

Online Carpooling Using Expander Decompositions

Anupam Gupta

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA
anupamg@cs.cmu.edu

Ravishankar Krishnaswamy

Microsoft Research, Bengaluru, India
rakri@microsoft.com

Amit Kumar

Department of Computer Science and Engineering, Indian Institute of Technology, Delhi, India
amitk@cse.iitd.ernet.in

Sahil Singla

Department of Computer Science, Princeton University, NJ, USA
singla@cs.princeton.edu

Abstract

We consider the online carpooling problem: given n vertices, a sequence of edges arrive over time. When an edge $e_t = (u_t, v_t)$ arrives at time step t , the algorithm must orient the edge either as $v_t \rightarrow u_t$ or $u_t \rightarrow v_t$, with the objective of minimizing the maximum discrepancy of any vertex, i.e., the absolute difference between its in-degree and out-degree. Edges correspond to pairs of persons wanting to ride together, and orienting denotes designating the driver. The discrepancy objective then corresponds to every person driving close to their fair share of rides they participate in.

In this paper, we design efficient algorithms which can maintain $\text{polylog}(n, T)$ maximum discrepancy (w.h.p) over any sequence of T arrivals, when the arriving edges are sampled independently and uniformly from any given graph G . This provides the first polylogarithmic bounds for the online (stochastic) carpooling problem. Prior to this work, the best known bounds were $O(\sqrt{n \log n})$ -discrepancy for any adversarial sequence of arrivals, or $O(\log \log n)$ -discrepancy bounds for the stochastic arrivals when G is the complete graph.

The technical crux of our paper is in showing that the simple greedy algorithm, which has provably good discrepancy bounds when the arriving edges are drawn uniformly at random from the complete graph, also has polylog discrepancy when G is an expander graph. We then combine this with known expander-decomposition results to design our overall algorithm.

2012 ACM Subject Classification Theory of computation \rightarrow Online algorithms

Keywords and phrases Online Algorithms, Discrepancy Minimization, Carpooling

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2020.23

Related Version <https://arxiv.org/abs/2007.10545>.

Funding This research was done under the auspices of the Indo-US Virtual Networked Joint Center IUSSTF/JC-017/2017.

Anupam Gupta: Supported in part by NSF awards CCF-1907820, CCF1955785, and CCF-2006953.

Sahil Singla: Supported in part by the Schmidt Foundation.

Acknowledgements We thank Thatchaphol Saranurak for explaining and pointing us to [7, Theorem 5.6]. The last author would like to thank Navin Goyal for introducing him to [1].



© Anupam Gupta, Ravishankar Krishnaswamy, Amit Kumar, and Sahil Singla;
licensed under Creative Commons License CC-BY

40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020).

Editors: Nitin Saxena and Sunil Simon; Article No. 23; pp. 23:1–23:14



Leibniz International Proceedings in Informatics

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

1 Introduction

Consider the following *edge orientation* problem: we are given a set V of n nodes, and undirected edges arrive online one-by-one. Upon arrival of an edge $\{u, v\}$, it has to be oriented as either $u \rightarrow v$ or $v \rightarrow u$, immediately and irrevocably. The goal is to minimize the *discrepancy* of this orientation at any time $t \in [T]$ during the arrival process, i.e., the maximum imbalance between the in-degree and out-degree of any node. Formally, if we let χ^t to denote the orientation at time t and $\delta_t^-(v)$ (resp. $\delta_t^+(v)$) to denote the number of in-edges (resp. out-edges) incident to v in χ^t , then we want to minimize

$$\max_t \text{disc}(\chi^t) := \max_t \max_v |\delta_t^-(v) - \delta_t^+(v)|.$$

If the entire sequence of edges is known up-front, one can use a simple cycle-and-path-peeling argument to show that any set of edges admit a discrepancy of at most 1. The main focus of this work is in understanding how much loss is caused by the presence of uncertainty, since we don't have knowledge of future arrivals when we irrevocably orient an edge.

This problem was proposed by Ajtai et al. [1] as a special case of the *carpooling problem* where hyperedges arrive online, each representing a carpool where one person must be designated as a driver. The “fair share” of driving for person i can be defined as $\sum_{e:i \in e} 1/|e|$, and we would like each person to drive approximately this many times. In the case of graphs where each carpool is of size $|e| = 2$, this carpooling problem is easily transformed into the edge-orientation problem.

Ajtai et al. showed that while deterministic algorithms cannot have an $o(n)$ discrepancy, they gave a randomized “local greedy” which has an expected discrepancy (for any $T \geq 1$) of $O(\sqrt{n \log n})$ for any online input sequence of T arrivals. Indeed, note that the discrepancy bound is independent of the length of the sequence T , and depends only on the number of nodes, thus giving a non-trivial improvement over the naive random assignment, which will incur a discrepancy of $O(\sqrt{T \log n})$. Intriguingly, the lower bound they show for online algorithms is only $\Omega((\log n)^{1/3})$ – leaving a large gap between the upper and lower bounds.

Given its apparent difficulty in the adversarial online model, Ajtai et al. proposed a stochastic model, where each edge is an independent draw from some underlying probability distribution over pairs of vertices. They considered the the uniform distribution, which is the same as presenting a uniformly random edge of the complete graph at each time. In this special case, they showed that the greedy algorithm (which orients each edge towards the endpoint with lower in-degree minus out-degree) has expected discrepancy $\Theta(\log \log n)$. Their analysis crucially relies on the structure and symmetry of the complete graph.

In this paper, we consider this stochastic version of the problem for general graphs: i.e., given an arbitrary simple graph G , the online input is a sequence of edges chosen independently and uniformly at random (with replacement) from the edges of this graph G^1 . Our main result is the following:

► **Theorem 1 (Main Theorem).** *There is an efficient algorithm for the edge-orientation problem that maintains, w.h.p, a maximum discrepancy of $O(\text{poly} \log(nT))$ on input sequences formed by i.i.d. draws from the edges of a given graph G .*

¹ It is possible to extend our results, by losing a $\log T$ factor, to edge-weighted distributions where an edge is drawn i.i.d. with probability proportional to its weight. Since this extension uses standard ideas like bucketing edges with similar weights, we restrict our attention to arrivals from a graph G for simplicity.

1.1 Our Techniques

Let us fix some notation. Given a (multi)graph $G = (V, E)$ with $|V| = n$, the algorithm is presented with a vector v^t at each time as follows. A uniformly random edge $(u, v) \in G$ is sampled, and the associated characteristic vector $v^t = \mathbf{e}_u - \mathbf{e}_v$ is presented to the algorithm, where $\mathbf{e}_u \in \mathbb{R}^n$ has all zeros except index u being 1. The algorithm must immediately sign v^t with $\chi^t \in \{-1, 1\}$, to keep the discrepancy bounded at all times t . Here the discrepancy of node u at time t is the u^{th} entry of the vector $\sum_{s \leq t} \chi^s v^s$ (which could be negative), and the discrepancy of the algorithm is the maximum absolute discrepancy over all vertices, i.e., $\left\| \sum_{s \leq t} \chi^s v^s \right\|_{\infty}$.

A natural algorithm is to pick a uniformly random orientation for each arriving edge. This maintains zero expected discrepancy at each node. However, the large variance may cause the maximum discrepancy over nodes to be as large as $\Omega(\sqrt{T})$, where T the total number of edges (which is the same as the number of time-steps). For example, this happens even on T parallel edges between two nodes. In this case, however, the *greedy algorithm* which orients the edge from the vertex of larger discrepancy to that of smaller discrepancy works well. Indeed it is not known to be bad for stochastic instances. (Since it is a deterministic algorithm, it can perform poorly on adversarial inputs due to known $o(n)$ lower bounds [1].)

Building on the work of Ajtai et al. who consider stochastic arrivals on complete graphs, the first step towards our overall algorithm is to consider the problem on *expander graphs*. At a high level, one hurdle to achieving low discrepancy in the stochastic case is that we reach states where both endpoints of a randomly-chosen edge already have equally high discrepancy. Then, no matter how we orient the edge, we increase the maximum discrepancy. But this should not happen in expander graphs: if S is the set of “high” discrepancy vertices, then the expansion of the graph implies that $|\partial S|$ must be a large fraction of the total number of edges incident to S . Therefore, intuitively, we have a good chance of reducing the discrepancy if we get edges that go from S to low-degree nodes. To make this idea formal, we relate the greedy process on expander graphs G to the so-called $(1 + \beta)$ -process over an *easier arrival* sequence where the end-points of a new edge are chosen from a *product distribution*, where the probability of choosing a vertex is proportional to its degree in G . However, in the $(1 + \beta)$ -process², the algorithm orients a new edge greedily with only probability β for some small value of β , and does a random orientation with the remaining probability $(1 - \beta)$.

Indeed, we compare these two processes by showing that (a) the expected increase of a natural potential $\Phi := \sum_v \cosh(\lambda \text{discrepancy}(v))$ – which can be thought of as a soft-max function – is lower for the greedy algorithm on expanders when compared to the $(1 + \beta)$ -process on the product distribution, and (b) the same potential increases very slowly (if at all) on the product distribution. A similar idea was used by Peres et al. [13] for a related stochastic load balancing problem; however, many of the technical details are different.

The second component of the algorithm is to decompose a general graph into expanders. This uses the (by-now commonly used) idea of expander decompositions. Loosely speaking, this says that the edges of any graph can be decomposed into some number of smaller graphs (each being defined on some subset of vertices), such that (a) each of these graphs is an expander, and (b) each vertex appears in only a poly-logarithmic number of these expanders. Our arguments for expanders require certain weak-regularity properties – namely the degrees of vertices should not be too small compared to the average degree – and hence some care is required in obtaining decompositions into such expanders. These details appear in the full version.

² The name $(1 + \beta)$ -process stems from the notion for an analogous load-balancing (or) balls-and-bins setting [13], this process would be like the $(1 + \beta)$ -fractional version of the power-of-two choices process.

Our overall algorithm can then be summarized in Algorithm 1.

■ **Algorithm 1 DivideAndGreedy** (graph $G = (V, E)$).

-
- 1: run the expander-decomposition algorithm in Theorem 19 (in Section 2.5) on G to obtain a collection $\mathcal{P} = \{G_1, G_2, \dots, G_k\}$ of edge-disjoint expander graphs.
 - 2: initialize $\mathcal{H} = \{H_1, H_2, \dots, H_k\}$ to be a collection of empty graphs, where H_i is the directed multi-graph consisting of all edges which have arrived corresponding to base graph G_i , along with their orientations assigned by the algorithm upon arrival.
 - 3: **for** each new edge $e \equiv \{u, v\}$ that arrives at time-step t **do**
 - 4: let i denote the index such that $e \in G_i$ according to our decomposition.
 - 5: add e to H_i , and orient e in a greedy manner w.r.t H_i , i.e., from u to v if $\text{disc}_{H_i}(u) \geq \text{disc}_{H_i}(v)$, where $\text{disc}_H(w) = \delta_{H_i}^{\text{in}}(w) - \delta_{H_i}^{\text{out}}(w)$ is the in-degree minus out-degree of any vertex w in the current sub-graph H_i maintained by the algorithm.
 - 6: **end for**
-

1.2 Related Work

The study of discrepancy problems has a long history; see the books [12, 8] for details on the classical work. The problem of online discrepancy minimization was studied by Spencer [14], who showed an $\Omega(\sqrt{T})$ lower bound for for adaptive adversarial arrivals. More refined lower bounds were given by Bárány [6]; see [4] for many other references. Much more recently, Bansal and Spencer [5] and Bansal et al. [4] consider a more general vector-balancing problem, where each request is a vector $v^t \in \mathbb{R}^n$ with $\|v^t\|_\infty \leq 1$, and the goal is to assign a sign $\chi^t \in \{-1, 1\}$ to each vector to minimize $\|\sum_t \chi^t v^t\|_\infty$, i.e., the largest coordinate of the signed sum. Imagining each edge $e_t = \{u, v\}$ to be the vector $\frac{1}{\sqrt{2}}(\mathbf{e}_u - \mathbf{e}_v)$ (where this initial sign is chosen arbitrarily) captures the edge-orientation problem up to constant factors. Bansal et al. gave an $O(n^2 \log(nT))$ -discrepancy algorithm for the natural stochastic version of the problem under general distributions. For some special geometric problems, they gave an algorithm that maintains $\text{poly}(s, \log T, \log n)$ discrepancy for sparse vectors that have only s non-zero coordinates. These improve on the work of Jiang et al. [11], who give a sub-polynomial discrepancy coloring for online arrivals of points on a line. A related variant of these geometric problems was also studied in Dwivedi et al. [9].

Very recently, an independent and exciting work of Alweiss, Liu, and Sawhney [2] gave a randomized algorithm that maintains a discrepancy of $O(\log(nT)/\delta)$ for any input sequence chosen by an oblivious adversary with probability $1 - \delta$, even for the more general vector-balancing problem for vectors of unit Euclidean norm (the so-called Kómlós setting). Instead of a potential based analysis like ours, they directly argue why a carefully chosen randomized greedy algorithm ensures w.h.p. that the discrepancy vector is always sub-Gaussian. A concurrent work of Bansal et al. [3] also obtains similar results for i.i.d. arrivals, but they use a very different potential than our expander-decomposition approach. It is an interesting open question to extend our approach to hypergraphs and re-derive their results.

1.3 Notation

We now define some graph-theoretic terms that are useful for the remainder of the paper.

► **Definition 2** (Volume and α -expansion). Given any graph $G = (V, E)$, and set $S \subseteq V$ its *volume* is defined to be $\text{vol}(S) := \sum_{v \in S} \text{degree}(v)$. We say G is an α -*expander* if

$$\min_{S \subseteq V} \frac{|E(S, V \setminus S)|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}} \geq \alpha.$$

We will also need the following definition of “weakly-regular” graphs, which are graphs where every vertex has degree *at least* a constant factor of the average degree. Note that the *maximum* degree can be arbitrarily larger than the average degree.

► **Definition 3** (γ -weakly-regular). For $\gamma \in [0, 1]$, a graph $G = (V, E)$ is called γ -*weakly-regular* if every vertex $v \in V$ has degree at least $\gamma \cdot \sum_{u \in V} \text{degree}(u) / |V|$.

► **Definition 4** (Discrepancy Vector). Given any directed graph $H = (V, A)$ (representing all the oriented edges until any particular time-step), let $\mathbf{d} \in \mathbb{Z}^{|V|}$ represent the discrepancy vector of the current graph, i.e. the v^{th} entry of \mathbf{d} , denoted by d_v is the difference between the number of in-edges incident at v and the number of out-edges incident at v in H .

2 The Greedy Algorithm on Expander Graphs

In this section, we consider the special case when the graph G is an expander. More formally, we show that the greedy algorithm is actually good for such graphs.

► **Definition 5** (Expander Greedy Process). The greedy algorithm maintains a current discrepancy d_v^t for each vertex v , which is the in-degree minus out-degree of every vertex among the previously arrived edges. Initially, $d_v^1 = 0$ for every vertex v at the beginning of time-step 1. At each time $t \geq 1$, a uniformly random edge $e \in G$ with end-points $\{u, v\}$ is presented to the algorithm, and suppose w.l.o.g. $d_u^t \geq d_v^t$, i.e., u has larger discrepancy (ties broken arbitrarily). Then, the algorithm orients the edge from u to v . The discrepancies of u and v become $d_u^{t+1} = d_u^t - 1$ and $d_v^{t+1} = d_v^t + 1$, and other vertices’ discrepancies are unchanged.

► **Theorem 6.** *Consider any γ -weakly-regular α -expander G , and suppose edges are arriving as independent samples from G over a horizon of T time-steps. Then, the greedy algorithm maintains a discrepancy d_v^t of $O(\log^5 nT)$ for every time t in $[0 \dots T]$ and every vertex v , as long as $\alpha \geq 6\lambda$, $\gamma \geq \lambda^{1/4}$, where $\lambda = O(\log^{-4} nT)$.*

For the sake of concreteness, it might be instructive to assume $\alpha \approx \gamma \approx O(\frac{1}{\log n})$, which is roughly what we will obtain from our expander-decomposition process.

2.1 Setting Up The Proof

Our main idea is to introduce *another random process* called the $(1+\beta)$ -process, and show that the $(1+\beta)$ -process stochastically dominates the expander-greedy process in a certain manner, and separately bound the behaviour of the $(1+\beta)$ -process subsequently. By combining these two, we get our overall analysis of the expander-greedy process.

To this end, we first define a random arrival sequence where the end-points of each new edge are actually sampled independently from a *product distribution*.

► **Definition 7** (Product Distribution). Given a set V of vertices with associated weights $\{w_v \geq 0 \mid v \in V\}$, at each time t , we select two vertices u, v as *two independent samples* from V , according to the distribution where any vertex $v \in V$ is chosen with probability $\frac{w_v}{\sum_{v' \in V} w_{v'}}$, and the vector $v^t := \chi_u - \chi_v$ is presented to the algorithm.

We next define the $(1 + \beta)$ -process, which will be crucial for the analysis.

► **Definition 8** ($(1 + \beta)$ -process on product distributions). Consider a product distribution over a set of vertices V . When presented with a vector $v^t := \chi_u - \chi_v$ from this product distribution at time t , the $(1 + \beta)$ -process assigns a sign to the vector v^t as follows: with probability $(1 - \beta)$, it assigns it uniformly ± 1 , and only with the remaining probability β it uses the greedy algorithm to sign this vector.

Note that setting $\beta = 1$ gives us back the greedy algorithm, and $\beta = 0$ gives an algorithm that assigns a random sign to each vector.

► **Remark 9.** The original $(1 + \beta)$ -process was in fact introduced in [13], where Peres et al. analyzed a general load-balancing process over n bins (corresponding to vertices), and balls arrive sequentially. Upon each arrival, the algorithm gets to sample a random edge from a k -regular expander³ G over the bins, and places the ball in the lighter loaded bin among the two end-points of the edge. They show that this process maintains a small maximum load, by relating it to an analogous $(1 + \beta)$ -process, where instead of sampling an edge from G , two bins are chosen uniformly at random, and the algorithm places the ball into a random bin with probability $1 - \beta$, and the lesser loaded bin with probability β . Note that their analysis inherently assumed that the two vertices are sampled from the uniform distribution where all weights w_u are equal. By considering arbitrary product distributions, we are able to handle arbitrary graphs with a non-trivial conductance, i.e., even those that *do not* satisfy the k -regularity property. This is crucial for us because the expander decomposition algorithms, which reduce general graphs to a collection of expanders, do not output regular expanders.

Our analysis will also involve a potential function (intuitively the soft-max of the vertex discrepancies) for both the expander-greedy process as well as the $(1 + \beta)$ -process.

► **Definition 10** (Potential Function). Given vertex discrepancies $\mathbf{d} \in \mathbb{Z}^{|V|}$, define

$$\Phi(\mathbf{d}) := \sum_v \cosh(\lambda d_v), \tag{1}$$

where $\lambda < 1$ is a suitable parameter to be optimized.

Following many prior works, we use the hyperbolic cosine function to symmetrize for positive and negative discrepancy values. When \mathbf{d} is clear from the context, we will write $\Phi(\mathbf{d})$ as Φ . We will also use \mathbf{d}^t to refer to the discrepancy vector at time t , and d_u^t to the discrepancy of u at time t . We will often ignore the superscript t if it is clear from the context.

We are now ready to define the appropriate parameters of the $(1 + \beta)$ -process. Indeed, given the expander-greedy process defined on graph G , we construct an associated $(1 + \beta)$ -process where for each vertex v , the probability of sampling any vertex in the product distribution is proportional to its degree in G , i.e., $w_v = \text{degree}_G(v)$ for all $v \in V$. We also set the β parameter equal to α , the conductance of the graph G .

2.2 One-Step Change in Potential

The main idea of the proof is to use a *majorization argument* to argue that *the expected one-step change* in potential of the expander process can be upper bounded by that of the $(1 + \beta)$ -process, if the two processes start at the same discrepancy configuration \mathbf{d}^t . Subsequently, we bound the one-step change for the $(1 + \beta)$ -process in Section 2.4.

³ Actually their proof works for a slightly more general notion of expanders, but which is still insufficient for our purpose.

To this end, consider a time-step t , where the current discrepancy vector of the expander process is \mathbf{d}^t . Suppose the next edge in the expander process is (i, j) , where $d_i^t > d_j^t$. Then the greedy algorithm will always choose a sign such that d_i decreases by 1, and d_j increases by 1. Indeed, this ensures the overall potential is non-increasing unless $d_i = d_j$. More importantly, the potential term for other vertices remains unchanged, and so we can express the expected change in potential as having contributions from precisely two terms, one due to $d_i \rightarrow d_i - 1$ (called the *decrease term*), and denoted as $\Delta_{-1}(t)$, and one due to $d_j \rightarrow d_j + 1$ (the *increase term*), denoted as $\Delta_{+1}(t)$:

$$\begin{aligned} \mathbb{E}_{(i,j) \sim G}[\Delta\Phi] &= \mathbb{E}_{(i,j) \sim G}[\Phi(\mathbf{d}^{t+1}) - \Phi(\mathbf{d}^t)] \\ &= \underbrace{\mathbb{E}_{(i,j)}[\cosh(\lambda(d_i - 1)) - \cosh(\lambda(d_i))]}_{=:\Delta_{-1}(\mathbf{d}^t)} + \underbrace{\mathbb{E}_{(i,j)}[\cosh(\lambda(d_j + 1)) - \cosh(\lambda(d_j))]}_{=:\Delta_{+1}(\mathbf{d}^t)}. \end{aligned}$$

Now, consider the $(1 + \beta)$ -process on the vertex set V , where the product distribution is given by weights $w_u = \deg(u)$ for each $u \in V$, starting with the same discrepancy vector \mathbf{d}^t as the expander process at time t . Then, if u and v are the two vertices sampled independently according to the product distribution, then by its definition, the $(1 + \beta)$ -process signs this pair randomly with probability $(1 - \beta)$, and greedily with probability β . For the sake of analysis, we define two terms analogous to $\Delta_{-1}(\mathbf{d}^t)$ and $\Delta_{+1}(\mathbf{d}^t)$ for the $(1 + \beta)$ -process. To this end, let $i \in \{u, v\}$ denote the identity of the random vertex to which the $(1 + \beta)$ -process assigns $+1$. Define

$$\tilde{\Delta}_{+1}(\mathbf{d}^t) := \mathbb{E}_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[\cosh(\lambda(d_i + 1)) - \cosh(\lambda(d_i))], \quad (2)$$

where $\mathbf{w} \times \mathbf{w}$ refers to two independent choices from the product distribution corresponding to w . Similarly let $j \in \{u, v\}$ denote the identity of the random vertex to which the $(1 + \beta)$ -process assigns -1 , and define

$$\tilde{\Delta}_{-1}(\mathbf{d}^t) := \mathbb{E}_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[\cosh(\lambda(d_j - 1)) - \cosh(\lambda(d_j))]. \quad (3)$$

In what follows, we bound $\Delta_{-1}(\mathbf{d}^t) \leq \tilde{\Delta}_{-1}(\mathbf{d}^t)$ through a coupling argument, and similarly bound $\Delta_{+1}(\mathbf{d}^t) \leq \tilde{\Delta}_{+1}(\mathbf{d}^t)$ using a separate coupling.

A subtlety: the expected one-step change in Φ in the expander process precisely equals $\Delta_{-1}(\mathbf{d}^t) + \Delta_{+1}(\mathbf{d}^t)$. However, if we define an analogous potential for the $(1 + \beta)$ -process, then the one-step change in potential there *does not* equal the sum $\tilde{\Delta}_{-1}(\mathbf{d}^t) + \tilde{\Delta}_{+1}(\mathbf{d}^t)$. Indeed, we sample u and v i.i.d. in the $(1 + \beta)$ -process, it is possible that $u = v$ and therefore the one-step change in potential is 0, while the sum $\tilde{\Delta}_{-1}(\mathbf{d}^t) + \tilde{\Delta}_{+1}(\mathbf{d}^t)$ will be non-zero. Hence the following lemma does not bound the expected potential change for the expander process by that for the $(1 + \beta)$ -process (both starting from the same state), but by this surrogate $\tilde{\Delta}_{-1}(\mathbf{d}^t) + \tilde{\Delta}_{+1}(\mathbf{d}^t)$, and it is this surrogate sum that we bound in Section 2.4.

2.3 The Coupling Argument

We now show a coupling between the expander-greedy process and the $(1 + \beta)$ -process defined in Section 2.1, to bound the expected one-step change in potential for the expander process.

► **Lemma 11.** *Given an α -expander $G = (V, E)$, let $\mathbf{d}^t \equiv (d_v : v \in V)$ denote the current discrepancies of the vertices at any time step t for the expander-greedy process. Consider a hypothetical $(1 + \beta)$ -process on vertex set V with $\beta = \alpha$, the weight of vertex $v \in V$ set to $w_v = \deg(v)$, and starting from the same discrepancy state \mathbf{d}^t . Then:*

23:8 Online Carpooling Using Expander Decompositions

(a) $\Delta_{-1}(\mathbf{d}^t) \leq \tilde{\Delta}_{-1}(\mathbf{d}^t)$, and (b) $\Delta_{+1}(\mathbf{d}^t) \leq \tilde{\Delta}_{+1}(\mathbf{d}^t)$.
Hence the expected one-step change in potential $\mathbb{E}[\Phi(\mathbf{d}^{t+1}) - \Phi(\mathbf{d}^t)] \leq \tilde{\Delta}_{-1}(\mathbf{d}^t) + \tilde{\Delta}_{+1}(\mathbf{d}^t)$.

Proof. We start by renaming the vertices in V such that $d_n \leq d_{n-1} \leq \dots \leq d_1$. Suppose the next edge in the expander process corresponds to indices i, j where $i < j$. We prove the lemma statement by two separate coupling arguments, which crucially depend on the following claim. Intuitively, this claim shows that a -1 is more likely to appear among the high discrepancy vertices of G in the expander process than the $(1 + \beta)$ -process (thereby having a lower potential), and similarly a $+1$ is more likely to appear among the low discrepancy vertices of G in the expander process than in the $(1 + \beta)$ -process. Peres et al. [13] also prove a similar claim for stochastic load balancing, but they only consider uniform distributions.

▷ **Claim 12.** For any $k \in [n]$, if S_k denotes the set of vertices with indices $k' \in [k]$ (the k highest discrepancy vertices) and T_k denotes $V \setminus S_k$, then

$$\Pr_{(i,j) \sim G}[-1 \in S_k] \geq \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[-1 \in S_k] \quad \text{and} \quad \Pr_{(i,j) \sim G}[+1 \in T_k] \geq \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[+1 \in T_k].$$

Above, we abuse notation and use the terminology “ $-1 \in S_k$ ” to denote that the vertex whose discrepancy decreases falls in the set S_k in the corresponding process.

Proof. Fix an index k , and let $\rho := \frac{\text{vol}(S_k)}{\text{vol}(V)}$ be the relative volume of S_k , i.e., the fraction of edges of G incident to the k nodes of highest degree. First we consider the $(1 + \beta)$ -process on V . With $(1 - \beta)$ probability we assign a sign to the input vector uniformly at random. Therefore, conditioned on this choice, a vertex in S_k will get a -1 sign with probability

$$\frac{1}{2} \cdot \Pr[u \in S_k] + \frac{1}{2} \Pr[v \in S_k] = \frac{\text{vol}(S_k)}{\text{vol}(V)} = \rho,$$

where u and v denote the two vertices chosen by the $(1 + \beta)$ -process. With probability β , we will use the greedy algorithm, and so -1 will appear on a vertex in S_k iff at least one of the two chosen vertices lie in S_k . Putting it together, we get

$$\begin{aligned} \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[-1 \in S_k] &= (1 - \beta) \cdot \frac{\text{vol}(S_k)}{\text{vol}(V)} + \beta \cdot \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[\{u, v\} \cap S_k \neq \emptyset] \\ &= (1 - \beta) \cdot \rho + \beta \cdot (1 - (1 - \rho)^2) = (1 + \beta - \beta \cdot \rho) \cdot \rho. \end{aligned} \quad (4)$$

Now we consider the expander process. A vertex in S_k gets -1 iff the chosen edge has at least one end-point in S_k . Therefore,

$$\begin{aligned} \Pr_{(i,j) \sim G}[-1 \in S_k] &= \Pr[i \in S_k] = \frac{|E(S_k, S_k)| + |E(S_k, V \setminus S_k)|}{|E|} \\ &= \frac{(2|E(S_k, S_k)| + |E(S_k, V \setminus S_k)|) + |E(S_k, V \setminus S_k)|}{2|E|} = \frac{\text{vol}(S_k) + |E(S_k, V \setminus S_k)|}{\text{vol}(V)}. \end{aligned}$$

Recalling that $\beta = \alpha$, and that G is an α -expander, we consider two cases:

Case 1: If $\text{vol}(S_k) \leq \text{vol}(V \setminus S_k)$, we use

$$\begin{aligned} \Pr_{(i,j) \sim G}[-1 \in S_k] &= \frac{\text{vol}(S_k) + |E(S_k, V \setminus S_k)|}{\text{vol}(V)} \\ &\geq (1 + \alpha) \frac{\text{vol}(S_k)}{\text{vol}(V)} = (1 + \beta) \rho \geq \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[-1 \in S_k]. \end{aligned}$$

Case 2: If $\text{vol}(S_k) > \text{vol}(V \setminus S_k)$, we use

$$\begin{aligned} \Pr_{(i,j) \sim G}[-1 \in S_k] &= \frac{\text{vol}(S_k) + |E(S_k, V \setminus S_k)|}{\text{vol}(V)} \geq \frac{\text{vol}(S_k) + \alpha \cdot \text{vol}(V \setminus S_k)}{\text{vol}(V)} \\ &\geq \left(1 + \beta \cdot \frac{\text{vol}(V \setminus S_k)}{\text{vol}(V)}\right) \cdot \rho = \Pr_{(i,j) \sim \mathbf{w} \times \mathbf{w}}[-1 \in S_k], \end{aligned}$$

where the last equality uses (4).

This completes the proof of $\Pr_{(i,j) \sim G}[-1 \in S_k] \geq \Pr_{(i,j) \sim \mathbf{w}}[-1 \in S_k]$. One can similarly show $\Pr_{(i,j) \sim G}[+1 \in T_k] \geq \Pr_{(u,v) \sim \mathbf{w} \times \mathbf{w}}[+1 \in T_k]$, which completes the proof of the claim. \triangleleft

Claim 12 shows that we can establish a coupling between the two processes such that if -1 belongs to S_k in $(1 + \beta)$ -process, then the same happens in the expander process. In other words, there is a joint sample space Ω such that for any outcome $\omega \in \Omega$, if vertices v_a and v_b get sign -1 in the expander process and the $(1 + \beta)$ -process respectively, then $a \leq b$.

Let \mathbf{d} and $\tilde{\mathbf{d}}$ denote the discrepancy vectors in the expander process and the $(1 + \beta)$ -process after the -1 sign has been assigned, respectively. Now, since both the processes start with the same discrepancy vector \mathbf{d}^t , we see that for any fixed outcome $\omega \in \Omega$, the vector $\tilde{\mathbf{d}}$ majorizes \mathbf{d} in the following sense.

► Definition 13 (Majorization). Let \mathbf{a} and \mathbf{b} be two real vectors of the same length n . Let $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ denote the vectors \mathbf{a} and \mathbf{b} with coordinates rearranged in descending order respectively. We say that \mathbf{a} majorizes \mathbf{b} , written $\mathbf{a} \succeq \mathbf{b}$, if for all i , $1 \leq i \leq n$, we have $\sum_{j=1}^i \vec{\mathbf{a}}_j \geq \sum_{j=1}^i \vec{\mathbf{b}}_j$.

One of the properties of majorization [10] is that any convex and symmetric function of the discrepancy vector (which Φ is) satisfies that $\Phi(\mathbf{d}) \leq \Phi(\tilde{\mathbf{d}})$. Thus, for any fixed outcome ω , the change in potential in the expander process is at most that of the surrogate potential in the $(1 + \beta)$ -process. Since $\Delta_{-1}(\mathbf{d}^t)$ and $\tilde{\Delta}_{-1}(\mathbf{d}^t)$ are just the expected change of these quantities in the two processes (due to assignment of -1 sign), the first statement of the lemma follows. Using an almost identical proof, we can also show the second statement. (Note that we may need to redefine the coupling between the two processes to ensure that if vertices v_a, v_b get sign $+1$ as above, then $b \leq a$.) \blacktriangleleft

2.4 Analyzing One-Step $\Delta\Phi$ of the $(1 + \beta)$ -process

Finally we bound the one-step change in (surrogate) potential of the $(1 + \beta)$ -process starting at discrepancy vector \mathbf{d}^t ; recall the definitions of $\tilde{\Delta}_{-1}(\mathbf{d}^t)$ and $\tilde{\Delta}_{+1}(\mathbf{d}^t)$ from Section 2.2.

► Lemma 14. *If $\Phi(\mathbf{d}^t) \leq (nT)^{10}$, and if the weights w_v are such that for all v , $\frac{w_v}{\sum_{v'} w_{v'}} \geq \frac{\gamma}{n}$ (i.e., the minimum weight is at least a γ fraction of the average weight), then we have that*

$$\tilde{\Delta}_{-1}(\mathbf{d}^t) + \tilde{\Delta}_{+1}(\mathbf{d}^t) \leq O(1),$$

as long as $\beta \geq 6\lambda$, $\gamma \geq 16\lambda^{1/4}$, and $\lambda = O(\log^{-4} nT)$.

Proof. Let u be an arbitrary vertex in V , and we condition on the fact that the first vertex chosen by the $(1 + \beta)$ -process is u . Then, we show that

$$\mathbb{E}_{v \sim \mathbf{w}} \left[\cosh(\lambda(d_i - 1)) - \cosh(\lambda(d_i)) + \cosh(\lambda(d_j + 1)) - \cosh(\lambda(d_j)) \mid u \text{ is sampled first} \right],$$

is $O(1)$ regardless of the choice of u , where we assume that i is the random vertex which is assigned -1 by the $(1 + \beta)$ -process, and j is the random vertex which is assigned $+1$. The proof of the lemma then follows by removing the conditioning on u .

Following [5, 4], we use the first two terms of the Taylor expansion of $\cosh(\cdot)$ to upper bound the difference terms of the form $\cosh(x + 1) - \cosh(x)$ and $\cosh(x - 1) - \cosh(x)$. To this end, note that, if $|\epsilon| \leq 1$ and $\lambda < 1$, we have that

$$\begin{aligned} \cosh(\lambda(x + \epsilon)) - \cosh(\lambda x) &\leq \epsilon \lambda \sinh(\lambda x) + \frac{\epsilon^2}{2!} \lambda^2 \cosh(\lambda x) + \frac{\epsilon^3}{3!} \lambda^3 \sinh(\lambda x) + \dots \\ &\leq \epsilon \lambda \sinh(\lambda x) + \epsilon^2 \lambda^2 \cosh(\lambda x). \end{aligned}$$

Using this, we proceed to bound the following quantity (by setting $\epsilon = -1$ and 1 respectively):

$$\mathbb{E}_{v \sim \mathbf{w}} \left[\underbrace{-\lambda(\sinh(\lambda d_i) - \sinh(\lambda d_j))}_{=: -L} + \underbrace{\lambda^2(\cosh(\lambda d_i) + \cosh(\lambda d_j))}_{=: Q} \mid u \text{ is sampled first} \right].$$

We refer to $L = \lambda(\sinh(\lambda d_i) - \sinh(\lambda d_j))$ and $Q = \lambda^2(\cosh(\lambda d_i) + \cosh(\lambda d_j))$ as the *linear* and *quadratic* terms, since they arise from the first- and second-order derivatives in the Taylor expansion.

To further simplify our exposition, we define the following random variables:

- (i) $u_>$ is the identity of the vertex among u, v with higher discrepancy, and $u_<$ is the other vertex. Hence we have that $d_{u_>} \geq d_{u_<}$.
- (ii) G denotes the random variable $\lambda(\sinh(\lambda d_{u_>}) - \sinh(\lambda d_{u_<}))$, which indicates an analogous term to L , but if we exclusively did a greedy signing always (recall that the greedy algorithm would always decrease the larger discrepancy, but the $(1 + \beta)$ -process follows a uniformly random signing with probability $(1 - \beta)$ and follows the greedy rule only with probability β).

Finally, for any vertex $w \in V$, we let $\text{Danger}(w) = \{v : |d_w - d_v| < \frac{2}{\lambda}\}$ to denote the set of vertices with discrepancy close to that of w , where the gains from the term corresponding to βG are insufficient to compensate for the increase due to Q .

We are now ready to proceed with the proof. Firstly, note that, since the $(1 + \beta)$ -process follows the greedy algorithm with probability β (independent of the choice of the sampled vertices u and v), we have that

$$\mathbb{E}_v[L \mid u \text{ is sampled first}] = (1 - \beta)0 + \beta \mathbb{E}_v[G \mid u \text{ is sampled first}]. \quad (5)$$

Intuitively, the remainder of the proof proceeds as follows: suppose $d_{u_>}$ and $d_{u_<}$ are both non-negative (the intuition for the other cases are similar). Then, Q is proportional to $\lambda^2 \cosh(\lambda d_{u_>})$. Now, if $d_{u_>} - d_{u_<}$ is sufficiently large, then G is proportional to $\lambda \sinh(\lambda d_{u_>})$, which in turn is close to $\lambda \cosh(\lambda d_{u_>})$. As a result, we get that as long as $\lambda = O(\beta)$, the term $-\beta G + Q$ can be bounded by 0 for each choice of v such that $d_{u_>} - d_{u_<}$ is large.

However, what happens when $d_{u_>} - d_{u_<}$ is small, i.e., when v falls in $\text{Danger}(u)$? Here, the Q term is proportional to $\lambda^2 \cosh(\lambda d_u)$, but the G term might be close to 0, and so we can't argue that $-\beta G + Q \leq O(1)$ in these events. Hence, we resort to an amortized analysis by showing that (i) when $v \notin \text{Danger}(u)$, $-\beta G$ can not just compensate for Q , it can in fact compensate for $\frac{1}{\sqrt{\lambda}} Q \geq \frac{1}{\sqrt{\lambda}} \cdot \lambda^2 \cosh(\lambda d_u)$, and secondly, (ii) the probability over a random choice of v of $v \notin \text{Danger}(u)$ is at least $\sqrt{\lambda}$, provided Φ is bounded to begin with. The overall proof then follows from taking an average over all v .

Hence, in what follows, we will show that in expectation the magnitude of βG can compensate for a suitably large multiple of Q when $v \notin \text{Danger}(u)$.

▷ **Claim 15.** Let $\beta \geq 6\lambda$. For any fixed choice of vertices u and v such that $v \notin \text{Danger}(u)$, we have $G := \lambda(\sinh(\lambda d_{u>}) - \sinh(\lambda d_{u<})) \geq \frac{\lambda}{3}(\cosh(\lambda d_u) + \cosh(\lambda d_v) - 4)$.

Proof. The proof is a simple convexity argument. To this end, suppose both $d_u, d_v \geq 0$. Then since $\sinh(x)$ is convex when $x \geq 0$ and its derivative is $\cosh(x)$, we get that

$$\sinh(\lambda d_{u>}) - \sinh(\lambda d_{u<}) \geq \lambda \cosh(\lambda d_{u<}) \cdot |d_u - d_v| \geq 2 \cosh(\lambda d_{u<}),$$

using $v \notin \text{Danger}(u)$. But since $|\sinh(x) - \cosh(x)| \leq 1$, we get that

$$\sinh(\lambda d_{u>}) - \sinh(\lambda d_{u<}) \geq 2 \sinh(\lambda d_{u<}) - 2.$$

Therefore, $\sinh(\lambda d_{u<}) \leq \frac{1}{3}(\sinh(\lambda d_{u>}) + 1)$. Now substituting, and using the monotonicity of \sinh and its closeness to \cosh , we get G is at least

$$\frac{2\lambda}{3}(\sinh(\lambda d_{u>}) - 1) \geq \frac{\lambda}{3}(\sinh(\lambda d_{u>}) + \sinh(\lambda d_{u<}) - 2) \geq \frac{\lambda}{3}(\cosh(\lambda d_u) + \cosh(\lambda d_v) - 4).$$

The case of $d_u, d_v \leq 0$ follows from setting $d'_u = |d_u|, d'_v = |d_v|$ and using the above calculations, keeping in mind that \sinh is an odd function but \cosh is even. Finally, when $d_{u<}$ is negative but $d_{u>}$ is positive,

$$\begin{aligned} G &= \lambda(\sinh(\lambda d_{u>}) - \sinh(\lambda d_{u<})) = \lambda(\sinh(\lambda d_{u>}) + \sinh(\lambda |d_{u<}|)) \\ &\geq \frac{\lambda}{3}(\cosh(\lambda d_{u>}) + \cosh(\lambda d_{u<}) - 2) \geq \frac{\lambda}{3}(\cosh(\lambda d_u) + \cosh(\lambda d_v) - 4). \quad \blacktriangleleft \end{aligned}$$

▷ **Claim 16.** Let $\beta \geq 6\lambda$. For any fixed choice of vertices u and v such that $v \notin \text{Danger}(u)$, we have $-\beta G + \left(1 + \frac{1}{\sqrt{\lambda}}\right) Q \leq O(1)$.

Proof. Recall that $G = \lambda(\sinh(\lambda d_{u>}) - \sinh(\lambda d_{u<}))$. Now, let A denote $\cosh(\lambda d_u) + \cosh(\lambda d_v)$. Then, by definition of Q and from Claim 15, we have that

$$-\beta G + \left(1 + \frac{1}{\sqrt{\lambda}}\right) Q \leq -\frac{\beta\lambda}{3}(A-4) + \left(1 + \frac{1}{\sqrt{\lambda}}\right) \lambda^2 A \leq \frac{4\lambda\beta}{3} + \left(\lambda^2 + \lambda^{\frac{3}{2}} - \frac{\lambda\beta}{3}\right) A \leq \lambda\beta$$

is at most $O(1)$, assuming $\beta \geq 6\lambda \geq 3(\lambda + \sqrt{\lambda})$, and recalling that λ, β are at most 1. \blacktriangleleft

We now proceed with our proof using two cases:

Case (i): $|d_u| \leq \frac{10}{\lambda}$. In this case, note that the Q term is

$$\begin{aligned} &\mathbb{E}_v[Q \mid u \text{ is sampled first}] \\ &= \mathbb{E}_v[Q \mid v \in \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \in \text{Danger}(u) \mid u \text{ is sampled first}] \\ &\quad + \mathbb{E}_v[Q \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}] \\ &\leq O(1) + \mathbb{E}_v[Q \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}]. \end{aligned}$$

Here the inequality uses $v \in \text{Danger}(u)$ and $|d_u| \leq \frac{10}{\lambda}$ to infer that both $|d_u|$ and $|d_v|$ are $\leq \frac{12}{\lambda}$. Hence the Q term in this scenario will simply be a constant.

Next we analyze the L term. For the following, we observe that the algorithm chooses a random ± 1 signing with probability $(1 - \beta)$, and chooses the greedy signing with probability β , and moreover, this choice is independent of the random choices of u and v . Hence, the expected L term conditioned on the algorithm choosing a random signing is simply 0, and the expected L term conditioned on the algorithm choosing the greedy signing is simply the term $\mathbb{E}[G]$. Hence, we can conclude that:

23:12 Online Carpooling Using Expander Decompositions

$$\begin{aligned}
& \mathbb{E}_v[-L \mid u \text{ is sampled first}] \\
&= \mathbb{E}_v[-L \mid v \in \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \in \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\quad + \mathbb{E}_v[-L \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\leq \mathbb{E}_v[-\beta G \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}].
\end{aligned}$$

Adding the inequalities and applying Claim 16, we get $\mathbb{E}_v[-L + Q \mid u \text{ is sampled first}] \leq O(1)$.

Case (ii): $|d_u| > \frac{10}{\lambda}$. We first prove two easy claims.

▷ **Claim 17.** Suppose $v \in \text{Danger}(u)$. Then $\cosh(\lambda d_v) \leq 8 \cosh(\lambda d_u)$.

Proof. Assume w.l.o.g. that $d_u, d_v \geq 0$. Also, assume that $d_v \geq d_u$, otherwise there is nothing to prove. Now $d_v \leq d_u + \frac{2}{\lambda}$. So $\frac{\cosh(\lambda d_v)}{\cosh(\lambda d_u)} \leq \sup_x \frac{\cosh(x+2)}{\cosh(x)}$. The supremum on the right happens when $x \rightarrow \infty$, and then the ratio approaches $e^2 < 8$. ◁

▷ **Claim 18.** For any discrepancy vector \mathbf{d}^t such that $\Phi(\mathbf{d}^t) \leq O((nT)^{10})$, and for any u such that $|d_u| > \frac{10}{\lambda}$, we have $\Pr[v \notin \text{Danger}(u)] \geq 8\sqrt{\lambda}$, as long as $\lambda = O(\log^{-4} nT)$.

Proof. We consider the case that $d_u > \frac{10}{\lambda}$; the case were $d_u < -\frac{10}{\lambda}$ is similar.

Assume for a contradiction that $\Pr[v \in \text{Danger}(u)] \geq 1 - 8\sqrt{\lambda}$, and so $\Pr[v \notin \text{Danger}(u)] \leq 8\sqrt{\lambda}$. We first show that the cardinality of the set $|w \notin \text{Danger}(u)|$ is small. Indeed, this follows immediately from our assumption on the minimum weight of any vertex in the statement of Lemma 14 being at least γ/n times the total weight. So we have that for every w , the probability of sampling w in the $(1 + \beta)$ -process is at least $\pi_w \geq \gamma/n$, implying that the total number of vertices not in $\text{Danger}(u)$ must be at most $\frac{8\sqrt{\lambda} \cdot n}{\gamma}$. This also means that the total number of vertices in $\text{Danger}(u) \geq \frac{n}{2}$ since $\gamma \geq \lambda^{1/4} \geq 16\sqrt{\lambda}$ for sufficiently small λ .

Since $d_u > \frac{10}{\lambda}$, we get that any vertex $v \in \text{Danger}(u)$ satisfies $d_v \geq d_u - \frac{2}{\lambda} \geq \frac{8}{\lambda}$. Moreover, since $\sum_v d_v = 0$, it must be that the negative discrepancies must in total compensate for the total sum of discrepancies of the vertices in $\text{Danger}(u)$. Hence, we have that $\sum_{w: d_w < 0} |d_w| \geq \sum_{v \in \text{Danger}(u)} d_v \geq |\{v : v \in \text{Danger}(u)\}| \cdot \frac{8}{\lambda} \geq 0.5n \cdot \frac{8}{\lambda}$.

From the last inequality, and since $|\{w : d_w < 0\}| \leq |\{w : w \notin \text{Danger}(u)\}| \leq \frac{8\sqrt{\lambda}n}{\gamma}$, we get that there exists a vertex \tilde{w} s.t $d_{\tilde{w}} < 0$ and $|d_{\tilde{w}}| \geq \frac{\gamma}{8\sqrt{\lambda}n} \cdot \frac{4n}{\lambda} = \frac{\gamma}{2\lambda^{3/2}}$. But this implies $\Phi(\mathbf{d}^t) \geq \cosh(\lambda d_{\tilde{w}}) \geq \cosh\left(\frac{\gamma}{2\sqrt{\lambda}}\right) > (nT)^{10}$, using that $\lambda = O(\log^{-4} nT)$ and that $\gamma \geq \lambda^{1/4}$. So we get a contradiction on the assumption that $\Phi(\mathbf{d}^t) \leq (nT)^{10}$. ◁

Returning to the proof for the case of $|d_u| \geq \frac{10}{\lambda}$, we get that

$$\begin{aligned}
& \mathbb{E}_v[Q \mid u \text{ is sampled first}] \\
&= \mathbb{E}_v[Q \mid v \in \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \in \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\quad + \mathbb{E}_v[Q \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\leq 8\lambda^2 \cosh(\lambda d_u) \\
&\quad + \mathbb{E}[Q \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}],
\end{aligned}$$

where the first term in inequality follows from Claim 17.

Next we analyze the L term similarly:

$$\begin{aligned}
& \mathbb{E}_v[-L \mid u \text{ is sampled first}] \\
&= \mathbb{E}_v[-L \mid v \in \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \in \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\quad + \mathbb{E}_v[-L \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}] \\
&\leq \mathbb{E}_v[-\beta G \mid v \notin \text{Danger}(u), u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}],
\end{aligned}$$

where the last inequality follows using the same arguments as in case (i). Adding these inequalities and applying Claim 16, we get that

$$\begin{aligned}
\mathbb{E}_v[-L + Q \mid u \text{ is sampled first}] &\leq O(1) + 8\lambda^2 \cosh(\lambda d_u) \\
&\quad - \frac{1}{\sqrt{\lambda}} \cdot \mathbb{E}_v[Q \mid u \text{ is sampled first}] \cdot \Pr[v \notin \text{Danger}(u) \mid u \text{ is sampled first}].
\end{aligned}$$

To complete the proof of Lemma 14, we note that $Q \geq \lambda^2 \cosh(\lambda d_u)$, and use Claim 18 to infer that $\Pr[v \notin \text{Danger}(u)] \geq 8\sqrt{\lambda}$. This implies

$$\mathbb{E}_v[-L + Q \mid u \text{ is sampled first}] \leq O(1) + 8\lambda^2 \cosh(\lambda d_u) - 8\lambda^2 \cosh(\lambda d_u) \leq O(1). \blacktriangleleft$$

We now can use this one-step expected potential change for the $(1 + \beta)$ -process to get the following result for the original expander process:

Proof of Theorem 6. Combining Lemma 14 and Lemma 11, we get that in the expander process, if we condition on the random choices made until time t , if $\Phi(\mathbf{d}^t) \leq (nT)^{10}$, then $\mathbb{E}[\Phi(\mathbf{d}^{t+1}) - \Phi(\mathbf{d}^t)] \leq C$ for some constant C . The potential starts off at n , so if it ever exceeds $CT(nT)^5$ in T steps, there must be a time t such that $\Phi(\mathbf{d}^t) \leq Ct(nT)^5$ and the increase is at least $C(nT)^5$. But the expected increase at this step is at most C , so by Markov's inequality the probability of increasing by $C(nT)^5$ is at most $1/(nT)^5$. Now a union bound over all times t gives that the potential exceeds $CT(nT)^5 \leq (nT)^{10}$ with probability at most $T/(nT)^5 = 1/\text{poly}(nT)$. But then $\cosh(\lambda d_v^t) \leq (nT)^{10}$, and therefore $d_v^t \leq O(\lambda \log(nT)^{10}) = O(\log^3 nT)$ for all vertices v and time t . \blacktriangleleft

In summary, if the underlying graph is γ -weakly-regular for $\gamma \geq \Omega(\log^{-1} nT)$, and has expansion $\alpha \geq \Omega(\log^{-2} nT)$, the greedy process maintains a poly-logarithmic discrepancy.

2.5 Putting It Together

We briefly describe the expander decomposition procedure and summarize the final algorithm.

► **Theorem 19** (Decomposition into Weakly-Regular Expanders). *Any graph $G = (V, E)$ can be decomposed into an edge-disjoint union of smaller graphs $G_1 \uplus G_2 \dots \uplus G_k$ such that each vertex appears in at most $O(\log^2 n)$ many smaller graphs, and (b) each of the smaller subgraphs G_i is a $\frac{\alpha}{4}$ -weakly regular α -expander, where $\alpha = O(1/\log n)$.*

The proof is in the full version. So, given a graph $G = (V, E)$, we use Theorem 19 to partition the edges into a union of $\frac{\alpha}{4}$ -weakly regular α -expanders, namely H_1, \dots, H_s , where $\alpha = O(1/\log n)$. Further, each vertex in V appears in at most $O(\log^2 n)$ of these expanders. For each graph H_i , we run the greedy algorithm independently. More formally, when an edge e arrives, it belongs to exactly one of the subgraphs H_i . We orient this edge with respect to the greedy algorithm running on H_i . Theorem 6 shows that the discrepancy of each vertex in H_i remains $O(\log^5(nT))$ for each time $t \in [0 \dots T]$ with high probability. Since each vertex in G appears in at most $O(\log^2 n)$ such expanders, it follows that the discrepancy of any vertex in G remains $O(\log^7 n + \log^5 T)$ with high probability. This proves Theorem 1.

References

- 1 Miklós Ajtai, James Aspnes, Moni Naor, Yuval Rabani, Leonard J. Schulman, and Orli Waarts. Fairness in scheduling. *J. Algorithms*, 29(2):306–357, 1998.
- 2 Ryan Alweiss, Yang P. Liu, and Mehtaab Sawhney. Discrepancy minimization via a self-balancing walk. *CoRR*, abs/2006.14009, 2020.
- 3 Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Online discrepancy minimization for stochastic arrivals. *CoRR*, abs/2007.10622, 2020.
- 4 Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. Online vector balancing and geometric discrepancy. In *Proceedings of STOC*, 2020.
- 5 Nikhil Bansal and Joel H. Spencer. On-line balancing of random inputs. *CoRR*, abs/1903.06898, 2019. URL: <http://arxiv.org/abs/1903.06898>.
- 6 Imre Bárány. On a Class of Balancing Games. *J. Comb. Theory, Ser. A*, 26(2):115–126, 1979.
- 7 Aaron Bernstein, Jan van den Brand, Maximilian Probst Gutenberg, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford, and He Sun. Fully-dynamic graph sparsifiers against an adaptive adversary. *CoRR*, abs/2004.08432, 2020. [arXiv:2004.08432](https://arxiv.org/abs/2004.08432).
- 8 Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2001.
- 9 Raaz Dwivedi, Ohad N. Feldheim, Ori Gurel-Gurevich, and Aaditya Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 174:103–131, 2019.
- 10 G. H. Hardy, J. E. Littlewood, and G. Polya. *Inequalities*. Cambridge University Press, 2 edition, 1952.
- 11 Haotian Jiang, Janardhan Kulkarni, and Sahil Singla. Online geometric discrepancy for stochastic arrivals with applications to envy minimization. *CoRR*, abs/1910.01073, 2019.
- 12 Jiri Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. Springer Science & Business Media, 2009.
- 13 Yuval Peres, Kunal Talwar, and Udi Wieder. Graphical balanced allocations and the $(1 + \beta)$ -choice process. *Random Struct. Algorithms*, 47(4):760–775, 2015.
- 14 Joel Spencer. Balancing games. *J. Comb. Theory, Ser. B*, 23(1):68–74, 1977.