

$$|\Pi_i| \geq (1 - 2\eta)k, \quad |N_i| \leq \eta|V_i| \quad \text{and} \quad |\mathcal{E}_{i,j}| \leq 2\eta|V_i| \quad \text{for all } i \in \Pi \setminus F, j \in \Pi_i. \tag{3.18}$$

Finally, define

$$\mathcal{H}_{i,j'} = N_i \cup \mathcal{E}_{i,j'}, \quad i \in \Pi \setminus F \text{ and } j \in \Pi_i$$

and

$$X_{t,i,j} = |I_t^{(pp)} \cap (V_i \setminus (N_i \cup \mathcal{E}_{i,j}))|, \quad i \in \Pi \setminus F \text{ and } j \in \Pi_i$$

as well as

$$X_{t,i} = \min_{j \in \Pi_i} X_{t,i,j}, \quad i \in \Pi \setminus F.$$

This definition guarantees that $|I_t^{(pp)}| \geq \|X_t\|_1$. The cornerstone of our proof is the following lemma, which bounds the growth of $X_t = (X_{t,i})_{i \in \Pi \setminus F}$ after one round.

Lemma 3.13. *Consider the situation as described above and assume additionally that $|X_{t,i}| \geq \log \log n$ for all $i \in \Pi \setminus F$ and that $|I_t^{(pp)}| \leq n/\log n$. Then, for all $\nu > 0$ and n large enough, there exists a symmetric matrix A with biggest eigenvalue $\lambda_{\max} \geq 1 + 2q - \nu$ and an error matrix ΔA with $\|\Delta A\|_F \leq \nu$ such that w.h.p.*

$$X_{t+1} \geq (A + \Delta A)X_t.$$

Proof. We set $I_t^{\mathcal{P},i} = I_t^{\mathcal{P}} \cap V_i, U_t^{\mathcal{P},i} = U_t^{\mathcal{P}} \cap V_i$ for $\mathcal{P} \in \{\text{push}, \text{pull}, \text{pp}\}$ and let

$$I_{t+1}^{\mathcal{P},i,j} \setminus I_t = \{u \in U_t \cap V_i \mid \text{there is } v \in I_t \cap V_j \text{ such that } u \text{ gets informed by } v \text{ using } \mathcal{P}\}$$

be the vertices in V_i newly informed in round $t + 1$ by operations involving only vertices from V_i and V_j . Let $(i, j) \in \Pi \setminus F$. For all $u \in U_t^i$ we know that $d(u) \geq \alpha n/2$. Moreover, $|I_t^i| \leq |I_t| \leq n/\log n$. Thus the probability of $u \in U_t^i$ being informed by vertices in I_t^j via *pull* is $q|N(u) \cap I_t^j|/|N(u)| = o(1)$. As the events of u being informed by *push* and *pull* are independent,

$$\mathbb{P}[u \in I_{t+1}^{(\text{push}),i,j} \cap I_{t+1}^{(\text{pull}),i,j}] = o(1)\mathbb{P}[u \in I_{t+1}^{(\text{push}),i,j}].$$

Thus, for any set $S \in V$,

$$\mathbb{E}[|(I_{t+1}^{(pp),i,j} \setminus I_t) \cap S|] = (1 - o(1))(\mathbb{E}[|(I_{t+1}^{(\text{pull}),i,j} \setminus I_t) \cap S|] + \mathbb{E}[|(I_{t+1}^{(\text{push}),i,j} \setminus I_t) \cap S|]). \tag{3.19}$$

Let $i \in \Pi \setminus F$ and $j \in \Pi_i$. We start by determining the expected number of vertices informed by *pull*. Further, set

$$D_i = (1 + \eta) \frac{n}{k} \sum_{1 \leq \ell \leq k} d_{i\ell}.$$

According to Lemma 3.12, all $v \in U_t^i \setminus N_i$ have degree less than D_i . Let $j' \in \Pi_i$. Then

$$\mathbb{E}_t[|I_{t+1}^{(pull),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] = \sum_{u \in U_t^i \setminus \mathcal{H}_{i,j'}} q \frac{|N(u) \cap I_t^j|}{|N(u)|} \geq q \frac{e^{(U_t^i \setminus \mathcal{H}_{i,j'}) \setminus I_t^j}}{D_i}.$$

Since $|I_t^j| \leq |I_t| \leq n/\log n$, we get with room to spare that $|U_t^i \setminus \mathcal{H}_{i,j'}| \geq (1 - 5\eta)n/k$ for n large enough and all $j' \in \Pi_i$. Applying (3.17), where we choose $W = U_t^i \setminus \mathcal{H}_{i,j'}$, yields $|d(U_t^i \setminus \mathcal{H}_{i,j'}, u) - d_{ij}| \leq 6\eta$ for all $u \in V_j \setminus \mathcal{E}_{j,i}$. Thus

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(pull),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] &\geq q \frac{(d_{ij} - 6\eta)|U_t^i \setminus \mathcal{H}_{i,j'}| |I_t^j \setminus \mathcal{E}_{j,i}|}{D_i} \\ &\geq (1 - 5\eta)q \frac{(d_{ij} - 6\eta)|I_t^j \setminus (\mathcal{E}_{j,i} \cup N_j)|}{D_i k/n}. \end{aligned}$$

As $D_i = (1 + \eta)n/k \sum_{1 \leq \ell \leq k} d_{i\ell}$, we get for

$$c_1 := (1 - 6\eta)(1 + \eta)^{-1}$$

with $X_{t,j,i} = |I_t^j \setminus (\mathcal{E}_{j,i} \cup N_j)|$ that

$$\mathbb{E}_t[|I_{t+1}^{(pull),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] \geq c_1 \cdot q \frac{(d_{ij} - 6\eta)X_{t,j,i}}{\sum_{1 \leq \ell \leq k} d_{i\ell}} \quad \text{for all } i \in \Pi \setminus F \text{ and } j, j' \in \Pi_i. \tag{3.20}$$

We continue with *push*. Let $i \in \Pi \setminus F$ and $j, j' \in \Pi_i$, and set (as before)

$$D_j = (1 + \eta) \frac{n}{k} \sum_{1 \leq \ell \leq k} d_{\ell j}.$$

Then

$$\mathbb{E}_t[|I_{t+1}^{(push),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] = \sum_{u \in U_t^i \setminus \mathcal{H}_{i,j'}} \left(1 - \prod_{v \in N(u) \cap I_t^j} \left(1 - \frac{q}{|N(v)|} \right) \right).$$

According to Lemma 3.12 all $v \in I_t^j \setminus N_j$ have degree less than D_j and furthermore $|I_t| = o(n) = o(D_j)$. Thus (3.3) yields the estimate

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(push),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] &\geq \sum_{u \in U_t^i \setminus \mathcal{H}_{i,j'}} \left(1 - \left(1 - \frac{q}{D_j} \right)^{|N(u) \cap (I_t^j \setminus N_j)|} \right) \\ &\geq (1 - o(1)) \sum_{u \in U_t^i \setminus \mathcal{H}_{i,j'}} q \frac{|N(u) \cap (I_t^j \setminus N_j)|}{D_j}. \end{aligned} \tag{3.21}$$

The remaining steps are similar to the previously considered case of *pull*. By assumption we have that $|I_t^j \setminus \mathcal{H}_{j,i}| = X_{t,j,i}$ and as $|I_t^j| \leq |I_t| \leq n/\log n$ we obtain that $|U_t^i \setminus \mathcal{H}_{i,j'}| \geq (1 - 5\eta)n/k$ for n large

enough and all $j' \in \Pi_i$. Using (3.17) we obtain that $|d(U_t^i \setminus \mathcal{H}_{i,j'}, u) - d_{ij}| \leq 6\eta$ for all $u \in V_j \setminus \mathcal{E}_{j,i}$. Thus

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(push),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] &\geq (q - o(1)) \frac{e(U_t^i \setminus \mathcal{H}_{i,j'}, I_t^i \setminus (N_j \cup \mathcal{E}_{j,i}))}{D_j} \\ &\geq (q - o(1)) \frac{(d_{ij} - 6\eta) |U_t^i \setminus \mathcal{H}_{i,j'}| X_{t,j,i}}{D_j}. \end{aligned}$$

Using the fact that $D_j = (1 + \eta)n/k \sum_{1 \leq \ell \leq k} d_{\ell j}$, we get for the same constant c_1 as in (3.20) and n large enough

$$\mathbb{E}_t[|I_{t+1}^{(push),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] \geq c_1 \cdot q \frac{(d_{ij} - 6\eta) X_{t,j,i}}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \quad \text{for all } i \in \Pi \setminus F \text{ and } j, j' \in \Pi_i. \tag{3.22}$$

With (3.19), we can combine the results for *pull*, (3.20), and *push*, (3.22), to get for $c_2 := c_1 - \eta$

$$\mathbb{E}_t[|I_{t+1}^{(pp),ij} \setminus (I_t^i \cup \mathcal{H}_{i,j'})|] \geq c_2 \cdot q \left(\frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right) X_{t,j,i} \quad \text{for all } i \in \Pi \setminus F, j, j' \in \Pi_i. \tag{3.23}$$

Next we will show how we can exploit (3.23) to obtain (a lower bound for) $\mathbb{E}_t[|(I_{t+1}^{(pp),i} \setminus I_t)|]$. Let $i \in \Pi \setminus F$ and $u \in U_t^i$. Using $|I_t| = o(n)$ and (3.3) we obtain

$$\mathbb{P}_t[u \in I_{t+1}^{(push),i} \setminus I_t] = 1 - \prod_{i \in N(u) \cap I_t} \left(1 - \frac{1}{|N(i)|} \right) = (1 - o(1)) \sum_{i \in N(u) \cap I_t} \frac{1}{|N(i)|}.$$

Let $W \subseteq V$. Using (3.19), the previous equation and that Π is a partition, we get

$$\begin{aligned} \mathbb{E}_t[|(I_{t+1}^{(pp),i} \setminus I_t) \cap W|] &= (1 - o(1)) \sum_{u \in U_t^i \cap W} \left(\frac{|N(u) \cap I_t|}{|N(u)|} + \sum_{i \in N(u) \cap I_t} \frac{1}{|N(i)|} \right) \\ &= (1 - o(1)) \sum_{u \in U_t^i \cap W} \left(\sum_{j \in [k]} \left(\frac{|N(u) \cap I_t \cap V_j|}{|N(u)|} + \sum_{i \in N(u) \cap I_t \cap V_j} \frac{1}{|N(i)|} \right) \right) \\ &= (1 - o(1)) \sum_{j \in [k]} \mathbb{E}_t[|(I_{t+1}^{(pp),ij} \setminus I_t) \cap W|]. \end{aligned}$$

Choose $W = V \setminus \mathcal{H}_{i,j'}$. Then the previous equation implies

$$\mathbb{E}_t[|I_{t+1}^{(pp),i} \setminus (I_t \cup \mathcal{H}_{i,j'})|] \geq (1 - o(1)) \sum_{j \in \Pi \setminus F} \mathbb{E}_t[|I_{t+1}^{(pp),ij} \setminus (I_t \cup \mathcal{H}_{i,j'})|] \quad \text{for all } i \in \Pi \setminus F, j' \in \Pi_i,$$

which in turn, using (3.23) and $X_{t,j,i} \geq X_{t,j}$ for all $j \in \Pi \setminus F$ and $i \in \Pi_j$, implies for $c := c_2 - \eta$

$$\mathbb{E}_t[X_{t+1,i,j'}] \geq X_{t,i} + c \cdot q \sum_{j \in \Pi_i} \left(\frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right) X_{t,j} \quad \text{for all } i \in \Pi \setminus F, j' \in \Pi_i. \tag{3.24}$$

Assume that (3.24) holds not only in expectation but also for a slightly smaller c , say $c - \eta$, with high probability. We are going to show this at the end of the proof. Using this assumption and a union bound over $j' \in \Pi_i$ gives w.h.p.

$$X_{t+1,i} = \min_{j' \in \Pi_i} X_{t+1,i,j'} \geq \langle a_i, (X_{t,j})_{j \in \Pi_i} \rangle \quad \text{for all } i \in \Pi \setminus F, \tag{3.25}$$

where for $i \in \Pi \setminus F$ and $j \in \Pi_i$ we have

$$a_{ij} = \mathbb{1}[i = j] + c \cdot q \left(\frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right). \tag{3.26}$$

Let A be the $|\Pi \setminus F| \times |\Pi \setminus F|$ matrix with entries as in the previous equation, that is, $A = (a_{ij})_{(i,j) \in (\Pi \setminus F)^2}$ is given by (3.26) for all $(i, j) \in (\Pi \setminus F)^2$. Note that A is symmetric. Then we obtain from (3.25)

$$X_{t+1} \geq B \cdot X_t,$$

with $B = A + \Delta A$, where

$$(\Delta A)_{ij} = \begin{cases} 0 & i \in \Pi \setminus F \text{ and } j \in \Pi_i, \\ -a_{ij} & i \in \Pi \setminus F \text{ and } j \in \Pi \setminus (F \cup \Pi_i). \end{cases}$$

Set

$$F' := \{(i, j) \in (\Pi \setminus F)^2 \mid j \in \Pi \setminus (F \cup \Pi_i)\}.$$

As $d(u) \geq \alpha n/2$ for all $u \in V$ and some $\alpha > 0$, we also know that $\sum_{1 \leq \ell \leq k} d_{\ell j} \geq k\alpha/2$. Together with $0 \leq d_{i,j} \leq 1$ for all $(i, j) \in [k]^2$, we get that

$$\left| \frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{i\ell}} \right| \leq \frac{2}{\alpha k}.$$

Using the fact that $|F'| \leq 2\eta k^2$ (see (3.18)), we obtain

$$\|\Delta A\|_F^2 = \sum_{(i,j) \in F'} a_{ij}^2 \leq \sum_{(i,j) \in F'} \left(\frac{4}{\alpha k}\right)^2 \leq 2\eta k^2 \left(\frac{4}{\alpha k}\right)^2 = \frac{4^2 \cdot 2 \cdot \eta}{\alpha^2}$$

and thus $\|\Delta A\|_F \leq 4\sqrt{2\eta}/\alpha$. This leaves us with bounding the biggest eigenvalue λ_{\max} of A . Using the well-known inequality for symmetric matrices,

$$\lambda_{\max} \geq \sum_{(i,j) \in (\Pi \setminus F)^2} A_{ij}/|\Pi \setminus F|,$$

we obtain

$$\begin{aligned} \lambda_{\max} &\geq \frac{1}{|\Pi \setminus F|} \sum_{(i,j) \in (\Pi \setminus F)^2} a_{ij} \\ &\geq \frac{1}{|\Pi \setminus F|} \left(\sum_{(i,i) \in (\Pi \setminus F)^2} 1 + \sum_{(i,j) \in [k]^2} \frac{cq(d_{ij} - 6\eta)}{\sum_{1 \leq \ell \leq k} d_{i\ell}} \right. \\ &\quad \left. + \sum_{(i,j) \in [k]^2} \frac{cq(d_{ij} - 6\eta)}{\sum_{1 \leq \ell \leq k} d_{\ell j}} - 2 \sum_{i \in [k] \setminus (\Pi \setminus F)} \sum_{j \in [k]} \frac{cq}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right). \end{aligned}$$

Note that $|\Pi \setminus F| \geq (1 - \eta)k$, $|[k] \setminus (\Pi \setminus F)| \leq \eta k$. Moreover, $\sum_{1 \leq \ell \leq k} d_{\ell j} \geq \alpha k/2$ for all $j \in [k]$. Thus

$$\lambda_{\max} \geq 1 + \frac{1}{k} \left(cqk + cqk - 12cq \sum_{(i,j) \in [k]^2} \frac{\eta}{\sum_{1 \leq \ell \leq k} d_{\ell j}} - 2cq \frac{\eta k^2}{\alpha k/2} \right) \geq 1 + 2cq(1 - 8\eta/\alpha).$$

Choosing η small enough such that $2q(1 - c(1 - 8\eta/\alpha))$, $4\sqrt{2\eta}/\alpha \leq \nu$ implies the claim of this lemma.

This leaves us with proving that (3.24) also holds with high probability. As $|I_{t+1}^{(pp)}|$ conditioned on I_t is a self-bounding function, so is $|I_{t+1}^{(pp),i} \setminus I_t^i|$ for all $i \in \Pi \setminus F$ and therefore also $|I_{t+1}^{(pp),i} \setminus (I_t \cup \mathcal{H}_{i,j'})| =: Y_{t+1,i,j'}$ for all $i \in \Pi \setminus F$ and $j' \in \Pi_i$. Note that $Y_{t+1,i,j'} = X_{t+1,i,j'} - X_{t,i,j'}$. Lemma 2.2 yields that

$$\mathbb{P}_t[Y_{t+1,i,j'} \geq (1 - \mathbb{E}_t[Y_{t+1,i,j'}])^{-1/3} \mathbb{E}_t[Y_{t+1,i,j'}]] \geq 1 - \mathbb{E}_t[Y_{t+1,i,j'}]^{-1/3}$$

and therefore, setting

$$Z_{t,i} = c \cdot q \sum_{j \in \Pi_i} \left(\frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{i\ell}} + \frac{d_{ij} - 6\eta}{\sum_{1 \leq \ell \leq k} d_{\ell j}} \right) X_{t,j} \quad \text{for all } i \in \Pi \setminus F,$$

and using (3.24), i.e. $\mathbb{E}_t[Y_{t+1,i,j'}] \geq Z_{t,i}$ for all $i \in \Pi \setminus F$ and $j' \in \Pi_i$, we get with probability at least $1 - k^3 Z_{t,i}^{-1/3}$

$$Y_{t+1,i,j'} \geq (1 - Z_{t,i}^{-1/3}) Z_{t,i} \quad \text{for all } i \in \Pi \setminus F \text{ and } j' \in \Pi_i.$$

This and $|I_t^i| \geq X_{t,i}$ for all $i \in \Pi \setminus F$ implies that (3.24) also holds with high probability for a marginally smaller c , as claimed. \square

Extension. We now solve the linear recurrence relation above and extend it to more than one round to get an upper bound on the runtime of *push&pull*. We first state a Chernoff bound that will be very useful in the next lemma.

Lemma 3.14 ([26]). *Let $\varepsilon, \delta > 0$. Suppose that X_1, \dots, X_n are independent geometric random variables with parameter δ , so $\mathbb{E}[X_i] = 1/\delta$ for each i . Let $X := \sum_{1 \leq i \leq n} X_i$, $\mu = \mathbb{E}[X] = n/\delta$. Then*

$$\mathbb{P}[X \geq (1 + \varepsilon)\mu] \leq e^{-n(\varepsilon - \log(1+\varepsilon))} \leq e^{-\varepsilon^2 n/2(1+\varepsilon)}$$

Together with Lemma 2.11 the following lemma implies Theorems 1.3(b) and 1.5(b).

Lemma 3.15. *Consider the setting of Theorems 1.3(b) and 1.5(b) and let $I_t = I_t^{(pp)}$. The following statements hold w.h.p.*

- (a) *Let $S \subseteq V_n$, $|S| = \Theta(n)$. Then there is $t = \Theta(\log \log n)$ such that w.h.p. $|I_t| \geq |I_t \cap S| \geq \log \log n$.*
- (b) *Let $\log \log n \leq |I_t| \leq n/\log n$. Then there is $\tau \leq \log_{1+2q}(n/|I_t|) + o(\log n)$ such that $|I_{t+\tau}| > n/\log n$.*
- (c) *Let $n/\log n \leq |I_t| \leq n - n/\log n$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| > n - n/\log n$.*
- (d) *Let $|I_t| \geq n - n/\log n$ and $q = 1$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| = n$.*

Proof. As $|I_t^{(pp)}| \geq |I_t^{(pull)}|$ clearly (c) and (d) follow from Lemma 3.2. We show (a) by determining a lower bound for the probability that an arbitrary vertex gets informed after a constant number of rounds. Set $\beta = \min\{\alpha, \epsilon\}$, let $S_0 = \{u\}$ and choose $w \in V$, $w \neq u$. By Lemma 2.9 there is $d \leq 8/\beta^2 + 2$ and $c = (\beta^4/64)^{8/\beta^2+3} \in (0, 1)$ such that there are at least cn^{d-1} paths of (edge) length d from u to w . Let $\gamma = (u, v_1, \dots, v_{d-1}, w)$ be such a path from u to w , and let A_γ denote the event that w is informed via γ after exactly d rounds performing only *push* operations, that is, A_γ is the event that in the first round the randomly selected neighbour of u is v_1 , in the second round the randomly selected neighbour of v_1 is v_2 and so forth, until in the d th round the randomly selected neighbour of v_{d-1} is w . Obviously, the probability of A_γ is bounded from below by n^{-d} . Further, let $\gamma' \neq \gamma$ be another path from u to w with length d . As γ and γ' differ by at least one edge we

readily obtain that $\mathbb{P}[A_\gamma \cap A_{\gamma'}] = 0$. Let Γ denote the set of all paths with length d from u to w . Having done these preparations we use them to conclude for all $w \in V$ and $t \geq 0$

$$\mathbb{P}_t[w \in I_{t+d}] \geq \mathbb{P}_t\left[\bigcup_{\gamma \in \Gamma} A_\gamma\right] \geq \sum_{\gamma \in \Gamma} \mathbb{P}_t[A_\gamma] \geq \sum_{\gamma \in \Gamma} n^{-d} \geq \frac{c}{n}. \tag{3.27}$$

We define a modified protocol as follows. Wait $d := \lceil 8/\beta^2 + 2 \rceil$ rounds, after that with probability c choose one uninformed vertex uniformly at random and set it as informed. Repeat. Call the vertices informed by this algorithm I_t^* . Then the probability of any vertex being informed after d rounds is

$$\mathbb{P}_t[v \in I_{t+d}^* | v \notin I_t^*] = c/n.$$

Thus, for any $t \geq 0$,

$$\mathbb{P}_t[v \in I_{t+d} | v \in U_t] \geq \mathbb{P}_t[v \in I_{t+d}^* | v \notin I_t^*] = c/n.$$

Note that for any $s \in \mathbb{N}$ the set I_{sd}^* is generated by a very simple procedure: s times independently, with probability c , we choose a random vertex and put it into I_{sd}^* . Thus $|I_{sd}^* \cap S|$ is binomially distributed with s trials, where each one has success probability $c|S|/n = \Theta(c)$; it follows readily that $|I_{sd}^* \cap S|$ concentrates around a multiple of s for large s , and the claim follows by choosing $s = \Theta(\log \log n)$.

This leaves (b) to be shown. Part (a) implies that there is some $t_0 = o(\log n)$ such that $X_{t_0,i} = \Theta(\log \log n)$ for all $i \in \Pi \setminus F$ by choosing $S = V_i \setminus (N_i \cup \mathcal{E}_{i,j})$, $j \in \Pi_i$ and applying a union bound over i and j . Thus we can apply Lemma 3.13. It gives w.h.p., say with probability $1 - g(n) = 1 - o(1)$, that $X_{t+1} \geq (A + \Delta A)X_t$, A has maximal eigenvalue $\lambda_{\max}(A) \geq 1 + 2q - \nu$ and $\|\Delta A\|_F \leq \nu$. Then $B := A + \Delta A$ has maximal eigenvalue

$$\lambda_{\max}(B) \geq \lambda_{\max}(A) - \|\Delta A\|_F \geq 1 + 2q - 2\nu$$

(Theorem of Wielandt and Hoffmann; see e.g. [24]).

Set $f(n) := (\log(n/\log n))^{2/3}$. Our assumptions guarantee that $f(n) = \omega(1)$ and $f(n) = o(\log n)$. Moreover, set

$$\tau := \frac{1}{1 - g(n)} \cdot \frac{\log(n/\log n)}{\log(\lambda_{\max}(B))} + f(n) = \frac{\log n}{\log(\lambda_{\max}(B))} + o(\log n).$$

Let $(X_i)_{i \in \mathbb{N}}$ be independent and identically distributed geometric random variables with expectation $1 - g(n)$. Set $X = X_1 + X_2 + \dots + X_{\mathcal{T}}$ with $\mathcal{T} = \log(n/\log n) / \log(\lambda_{\max}(B))$. We show that $\mathbb{P}[X \leq \tau] = 1 - o(1)$. To see this, note first that by linearity of expectation $\mathbb{E}[X] = \tau - f(n)$. Then, by Lemma 3.14,

$$\mathbb{P}[X \leq \tau] = \mathbb{P}\left[X \leq \left(1 + \frac{f(n)}{\tau + f(n)}\right)\mathbb{E}[X]\right] \geq 1 - \exp\left(-\Theta\left(\frac{f(n)^2}{\tau}\right)\right) = 1 - o(1).$$

Thus we have w.h.p.

$$|I_{t+\tau}| \geq \|X_{t+\tau}\|_1 \geq \|B^{\mathcal{T}} X_{t_0}\|_1.$$

Let v be an eigenvector of B to $\lambda_{\max}(B)$. As $v \neq 0$ there is an index ℓ such that $v_\ell \neq 0$. Without loss of generality we can assume that $v_\ell = 1$, as v/v_ℓ is also an eigenvector to $\lambda_{\max}(B)$. Thus $(B^{\mathcal{T}} v)_\ell = \lambda_{\max}(B)^{\mathcal{T}}$, $(B^{\mathcal{T}}(X_{t_0} - v))_i \geq 0$ for all $1 \leq i \leq k$ and therefore

$$|I_{t+\tau}| \geq (B^{\mathcal{T}} X_{t_0})_\ell \geq (B^{\mathcal{T}}(v + X_{t_0} - v))_\ell = (B^{\mathcal{T}} v)_\ell + (B^{\mathcal{T}}(X_{t_0} - v))_\ell \geq (B^{\mathcal{T}} v)_\ell \geq \lambda_{\max}(B)^{\mathcal{T}}.$$

Our choice of \mathcal{T} yields w.h.p. $|I_{t+\tau}| \geq \lambda_{\max}(B)^{\mathcal{T}} \geq n/\log n$. Note that, since $\nu > 0$ was chosen arbitrarily, we in fact have that $\tau \leq \log_{1+2q}(n) + o(\log n)$, and the proof is completed. \square




Figure 1. Plotted values of Δ in $T_{pp}(G_n(\varepsilon), q) - c_{pp} \log n = \Delta \log n + o(\log n)$, for $0.9 < q < 1$ and $0 < \varepsilon < 1/2$.

3.6 Proof of Theorem 1.4(b) – edge deletions may slow down push&pull

For any $0 < \varepsilon < 1/2$, $q \in (0, 1)$ we consider a sequence of graphs $(G_n(\varepsilon))_{n \in \mathbb{N}} = ((V_n, E_n))_{n \in \mathbb{N}}$ that is similar to the one studied in the proof of Theorem 1.4(a). Let $V_n = A_n \cup B_n$ with $A_n := \{1, \dots, \lfloor n/2 \rfloor\}$, $B_n := \{\lfloor n/2 \rfloor + 1, \dots, n\}$ and $\deg(v) = n - 1$ for all $v \in A_n$. Let the induced subgraph of B_n be a random graph in which each edge is included independently with probability $p = 1 - 2\varepsilon$. We know and it is easy to show (see e.g. [15, Section IV]) that w.h.p. this subgraph is almost regular, that is,

$$d_{B_n}(v) = (1 + o(1))(1 - 2\varepsilon)n/2 \quad \text{for all } v \in B_n, \tag{3.28}$$

and is an expander, which means that for every $S_n \subseteq B_n$, $1 \leq |S_n| \leq n/4$ and $d_{B_n} := (1 - 2\varepsilon)n/2$ we have

$$e(S_n, B_n \setminus S_n) = (1 + o(1)) \frac{d_{B_n} |S_n| |B_n \setminus S_n|}{|B_n|} = (1 - 2\varepsilon + o(1)) |S_n| |B_n \setminus S_n|. \tag{3.29}$$

First we give a statement that describes the expected number of informed vertices after performing one round of *push&pull*.

Lemma 3.16. *Let $G_n(\varepsilon) = (A_n \cup B_n, E_n)$ be as above.*

(a) *Let $\sqrt{\log n} \leq |I_t| \leq n / \log n$ and set*

$$X_t = (|I_t^{(pp),(A)}|, |I_t^{(pp),(B)}|) := (|I_t^{(pp)} \cap A_n|, |I_t^{(pp)} \cap B_n|).$$

Then $\mathbb{E}_t[X_{t+1}] = (1 + o(1))MX_t$, where

$$M = \begin{pmatrix} 1 + q & q(1 + \varepsilon/(2 - 2\varepsilon)) \\ q(1 + \varepsilon/(2 - 2\varepsilon)) & 1 + q(1 - 2\varepsilon/(2 - 2\varepsilon)) \end{pmatrix}.$$

(b) *Let $|U_t^{(pp)}| \leq n / \log n$. Then*

$$\mathbb{E}_t[|U_{t+1}^{(pp)}|] \leq (1 + o(1))e^{-q(1/2+(1/2-\varepsilon)/(1-\varepsilon))}(1 - q)|U_t|.$$

Proof. For $J \in \{A, B\}$, $J_n \in \{A_n, B_n\}$ set $U_t^{(J)} := U_t \cap J_n$, $I_t^{(J)} := I_t \cap J_n$ and $I_{t+1}^{(pp),(J)} = I_{t+1}^{(pp)} \cap J_n$. We first prove (a) by computing the expected number of informed vertices after a single round. Since

$d(u) = \Theta(n)$ for all $u \in V_n$ and $|I_t| \leq n/\log n$, the probability of $u \in U_t$ being informed by *pull* is

$$\mathbb{P}_t[u \in I_{t+1}^{(pull)} \setminus I_t] = \frac{q|N(u) \cap I_t|}{|N(u)|} = o(1).$$

As the events of u being informed by *push* and *pull* are independent, we have

$$\mathbb{P}_t[u \in (I_{t+1}^{(push)} \cap I_{t+1}^{(pull)}) \setminus I_t] = o(1)\mathbb{P}_t[u \in I_{t+1}^{(push)} \setminus I_t].$$

Thus

$$\mathbb{E}_t[|I_{t+1}^{(pp)} \setminus I_t|] = (1 + o(1))(\mathbb{E}_t[|I_{t+1}^{(push)} \setminus I_t|] + \mathbb{E}_t[|I_{t+1}^{(pull)} \setminus I_t|]).$$

We look at *pull* in detail first. Recall that

$$\deg(v) = n - 1 \text{ for all } v \in A_n \quad \text{and} \quad \deg(v) = (1 + o(1))(1 - \varepsilon)n \text{ for all } v \in B_n.$$

Moreover, using (3.29), we obtain

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(pull)} \setminus I_t|] &= \sum_{u \in U_t} q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &= \sum_{u \in U_t^{(A)}} q \frac{|N(u) \cap I_t|}{|N(u)|} + \sum_{u \in U_t^{(B)}} q \frac{|N(u) \cap I_t|}{|N(u)|} \\ &= (q + o(1)) \frac{n}{2} \left(\frac{|I_t^{(A)}| + |I_t^{(B)}|}{n} + \frac{|I_t^{(A)}| + (1 - 2\varepsilon)|I_t^{(B)}|}{(1 - \varepsilon)n} \right) \end{aligned}$$

and thus

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(pull),(A)} \setminus I_t|] &= (q + o(1)) \frac{|I_t^{(A)}| + |I_t^{(B)}|}{2}, \\ \mathbb{E}_t[|I_{t+1}^{(pull),(B)} \setminus I_t|] &= (q + o(1)) \frac{|I_t^{(A)}| + (1 - 2\varepsilon)|I_t^{(B)}|}{2(1 - \varepsilon)}. \end{aligned}$$

Next we consider *push*. By using $|I_t| = o(n) = o(\delta_{G_n(\varepsilon)})$ and (3.3), we obtain

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(push)} \setminus I_t|] &= \sum_{u \in U_t} 1 - \prod_{i \in N(u) \cap I_t} \left(1 - \frac{q}{|N(i)|} \right) \\ &= \sum_{u \in U_t} (1 + o(1)) \sum_{i \in N(u) \cap I_t} \frac{q}{|N(i)|} \\ &= (q + o(1)) \sum_{u \in U_t} \left(\frac{|I_t^{(A)}|}{n} + \frac{\mathbb{1}[u \in U_t^{(A)}]|I_t^{(B)}| + \mathbb{1}[u \in U_t^{(B)}]|N(u) \cap I_t^{(B)}|}{(1 - \varepsilon)n} \right) \end{aligned}$$

and thus, with $|U_t^{(A)}|, |U_t^{(B)}| = (1 - o(1))n/2$ and (3.29),

$$\begin{aligned} \mathbb{E}_t[|I_{t+1}^{(push),(A)} \setminus I_t|] &= (q + o(1)) \left(\frac{|I_t^{(A)}|}{2} + \frac{|I_t^{(B)}|}{2} + \frac{\varepsilon|I_t^{(B)}|}{2(1 - \varepsilon)} \right), \\ \mathbb{E}_t[|I_{t+1}^{(push),(B)} \setminus I_t|] &= (q + o(1)) \left(\frac{|I_t^{(A)}|}{2} + \frac{|I_t^{(B)}|}{2} - \frac{\varepsilon|I_t^{(B)}|}{2(1 - \varepsilon)} \right). \end{aligned}$$

Accumulating the calculated expectations for *pull* and *push* yields the claim.

Next we show (b). The assumption implies that $|I_t| = (1 - o(1))n$ and therefore $|I_t^{(A)}| = |I_t^{(B)}| = (1 - o(1))n/2$. Let D_u be the event that an uninformed vertex u does not get informed by the

push algorithm; let E_u be the corresponding event for pull. Then D_u and E_u are independent and $D_u \cap E_u$ is the event that u does not get informed in the current round. Let $u \in U_t^{(A)}$, and then

$$\begin{aligned} \mathbb{P}_t[D_u] &= \prod_{v \in I_t^{(A)}} \left(1 - \frac{q}{|N(v)|}\right) \prod_{v \in I_t^{(B)}} \left(1 - \frac{q}{|N(v)|}\right) \\ &= (1 - o(1)) \left(1 - \frac{q}{n}\right)^{|I_t^{(A)}|} \left(1 - \frac{q}{(1 - \varepsilon)n}\right)^{|I_t^{(B)}|} \\ &= e^{-q(1/2+1/(2(1-\varepsilon)))} + o(1) \\ &\leq e^{-q(1/2+(1-2\varepsilon)/(2(1-\varepsilon)))} + o(1) \end{aligned}$$

and

$$\mathbb{P}_t[E_u] = 1 - \frac{q|N(u) \cap I_t|}{|N(u)|} = 1 - \frac{q|I_t|}{n-1} = 1 - q + o(1).$$

Now consider $u \in U_t^{(B)}$; then according to (3.28) we have

$$|N(u) \cap I_t^{(B)}| = |N(u) \cap B_n| - |N(u) \cap U_t^{(B)}| = (1 + o(1))(1 - 2\varepsilon)n/2.$$

Therefore

$$\begin{aligned} \mathbb{P}_t[D_u] &= \prod_{v \in I_t^{(A)}} \left(1 - \frac{q}{|N(v)|}\right) \prod_{v \in N(u) \cap I_t^{(B)}} \left(1 - \frac{q}{|N(v)|}\right) \\ &= (1 - o(1))e^{-q/2} \left(1 - \frac{q}{(1 - \varepsilon)n}\right)^{|N(u) \cap I_t^{(B)}|} \\ &= e^{-q(1/2+(1-2\varepsilon)/(2(1-\varepsilon)))} + o(1) \end{aligned}$$

and

$$\mathbb{P}_t[E_u] = 1 - \frac{q|N(u) \cap I_t|}{|N(u)|} = 1 - (1 + o(1)) \frac{q(|I_t^{(A)}| + |N(u) \cap I_t^{(B)}|)}{(1 - \varepsilon)n} = 1 - q + o(1).$$

Combining the results for $u \in U_t^{(A)}$ and $u \in U_t^{(B)}$, we get

$$\mathbb{E}_t[|U_{t+1}|] = \sum_{u \in U_t} \mathbb{P}_t[D_u] \mathbb{P}_t[E_u] \leq (1 + o(1))e^{-q(1/2+(1/2-\varepsilon)/(1-\varepsilon))} (1 - q)|U_t|. \quad \square$$

Remark 3.1. Let λ_{\max} be the greatest eigenvalue of M as defined in Lemma 3.16(a). Then

$$\lambda_{\max} = 1 + 2q + \left(2q \left(\sqrt{\varepsilon^2/2 - \varepsilon + 1} - 1\right) + q\varepsilon\right) / (2 - 2\varepsilon) > 1 + 2q.$$

Next comes a lemma that bounds the runtime of push-pull on $G_n(\varepsilon)$. In particular, (a) and (c) of Lemma 3.17 provide a lower bound on the runtime, and (a), (b) and (d) of Lemma 3.17 together with Lemma 3.15(a) provide an upper bound.

Lemma 3.17. Let $I_t = I_t^{(pp)}$, $\varepsilon > 0$ and $\lambda = \lambda_{\max}(M)$ be the greatest eigenvalue of M as given in Lemma 3.16(a). Consider $G_n(\varepsilon)$.

- (a) Let $\sqrt{\log n} \leq |I_t| \leq n/\log n$. Then there are $\tau_1, \tau_2 = \log_\lambda(n/|I_t|) + o(\log n)$ such that $|I_{t+\tau_1}| < n/\log n < |I_{t+\tau_2}|$.

- (b) Let $n/\log n \leq |I_t| \leq n - n/\log n$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| > n - n/\log n$.
- (c) Let $|I_t| \leq n/\log n$. Then there is

$$\tau \geq \log n / \log ((1 - q)^{-1} \exp (q(1/2 + (1/2 - \varepsilon)/(1 - \varepsilon)))) - o(\log n)$$

such that $|I_{t+\tau}| < n$.

- (d) Let $|I_t| \geq n - n/\log n$ and $q \in (0, 1)$. Then there is

$$\tau \leq \log n / \log ((1 - q)^{-1} \exp (q(1/2 + (1/2 - \varepsilon)/(1 - \varepsilon)))) + o(\log n)$$

such that $|I_{t+\tau}| = n$.

Proof. We do not give a proof for (b) as it follows immediately from Lemma 3.15(a). For $J \in \{A, B\}$ set $U_t^{(J)} := U_t \cap J_n, I_t^{(J)} := I_t \cap J_n$. We prove (a) first. Let $t_0 > 0$ be the first round such that $|I_{t_0}| \geq \log \log n$ and set X_t and M as in Lemma 3.16(a); note that Lemma 3.15(a) also gives that $(X_{t_0})_i \geq \log \log n/2$ for $i \in \{1, 2\}$. Then, for all $t \geq t_0$ such that $|I_t| \leq n/\log n$, we obtain from Lemma 3.16(a) that $\mathbb{E}_t[X_{t+1}] = (1 + o(1))MX_t$ and, in particular, $\mathbb{E}_t[(X_{t+1})_i] = \Theta(|I_t|)$ for $i \in \{1, 2\}$. As every component of X_t is self-bounding, Lemma 2.1 applies and we get for $i \in \{1, 2\}$

$$\mathbb{P}_t[|(X_{t+1})_i - \mathbb{E}_t[(X_{t+1})_i]| \geq \mathbb{E}_t[(X_{t+1})_i]^{2/3}] = O(|I_t|^{-1/3})$$

and by the union bound, provided that $|I_t| \leq n/\log n$,

$$\mathbb{P}_t \left[\bigcap_{i \in \{1,2\}} (|(X_{t+1})_i - \mathbb{E}_t[(X_{t+1})_i]| \leq \mathbb{E}_t[(X_{t+1})_i]^{2/3}) \right] = 1 - O(|I_t|^{-1/3}). \tag{3.30}$$

Using (3.30) we want to find a bound on $|I_{t+1}|$. As long as $|I_t| \leq n/\log n$, we get

$$((1 - O(|I_{t_0}|^{-1/3}))M)^{t+1-t_0} X_{t_0} \leq X_{t+1} \leq ((1 + O(|I_{t_0}|^{-1/3}))M)^{t+1-t_0} X_{t_0}.$$

As seen in Remark 3.1, M has maximal eigenvalue $\lambda_{\max} > 1$, and as M is a positive matrix there is a positive eigenvector v to λ_{\max} ; see [32]. This gives constants $c_1, c_2 > 0$ such that $c_1 v \log \log n \leq X_{t_0} \leq c_2 v \log \log n$, and for t large enough

$$\frac{c_1}{c_2} ((1 - O(|I_{t_0}|^{-1/3}))\lambda_{\max})^{t+1-t_0} X_{t_0} \leq X_{t+1} \leq \frac{c_2}{c_1} ((1 + O(|I_{t_0}|^{-1/3}))\lambda_{\max})^{t+1-t_0} X_{t_0},$$

and therefore

$$|I_{t+1}| \leq \frac{c_1}{c_2} ((1 + o(1))\lambda_{\max})^{t-t_0} |I_{t_0}|,$$

as long as the right-hand side is bounded by $n/\log n$. For all these t we also get

$$|I_{t+1}| \geq \frac{c_2}{c_1} ((1 - o(1))\lambda_{\max})^{t-t_0} |I_{t_0}|.$$

Proceeding as in Lemmas 2.5 and 2.6 and their proofs, where we replace the events

$$|I_t| \geq \mathbb{E}_{t-1}[|I_t|] - \mathbb{E}_{t-1}[|I_t|]^{2/3} \quad \text{or} \quad |I_t| \geq n/g(n)$$

and

$$||I_t| - \mathbb{E}_{t-1}[|I_t|]| \leq \mathbb{E}_{t-1}[|I_t|]^{2/3}$$

with

$$\bigcap_{i \in \{1,2\}} ((X_{t+1})_i \geq (1 - \mathbb{E}_t[(X_{t+1})_i]^{-1/3})\mathbb{E}_t[(X_{t+1})_i]) \quad \text{or} \quad |I_t| \geq n/\log n$$

and

$$\bigcap_{i \in \{1,2\}} (|(X_{t+1})_i - \mathbb{E}_t[(X_{t+1})_i]| \leq \mathbb{E}_t[(X_{t+1})_i]^{2/3})$$

we obtain the statement. Next we show (c). The assumption guarantees that less than $n/\log n$ vertices are informed. Thus $|U_t^{(B)}| \geq n/2 - |I_t| \geq (1/2 - 1/\log n)n$. We consider a modified dissemination process, where in each round, each uninformed vertex always chooses an informed neighbour (but does not necessarily get informed as the message transmission may fail), and additionally each vertex chooses a neighbour independently and uniformly at random and after this round the chosen vertex is informed with probability q ; in other words, we assume that uninformed vertices can inform other vertices. In this modified process the probability of an uninformed vertex $u \in U_t^{(B)}$ staying uninformed after performing one round is given by the product of the probabilities of not being informed by *pull* or via *push* by a vertex in A_n or B_n . Using (3.29) and $(1 - 1/n)^n = e^{-1+o(1)}$, we get $g(n) = o(1)$ such that

$$\begin{aligned} \mathbb{P}_t[u \in U_{t+1}^{(B)}] &= (1 - q) \left(1 - \frac{q}{n}\right)^{n/2} \left(1 - \frac{q}{(1 - \varepsilon)n}\right)^{|N(u) \cap B_n|} \\ &= (1 - q) \exp\left(-q\left(\frac{1}{2} + \frac{1/2 - \varepsilon}{1 - \varepsilon}\right) + g(n)\right). \end{aligned}$$

As we have seen in the proof of Lemma 3.16(b), the probability of being informed by *push&pull* is greater for a vertex in A_n than for a vertex in B_n . Therefore it is sensible to expect that some vertices in B_n will be the last to be informed. Consequently let E_u denote the event that a currently uninformed vertex $u \in U_t^{(B)}$ does not get informed in this modified version within the next

$$\tau := \frac{1}{\log((1 - q)^{-1} \exp(q(1/2 + (1/2 - \varepsilon)/(1 - \varepsilon) - g(n))))} \log(n) - h(n)$$

rounds, where $h = o(\log n)$ and $h = \omega(1)$. Therefore we have

$$\mathbb{P}_t[E_u] = \left((1 - q) \exp\left(-q\left(\frac{1}{2} + \frac{1/2 - \varepsilon}{1 - \varepsilon}\right) + g(n)\right) \right)^\tau = \frac{1}{n} e^{\omega(1)}.$$

In this modified model the events $\{E_u \mid u \in U_t^{(B)}\}$ satisfy that there is $p = \omega(n^{-1})$ such that

$$\mathbb{P}_t[E_u \mid \{\overline{E}_v : v \in U\}] \geq p \quad \text{for all } u \in B_n \text{ and } U \subseteq V \setminus \{u\}.$$

This follows immediately by the above calculations. Thus, as $|U_t^{(B)}| = \Theta(n)$,

$$\mathbb{P}_t\left[\bigwedge_{u \in U_t^{(B)}} \overline{E}_u\right] \leq \prod_{u \in U_t^{(B)}} (1 - p) \leq \exp\left(-\sum_{u \in U_t^{(B)}} p\right) = o(1).$$

Finally we show (d). By Lemma 3.16(b), we obtain that for any $\tau \in \mathbb{N}$,

$$\mathbb{E}_t[|U_{t+\tau}|] \leq ((1 + o(1))e^{-q(1/2+(1/2-\varepsilon)/(1-\varepsilon))} (1 - q))^\tau |U_t|.$$

Then, for some

$$\tau := \frac{\log(n)}{\log((1 - q)^{-1} \exp(q(1/2 + (1/2 - \varepsilon)/(1 - \varepsilon))))} + o(\log n),$$

we obtain that, say, $\mathbb{E}_t[|U_{t+\tau}|] \leq |U_t|/n \leq 1/\log n$. Thus $\mathbb{P}_t[|U_{t+\tau}| \geq 1] \leq o(1)$ by Markov's inequality. □

Lemma 3.17 together with Lemma 2.11 gives that

$$T_{pp}(G_n(\varepsilon), q) = \log_\lambda n + \frac{1}{q(1 - 1.5\varepsilon)/(1 - \varepsilon) - \log(1 - q)} \log n + o(\log n),$$

where

$$\lambda = 1 + 2q + \left(2q \left(\sqrt{(\varepsilon^2/2 - \varepsilon + 1)} - 1 \right) + q\varepsilon \right) / (2 - 2\varepsilon) > 1 + 2q.$$

To see whether *push&pull* actually slowed down (in terms of order $\log n$) one has to compare the runtime on this sequence of graphs to $c_{pp} \log n$, the runtime on expander sequences. In Figure 1 we can see that it slows down for nearly all values of ε and q in question; however, there are admissible values of ε and q such that the process even speeds up.

References

- [1] Acan, H., Collecchio, A., Mehrabian, A. and Wormald, N. (2017) On the push&pull protocol for rumour spreading. In *Extended Abstracts Summer 2015* (J. Díaz *et al.*, eds), pp. 3–10. Springer.
- [2] Angel, O., Mehrabian, A. and Peres, Y. (2017) The string of diamonds is tight for rumor spreading. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques (APPROX/RANDOM 2017)*, pp. 26:1–26:9. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [3] Boucheron, S., Lugosi, G. and Bousquet, O. (2004) Concentration inequalities. In *Advanced Lectures on Machine Learning*, Vol. 3176 of Lecture Notes in Computer Science, pp. 208–240. Springer.
- [4] Boyd, S., Ghosh, A., Prabhakar, B. and Shah, D. (2006) Randomized gossip algorithms. *IEEE/ACM Trans. Inform. Theory* **52** 2508–2530.
- [5] Censor-Hillel, K., Haeupler, B., Kelner, J. and Maymounkov, P. (2012) Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing (STOC '12)*, pp. 961–970. ACM.
- [6] Chierichetti, F., Giakkoupis, G., Lattanzi, S. and Panconesi, A. (2018) Rumor spreading and conductance. *J. Assoc. Comput. Mach.* **65** 17.
- [7] Daum, S., Kuhn, F. and Maus, Y. (2016) Rumor spreading with bounded in-degree. In *Structural Information and Communication Complexity: 23rd International Colloquium (SIROCCO 2016)*, pp. 323–339.
- [8] Dellamonica, D., Kohayakawa, Y., Marcinişzyn, M. and Steger, A. (2008) On the resilience of long cycles in random graphs. *Electron. J. Combin.* **15** #R321.
- [9] Demers, A., Greene, D., Houser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D. and Terry, D. (1988) Epidemic algorithms for replicated database maintenance. *ACM SIGOPS Operating Systems Review* **22** 8–32.
- [10] Doerr, B., Fouz, M. and Friedrich, T. (2011) Social networks spread rumors in sublogarithmic time. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing (STOC '11)*, pp. 21–30. ACM.
- [11] Doerr, B. and Kozrygin, A. (2017) Randomized rumor spreading revisited. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, pp. 138:1–138:14.
- [12] Doerr, B. and Künemann, M. (2014) Tight analysis of randomized rumor spreading in complete graphs. In *Proceedings of the Meeting on Analytic Algorithmics and Combinatorics (ANALCO 2014)*, pp. 82–91. SIAM.
- [13] Elsässer, R. and Sauerwald, T. (2009) On the runtime and robustness of randomized broadcasting. *Theoret. Comput. Sci.* **410** 3414–3427.
- [14] Feige, U., Peleg, D., Raghavan, P. and Upfal, E. (1990) Randomized broadcast in networks. *Random Struct. Algorithms* **1** 447–460.
- [15] Fountoulakis, N., Huber, A. and Panagiotou, K. (2010) Reliable broadcasting in random networks and the effect of density. In *2010 Proceedings IEEE INFOCOM*, pp. 1–9. IEEE.
- [16] Fountoulakis, N. and Panagiotou, K. (2010) Rumor spreading on random regular graphs and expanders. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques (RANDOM 2010)*, Vol. 6302 of Lecture Notes in Computer Science, pp. 560–573. Springer.
- [17] Fountoulakis, N., Panagiotou, K. and Sauerwald, T. (2012) Ultra-fast rumor spreading in social networks. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*, pp. 1642–1660. SIAM.
- [18] Friedrich, T., Sauerwald, T. and Stauffer, A. (2013) Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica* **67** 65–88.
- [19] Frieze, A. M. and Grimmett, G. R. (1985) The shortest-path problem for graphs with random arc-lengths. *Discrete Appl. Math.* **10** 57–77.
- [20] Giakkoupis, G. (2011) Tight bounds for rumor spreading in graphs of a given conductance. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)*, pp. 57–68. Schloss Dagstuhl-Leibniz-Zentrum für Informatik
- [21] Giakkoupis, G. (2014) Tight bounds for rumor spreading with vertex expansion. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*, pp. 801–815.

- [22] Greenberg, S. and Mohri, M. (2014) Tight lower bound on the probability of a binomial exceeding its expectation. *Statist. Probab. Lett.* **86** 91–98.
- [23] Haeupler, B. (2015) Simple, fast and deterministic gossip and rumor spreading. *J. Assoc. Comput. Mach.* **62** 47.
- [24] Hoffman, A. J. and Wielandt, H. W. (1953) The variation of the spectrum of a normal matrix. *Duke Math. J.* **20** 37–39.
- [25] Hoory, S., Linial, N. and Wigderson, A. (2006) Expander graphs and their applications. *Bull. Amer. Math. Soc.* **43** 439–561.
- [26] Janson, S. (2018) Tail bounds for sums of geometric and exponential variables. *Statistics & Probability Letters*, **135** 1–6.
- [27] Karp, R. M., Schindelhauer, C., Shenker, S. and Vöcking, B. (2000) Randomized rumor spreading. In *41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, pp. 565–574.
- [28] Panagiotou, K., Pérez-Giménez, X., Sauerwald, T. and Sun, H. (2015) Randomized rumour spreading: the effect of the network topology. *Combin. Probab. Comput.* **24** 457–479.
- [29] Panagiotou, K. and Speidel, L. (2017) Asynchronous rumor spreading on random graphs. *Algorithmica* **78** 968–989.
- [30] Rödl, V. and Schacht, M. (2010) Regularity lemmas for graphs. In *Fete of Combinatorics and Computer Science*, Vol. 20 of Bolyai Society Mathematical Studies, pp. 287–325. Springer.
- [31] Sudakov, B. and Vu, V. H. (2008) Local resilience of graphs. *Random Struct. Algorithms* **33** 409–433.
- [32] Wielandt, H. (1950) Unzerlegbare, nicht negative Matrizen. *Math. Z.* **52** 642–648.