Stochastic Kronecker Graphs

Mohammad Mahdian* Ying Xu†

Abstract

A random graph model based on Kronecker products of probability matrices has been recently proposed as a generative model for large-scale real-world networks such as the web. This model simultaneously captures several well-known properties of real-world networks; in particular, it gives rise to a heavy-tailed degree distribution, has a low diameter, and obeys the densification power law. Most properties of Kronecker products of graphs (such as connectivity and diameter) are only rigorously analyzed in the deterministic case. In this paper, we study the basic properties of stochastic Kronecker products based on an initiator matrix of size two (which is the case that is shown to provide the best fit to many real-world networks). We will show a phase transition for the emergence of the giant component and another phase transition for connectivity, and prove that such graphs have constant diameters beyond the connectivity threshold, but are not searchable using a decentralized algorithm.

1 Introduction

A generative model based on Kronecker matrix multiplication was recently proposed by Leskovec et al. [10] as a model that captures many properties of real-world networks. In particular, they observe that this model exhibits a heavy-tailed degree distribution, and has an average degree that grows as a power law with the size of the graph, leading to a diameter that stays bounded by a constant (the so-called *densification power law* [12]). Furthermore, Leskovec and Faloutsos [11] fit the stochastic model to some real world graphs, such as Internet Autonomous Systems graph and Epinion trust graphs, and find that Kronecker graphs with appropriate 2×2 initiator matrices mimic very well many properties of the target graphs.

Most properties of the Kronecker model (such as connectivity and diameter) are only rigorously analyzed in the deterministic case (i.e., when the initiator matrix is a binary matrix, generating a single graph, as opposed to a distribution over graphs), and empirically shown in the general stochastic case [10]. In this paper we analyze some basic graph properties of stochastic Kronecker graphs with an initiator matrix of size 2. This is the case that is shown by Leskovec and Faloutsos [11] to provide the best fit to many real-world networks. We give necessary and sufficient conditions for Kronecker graphs to be connected or to have giant components of size $\Theta(n)$ with high probability[1]. Our analysis of the connectivity of Kronecker graphs is based on a general lemma (Theorem 1) that might be of independent interest. We prove that under the parameters that the graph is connected with high probability, it also has a constant diameter with high probability. This unusual property is consistent with the observation of Leskovec et al. [12] that in many real-world graphs the effective diameters do not increase, or even shrink, as the sizes of the graphs increase,

*Yahoo! Research. E-mail: mahdian@yahoo-inc.com.

†Stanford University. E-mail: xuying@cs.stanford.edu. Work performed in part while visiting Yahoo! Research.

[1] Throughout the paper by "with high probability" we mean with probability $1 - o(1)$.

which is violated by many other random graph models with increasing diameters. Finally we show that Kronecker graphs do not admit short (poly-logarithmic) routing paths by decentralized routing algorithms based on only local information.

1.1 The Model and Overview of Results

In this paper we analyze stochastic Kronecker graphs with an initiator matrix of size 2, as defined below:

Definition 1. *A* (stochastic) Kronecker graph *is defined by*

(i) an integer k, and

(ii) a symmetric 2×2 matrix θ: $\theta[1,1] = \alpha, \theta[1,0] = \theta[0,1] = \beta, \theta[0,0] = \gamma$, where $0 \leq \gamma \leq \beta \leq \alpha \leq 1$. We call θ the base *matrix or the* initiator *matrix.*

The graph has $n = 2^k$ vertices, each vertex labeled by a unique bit vector of length k; given two vertices u with label $u_1 u_2 \ldots u_k$ and v with label $v_1 v_2 \ldots v_k$, the probability of edge (u, v) existing, denoted by $P[u, v]$, is $\prod_i \theta[u_i, v_i]$, independent on the presence of other edges.

In particular, when $\alpha = \beta = \gamma$, the Kronecker graph becomes the well studied random graph $G(n, p)$ with $p = \alpha^k$. Leskovec and Faloutsos [11] showed that the Kronecker graph model with 2×2 initiator matrices satisfying the above conditions is already very powerful in simulating real world graphs. In fact, their experiment shows that the matrix $[.98, .58; .58, .06]$ is a good fit for the Internet AS graph. When the base matrix does not satisfy the condition stated in the above definition (i.e., when $\alpha \geq \gamma \geq \beta$ or $\beta \geq \alpha \geq \gamma$), Kronecker graphs appear to have different structural properties, and require different analytic techniques. We can prove some of our results in these regimes as well; however, due to lack of space and since this setting of parameters is less appealing as a generative model for the web, these results are deferred to Appendix A.

We analyze basic graph properties of the stochastic Kronecker graph model. In particular, we prove that the necessary and sufficient condition for Kronecker graphs to be connected with high probability is $\beta + \gamma > 1$ or $\alpha = \beta = 1, \gamma = 0$ (Section 2.2); the necessary and sufficient condition for Kronecker graphs to have a giant component of size $\Theta(n)$ with high probability is $(\alpha+\beta)(\beta+\gamma) > 1$, or $(\alpha + \beta)(\beta + \gamma) = 1$ and $\alpha + \beta > \beta + \gamma$ (Section 2.3); if $\beta + \gamma > 1$, the diameters of Kronecker graphs are constant with high probability (Section 3); and that no decentralized search algorithm can find a path of length $o(n^{(1-\alpha)\log_2 e})$ between a given pair of vertices in Kronecker graphs, unless the graph is deterministic (Section 4).

Besides Kronecker graphs, we also define a general family of random graphs $G(n, P)$, which generalizes all random graph models where edges are independent, including Kronecker graphs and $G(n, p)$.

Definition 2. *A random graph $G(n, P)$, where n is an integer and P is an $n \times n$ matrix with elements in $[0, 1]$, has n vertices and includes each edge (i, j) independently with probability $P[i, j]$.*

Throughout this paper we consider undirected $G(n, P)$: P is symmetric and edges are undirected. We prove two useful theorems about connectivity and searchability in this model, which may be of independent interest; namely, we show that if the min-cut size of the weighted graph defined by P is at least $c \ln n$ (c is a sufficiently large constant), then with high probability $G(n, P)$ is connected (Section 2.1); we also prove a monotonicity property for searchability in this model (Section 4).

2 Connectivity and Giant Components

We first state a sufficient condition for connectivity of general random graphs $G(n, P)$ (Section 2.1), then use this condition to analyze connectivity and giant components of Kronecker graphs (Section 2.2, 2.3).

2.1 Connectivity of $G(n, P)$

We give a sufficient condition of the matrix P for $G(n, P)$ graphs to be connected. Let V be the set of all vertices. For any $S, S' \subseteq V$, define $P(i, S) = \sum_{j \in S} P[i, j]; P(S, S') = \sum_{i \in S, j \in S'} P[i, j]$.

Theorem 1. *If the min-cut size of the weighted graph defined by P is $c \ln n$ (c is a sufficiently large constant), i.e. $\forall S \subset V, P(S, V \setminus S) \geq c \ln n$, then with high probability $G(n, P)$ is connected.*

Proof. A *k-minimal cut* is a cut whose size is at most k times the min-cut size. We use the following result about the number of k-minimal cuts due to Karger and Stein [5]: In any weighted graph, the number of k-minimal cuts is at most $O((2n)^{2k})$.

Consider the weighted graph defined by P. Denote its min-cut size by t. We say a cut is a k-cut if its size is between kt and $(k+1)t$. By the above result there are at most $O((2n)^{2k+2})$ k-cuts. Now consider a fixed k-cut in a random realization of $G(n, P)$: the expected size of the cut is at least kt, so by Chernoff bound the probability that the cut has size 0 in the realization is at most $e^{-kt/2}$. Taking the union bound over all k-cuts, for all $k = 1, 2, \ldots, n^2$, the probability that at least one cut has size 0 is bounded by

$$\sum_{k=1,\ldots,n^2} e^{-kt/2} O((2n)^{2k+2})$$

For $t = c \ln n$ where c is a sufficiently large constant, this probability is $o(1)$. Therefore with high probability $G(n, P)$ is connected. □

Note that $G(n, p)$ is known to be disconnected with high probability when $p \leq (1 - \epsilon) \ln n / n$, i.e., when the min-cut size is $(1 - \epsilon) \ln n$. Therefore the condition in the above theorem is tight up to a constant factor. Also, extrapolating from $G(n, p)$, one might hope to prove a result similar to the above for the emergence of the giant component; namely, if the size of the min-cut in the weighted graph defined by P is at least a constant, $G(n, P)$ has a giant component. However, this result is false, as can be seen from this example: n vertices are arranged on a cycle, and P assigns a probability of 0.5 to all pairs that are within distance c (a constant) on the cycle, and 0 to all other pairs. It is not hard to prove that with high probability $G(n, P)$ does not contain any connected component of size larger than $O(\log n)$.

2.2 Connectivity of Kronecker Graphs

We define the *weight* of a vertex to be the number of 1's in its label; denote the vertex with weight 0 by $\vec{0}$, and the vertex with weight k by $\vec{1}$. We say a vertex u is *dominated* by vertex u', denoted by $u \leq u'$, if for any bit i, $u_i \leq u'_i$. Recall that $P[u, v]$ is as defined in Definition 1.

The following lemmas state some simple facts about Kronecker graphs. Lemma 2 is trivially true given the condition $\alpha \geq \beta \geq \gamma$. The proof of Lemma 3 is presented in Appendix B.

Lemma 2. *For any vertex u, $\forall v, P[u, v] \geq P[\vec{0}, v]; \forall S, P(u, S) \geq P(\vec{0}, S)$. Generally, for any vertices $u \leq u'$, $\forall v, P[u, v] \leq P[u', v]; \forall S, P(u, S) \leq P(u', S)$.*

Lemma 3. *The expected degree of a vertex u with weight l is $(\alpha+\beta)^l(\beta+\gamma)^{k-l}$.*

Theorem 4. *The necessary and sufficient condition for Kronecker graphs to be connected with high probability (for large k) is $\beta+\gamma>1$ or $\alpha=\beta=1,\gamma=0$.*

Proof. We first show that this is a necessary condition for connectivity.
Case 1. If $\beta+\gamma<1$, the expected degree of vertex $\vec{0}$ is $(\beta+\gamma)^k=o(1)$, with high probability vertex $\vec{0}$ is isolated and the graph is thus disconnected.
Case 2. If $\beta+\gamma=1$ but $\beta<1$, we again prove that with constant probability vertex $\vec{0}$ is isolated:

$$\begin{aligned} Pr[\vec{0} \text{ has no edge}] &= \prod_v (1-P[\vec{0},v]) = \prod_{w=0}^{k}(1-\beta^w\gamma^{k-w})^{\binom{k}{w}} \geq \prod_{w=0}^{k} e^{-2\binom{k}{w}\beta^w\gamma^{k-w}} \\ &= e^{-2\sum_{w=0}^{k}\binom{k}{w}\beta^w\gamma^{k-w}} = e^{-2(\beta+\gamma)^k} = e^{-2} \end{aligned}$$

Now we prove it is also a sufficient condition. When $\alpha=\beta=1,\gamma=0$, the graph embeds a deterministic star centered at vertex $\vec{1}$, and is hence connected. To prove $\beta+\gamma>1$ implies connectivity, we only need to show the min-cut has size at least $c\ln n$ and apply Theorem 1. The expected degree of vertex $\vec{0}$ excluding self-loop is $(\beta+\gamma)^k-\gamma^k>2ck=2c\ln n$ given that β and γ are constants independent on k satisfying $\beta+\gamma>1$, therefore the cut $(\{\vec{0}\},V\setminus\{\vec{0}\})$ has size at least $2c\ln n$. Remove $\vec{0}$ and consider any cut $(S,V\setminus S)$ of the remaining graph, at least one side of the cut gets at least half of the expected degree of vertex $\vec{0}$; without loss of generality assume it is S i.e. $P(\vec{0},S)>c\ln n$. Take any node u in $V\setminus S$, by Lemma 2, $P(u,S)\geq P(\vec{0},S)>c\ln n$. Therefore the cut size $P(S,V\setminus S)\geq P(u,S)>c\ln n$. □

2.3 Giant Components

Lemma 5. *Let H denote the set of vertices with weight at least $k/2$, then for any vertex u, $P(u,H)\geq P(u,V)/4$.*

Proof. Given u, let l be the weight of u. For a vertex v let $i(v)$ be the number of bits where $u_b=v_b=1$, and let $j(v)$ be the number of bits where $u_b=0,v_b=1$. we partition the vertices in $V\setminus H$ into 3 subsets: $S_1=\{v: i(v)\geq l/2, j(v)<(k-l)/2\}, S_2=\{v: i(v)<l/2, j(v)\geq (k-l)/2\}, S_3=\{v: i(v)<l/2, j(v)<(k-l)/2\}$.

First consider S_1. For a vertex $v\in S_1$, we flip the bits of v where the corresponding bits of u is 0 to get v'. Then $i(v')=i(v)$ and $j(v')\geq(k-l)/2>j(v)$. It is easy to check that $P[u,v']\geq P[u,v]$, $v'\in H$, and different $v\in S_1$ maps to different v'. Therefore $P(u,H)\geq P(u,S_1)$.

Similarly we can prove $P(u,H)\geq P(u,S_2)$ by flipping the bits corresponding to 1s in u, and $P(u,H)\geq P(u,S_3)$ by flipping all the bits. Adding up the three subsets, we get $P(u,V\setminus H)\leq 3P(u,H)$. Thus, $P(u,H)\geq P(u,V)/4$. □

Theorem 6. *The necessary and sufficient condition for Kronecker graphs to have a giant component of size $\Theta(n)$ with high probability is $(\alpha+\beta)(\beta+\gamma)>1$, or $(\alpha+\beta)(\beta+\gamma)=1$ and $\alpha+\beta>\beta+\gamma$.*

Proof. When $(\alpha+\beta)(\beta+\gamma)<1$, we prove that the expected number of non-isolated nodes are $o(n)$. Let $(\alpha+\beta)(\beta+\gamma)=1-\epsilon$. Consider vertices with weight at least $k/2+k^{2/3}$, by Chernoff bound the fraction of such vertices is at most $\exp(-ck^{4/3}/k)=\exp(-ck^{1/3})=o(1)$, therefore the number of non-isolated vertices in this category is $o(n)$; on the other hand, for a vertex with weight less than $k/2+k^{2/3}$, by Lemma 3 its expected degree is at most

$$(\alpha+\beta)^{k/2+k^{2/3}}(\beta+\gamma)^{k/2-k^{2/3}} = (1-\epsilon)^{k/2}(\frac{\alpha+\beta}{\beta+\gamma})^{k^{2/3}} = n^{-\epsilon'}c^{o(\log n)} = o(1)$$

Therefore overall there are $o(n)$ non-isolated vertices.

When $\alpha+\beta=\beta+\gamma=1$, i.e. $\alpha=\beta=\gamma=1/2$, the Kronecker graph is equivalent to $G(n,1/n)$, which has no giant component of size $\Theta(n)$ [4].

When $(\alpha+\beta)(\beta+\gamma)>1$, we prove that the subgraph induced by $H=\{v: weight(v)\geq k/2\}$ is connected with high probability, hence forms a giant connected component of size at least $n/2$. Again we prove that the min-cut size of H is $c\ln n$ and apply Theorem 1. For any vertex u in H, its expected degree is at least $((\alpha+\beta)(\beta+\gamma))^{k/2}=\omega(\ln n)$; by Lemma 5 $P(u,H)\geq P(u,V)/4>2c\ln n$. Now given any cut $(S,H\setminus S)$ of H, we prove $P(S,H\setminus S)>c\ln n$. Without loss of generality assume vertex $\vec{1}$ is in S. For any vertex $u\in H$, either $P(u,S)$ or $P(u,H\setminus S)$ is at least $c\ln n$. If $\exists u$ such that $P(u,H\setminus S)>c\ln n$, then since $u\leq\vec{1}$, by Lemma 2 $P(S,H\setminus S)\geq P(\vec{1},H\setminus S)\geq P(u,H\setminus S)>c\ln n$; otherwise $\forall u\in H, P(u,S)>c\ln n$, since at least one of the vertex is in $H\setminus S$, $P(S,H\setminus S)>c\ln n$.

Finally, when $(\alpha+\beta)(\beta+\gamma)=1$ and $\alpha+\beta>\beta+\gamma$, let $H_1=\{v: weight(v)\geq k/2+k^{1/6}\}$, and we will prove that the subgraph induced by H_1 is connected with high probability by proving its min-cut size is at least $c\ln n$ (Claim 4), and that $|H_1|=\Theta(n)$ (Claim 5), therefore with high probability H_1 forms a giant connected component of size $\Theta(n)$. The proofs of these claims are presented in Appendix B. $\square$

3 Diameter

We analyze the diameter of a Kronecker graph under the condition that the graph is connected with high probability. When $\alpha=\beta=1,\gamma=0$, every vertex links to $\vec{1}$ so the graph has diameter 2; below we analyze the case where $\beta+\gamma>1$. We will use the following result about the diameter of $G(n,p)$, which has been extensively studied in for example [6, 2, 3].

Theorem 7. *[6, 2] If $(pn)^{d-1}/n\to 0$ and $(pn)^d/n\to\infty$ for a fixed integer d, then $G(n,p)$ has diameter d with probability approaching 1 as n goes to infinity.*

Theorem 8. *If $\beta+\gamma>1$, the diameters of Kronecker graphs are constant with high probability.*

Proof. Let S be the subset of vertices with weight at least $\frac{\beta}{\beta+\gamma}k$. We will prove that the subgraph induced by S has a constant diameter, and any other vertex directly connects to S with high probability.

Claim 1. *With high probability, any vertex u has a neighbor in S.*

Proof of Claim 1. We compute the expected degree of u to S:

$$P(u,S)\geq\sum_{j\geq\frac{\beta}{\beta+\gamma}k}\binom{k}{j}\beta^j\gamma^{k-j}=(\beta+\gamma)^k\sum_{j\geq\frac{\beta}{\beta+\gamma}k}\binom{k}{j}(\frac{\beta}{\beta+\gamma})^j(\frac{\gamma}{\beta+\gamma})^{k-j}$$

The summation is exactly the probability of getting at least $\frac{\beta}{\beta+\gamma}k$ HEADs in k coin flips where the probability of getting HEAD in one trial is $\frac{\beta}{\beta+\gamma}$, so this probability is at least a constant. Therefore $P(u,S)\geq(\beta+\gamma)^k/2>c\ln n$ for any u; by Chernoff bound any u has a neighbor in S with high probability. $\square$

Claim 2. $|S|\cdot\min_{u,v\in S}P[u,v]\geq(\beta+\gamma)^k$.

Proof of Claim 2. We have

$$\min_{u,v\in S} P[u,v] \geq \beta^{\frac{\beta}{\beta+\gamma}k}\gamma^{\frac{\gamma}{\beta+\gamma}k}$$

and

$$|S| \geq \binom{k}{\frac{\beta}{\beta+\gamma}k} \approx \frac{(\frac{k}{e})^k}{(\frac{\beta k}{(\beta+\gamma)e})^{\frac{\beta}{\beta+\gamma}k}(\frac{\gamma k}{(\beta+\gamma)e})^{\frac{\gamma}{\beta+\gamma}k}} = \frac{(\beta+\gamma)^k}{\beta^{\frac{\beta}{\beta+\gamma}k}\gamma^{\frac{\gamma}{\beta+\gamma}k}}$$

Therefore $|S| \cdot \min_{u,v\in S} P[u,v] \geq (\beta+\gamma)^k$. □

Given Claim 2, it follows easily that the diameter of the subgraph induced by S is constant: let $\beta + \gamma = 1 + \epsilon$ where ϵ is a constant, the diameter of $G(|S|, (\beta+\gamma)^k/|S|)$ is at most $d = 1/\epsilon$ by Theorem 7; since by increasing the edge probabilities of $G(n, P)$ the diameter cannot increase, the diameter of the subgraph of the Kronecker graph induced by S is no larger than that of $G(|S|, (\beta+\gamma)^k/|S|)$. Therefore, by Claim 1, for every two vertices u and v in the Kronecker graph, there is a path of length at most $2 + 1/\epsilon$ between them. □

4 Searchability

In Section 3 we showed that the diameter of a Kronecker graph is constant with high probability, given that the graph is connected. However it is yet a question whether a short path can be found by a decentralized algorithm where each individual only has access to local information. We use a similar definition as used by Kleinberg [7, 8, 9].

Definition 3. *In a* decentralized routing algorithm for $G(n, P)$*, the message is passed sequentially from a current message holder to one of its neighbors until reach the destination t, using only local information. In particular, the message holder u at a given step has knowledge of:*
(i) the probability matrix P;
(ii) the label of destination t;
(iii) edges incident to all visited vertices.
A $G(n, P)$ graph is d-searchable *if there exists a decentralized routing algorithm such that for any destination t, source s, with high probability the algorithm can find an s-t path no longer than d.*

We first give a monotonicity result on general random graphs $G(n, P)$, then use it to prove Kronecker graphs with $\alpha < 1$ is not poly-logarithmic searchable. It is possible to directly prove our result on Kronecker graphs, but we believe the monotonicity theorem might be of independent interests. More results on searchability in $G(n, P)$ using deterministic memoryless algorithms can be found in [1]. The proof of Theorem 9 is left to Appendix B.

Theorem 9. *If $G(n, P)$ is d-searchable, and $P \leq P'$ $(\forall i, j, P[i, j] \leq P'[i, j])$, then $G(n, P')$ is d-searchable.*

Theorem 10. *Kronecker graphs are not $n^{(1-\alpha)\log_2 e}$-searchable.*

Proof. Let P be the probability matrix of the Kronecker graph, and P' be the matrix where each element is $p = n^{-(1-\alpha)\log_2 e}$. We have $P \leq P'$ because $max_{i,j} P[i, j] \leq \alpha^k \leq e^{-(1-\alpha)k} = n^{-(1-\alpha)\log_2 e} = p$. If the Kronecker graph is $n^{(1-\alpha)\log_2 e}$-searchable, then by Theorem 9 $G(n, p)$ where $p = n^{-(1-\alpha)\log_2 e}$ is also $n^{(1-\alpha)\log_2 e}$-searchable. However, $G(n, p)$ is not $\frac{1}{p}$-searchable. This is because given any decentralized algorithm, whenever we first visit a vertex u, independent on the routing history, the probability that u has a direct link to t is no more than p, hence the routing path is longer than the geometry distribution with parameter p, i.e. with constant probability the algorithm cannot reach t in $1/p$ steps. □

References

[1] E. Arcaute, N. Chen, R. Kumar, D. Liben-Nowell, M. Mahdian, H. Nazerzadeh, and Y. Xu. Searchability in random graphs. *Proceedings of the 5th Workshop on Algorithms and Models for the Web-Graph*, 2007.

[2] B. Bollobás. The diameter of random graphs. *IEEE Trans. Inform. Theory, 36(2):285-288*, 1990.

[3] F. Chung and L. Lu. The diameter of random sparse graphs. *Advances in Applied Math. 26, 257–279*, 2001.

[4] P. Erdős and A. Rényi. On random graphs I. *Publications Mathematics, Debrecen*, 6:290–297, 1959.

[5] D. Karger and C. Stein. A new approach to the minimum cut problem. *Journal of the ACM 43(4)*, 1996.

[6] V. Klee and D. Larmann. Diameters of random graphs. *Canad. J. Math. 33:618-640*, 1981.

[7] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. *Proc. 32nd ACM Symposium on Theory of Computing*, 2000.

[8] J. Kleinberg. Small-world phenomena and the dynamics of information. *Advances in Neural Information Processing Systems (NIPS)*, 2001.

[9] J. Kleinberg. Complex networks and decentralized search algorithms. *Proc. the International Congress of Mathematicians (ICM)*, 2006.

[10] J. Leskovec, D. Chakrabarti, J. Kleinberg, and C. Faloutsos. Realistic, mathematically tractable graph generation and evolution, using kronecker multiplication. *PKDD*, 2005.

[11] J. Leskovec and C. Faloutsos. Scalable modeling of real graphs using kronecker multiplication. *International Conference on Machine Learning (ICML)*, 2007.

[12] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data 1(1)*, 2007.

A Other settings of parameters

We study connectivity of Kronecker graphs with 2×2 initiator matrices without the constraint $\alpha \geq \beta \geq \gamma$. We prove that if $\beta + max(\alpha, \gamma) > 1$, then Kronecker graphs are connected with high probability, for arbitrary $\alpha, \beta, \gamma \in [0, 1]$.

Theorem 11. *If $\beta + max(\alpha, \gamma) > 1$, then Kronecker graphs are connected with high probability.*

Proof. We first prove for the case where $\alpha \geq \gamma \geq \beta$. We may assume $\alpha = \gamma$. Let $d(u, v)$ be the Hamming-distance between the labels of u and v. Then $P[u, v] = \gamma^k (\frac{\beta}{\gamma})^{d(u,v)}$.

Claim 3. *In λk-th power of hypercube (i.e. a vertex u is connected to all vertices v with $d(u, v) \leq \lambda k$), the min-cut size is at least $\binom{k}{\lambda k-1}/2$.*

Proof of Claim 3. Given a cut $(S, \bar{S})$, take any vertex $u \in S$. For any vertex u' with $d(u, u') = 1$, the number of common neighbors of u and u' is at least $\binom{k}{\lambda k-1}$, because all vertices within distance $\lambda k - 1$ to u are within distance λk to u' and are thus their common neighbors. If the cut size is less than $\binom{k}{\lambda k-1}/2$, then at least half of those common neighbors must be in S, and u' must also be in S. Apply the same argument iteratively and all vertices will end up being in S. □

Now given a cut $(S, \bar{S})$ in the Kronecker graph, for any edge of this cut in λk-th power of hypercube, the Hamming-distance of the two endpoints is at most λk, and the corresponding edge presents in the Kronecker graph with probability at least $\gamma^k(\frac{\beta}{\gamma})^{\lambda k}$, therefore

$$P(S, \bar{S}) \geq \frac{1}{2}\binom{k}{\lambda k - 1}\gamma^k(\frac{\beta}{\gamma})^{\lambda k} \geq \frac{1}{2k}((\frac{1}{\lambda})^{\lambda}(\frac{1}{1-\lambda})^{1-\lambda}\gamma(\frac{\beta}{\gamma})^{\lambda})^k$$

Let $\lambda = \frac{\beta}{\beta+\gamma}$, $P(S, \bar{S}) \geq \frac{1}{2k}(\beta+\gamma)^k = \omega(\ln n)$. According to Theorem 1, the graph is connected with high probability.

For the case where $\beta \geq max(\alpha, \gamma)$, the proof is similar as above. Instead of λk-th power of hypercube, we can prove that in the graph where a vertex u is connected to all vertices v with $d(u, v) \geq \lambda k$, the min-cut size is at least $\binom{k}{\lambda k-1}/2$, and use it to bound the cut size in Kronecker graphs. □

When $\beta + max(\alpha, \gamma) < 1$ or $\beta + max(\alpha, \gamma) = 1$ but all of them are strictly less than 1, then the graph is not connected with at least constant probability because vertex $\vec{0}$ is isolated with probability at least e^{-2} as the proof in Theorem 4.

B Proofs

Proof of Lemma 3. For any vertex v, let i be the number of bits where $u_b = v_b = 1$, and let j be the number of bits where $u_b = 1, v_b = 0$, then $P[u, v] = \alpha^i\beta^{j+l-i}\gamma^{k-l-j}$. Summing $P[u, v]$ over all v, the expected degree of u is

$$\sum_{i=0}^{l}\sum_{j=0}^{k-l}\binom{l}{i}\binom{k-l}{j}\alpha^i\beta^{j+l-i}\gamma^{k-l-j} = \sum_{i=0}^{l}\binom{l}{i}\alpha^i\beta^{l-i}\sum_{j=0}^{k-l}\binom{k-l}{j}\beta^j\gamma^{k-l-j} = (\alpha+\beta)^l(\beta+\gamma)^{k-l}$$

□

Claim 4. *For any cut $(S, H_1 \setminus S)$ of H_1, $P(S, H_1 \setminus S) > c \ln n$.*

Proof of Claim 4. First, for any $u \in H_1$,

$$P(u, V) \geq (\alpha+\beta)^{k/2+k^{1/6}}(\beta+\gamma)^{k/2-k^{1/6}} = ((\alpha+\beta)/(\beta+\gamma))^{k^{1/6}} = \omega(\ln n).$$

By Lemma 5, $P(u, H) = \omega(\ln n)$. We will prove $P(u, H_1) \geq P(u, H \setminus H_1)/2$, and it follows that $P(u, H_1) = \omega(\ln n)$. Then we can apply the same argument as in case $(\alpha+\beta)(\beta+\gamma) > 1$ and prove that for any cut $(S, H_1 \setminus S)$ of H_1, $P(S, H_1 \setminus S) > c \ln n$: assume vertex $\vec{1}$ is in S; for any vertex $u \in H_1$, either $P(u, S)$ or $P(u, H_1 \setminus S)$ is at least $c \ln n$; if $\exists u$ such that $P(u, H_1 \setminus S) > c \ln n$, then $P(S, H_1 \setminus S) \geq P(\vec{1}, H_1 \setminus S) \geq P(u, H_1 \setminus S) > c \ln n$; otherwise $\forall u \in H_1, P(u, S) > c \ln n$, since at least one vertex is in $H_1 \setminus S$, we have $P(S, H_1 \setminus S) > c \ln n$.

It remains to prove $P(u, H_1) \geq P(u, H \setminus H_1)/2$. We will map each vertex $v \in H \setminus H_1$ to a vertex $f(v) = v' \in H_1$ such that $v \leq v'$ (and hence $P[u, v] \leq P[u, v']$), and each vertex in H_1 is mapped to at most twice. Once we have such a mapping, then $P(u, H \setminus H1) = \sum_{v \in H \setminus H_1} P[u, v] \leq \sum_{v \in H \setminus H_1} P[u, f(v)] \leq \sum_{v' \in H_1} 2P[u, v'] = 2P(u, H_1)$. The mapping is as follows: for each i in $[k/2, k/2 + k^{1/6})$, construct a bipartite graph G_i where the left nodes L_i are vertices with weight i, and make two copies of all vertices with weight $i + k^{1/6}$ to form the right nodes R_i, and add an edge if a right node dominates a left node. It is easy to see that the union of L_is forms exactly $H \setminus H_1$, while all right nodes are in H_1 and each node appears at most twice. The bipartite graph G_i has a maximum matching of size $|L_i|$, because all left (right) nodes have the same degree by symmetry and $|L_i| < |R_i|$ (proved below). We take any such maximum matching to define the mapping and it satisfies that $v \le f(v)$, $f(v) \in H_1$ and each $v' \in H_1$ is mapped to at most twice. Finally we prove $|L_i| < |R_i|$ for $k/2 \le i < k/2 + k^{1/6}$:

$$\begin{aligned}\frac{|R_i|}{|L_i|} &= \frac{2\binom{k}{i+k^{1/6}}}{\binom{k}{i}} \ge \frac{2\binom{k}{k/2+2k^{1/6}}}{\binom{k}{k/2}} = \frac{2\frac{k}{2}\dots(\frac{k}{2}-2k^{1/6}+1)}{(\frac{k}{2}+2k^{1/6})\dots(\frac{k}{2}+1)} \\ &\ge 2(\frac{\frac{k}{2}-2k^{1/6}}{\frac{k}{2}})^{2k^{1/6}} \ge 2e^{-ck^{-5/6}*k^{1/6}} = 2e^{-o(1)} = 2(1-o(1)) > 1\end{aligned}$$

□

Claim 5. $|H_1| = \Theta(n)$.

Proof of Claim 5. We count the number of vertices with weight $k/2 + i$:

$$\binom{k}{k/2+i} \le \binom{k}{k/2} = \Theta(\frac{\sqrt{2\pi k}(k/e)^k}{(\sqrt{2\pi k/2}(k/2e)^{k/2})^2}) = \Theta(\frac{2^k}{\sqrt{k}})$$

Therefore the size of $H \setminus H_1$ is at most $\sum_{i=0}^{k^{1/6}} \binom{k}{k/2+i} \le k^{1/6} * 2^k/\sqrt{k} = o(n)$. It is easy to see $|H| > n/2$, thus $|H_1| > n/2 - o(n)$. □

Proof of Theorem 9. Given $G(n, P')$ we simulate $G(n, P)$ by ignoring some edges. Given a realization G of $G(n, P')$, we keep an edge (i, j) in G with probability $P[i, j]/P'[i, j]$, and delete the edge otherwise; do so for each edge independently. We claim that random graphs generated by the above process is equivalent to $G(n, P)$: the probability that edge (i, j) presents is $P'[i, j] * (P[i, j]/P'[i, j]) = P[i, j]$, independent on other edges. Now we have a $G(n, P)$ graph, we use its decentralized routing algorithm, which will find a path with length at most d with high probability for any s and t.

Note that we cannot process all edges in the beginning, because there is no global data structure to remember which edges are deleted. Instead we will decide whether to delete an edge the first time we visit one of its endpoints and this information will be available to all vertices visited later. □