

Sparse Reliable Graph Backbones

Shiri Chechik

Department of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot 76100, Israel

Yuval Emek

Computer Engineering and Networks Laboratory
ETH Zurich
CH-8092 Zurich, Switzerland

Boaz Patt-Shamir*

School of Electrical Engineering
Tel Aviv University
Tel Aviv 69978, Israel

David Peleg

Department of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot 76100, Israel

Abstract

Given a connected graph G and a failure probability $p(e)$ for each edge e in G , the *reliability* of G is the probability that G remains connected when each edge e is removed independently with probability $p(e)$. In this paper it is shown that every n -vertex graph contains a sparse *backbone*, i.e., a spanning subgraph with $O(n \log n)$ edges whose reliability is at least $(1 - n^{-\Omega(1)})$ times that of G . Moreover, for any pair of vertices s, t in G , the (s, t) -*reliability* of the backbone, namely, the probability that s and t remain connected, is also at least $(1 - n^{-\Omega(1)})$ times that of G . Our proof is based on a polynomial time randomized algorithm for constructing the backbone. In addition, it is shown that the constructed backbone has nearly the same *Tutte polynomial* as the original graph (in the quarter-plane $x \geq 1, y > 1$), and hence the graph and its backbone share many additional features encoded by the Tutte polynomial.

Keywords: network reliability, sparse subgraphs, Tutte polynomial.

1 Introduction

Finding a sparse subgraph that approximately preserves some key attribute of the original graph is fundamental to network algorithms: any lazy network manager would find the capability to maintain fewer links in a large network a precious gift. This can also be considered from the perspective of identifying a set of redundant edges in a graph. Whether an edge is redundant or not depends of course on the attributes that should be preserved. Spanners [15, 16] for example, approximately preserve pairwise distances in graphs, with a trade-off spectrum

*Supported in part by Israel Science Foundation (grant 1372/09) and by the Israel Ministry of Science and Technology.

between the quality of approximation and the number of edges in the spanner. The general graph attribute we focus on in the current paper is connectivity under random edge failures.

Specifically, we consider the classical setting of *network reliability*, defined over a graph G whose edges e are associated with *failure* probabilities $p(e)$. The *reliability* of G is the probability that G remains connected when each edge e of G is removed independently with probability $p(e)$. Clearly, the reliability of a graph is monotone non-increasing with respect to edge removal. We seek a sparse spanning subgraph (containing all vertices and only a small subset of the edges) of G , referred to henceforth as a *backbone*, whose reliability is almost as good as that of G .

Our main result is a randomized algorithm for constructing a backbone with $O(n \log n)$ edges that approximates the reliability of G to within a (multiplicative) factor of $1 - n^{-\Omega(1)}$, where n denotes the number of vertices. The randomized algorithm allows edge multiplicities, so the original graph G may have significantly more than $\binom{n}{2}$ edges. This construction is tight: we show that there are graphs whose reliability cannot be approximated to within any positive factor by any subgraph with significantly less than $n \log n$ edges. Moreover, the backbone graph approximates not only the *all-terminal* variant of the reliability (the probability that the whole graph remains connected), but also the (s, t) -reliability of G for any two vertices s and t , defined as the probability that s and t remain in the same connected component. Our construction is presented first for the *homogeneous* case, where the failure probability of every edge is some constant $0 < p < 1$, and then extended to the general *heterogeneous* case, assuming that there aren't "too many" edges whose failure probabilities are very close to 1 (see Sect. 3.2 for a precise statement).

It turns out that our backbone also provides a good approximation for the *Tutte polynomial*¹. Specifically, in the quarter-plane $x \geq 1, y > 1$ the Tutte polynomial of the backbone approximates the Tutte polynomial of the original graph to within a factor of $1 \pm n^{-\Omega(1)}$ after multiplying by a (trivially calculated) normalizing factor that accounts for the different number of edges. Since the Tutte polynomial encodes many interesting features of the graph (including its reliability), this result seems to indicate that our backbone construction provides a good representation of the graph in some deeper sense.

Related work. Network reliability is a fundamental problem in operations research since the early days of that discipline [13]; see the survey [2] for a comprehensive account. It is also well-known in the area of computational complexity; various versions of the network reliability problem are listed among the 14 basic #P-complete problems² presented in [19]. In particular, both the all-terminal reliability problem and the (s, t) -reliability problem are known to be #P-hard even when the failure probabilities $p(e)$ are homogeneous. [10] establishes a fully polynomial time randomized approximation scheme (FPRAS) for the problem of evaluating the probability that the graph disconnects under random edge failures. Although this disconnection probability is simply one minus the reliability of the graph, the algorithm of [10] does not translate to a (multiplicative) approximation for the problem of evaluating the reliability. In fact, the approximability of the all-terminal reliability and the (s, t) -reliability problems is still an open question.

A notion somewhat related to ours is that of graph *sparsifiers* [17, 18]: An n vertex weighted graph H is said to be a κ -*sparsifier* of an n vertex weighted graph G if $x^T L_G x \leq x^T L_H x \leq \kappa \cdot x^T L_G x$ for every vector $x \in \mathbb{R}^n$, where L_H and L_G are the Laplacian matrices of H and G , respectively. Sparsifiers are a generalization of the *compressed graphs* of [3], that approximately preserve the total weight of edges crossing any cut in the

¹ The Tutte polynomial $T_G(x, y)$ is a bivariate polynomial whose coefficients are determined by the graph G . See Sect. 5 for details.

² The complexity class #P consists of the counting problems whose decision versions are in NP.

original graph. Indeed, the graph compression condition corresponds to the sparsifier condition restricted to vectors $x \in \{0, 1\}^n$.

One is interested in constructing sparse sparsifiers (hence the name) and the state of the art in that context is the recent construction of $(1 + \epsilon)$ -sparsifiers with $O(n/\epsilon^2)$ edges presented in [6]. Note that unlike the backbone constructed in the current paper, sparsifiers are not required to be subgraphs of the original graph. Furthermore, even if a sparsifier edge is present in the original graph, its weight may be different. In fact, there exist unweighted graphs for which every good sparsifier must introduce edges of widely varying weights [18].

A brief overview of the Tutte polynomial is given in Sect. 5. Here we comment that the computational complexity of evaluating the Tutte polynomial on various points $(x, y) \in \mathbb{R}^2$ is almost completely understood. The problem admits an efficient algorithm if $(x, y) \in \{(1, 1), (-1, -1), (0, 1), (-1, 0)\}$ or if $(x-1)(y-1) = 1$; otherwise it is #P-hard [8]. An FPRAS exists for the $y > 0$ portion of the “Ising” hyperbola $(x-1)(y-1) = 2$ [9]; and unless $\text{RP} = \text{NP}$, an FPRAS does not exist if $x < -1$ or if $y < -1$ except for the aforementioned easy-to-compute points, the ray $x < -1, y = 1$, and the $y < -1$ portion of the hyperbola $(x-1)(y-1) = 2$ [7]. An FPRAS also exists for the quarter-plane $x \geq 1, y \geq 1$ if the minimum degree in G is $\Omega(n)$ [1] and for the half-plane $y > 1$ if the size of a minimum cut in G is $\Omega(\log n)$ [10].

Technique. Our backbone construction samples each edge with probability inverse proportional to its *strength*, a parameter closely related to edge connectivity. This technique was introduced in [3] for the construction of compressed graphs. In [3], the weights of the selected edges are then modified to meet the graph compression condition. This cannot be done when constructing a backbone: we can only remove edges, and are not allowed to change intrinsic attributes (namely failure probabilities) of the remaining ones. Nevertheless, we show that with high probability, the resulting backbone approximately preserves the reliability of the original graph. The main ingredient in our analysis is the fact that graphs with logarithmic edge connectivity are highly reliable [12, 10]. (Note that we do not make any assumptions on the connectivity of the original graph.) The Tutte polynomial analysis is slightly more involved and it essentially relies on an observation of [1] combined with a theorem of [10].

Paper organization. The remainder of this paper is organized as follows. Sect. 2 includes the preliminaries used throughout the paper. The backbone construction is presented in Sect. 3 and the matching lower bound is established in Sect. 4. In Sect. 5 we prove that our backbone also provides a good approximation for the Tutte polynomial.

2 Preliminaries

Unless stated otherwise, all graphs mentioned in this paper are undirected and not necessarily simple (i.e., they may contain parallel edges and self loops). We denote the vertex set and edge set of a graph G by $V(G)$ and $E(G)$, respectively. The graph *induced* on G by a vertex subset $U \subseteq V(G)$ is $G(U) = (U, E(G) \cap (U \times U))$. The graph *induced* on G by an edge subset $F \subseteq E(G)$ is simply $G(F) = (V(G), F)$. Consider some partition of $V(G)$ into $V(G) = U_1 \cup \dots \cup U_r$ and let $\mathcal{U} = \{U_1, \dots, U_r\}$. We refer to the edges in $E(G) \cap \bigcup_{i=1}^r U_i \times U_i$ as the *internal* edges of \mathcal{U} and to the edges in $E(G) \cap \bigcup_{i \neq j} U_i \times U_j$ as the *external* edges of \mathcal{U} .

A *cut* C of a graph G is a partition of $V(G)$ into two non-empty subsets, that is, $C = \{U_1, U_2\}$, where $U_1 \cap U_2 = \emptyset$ and $U_1 \cup U_2 = V(G)$. We say that an edge $e \in E(G)$ *crosses* C if $e \in U_1 \times U_2$. The set of

edges crossing C is denoted by $E(C)$. The cardinality $|E(C)|$ is referred to as the *size* of C ; if the edges of G are associated with weights, then the total weight of all edges in $E(C)$ is referred to as the *weight* of C . A *min cut* (respectively, *min weight cut*) is a cut of minimum size (resp., weight).

3 Backbone Construction and Reliability Analysis

A *network reliability* instance consists of a connected graph G and a *failure* probability $0 < p(e) < 1$ associated with each edge $e \in E(G)$. The network is assumed to occasionally undergo an *edge failure event* \mathcal{F} . Upon such an event, each edge $e \in E(G)$ *fails*, i.e., is removed from the graph, with probability $p(e)$ independently of all other edges. In the *all terminal network reliability* problem, one is interested in the probability that G remains connected following the failure event \mathcal{F} , whereas in the *two terminal network reliability* problem one is interested in the probability that two designated vertices s and t remain in the same connected component of G following the event \mathcal{F} . The former probability, denoted $\text{REL}(G, p)$, is referred to as the *reliability* of G and the latter, denoted $\text{REL}(G, s, t, p)$, is referred to as the *reliability* of s and t in G . Our goal in this section is to establish the existence of a backbone with $O(n \log n)$ edges that approximates the reliability of the original graph.

3.1 Homogeneous failure probabilities

We first focus on the homogeneous case, proving Theorem 3.1; the extension to heterogeneous failure probabilities is discussed in Sect. 3.2.

Theorem 3.1. *There exists an efficient randomized algorithm that given a connected graph G , failure probability $0 < p < 1$, and performance parameters $\delta_1, \delta_2 \geq 1$, outputs a backbone G' of G that satisfies the following three requirements with probability $1 - O(n^{-\delta_1})$:*

- (1) $|E(G')| = O\left(n \log n \cdot \left(\delta_1 + \frac{\delta_2}{1-p}\right)\right)$;
- (2) $\text{REL}(G', p) \geq \text{REL}(G, p) \cdot (1 - O(n^{-\delta_2}))$; and
- (3) $\text{REL}(G', s, t, p) \geq \text{REL}(G, s, t, p) \cdot (1 - O(n^{-\delta_2}))$ for every $s, t \in V(G)$.

Our technique derives from that presented in [3]; for completeness, we describe some ingredients in detail.

Strong components. A graph G is said to be *k-connected* if the size of every cut in G is at least k . Fix some vertex subset $U \subseteq V(G)$. The vertex induced subgraph $G(U)$ is called a *k-strong component* of G if it is k -connected and $G(U')$ is not k -connected for any vertex subset $U' \subseteq V(G)$ such that $U' \supsetneq U$. If $G(U_1)$ and $G(U_2)$, $U_1 \neq U_2$, are k -strong components of G , then U_1 and U_2 must be disjoint, as otherwise $G(U_1 \cup U_2)$ is k -connected. Therefore, if the size of a minimum cut in G is c , then the k -strong components of G for $k = c, c+1, \dots$ define a unique laminar family over $V(G)$, that is, G itself is the sole c -strong component, and for every $k \geq c$, the collection \mathcal{U}_k of vertex sets of the k -strong components forms a partition of $V(G)$, refined by the partition \mathcal{U}_{k+1} .

The *strength* of an edge $e = (u, v) \in E(G)$, denoted k_e , is defined to be the maximum k such that u and v belong to the same k -strong component of G . Note that $k_e \geq k$ for every internal edge of \mathcal{U}_k and $k_e < k$

for every external edge of \mathcal{U}_k . Moreover, if $G(U)$ is a k -strong component, then the strength in $G(U)$ of every edge $e \in E(G) \cap (U \times U)$ is equal to its original strength k_e in G .

Edge sampling. Consider some n -vertex graph G and let $q : E(G) \rightarrow [0, 1]$ be a mapping that assigns some *sampling probability* $q(e)$ to each edge $e \in E(G)$. Given some edge subset $F \subseteq E(G)$, let F^q be a random subset of F that contains each edge $e \in F$ with probability $q(e)$ independently of all other edges and let $G^q = (V(G), E(G)^q)$ be the random graph obtained from G by selecting each edge $e \in E(G)$ in that manner. The *expected graph* \bar{G}^q of G^q is the weighted graph obtained from G by associating a weight $q(e)$ with each edge $e \in E(G)$. As the name implies, for each cut C in G , the weight of C in \bar{G}^q reflects the expected size of C in G^q . The following theorem, established in [11], guarantees that if every cut in the expected graph is sufficiently heavy, then the sizes of cuts in G^q can be “predicted” with high probability.

Theorem 3.2 ([11]). *Let \bar{c} be the weight of a min weight cut in \bar{G}^q and fix some $0 < \epsilon < 1$ and $d > 0$. If $\bar{c} \geq 3(d+2) \ln(n)/\epsilon^2$, then with probability $1 - O(n^{-d})$, every cut in G^q has size between $1 - \epsilon$ and $1 + \epsilon$ times its expected size (i.e., its weight in \bar{G}^q).*

Consider some r disjoint graphs G_1, \dots, G_r . Let $n_i = |V(G_i)|$ for every $1 \leq i \leq r$ and let $n = \sum_{i=1}^r n_i$. For $i = 1, \dots, r$, let $q_i : E(G_i) \rightarrow [0, 1]$ be a mapping that assigns some probability $q_i(e)$ to each edge $e \in E(G_i)$. The statement of Theorem 3.2 can be extended to hold for all graphs G_i simultaneously. This extension can be established by a careful examination of the proof in [11]; for completeness, we provide here a “black-box” proof for this extension.

Corollary 3.3. *Let \bar{c}_i be the weight of a min weight cut in $\bar{G}_i^{q_i}$ for $i = 1, \dots, r$ and fix some $0 < \epsilon < 1$ and $d > 0$. If $\min_{1 \leq i \leq r} \bar{c}_i \geq 3(d+2) \ln(n)/\epsilon^2$, then with probability $1 - O(n^{-d})$, every cut in $G_i^{q_i}$ has size between $1 - \epsilon$ and $1 + \epsilon$ times its expected size (i.e., its weight in $\bar{G}_i^{q_i}$) for all $1 \leq i \leq r$.*

Proof. Let v_i be an arbitrary vertex in $V(G_i)$ for every $1 \leq i \leq r$. Consider the graph G obtained by augmenting the union of G_1, \dots, G_r with sufficiently many sturdy “connector edges” connecting v_i to v_{i+1} for every $1 \leq i \leq r-1$, that is, $V(G) = \bigcup_{i=1}^r V(G_i)$ and $E(G) = \bigcup_{i=1}^r E(G_i) \cup \bigcup_{i=1}^{r-1} F_i$, where F_i consists of m parallel (v_i, v_{i+1}) connector edges for some sufficiently large m . Let $q : E(G) \rightarrow [0, 1]$ be a mapping that agrees with $q_i(e)$ on every edge $e \in E(G_i)$, $1 \leq i \leq r$, and assigns sampling probability $q(e) = 1$ to every edge $e \in F_i$, $1 \leq i \leq r-1$.

We argue that the weight of every cut in \bar{G}^q is at least $3(d+2) \ln(n)/\epsilon^2$; the assertion follows by applying Theorem 3.2 to G and q . To that end, consider some cut $C = \{U_1, U_2\}$ of \bar{G}^q . If there exists some $1 \leq i \leq r-1$ such that $v_i \in U_j$ and $v_{i+1} \in U_{3-j}$, $j \in \{1, 2\}$, then $E(C)$ contains at least m connector edges and the weight of C is at least m . Otherwise, there must exist some $1 \leq i \leq r$ and some cut C_i of $\bar{G}_i^{q_i}$ such that $E(C_i) \subseteq E(C)$, hence the weight of C is at least $\bar{c}_i \geq 3(d+2) \ln(n)/\epsilon^2$. \square

Sampling edges by their strength. We now turn to describe Algorithm **SRGB** (acronym for the paper’s title), performing the actual construction of the sparse reliable backbone. The algorithm is given an n -vertex graph G with edge failure probability p and two performance parameters $\delta_1, \delta_2 \geq 1$. Let

$$\rho = \left\lceil 12 \ln n \cdot \max \left\{ \delta_1 + 2, 2 \frac{\delta_2 + 2}{1 - p} \right\} \right\rceil \quad (1)$$

and define $q(e) = \min\{1, \rho/k_e\}$ for every $e \in E(G)$, where k_e is the strength of e in G . The algorithm constructs the backbone G' of G by selecting each edge $e \in E(G)$ independently with probability $q(e)$, namely, $G' \leftarrow G^q$.

We need to show that Algorithm **SRGB** guarantees the requirements of Theorem 3.1. The authors of [3] analyze a similar construction³ and establish, among other things, the following lemma whose proof is included here for completeness.

Lemma 3.4 ([3]). *The edge strengths satisfy $\sum_{e \in E(G)} 1/k_e \leq n - 1$.*

Proof. Consider some vertex subset $U \subseteq V(G)$. Let C be a min cut of the subgraph $G(U)$ induced by U on G and assume that $|E(C)| = k$. Since the strength of every edge in $G(U)$ is at least k , it follows that $\sum_{e \in E(C)} 1/k_e \leq 1$. On the other hand, as C is a cut, by removing the edges crossing C , the subgraph $G(U)$ breaks down into several connected components. Therefore, the edges in $E(C)$ contributes at most 1 to $\sum_{e \in E(G)} 1/k_e$, whereas their removal increases the number of connected components by at least 1. This gives rise to the following recursive process: find a min cut in G and remove its edges; continue recursively with the resulting connected components. As every application of this recursive process increases the number of connected components by at least 1, it cannot be applied more than $n - 1$ times. The assertion follows since every application removes a subset of the edges that contributes at most 1 to $\sum_{e \in E(G)} 1/k_e$. \square

Since Algorithm **SRGB** takes each edge $e \in E(E)$ into G' with probability at most ρ/k_e , Lemma 3.4 implies that the expected number of edges in G' is $\mathbb{E}[|E(G')|] \leq \rho(n - 1)$; as these random experiments are independent, a standard Chernoff bound argument (see, e.g., [14]) shows that the probability that $|E(G')|$ is greater than, say, twice its expected value is exponentially small. Part (1) of Theorem 3.1 follows. Our goal in the remainder of this section is to prove that with probability $1 - O(n^{-\delta_1})$, the random graph G' satisfies $\text{REL}(G', p) \geq \text{REL}(G, p) \cdot (1 - O(n^{-\delta_2}))$. Proving Part (3) of the theorem, namely, showing that with probability $1 - O(n^{-\delta_1})$ the random graph G' satisfies $\text{REL}(G', s, t, p) \geq \text{REL}(G, s, t, p) \cdot (1 - O(n^{-\delta_2}))$ for every $s, t \in V(G)$, is analogous.

Let $G(U_1), \dots, G(U_r)$ be the ρ -strong components of G and consider some $G(U_i)$, $1 \leq i \leq r$. Let C be a cut in $G(U_i)$ and let e be some edge in $E(C)$. Recall that the strength of e in $G(U_i)$ is equal to its strength in G , denoted k_e . Since e crosses a cut of size $|C|$ in $G(U_i)$, it follows that $k_e \leq |C|$, thus $\sum_{e \in E(C)} 1/k_e \geq 1$. On the other hand, $G(U_i)$ is ρ -connected, hence $k_e \geq \rho$ and $q(e) = \rho/k_e$. Therefore, the weight of C in the expected graph \bar{G}^q is

$$\sum_{e \in E(C)} q(e) = \rho \sum_{e \in E(C)} 1/k_e \geq \rho.$$

By Eq. (1), $\rho \geq 12(\delta_1 + 2) \ln n$, so Corollary 3.3 can be applied to $G(U_1), \dots, G(U_r)$ to conclude that with probability $1 - O(n^{-\delta_1})$, every cut in $G'(U_i)$, $1 \leq i \leq r$, has size at least $\rho/2$ (this probability is with respect to the random choices of Algorithm **SRGB**) — condition the subsequent analysis on that event. Since Eq. (1) also implies that $(1 - p)\rho/2 \geq 12(\delta_2 + 2) \ln n$, an application of Corollary 3.3 to $G'(U_1), \dots, G'(U_r)$ derives⁴ the following corollary.

Corollary 3.5. *By setting $\rho = \left\lceil 12 \ln n \cdot \max \left\{ \delta_1 + 2, 2 \frac{\delta_2 + 2}{1 - p} \right\} \right\rceil$, we ensure that with probability $1 - O(n^{-\delta_2})$, all the components $G'(U_1), \dots, G'(U_r)$ remain connected following an edge failure event \mathcal{F} (in fact, the size of every cut in these components decreases by at most half).*

³ The construction in [3] assigns (new) weights to the edges of the random graph, and hence its analysis follows a different path that requires some additional complications.

⁴ The fact that components of large edge connectivity admit high reliability was originally discovered by [12] and later on restated in [10]. Using their frameworks instead of Corollary 3.3 would have resulted in slightly better constants.

Let A (respectively, A') denote the event that G (resp., G') remains connected after an edge failure event \mathcal{F} and let B (resp., B') denote the event that all the components $G(U_1), \dots, G(U_r)$ (resp., $G'(U_1), \dots, G'(U_r)$) remain connected after an edge failure event \mathcal{F} . We argue that $\mathbb{P}(A') \geq \mathbb{P}(A) \cdot (1 - O(n^{-\delta_2}))$. Corollary 3.5 implies that $\mathbb{P}(B') \geq 1 - O(n^{-\delta_2})$ and by definition, $\mathbb{P}(B') \leq \mathbb{P}(B) \leq 1$. Let $\mathcal{E}_X \subseteq E(G)$ be the set of all edges external to $\{U_1, \dots, U_r\}$. Note that every edge $e \in \mathcal{E}_X$ has strength $k_e < \rho$ in G , and therefore was selected by Algorithm **SRGB** with probability 1. It follows that all those edges are included in G' , i.e., $\mathcal{E}_X \subseteq E(G')$, and thus $\mathbb{P}(A' | B') = \mathbb{P}(A | B) \geq \mathbb{P}(A | \neg B)$. The argument follows by observing that

$$\mathbb{P}(A') \geq \mathbb{P}(A' | B') \cdot \mathbb{P}(B') \geq \mathbb{P}(A | B) \cdot (1 - O(n^{-\delta_2}))$$

and

$$\mathbb{P}(A) \leq \mathbb{P}(A | B) + \mathbb{P}(A | \neg B) \cdot \mathbb{P}(\neg B) \leq \mathbb{P}(A | B) \cdot (1 + O(n^{-\delta_2})).$$

This completes the proof of part (2) of Theorem 3.1 as $\text{REL}(G, p) = \mathbb{P}(A)$ and $\text{REL}(G', p) = \mathbb{P}(A')$.

Las-Vegas implementation. As discussed above, our algorithm satisfies all three requirements with high probability. However, once invoking the algorithm on some instance graph G , one may wish to *ensure* that indeed all three requirements are satisfied. As stated above, the approximability of the all-terminal reliability and (s, t) -reliability problems is still an open question. So, it may seem hopeless to be able to check if requirements (2) and (3) indeed hold for a specific invocation of our algorithm. However, following our line of arguments, one can see that to guarantee that requirements (2) and (3) hold, it suffices to check that the minimal cut in all ρ -strong components $G'(U_1), \dots, G'(U_r)$ is at least $\rho/2$. This, of course, can be done in polynomial time.

Running Time. The running time of our algorithm is dominated by finding the strength of the edges. It is not hard to see that this can be done in polynomial time (by hierarchically decomposing the graph via n minimum cut computations). However, this could be too slow for certain applications. Luckily, our algorithm does not require the exact values k_e ; rather, one can settle for approximate values \tilde{k}_e satisfying some desired properties. This can be done, using some ideas presented in [3], so as to improve the overall running time to $O(m \log^2 n)$.

Specifically, it is shown in [3] (Section 4) how to find in $O(m \log^2 n)$ time approximate values \tilde{k}_e obeying the following two requirements: (R1) $\tilde{k}_e \leq k_e$; and (R2) $\sum 1/\tilde{k}_e = O(n)$. Using the estimates \tilde{k}_e rather than k_e in our construction can be implemented to run in time $O(m \log^2 n)$. Moreover, observe that by (R1), each edge e is now taken with higher probability, therefore the probability that requirements (2) and (3) of Theorem 3.1 still hold may only increase. In addition, by (R2), the number of edges in our resulting subgraph may only increase by a constant factor, hence requirement (1) of Theorem 3.1 is also satisfied.

3.2 Heterogeneous failure probabilities

We now turn to discuss the heterogeneous case, where each edge e has a different failure probability $p(e)$. It's not hard to verify that setting $\rho = \left\lceil 12 \ln n \cdot \max \left\{ \delta_1 + 2, 2 \frac{\delta_2 + 2}{1 - \hat{p}} \right\} \right\rceil$, where \hat{p} is the highest failure probability in G , yields the same analysis and results as for the homogeneous case. However, if \hat{p} is close to 1, then this would result in a backbone G' with too many edges. Consider, for example, an arbitrary graph G^- where all edges have the same (constant) failure probability $0 < p < 1$, and augment it into a graph G by adding a single new edge with very high failure probability, say, $\hat{p} = 1 - 1/n^2$. Clearly, applying Algorithm **SRGB** to G^- will generate, with probability at least $1 - O(n^{-\delta_1})$, a backbone G'^- with $O(n \log n)$ edges such that

$\text{REL}(G'^-, p) \geq \text{REL}(G, p) \cdot (1 - O(n^{-\delta_2}))$. Using the algorithm with \hat{p} , however, will yield a very high value for ρ , and the resulting backbone G' is likely to contain $\Omega(n^2)$ edges.

Hence, we are interested in constructing a backbone G'' with $O(n \log n)$ edges that approximates the reliability of G even when some of the failure probabilities are close to 1. Define the *average failure probability* of cut C in G as $\sum_{e \in E(C)} p(e) / |E(C)|$. We show that if the average failure probability of every cut in G is at most \bar{p} , then it is possible to construct a backbone G'' such that with probability at least $1 - O(n^{-\delta_1})$, G'' has $O\left(\frac{n \log n}{1-\bar{p}} \left(\delta_1 + \frac{\delta_2}{1-\bar{p}}\right)\right)$ edges and $\text{REL}(G'', p) \geq \text{REL}(G, p) \cdot (1 - O(n^{-\delta_2}))$.

Let \hat{G} be the graph obtained from G by erasing all edges with failure probability greater than $1/2 + \bar{p}/2$. Set

$$\rho = \left\lceil 12 \ln n \cdot \max \left\{ \delta_1 + 2, 2 \frac{\delta_2 + 2}{1/2 - \bar{p}/2} \right\} \right\rceil \quad (2)$$

and construct the backbone G' by applying Algorithm **SRGB** (with ρ as defined in Eq. (2)) to \hat{G} . Let $\hat{G}(U_1), \dots, \hat{G}(U_r)$ be the ρ -strong components of \hat{G} and fix $\hat{U} = \{U_1, \dots, U_r\}$. Denote the set of external edges of \hat{U} in the graph G by \mathcal{E}_X . Enhance G' by augmenting it with all edges in \mathcal{E}_X that are not already in $E(G')$ — set G'' to be the resulting graph.

Let A (respectively, A'') denote the event that G (resp., G'') remains connected following an edge failure event \mathcal{F} and let B (resp., B'') denote the event that all the components $G(U_1), \dots, G(U_r)$ (resp., $G''(U_1), \dots, G''(U_r)$) remain connected following an edge failure event \mathcal{F} . Since every edge in \hat{G} has failure probability at most $1/2 + \bar{p}/2$, Corollary 3.5 guarantees that $1 - O(n^{-\delta_2}) \leq \mathbb{P}(B'') \leq \mathbb{P}(B) \leq 1$. Since $\mathcal{E}_X \subseteq E(G'')$, it follows that $\mathbb{P}(A'' \mid B'') = \mathbb{P}(A \mid B)$. Therefore, by the line of arguments used in Sect. 3.1, we conclude that $\mathbb{P}(A'') \geq \mathbb{P}(A) \cdot (1 - O(n^{-\delta_2}))$. So, it remains to show that $\mathbb{E}[|E(G'')|] = O(\rho(n-1)/(1-\bar{p}))$.

Denote the set of external edges of \hat{U} in the graph G' by E_1 and the set of external edges of \hat{U} in G that were subsequently added to G'' by $E_2 = \mathcal{E}_X \setminus E_1 = E(G'') \setminus E(G')$. By the line of arguments used in Sect. 3.1, we get $|E_1| \leq \rho(n-1)$. Note that the removal of $\mathcal{E}_X = E_1 \cup E_2$ disconnects G . This does not mean that $E_1 \cup E_2$ are the crossing edges of some cut in G , as its removal may disconnect G into more than two connected components. Nevertheless, we argue that the average failure probability over all edges in $E_1 \cup E_2$ is at most \bar{p} . To see this, let $C_i = \{U_i, V(G) - U_i\}$ be the cut that disconnects U_i from the rest of the graph. Then $\bigcup_i E(C_i) = \mathcal{E}_X$, where each edge of \mathcal{E}_X appears exactly twice in $\bigcup_i E(C_i)$. As the average failure probability on each cut C_i separately is at most \bar{p} , we get the same bound also for the average over $\mathcal{E}_X = E_1 \cup E_2$.

So, we know that $\frac{\sum_{e \in E_1} p(e) + \sum_{e \in E_2} p(e)}{|E_1| + |E_2|} \leq \bar{p}$ and recall that $p(e) > 1/2 + \bar{p}/2$ for every $e \in E_2$. Therefore, we can apply a Markov type argument to conclude that $\frac{|E_2|}{|E_1| + |E_2|} < \frac{\bar{p}}{1/2 + \bar{p}/2} = \frac{2\bar{p}}{1+\bar{p}}$. Plugging in the fact that $|E_1| \leq \rho(n-1)$, we get that $|E_2| < \frac{2\bar{p}/(1+\bar{p})}{1-2\bar{p}/(1+\bar{p})} \rho(n-1) = \frac{2\bar{p}}{1-\bar{p}} \rho(n-1)$. We summarize as follows.

Theorem 3.6. *There exists an efficient randomized algorithm that given a connected graph G , failure probability $p(e)$ for each $e \in E(G)$, where the average failure probability of every cut in G is at most \bar{p} , and performance parameters $\delta_1, \delta_2 \geq 1$, outputs a backbone G' of G that satisfies the following three requirements with probability $1 - O(n^{-\delta_1})$:*

- (1) $|E(G')| = O\left(\frac{n \log(n)}{1-\bar{p}} \left(\delta_1 + \frac{\delta_2}{1-\bar{p}}\right)\right)$;
- (2) $\text{REL}(G', p) \geq \text{REL}(G, p) \cdot (1 - O(n^{-\delta_2}))$; and

(3) $\text{REL}(G', s, t, p) \geq \text{REL}(G, s, t, p) \cdot (1 - O(n^{-\delta_2}))$ for every choice of $s, t \in V(G)$.

4 A tight lower bound

We now turn to show that the $O(n \log n)$ upper bound on the number of edges is indeed tight. Consider some graph G and let \mathcal{S}_G be the collection of all spanning subgraphs of G . Given some failure probability $0 < p < 1$ and some real $\epsilon > 0$, let

$$\psi_{p,\epsilon}(G) = \max \left\{ \text{REL}(H, p) \mid H \in \mathcal{S}_G, |E(H)| \leq (1 - \epsilon)n \log_{1/p} n \right\}.$$

We establish the following theorem.

Theorem 4.1. *For every failure probability $0 < p < 1$, the family $\{K_{n,n}\}_{n=1}^\infty$ of complete bipartite graphs with n vertices on each side satisfies*

- (1) $\lim_{n \rightarrow \infty} \text{REL}(K_{n,n}, p) = 1$; and
- (2) for every constant $\epsilon > 0$, $\lim_{n \rightarrow \infty} \psi_{p,\epsilon}(K_{n,n}) = 0$.

Proof. Requirement (1) is immediately satisfied by Theorem 3.2, so it remains to establish requirement (2). To that end, fix some n and consider some constant $\epsilon > 0$ and some spanning subgraph H of $K_{n,n}$ such that $|E(H)| \leq (1 - \epsilon)n \log_{1/p} n$. The subgraph H is bipartite as well; let $Z = \{v_1, \dots, v_k\}$ be the set of vertices of degree at most $(1 - \epsilon/2) \log_{1/p} n$ on its left side. By a straightforward counting argument, $k \geq n \left(1 - \frac{1-\epsilon}{1-\epsilon/2}\right) > \epsilon n/2$.

Let A_i be the event that v_i becomes an isolated vertex under an edge failure event \mathcal{F} . By definition, $\mathbb{P}(A_i) \geq p^{(1-\epsilon/2) \log_{1/p} n} = n^{-(1-\epsilon/2)}$. Since H is bipartite, the events A_1, \dots, A_k are independent (each determined by a disjoint set of edges), hence the probability that none of them occurs is at most

$$\left(1 - n^{-(1-\epsilon/2)}\right)^k \leq \left(1 - n^{-(1-\epsilon/2)}\right)^{\epsilon n/2} \leq e^{-\epsilon n^{\epsilon/2}/2},$$

which tends to 0 as $n \rightarrow \infty$. The assertion follows as $\text{REL}(H, p) \leq \mathbb{P}(\neg A_1 \wedge \dots \wedge \neg A_k)$. \square

5 The Tutte Polynomial of the Backbone

The Tutte polynomial, introduced by W.T. Tutte, is a bivariate polynomial whose coefficients are determined by a given graph. The Tutte polynomial is a central concept in algebraic graph theory, as it captures many interesting properties of the graph from which it is derived. [4] gives a relatively updated treatment of the concept. Below, we only review the basic definitions and some key results.

Let G be a graph. The Tutte polynomial of G at point $(x, y) \in \mathbb{R}^2$, denoted $T_G(x, y)$, is defined by

$$T_G(x, y) = \sum_{F \subseteq E(G)} (x-1)^{K(F)-K(G)} (y-1)^{K(F)+|F|-n},$$

where $n = |V(G)|$, and for $F \subseteq E(G)$, $K(F)$ denotes the number of connected components in the graph $(V(G), F)$, and $K(G) = K(E(G))$. The Tutte polynomial contains many interesting points and lines that capture combinatorial features of the graph G , including:

- $T_G(1, 1)$ counts the number of spanning trees of G .
- $T_G(2, 1)$ counts the number of spanning forests of G .
- $T_G(1, 2)$ counts the number of connected spanning subgraphs of G .
- At $y = 0$ and $x = 1 - \lambda$ for positive integer λ , the Tutte polynomial specializes to yield the *chromatic polynomial* $\chi_G(\lambda) = (-1)^{n-K(G)} \lambda^{K(G)} T_G(1 - \lambda, 0)$ that counts the number of legal vertex colorings of G using λ colors.
- At $x = 1$ and $y = 1/p$ for $0 < p < 1$, the Tutte polynomial specializes to yield the reliability of G , $\text{REL}(G, p) = (1 - p)^{n-1} p^{|E(G)|-n+1} T_G(1, 1/(1 - p))$.
- Along the hyperbolas $(x - 1)(y - 1) = s$ for any positive integer s , the Tutte polynomial specializes to the partition function of the s -state *Potts model* of statistical mechanics.

The reader is referred to the survey [5] for more interpretations.

Our goal in this section is to prove the following theorem.

Theorem 5.1. *For every point (x, y) in the quarter-plane $x \geq 1, y > 1$, there exists an efficient randomized algorithm that given a connected graph G and performance parameters $\delta_1, \delta_2 \geq 1$, outputs a backbone G' of G that satisfies the following two requirements with probability $1 - O(n^{-\delta_1})$:*

- (1) $|E(G')| = O\left(n \log(n) \left(\delta_1 + \frac{\delta_2}{1-1/y}\right)\right)$; and
- (2) the evaluations of $T_G(\cdot, \cdot)$ and $T_{G'}(\cdot, \cdot)$ at (x, y) satisfy

$$T_G(x, y) \cdot \left(1 - O\left(n^{-\delta_2}\right)\right) \leq y^{|E(G)|-|E(G')|} \cdot T_{G'}(x, y) \leq T_G(x, y) \cdot \left(1 + O\left(n^{-\delta_2}\right)\right).$$

It is important to point out that the role of the $y^{|E(G)|-|E(G')|}$ normalizing factor in part (2) of Theorem 5.1 is to compensate for the fact that the number of edges in the backbone G' is smaller than that of the original graph G . This can not be avoided since, in general, the more edges a graph has, the larger value its Tutte polynomial evaluates to at points (x, y) in the quarter-plane $x \geq 1, y > 1$. This is best demonstrated at point $(x = 2, y = 2)$, where the Tutte polynomial merely counts the number of edge subsets, i.e., $T_G(2, 2) = 2^{|E(G)|}$. The key point here is that this normalizing factor depends only on the number of edges in G and G' and not on the topologies of these graphs. In particular, since $y^{|E(G)|-|E(G')|}$ can obviously be calculated in polynomial time, the problem of approximating the Tutte polynomial of graphs in the quarter-plane $x \geq 1, y > 1$ (which is still open to the most part) reduces to the case of graphs with $O(n \log n)$ edges.

Note first that along the ray $x = 1, y > 1$, the Tutte polynomial of G specializes to the reliability of G following the identity

$$\text{REL}(G, p) = (1 - p)^{n-1} p^{|E(G)|-n+1} T_G(1, 1/p).$$

Therefore, when $x = 1$, Theorem 5.1 follows directly from Theorem 3.1. Assume hereafter that $x > 1$.

Fix $q = 1 - 1/y$. The construction of G' is identical to that described in Sect. 3.1 when setting $p = 1 - q$. In Sect. 3.1 we argued that with very high probability, $|E(G')| = O(n\rho)$, which implies requirement (1) of Theorem 5.1 by the choice of ρ . Our goal in the remainder of this section is to prove that requirement (2) holds with probability $1 - O(n^{-\delta_1})$.

The authors of [1] observe that in the quarter-plane $x > 1, y > 1$, the Tutte polynomial of a connected graph G with n vertices and m edges can be expressed as

$$T_G(x, y) = \frac{y^m}{(x-1)(y-1)^n} \mathbb{E} \left[z^{K(G^q)} \right],$$

where $z = (x-1)(y-1)$. Theorem 5.1 will be established by showing that $\mathbb{E}[z^{K(G^q)}] \approx \mathbb{E}[z^{K(G'^q)}]$.

Let $G(U_1), \dots, G(U_r)$ be the ρ -strong components of G and Let $\mathcal{E}_X \subseteq E(G)$ be the set of all edges external to $\{U_1, \dots, U_r\}$. Consider the collection \mathcal{H} of all spanning subgraphs H of G such that

- (1) $\mathcal{E}_X \subseteq E(H)$; and
- (2) $H(U_i)$ is $(\rho/2)$ -connected for every $1 \leq i \leq r$.

By definition, G itself is in \mathcal{H} . Recall that G' contains all edges whose strength in G is smaller than ρ . Eq. (1) implies that $\rho \geq 12(\delta_1 + 2) \ln n$, thus we can follow the line of arguments used in Sect. 3.1 and apply Corollary 3.3 to $G(U_1), \dots, G(U_r)$ to conclude that with probability $1 - O(n^{-\delta_1})$, G' is also in \mathcal{H} , where the probability is taken with respect to the random choices of Algorithm **SRGB**. Our analysis relies on showing that $\mathbb{E}[z^{K(H^q)}]$ is approximately the same for all graphs $H \in \mathcal{H}$.

Consider an arbitrary graph $H \in \mathcal{H}$. Partition the edges of H into $E(H) = \mathcal{E}_I \cup \mathcal{E}_X$, where $\mathcal{E}_I = \bigcup_{i=1}^r E(H) \cap (U_i \times U_i)$ and $\mathcal{E}_X = E(H) - \mathcal{E}_I$. We express $\mathbb{E}[z^{K(H^q)}]$ as

$$\mathbb{E} \left[z^{K(H^q)} \right] = \sum_{F \subseteq \mathcal{E}_X} \mathbb{E} \left[z^{K(H^q)} \mid \mathcal{E}_X^q = F \right] \cdot \mathbb{P}(\mathcal{E}_X^q = F)$$

and establish Theorem 5.1 by proving that

$$\mathbb{E} \left[z^{K(H^q)} \mid \mathcal{E}_X^q = F \right] = z^{K_F} \left(1 \pm O \left(n^{-\delta_2} \right) \right)$$

for every $F \subseteq \mathcal{E}_X$, where $K_F = K(V(H), \mathcal{E}_I \cup F)$ denotes the number of connected components in the graph induced on H by the edges in $\mathcal{E}_I \cup F$.

Assume first that $0 < z \leq 1$ and fix some edge subset $F \subseteq \mathcal{E}_X$. By Eq. (1), $q\rho/2 \geq 12(\delta_2 + 2) \ln n$, thus an application of Corollary 3.3 to $H(U_1), \dots, H(U_r)$ implies that with probability $1 - O(n^{-\delta_2})$, all these components remain connected, where the probability is taken with respect to the experiment H^q . Therefore,

$$z^{K_F} \left(1 - O \left(n^{-\delta_2} \right) \right) \leq \mathbb{E} \left[z^{K(H^q)} \mid \mathcal{E}_X^q = F \right] \leq z^{K_F}$$

which establishes the assertion.

Now, assume that $z > 1$ and fix some edge subset $F \subseteq \mathcal{E}_X$. Let $\Gamma = (V(H), \mathcal{E}_I^q \cup F)$ be the random graph obtained from H by taking the edges in F and selecting each edge $e \in \mathcal{E}_I$ independently with probability q . Let $H_I = (V(H), \mathcal{E}_I)$ be the graph induced on H by the edges in \mathcal{E}_I and let $\kappa = K(H_I^q) - K(H_I)$ be a random variable that takes on the number of connected components ‘‘added’’ to H_I due to the experiment H_I^q . We have

$$\begin{aligned} z^{K_F} &\leq \mathbb{E} \left[z^{K(H^q)} \mid \mathcal{E}_X^q = F \right] = \sum_{j \geq 0} \mathbb{P}(K(\Gamma) = K_F + j) \cdot z^{K_F + j} \\ &= z^{K_F} \cdot \sum_{j \geq 0} \mathbb{P}(K(\Gamma) = K_F + j) \cdot z^j \leq z^{K_F} \cdot \sum_{j \geq 0} \mathbb{P}(K(\Gamma) \geq K_F + j) \cdot z^j \\ &\leq z^{K_F} \cdot \sum_{j \geq 0} \mathbb{P}(\kappa \geq j) \cdot z^j = z^{K_F} \left(1 + \sum_{j \geq 1} \mathbb{P}(\kappa \geq j) \cdot z^j \right), \end{aligned}$$

where the last inequality follows from the definition of κ as the event $K(\Gamma) \geq K_F + j$ cannot occur unless $\kappa \geq j$. It remains to show that $\sum_{j \geq 1} \mathbb{P}(\kappa \geq j) \cdot z^j = O(n^{-\delta_2})$. The following theorem is established in [10].

Theorem 5.2 ([10]). *Let G be a connected n -vertex graph and let c be the size of a minimum cut in G . Fix some reals $d > 1$ and $q \in [0, 1]$ and integer $t \geq 2$. If $c \geq (d + 2) \log_{1/(1-q)} n$, then $\mathbb{P}(K(G^q) \geq t) < n^{-dt/2}$.*

Theorem 5.2 can be extended to yield the following corollary by following the same “black-box” type of argument employed in the proof of Corollary 3.3.

Corollary 5.3. *Consider some r disjoint graphs G_1, \dots, G_r . Let $n_i = |V(G_i)|$ for every $1 \leq i \leq r$ and let $n = \sum_{i=1}^r n_i$. Let c_i be the size of a minimum cut in G_i for $i = 1, \dots, r$. Set $\tilde{G} = (\bigcup_{i=1}^r V(G_i), \bigcup_{i=1}^r E(G_i))$. Fix some reals $d > 1$ and $q \in [0, 1]$ and integer $t \geq 2$. If $\min_{1 \leq i \leq r} c_i \geq (d + 2) \log_{1/(1-q)} n$, then $\mathbb{P}(K(\tilde{G}^q) \geq r + t - 1) < n^{-dt/2}$.*

Recall that we wish to show that $\sum_{j \geq 1} \mathbb{P}(\kappa \geq j) \cdot z^j = O(n^{-\delta_2})$. Eq. (1) yields $\rho/2 \geq 12(\delta_2 + 2) \ln(n)/q > (\delta_2 + 2) \log_{1/(1-q)} n$, so we can use Corollary 5.3 to deduce that $\mathbb{P}(\kappa \geq j) < n^{-\delta_2(j+1)/2}$. Therefore,

$$\begin{aligned} \sum_{j \geq 1} \mathbb{P}(\kappa \geq j) \cdot z^j &< \sum_{j \geq 1} n^{-\delta_2(j+1)/2} \cdot z^j = z^{-1} \cdot \sum_{j \geq 2} \left(z n^{-\delta_2/2} \right)^j \\ &= z^{-1} \left(z n^{-\delta_2/2} \right)^2 \cdot \sum_{j \geq 0} \left(z n^{-\delta_2/2} \right)^j \leq 2z n^{-\delta_2}, \end{aligned}$$

where the last inequality follows by assuming that n is sufficiently large so that $z n^{-\delta_2/2} \leq 1/2$. This completes the proof of Theorem 5.1.

References

- [1] N. Alon, A. Frieze, and D. Welsh. Polynomial time randomized approximation schemes for Tutte–Gröthendieck invariants: the dense case. *Random Struct. Algorithms*, 6(4):459–478, 1995.
- [2] M.O. Ball, C.J. Colbourn, and J.S. Provan. Network reliability. *Handbook of Operations Research: Network Models*, pp. 673–762, Elsevier North-Holland, 1995.
- [3] A. Benczúr and D.R. Karger. Approximating $s - t$ minimum cuts in $\tilde{O}(n^2)$ time. In *Proc. 28th ACM Symp. on the Theory of Computing (STOC)*, pages 47–55, 1996.
- [4] B. Bollobás. *Modern Graph Theory*. Graduate texts in mathematics, Springer-Verlag, Berlin, New York, 1998.
- [5] T.H. Brylawski and J.G. Oxley. The Tutte polynomial and its applications. *Matroid Applications*, N. White, Ed., Cambridge Univ. Press, Cambridge, pp. 123–225, 1992.
- [6] J.D. Batson, D.A. Spielman, and N. Srivastava. Twice-Ramanujan sparsifiers. In *Proc. of the 41st ACM Symp. on Theory of Computing (STOC)*, pages 255–262, 2009.
- [7] L.A. Goldberg and M. Jerrum. Inapproximability of the Tutte polynomial. *Inf. Comput.*, 206(7):908–929, 2008.
- [8] F. Jaeger, D.L. Vertigan, and D.J.A. Welsh. On the computational complexity of the Jones and Tutte polynomials. *Math. Proc. Camb. Phil. SOC.*, 108:35–53, 1990.
- [9] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput. (SICOMP)*, 22(5):1087–1116, 1993.
- [10] D.R. Karger. A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem. *SIAM J. Comput. (SICOMP)*, 29(2):492–514, 1999.

- [11] D.R. Karger. Random sampling in cut, flow, and network design problems. *Mathematics of Operations Research (MOR)*, 24(2):383–413, 1999.
- [12] M.V. Lomonosov and V.P. Polesskii. Lower bound of network reliability. *Probl. Peredachi Inf.*, 8(2):47–53, 1972.
- [13] E.F. Moore and C.E. Shannon. Reliable circuits using less reliable relays. *J. Franklin Inst.*, 262:191–208; 263:281–297, 1956.
- [14] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.
- [15] D. Peleg and A.A. Schäffer. Graph spanners. *J. of Graph Theory*, 13:99–116, 1989.
- [16] D. Peleg and J.D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. on Computing*, 18:740–747, 1989.
- [17] D.A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proc. of the 36th ACM Symp. on Theory of Computing (STOC)*, pages 81–90, 2004.
- [18] D.A. Spielman and S.-H. Teng. Spectral sparsification of graphs. Available at <http://arxiv.org/abs/0808.4134>, 2008.
- [19] L.G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput. (SICOMP)*, 8(3):410–421, 1979.