Linked Decompositions of Networks and the Power of Choice in Polya Urns

Henry Lin* Christos Amanatidis† Martha Sideri‡ Richard M. Karp§
Christos H. Papadimitriou¶

December 8, 2007

Abstract

A *linked decomposition* of a graph with n nodes is a set of subgraphs covering the n nodes such that all pairs of subgraphs intersect; we seek linked decompositions such that all subgraphs have about $\sqrt{n}$ vertices, logarithmic diameter, and each vertex of the graph belongs to either one or two subgraphs. A linked decomposition enables many control and management functions to be implemented locally, such as resource sharing, maintenance of distributed directory structures, deadlock-free routing, failure recovery and load balancing, without requiring any node to maintain information about the state of the network outside the subgraphs to which it belongs. Linked decompositions also enable efficient routing schemes with small routing tables, which we describe in Section 5. Our main contribution is to show that "Internet-like graphs" (e.g. the preferential attachment model proposed by Barabasi et al. [11] and other similar models) have linked decompositions with the parameters described above with high probability; moreover, our experiments show that the Internet topology itself can be so decomposed. Our proof proceeds by analyzing a novel process, which we call *Polya urns with the power of choice*, which may be of great independent interest. In this new process, we start with n nonempty bins containing $O(n)$ balls total, and each arriving ball is placed in the least loaded of m bins, drawn independently at random *with probability proportional to load*. Our analysis shows that in our new process, with high probability the bin loads become roughly balanced some time before $O(n^{2+\epsilon})$ further balls have arrived and stay roughly balanced, regardless of how the initial $O(n)$ balls were distributed, where $\epsilon > 0$ can be arbitrarily small, provided m is large enough.

*Computer Science Division, University of California at Berkeley. Supported by NSF grant CCF-0635319 and NSF grant CCF-0515259. E-mail: henrylin@eecs.berkeley.edu

†Univ. of Economics and Business, Athens, Greece. Supported by a project co-financed by E.U. European Social Fund and the Greek Ministry of Development - GSRT. E-mail: amanatidis_c@yahoo.gr

‡Univ. of Economics and Business, Athens, Greece. Supported by a project co-financed by E.U. European Social Fund and the Greek Ministry of Development - GSRT. E-mail: sideri@aueb.gr

§Computer Science Division, University of California at Berkeley. Supported by NSF grant CCF-0515259. E-mail: karp@cs.berkeley.edu

¶Computer Science Division, University of California at Berkeley. Supported by NSF grant CCF-0635319 and a gift from Yahoo! Research. E-mail: christos@cs.berkeley.edu

1 Introduction

Throwing balls into bins uniformly at random is well-known to produce relatively well-balanced occupancies with high probability. In contrast, when each new ball is added to a bin selected with probability proportional to its current occupancy, we get the much-studied model of *Polya urns,* which is known to produce large imbalances: with n urns, the largest urn is expected to have $\log n$ times more balls, and the smallest urn n times fewer balls, than the average one. It is a celebrated result in our field that if, while throwing balls into bins, we choose two random bins and add a ball only to the smallest one, then this *power of two choices* makes the already small imbalances significantly smaller [10, 16, 20]. But what about utilizing the power of choice in the Polya urns model? Does selecting $m \geq 2$ bins according to the Polya urn distribution, and adding a ball to the least loaded bin, result in balanced loads?

In this paper, we show that *the power of choice in the Polya urn model does balance bin loads.* In particular, we show (Theorem 3) that if n nonempty bins start with $O(n)$ balls total, then throwing $O(n^{2+\epsilon})$ more balls according to our new Polya urns process with the power of choice, balances all bins within a multiplicative factor of $(1+\epsilon)$ with high probability, where $\epsilon > 0$ can be made arbitrarily small, provided we make the number of choices m large enough. The exact dependence of ϵ on m is a bit messy (see Section 3), but essentially ϵ decreases exponentially as m increases. Moreover, once the bins become roughly balanced, they stay roughly balanced as more balls are added. As m increases, our result becomes essentially tight, since $\Omega(n^2)$ balls are required to balance an initial distribution that starts with $\Omega(n)$ balls in one bin, and 1 ball in every other bin. We can also show the bin loads balance when $m = 2$ choices are used and poly(n) balls are thrown, but it remains an interesting open question whether or not the loads also balance when $O(n^{2+\epsilon})$ new balls are thrown with $m = 2$ choices.

The motivation for this investigation comes from a graph-theoretic problem related to the Internet. Let G be a graph, and a, b, c, d be integer parameters. We say that a graph G has an (a, b, c, d)*-linked decomposition* if there are c connected subgraphs of G, which we call *components*, such that

1. each component has between $\frac{a}{4}$ and a nodes
2. each node belongs to at least one and at most b components
3. the diameter of each component is at most d
4. most importantly, *any two components have a node in common.*

Our notion of linked decomposition is related to the notion of *sparse covers* defined in [9], which decomposes the graph into connected subgraphs satisfying properties (2) and (3). Other related concepts of network decomposition are introduced in [8] and [21].

For our purposes, we would like $a, c \approx \sqrt{n}$, $b = 2$, and $d = O(\log n)$, where n is the number of nodes in G. It is not hard to see that a linked decomposition, if it exists for graph G, would yield a routing algorithm with unstructured node names with $O(d)$ delay (worst-case number of hops) and $O(d)$ traffic (total messages sent per packet), and $\tilde{O}(ab + n/a)$ storage per node (see Section 5 for details), which is about $\sqrt{n}$ for the parameters discussed. Although there is an existing routing scheme [5] known to achieve $\tilde{O}(\sqrt{n})$ size routing tables while maintaining short routing paths (stretch 3), our new routing scheme is a bit simpler, and does not require the use of landmark nodes, which may provide benefits in terms of robustness and maintaining balanced congestion (Section 5 for more details). In particular, our routing scheme decomposes the network into connected subnetworks of roughly balanced size, which can be administered relatively independently, with very little harm to communication. Moreover, using a more relaxed form of linked decomposition and structured node names, we can achieve about $O(n^{1/3})$ storage per node and $\tilde{O}(\log n)$ delay/messages sent as well. For more details regarding our routing schemes and their implications for network routing, see Section 5.

But why should we expect that we can decompose networks in this way? Rather surprisingly, our experiments (in Appendix L) indicate that *the existing Internet graph does have linked decompositions with these approximate parameters* (both at the router and autonomous system level).

A linked decomposition with parameters $a, c \approx \sqrt{n}$, $b = 2$, and $d = O(\log n)$ is a very demanding requirement, and hence an unlikely property of graphs; on the other hand the Internet seems to have it. We suspect that this is not a coincidence, and that "any Internet-like graph" is very likely to have a linked decomposition with these parameters. Towards this goal, we turn to the simplest and perhaps most influential model of Internet-like graphs, namely the *preferential attachment model* $\mathcal{PA}(m)$ proposed by [11] in 1999, and since then studied extensively (see [12] for a survey). In this model, nodes arrive one after the other, and when a node arrives, it is connected via m edges to previously arrived nodes. In particular, for each new node t, we pick m previous nodes, i.i.d. at random and with replacement, with probability proportional to the degree of the node at the current time, and create an edge from node t to each of the m nodes picked with parallel edges possible.

Our main result is Theorem 1 which states that *for any $\epsilon > 0$, there is a $m \geq 2$, such that $\mathcal{PA}(m)$ produces a graph which has a linked decomposition with parameters $a = \Theta(n^{1/2+\epsilon})$, $b = 2$, $c = \Theta(n^{1/2-\epsilon})$, and $d = O(\log n)$, with probability $1 - n^{-\epsilon}$.*

The proof of our main result is based on the Polya urns insight from the beginning of the Introduction. To find our decomposition, the idea is to follow the preferential attachment process, and assign the first $c = \Theta(n^{1/2-\epsilon})$ nodes to their own component. Then as the next $\frac{n}{2} - c$ nodes arrive, we look at the nodes/components to which the m edges of arriving node point, and assign the new node to the component whose total degree (the sum of the degrees of the nodes in the component) is lowest. Now, provided the components are roughly balanced in terms of total degree after these first $\frac{n}{2}$ nodes have arrived, we can assign the last $\frac{n}{2}$ nodes to two components each in a simple way such that a coupon collect argument can be used to show that each pair of components intersect at some node with high probability (think of the intersections as coupons being collected). The coupon collector argument shows that property (4) of linked decomposition holds with high probability, and as we show in Section 2, it is not hard to prove that the other three properties of linked decomposition hold with high probability as well.

The main challenge in proving Theorem 1 is to show that the components obtain roughly balanced total degree after $\frac{n}{2}$ nodes arrive, Theorem 2, which is needed for property (4) of linked decomposition to hold. To prove that the degrees of the components become roughly balanced when running $\mathcal{PA}(m)$, note that if we think of the total degrees of each component as the occupancy of a Polya urn, then our random process is very much like the Polya urns process with the power of choice, defined previously, with c bins and m choices. Although Polya urns with the power of choice is not quite the same process as the one we would like to analyze, the analysis used to prove that Polya urns with the power of choice produces roughly balanced bin loads when $O(c^{2+\epsilon})$ balls are thrown (Theorem 3), can be modified to prove that the degrees of the components balance when

$O(c^{2+\epsilon})$ new nodes arrive, Theorem 2. Here again, in order to make ϵ arbitrarily small, we need to make m sufficiently large. We expect that the dependence of ϵ on m to be similar for Theorem 2, as the dependence given in Theorem 3 (see Section 3), but some details remain to be checked. By proving Theorem 2, we complete the last step needed to prove that a linked decomposition with the required parameters exists with high probability.

One negative aspect of the proof above is the fact that we need to increase m in order to make ϵ arbitrarily small. However, if we relax the definition of linked decomposition so that each component may contain at most one special node violating property (2) of the linked decomposition — that is, some nodes may be allowed to belong to more than two components — then we can find a linked decomposition even when $m = 2$ (Theorem 4). Moreover, it is easy to see that routing is not harmed by this exception, and we obtain a slightly lower value for $a = \Theta(\sqrt{n \log n})$. The proof proceeds by using a dynamic programming algorithm to decompose the network formed by the first $\Theta(n/\log n)$ nodes into $c = \Theta(\sqrt{n/\log n})$ components of approximately equal size. Once we have c balanced components, it is not very hard to use a coupon collector argument to show that the remaining $\Theta(c^2 \log c)$ nodes are enough to link the components together as required by property (4) of linked decomposition, and one can show the other three properties of linked decomposition hold as well.

2 Linked Decompositions in $\mathcal{PA}(m)$ Graphs

The *preferential attachment model* of random graphs $\mathcal{PA}(m)$ [11, 12] creates a sequence of random graphs $G_1, G_2, \ldots, G_t, \ldots$ with multiple edges, where G_t has t nodes. G_1 is a single node with no edges, and G_2 is a graph with m edges between two nodes. To generate G_{t+1} from G_t for $t \geq 2$, a new node $t+1$ is added, and then m nodes are selected i.i.d. at random with replacement, from among nodes $1, \ldots, t$, where each node $i \leq t$ is selected with probability $\frac{d_t(i)}{2mt}$ and $d_t(i)$ is the degree of node i in G_t. Then edges are added in G_{t+1} from node $t+1$ to the m selected nodes. This is a very influential graph-theoretic model in the context of the Internet (even though it is understood that it does not satisfy all known properties of the Internet graph).

By the phrase "with high probability" (whp) in this paper we shall mean with probability greater than or equal to $1 - n^{-\alpha}$, for some $\alpha > 0$. In most cases, α can be made arbitrarily large by deteriorating the other parameters (e.g. by increasing m in Theorem 1). Following the definition of linked decomposition described in the Introduction, we can prove one of our main theorems:

Theorem 1 *For any $\epsilon > 0$, there exists a $m \geq 2$, such that a graph G_n generated by the $\mathcal{PA}(m)$ has a linked decomposition (whp) with parameters $a = \Theta(n^{1/2+\epsilon})$, $b = 2$, $c = \Theta(n^{1/2-\epsilon})$, and $d = O(\log n)$.*

Proof. To find a linked decomposition in a $\mathcal{PA}(m)$ graph with n nodes, we follow the preferential attachment process, and as nodes arrive, we use a very simple procedure to assign each node t to one or two components.

(1) For node $t \in \{1, \ldots, n^{1/2-\epsilon}\}$, we assign node t to its own component.

(2) For node $t \in \{n^{1/2-\epsilon} + 1, \ldots, \frac{n}{2}\}$, we look at the nodes/components to which the m edges of node t point, and assign node t to the component whose total degree (i.e. the sum of the degrees of its nodes) is lowest so far.

(3) For node $t \in \{\frac{n}{2}+1, \ldots, n\}$, we look at where the first two edges of node t point, say to nodes u and v, and assign node t to a component to which u belongs, and a component to which v belongs. (If u or v belongs to two components, then we assign t to a random component of u or v).

It is easy to see that we have decomposed the graph into connected components, and that we have satisfied property 2 of linked decompositions with $b = 2$. Furthermore, since each node t has a constant probability (dependent on m) of joining the component of a node $v \leq \frac{t}{2}$, it is easy to prove that each component has diameter $O(\log n)$ (whp).

To prove the final two properties of our linked decomposition, the key step is to show that at the end of step (2), the total degree of each component is roughly balanced (whp). In the next section, we define a Polya urns process, which models the degrees of the components, and analyze it to prove one of our main theorems, Theorem 2, which states that at the end of step (2), the total degrees of any two components differ by at most a multiplicative factor of $(1 + \epsilon)$ (whp).

Theorem 2 can be used to prove the final two properties hold (whp), because if we can show that the total degree is roughly the same for each component at the end of step (2), then it is not hard to apply Azuma's inequality (see the analysis in Appendix C or Appendix I) to show that the total degree of each component remains roughly balanced within a constant (whp) throughout step (3). Furthermore, if the components remain balanced within a multiplicative constant throughout step (3), then at the end of step (3), each component must have total degree $\Theta(n^{1/2+\epsilon})$, and each component contains at most $O(n^{1/2+\epsilon})$ nodes. Moreover, it is not hard to see that each component obtains $\Omega(n^{1/2+\epsilon})$ nodes (whp) in step (3), and thus property 1 holds (whp) with $a = \Theta(n^{1/2+\epsilon})$, assuming Theorem 2 is true.

Moreover, Theorem 2 can also be used to show that property 4 holds (whp). Recall that property 4 states that each pair of components must intersect at some node. If we view these $\Theta(n^{1-2\epsilon})$ pairs of intersections as coupons we wish to collect, then step (3) provides us with $\Theta(n)$ opportunities to collect $\Theta(n^{1-2\epsilon})$ coupons. Even though we are not collecting coupons uniformly at random, Theorem 2 (and Azuma's inequality) shows that the degrees of the components remain balanced throughout step (3) (whp), which implies that each coupon is still collected with probability at least $\Omega(n^{-(1-2\epsilon)})$. Thus, we can conclude that given $(\frac{n}{2})$ opportunities, our process collects all the coupons (intersections) within the required steps (whp). Therefore, provided that we can prove Theorem 2, all four properties of our linked decomposition hold with high probability. □

3 Polya Urns with the Power of Choice

To prove that the total degrees of the components become balanced, we analyze an equivalent random process $\mathcal{P}$ on n bins, where each bin represents a component and each bin's load represents the size of a component. (Here, n should be thought of as the number of components, called c in the definition of linked decomposition, and should not be confused with the number of nodes in our graph). Analogous to the preferential attachment process, our random process starts with $2mn$ balls distributed arbitrarily among n bins, such that each bin contains at least m balls each. Furthermore, it is known that with high probability our process starts with at most $O(n^{1/2+\epsilon})$ balls in each bin [15], where $\epsilon > 0$ can be arbitrarily small, a fact which will be useful later on. At each step of our random process $\mathcal{P}$, $2m$ new balls are thrown into the n bins, according to the following rule:

- Pick m bins i.i.d. at random, with replacement, with probability proportional to bin load
    - (I) Throw 1 ball into each of the m random bins picked
    - (II) Throw m more balls into the least loaded of the m random bins picked (breaking ties arbitrarily).

We can prove the following theorem about our random process $\mathcal{P}$, for any arbitrarily small $\epsilon > 0$, provided m is a sufficiently large constant:

Theorem 2 *Given the starting conditions described above, any time after $\Omega(n^{2+\epsilon})$ iterations of our process $\mathcal{P}$, the loads of any two bins differ by a multiplicative factor of at most $(1+\epsilon)$ (whp).*

Remark: Note that proving Theorem 2 is sufficient for proving that properties (1) and (4) of linked decomposition hold (whp), and completes the proof of Theorem 1.

Even though step (II) of our random process tends to balance bin loads, proving Theorem 2 is a bit challenging since step (I) tends to preserve load imbalances. For this reason, we study a simpler random process $\bar{\mathcal{P}}$ on n bins, which we call *Polya urns with the power of choice.* By proving $\bar{\mathcal{P}}$ balances bin loads (whp), we illustrate the main techniques needed to prove $\mathcal{P}$ produces balanced bin loads (Theorem 2). Our new process, Polya urns with the power of choice $\bar{\mathcal{P}}$, starts with n nonempty bins containing $N_0 = \hat{c}n$ balls total. At each step of our random process $\bar{\mathcal{P}}$, we throw one more ball into one of the n bins according to the following rule:

(a) Pick m bins i.i.d. at random, with replacement, with probability proportional to bin load

(b) Throw 1 ball into the least loaded of the m random bins picked

We can prove the following theorem about Polya urns with the power of choice, for any $\epsilon > 0$, provided m is a sufficiently large constant:

Theorem 3 *Given the starting conditions described above, any time after $\Omega(n^{2+\epsilon})$ balls have been thrown, the loads of any two bins differ by a multiplicative factor of at most $(1+\epsilon)$ (whp).*

To be more precise, we show that if each bin starts with fractional load at least $\frac{1}{\hat{c}n}$ for $\hat{c} > 1$, then all bins become balanced within $(1+\hat{\epsilon})$ (whp) some time before $O(n^{1/(1-(1-1/\hat{c})^{m-1})+1+\hat{\epsilon}})$ new balls arrive, and they stay balanced within $(1+\hat{\epsilon})$ (whp) at all later times, where $\hat{\epsilon} > 0$ can be arbitrarily small. Rewriting the above expression, note that the bin loads balance (whp) sometime before $O(n^{2+1/(\mathcal{C}^{m-1}-1)+\hat{\epsilon}})$ new balls have been thrown, where $\mathcal{C} = \frac{\hat{c}}{\hat{c}-1}$, and thus, for sufficiently large m or sufficiently small $\hat{c}$, we can conclude that the bin loads become balanced within $(1+\epsilon)$ in $O(n^{2+\epsilon})$ steps (whp) for any $\epsilon > 0$. Moreover, as we increase m, the parameter ϵ decreases exponentially fast.

Due to space limitations, we only provide a sketch of Theorem 3 in the two next subsections, and details are deferred to Appendix A to E. Additionally, we defer the full proof of Theorem 2 to Appendix F to K. The proof for Theorem 2 follows roughly the same steps as the proof for Theorem 3, although the details are slightly different, since our random process is slightly different.

3.1 An Alternate Random Process

To prove Theorem 3, we start by defining a new random process $\bar{\mathcal{P}}_0$ which is somewhat easier to analyze than the Polya urn process with the power of choice $\bar{\mathcal{P}}$. Our new random process $\bar{\mathcal{P}}_0$, defined below, also throws balls into bins one at a time, and for each ball thrown, m random bins are generated in the same manner as described for $\bar{\mathcal{P}}$. However, our new process $\bar{\mathcal{P}}_0$ sometimes ignores the power of m choices and throws a ball into a bin that is not the least loaded of the m bins.

Eventually, our goal will be to show that for any arbitrarily small $\epsilon > 0$, after $N_F = \Theta(n^{2+\epsilon})$ balls are thrown according to $\bar{\mathcal{P}}_0$, the number of balls in the bins are within a factor of $(1+\epsilon)$ from one another, with high probability. Note that the previous statement implies our main theorem by making the following observation: if $\vec{a}(t)$ represents the bin loads of the Polya urns process with the power of choice $\bar{\mathcal{P}}$ at time t and $\vec{b}(t)$ represents the bin loads of the new process $\bar{\mathcal{P}}_0$ (or any process that sometimes ignores the power of m choices) at time t, then there exists a coupling of the two processes such that $\vec{a}(t)$ always majorizes $\vec{b}(t)$. We omit the formal details of the coupling, which are not difficult, but note that in order to prove the theorem, the existence of the coupling implies that we only need to analyze our new random process $\bar{\mathcal{P}}_0$, instead of $\bar{\mathcal{P}}$.

Our new process $\bar{\mathcal{P}}_0$ throws balls into bins in two phases. In phase A, we throw $N_A = \Theta(N_0^{2+\epsilon_0}) = \Theta(n^{2+\epsilon_0})$ balls into bins, where N_0 is the number of balls at the start of the random process, and $\epsilon_0 > 0$ is an arbitrarily small constant. At any time in phase A, we classify the bins into two types of bins, which are used to determine which bin each ball is thrown into. A bin is considered to be a *low* bin if it contains less than $\frac{c_1}{n}$ fraction of the balls, and otherwise it is consider to be a *high* bin. Here, c_1 is a constant greater than 1 to be defined later, dependent on ϵ_0. Given these two definitions, we throw each new ball into a bin as follows:

- If at least one of the random bins generated for the new ball is a low bin, then we throw the new ball into the first low bin generated.
- If all the bins generated are high bins, then we throw the new ball into the first (high) bin generated.

In phase B, we throw balls into bins until the bins contain a total of $N_B = \Theta(N_A^{1+\epsilon_0})$ balls. In phase B, we define M be the $(\frac{n}{c_2})$th smallest bin at the end of phase A, where $c_2 > 1$ is a constant dependent on ϵ_0, and we will use bin M to classify the bins into three types of bins. A bin is considered to be a *middle* bin at the current time if contains same load as bin M, or it has obtained the same load as bin M at some prior time in phase B. A bin is *low* if it currently contains strictly less load then bin M, and it has never obtained the same load as bin M at any prior time in phase B. The remaining bins, those that have so far always contained more balls than bin M in phase B, are *high* bins. In Phase B, we throw each new ball into a bin as follows:

- We look at the random bins generated for the new ball, and we let $T \in \{\text{low, middle, high}\}$ be the lowest bin type generated.
- We throw a ball into the first bin generated of type T.

To complete the proof, we prove the following five claims about the two phases:

- At the end of phase A:

Claim 1: Each bin contains at least $\omega(\log n)$ balls with high probability.

Claim 2: Each bin contains at most $1.01(\frac{c_1}{n})$ fractional load with high probability.

- If each bin contains at least $\omega(\log n)$ balls and at most $1.01(\frac{c_1}{n})$ fractional load at the end of phase A, then with high probability:

Claim 3: Every high bin becomes a middle bin sometime during phase B.

Claim 4: Every low bin becomes a middle bin sometime during phase B.

Claim 5: If a bin becomes a middle bin (i.e. obtains the same load as bin M) sometime during phase B, then at the end of phase B, the load of that bin is at most $(1+\epsilon_1)$ times the load of bin M, and at least $(1-\epsilon_1)$ times the load of bin M, where $\epsilon_1 > 0$ can be arbitrarily small, provided n is sufficiently large.

Note that the five claims imply Theorem 3, and by making ϵ_1 and ϵ_0 arbitrarily small, we also make ϵ arbitrarily small. Proving the five claims requires applying a variety of probability theory techniques, but we only provide a rough proof sketch for each claim in the next subsection. The full details are in Appendix A to E.

3.2 A Proof Sketch for Each Claim

To prove Claim 1, we first note that a coupling can be defined such that bin loads arising from the $\bar{\mathcal{P}}_0$ process majorize the bin loads arising from the standard Polya process, which throws balls into bins with probability proportional to load, but does not utilize the power of choice. We can then utilize a known result, which lower bounds the load of the lowest bin arising from the standard Polya urn process [14]. This lower bound states that the lowest bin will contain $\omega(\log n)$ balls (whp) after N_A balls have been thrown according to the standard Polya urns process. With a coupling argument, we can then conclude that at the end of phase A of process $\bar{\mathcal{P}}_0$, each bin also must contain $\omega(\log n)$ balls (whp), and thus Claim 1 holds.

To prove Claims 2, we first note that the probability of throwing a ball into a high bin is $(h_t)^m$, where h_t is the total fractional load of all high bins at time t, and the probability of throwing a ball into a particular high bin with fractional load p_t at time t is equal to $(h_t)^m \cdot \frac{p_t}{h_t} = (h_t)^{m-1} p_t$. Provided each bin starts with fractional load at least $\frac{1}{\hat{c}n}$, we can then do some work to upper bound $h_t \leq (1 - 1/\hat{c}) + \epsilon_1$ at any time t in phase A (whp), for any arbitrarily small $\epsilon_1 > 0$, provided we set c_1 large enough. Thus, the probability of throwing a ball into any high bin is $\leq \gamma p_t$ (whp), where γ can be made arbitrarily close to $(1 - 1/\hat{c})^{m-1}$ by setting ϵ_1 small enough. Furthermore, since $\gamma < 1$ for small ϵ_1, the high bins are in essence suffering a *shrinkage condition*, which causes their fractional load to decrease (whp). Given that the probability of throwing a ball into a high bin is bounded by γp_t, we can then prove a lemma, which states that any time after at least $\Omega(N_0 \cdot n^{1/(1-(1+\epsilon_1)\gamma)}) = \Omega(n^{1+1/(1-(1-1/\hat{c})^{m-1})+\hat{\epsilon}_1}) = \Omega(n^{2+1/((\hat{c}/(\hat{c}-1))^{m-1}-1)+\hat{\epsilon}_1})$ new balls have been thrown, each high bin contains at most $1.01\frac{c_1}{n}$ fractional load (whp), where $\hat{\epsilon}_1$ can be made arbitrarily small, by setting c_1 large enough. Note that our last lemma then implies that Claim 2 holds (whp), since only $N_A = \Theta(n^{2+\epsilon_0})$ new balls need to be thrown, where $\epsilon_0 > 0$ can be arbitrarily small, provided m is sufficiently large.

To prove Claim 3, we follow the same steps as Claim 2, and show that (whp) at any time in phase B the probability that a high bin receives a new ball is $\leq \gamma p_t$, where γ is a constant strictly less than 1. This *shrinkage condition* can then be used with the same lemma in Claim 2 to show that the fractional load of each high bin decreases (whp), and subsequently each high bin must become a middle bin (whp) after only $O(N_A)$ new balls have been thrown, and therefore, Claim 3 holds.

To prove Claim 4, we first note that the probability of throwing a ball into a low bin at time t is $1-(1-l_t)^m$, where l_t is the total fractional load of the low bins at time t, and the probability of throwing a ball into a particular low bin with fractional load p_t at time t is $(1-(1-l_t)^m)\cdot\frac{p_t}{l_t} = \frac{1-(1-l_t)^m}{l_t}\cdot p_t$. Furthermore, we can do some work to prove that at any time in phase B, $\frac{1-(1-l_t)^m}{l_t} \geq \gamma$ (whp), where γ can be arbitrarily large provided we set c_2 large enough, and thus, each low bin maintains a *growth condition* (whp). Given that this *growth condition* holds for any low bin at any time t in phase B, we can then prove a lemma which states that (whp) after $O(N_A^{1+1/(.99\gamma-1)})$ more balls have been thrown in phase B, each low bin must become a middle bin (whp). Therefore, all low bins must become middle bins (whp) before the end of phase B when $N_B = \Theta(N_A^{1+\epsilon_0})$ balls have been thrown, where $\epsilon_0 > 0$ can be arbitrarily small provided we set c_2 large enough, and thus, Claim 4 holds.

To prove Claim 5, we just need to apply Azuma's inequality to an appropriately defined martingale. Consider defining the random variable a_t to be the number of balls in the marker bin M at time t, and b_t to be the number of balls in any other middle bin M' at time t. It is not hard to show that if bin M' first becomes a middle bin at time t_0 in phase B, then the sequence of random variables $X_t \equiv \frac{a_t}{a_t+b_t}$ is a martingale, starting from time t_0. Now, to prove Claim 5, we just need to note that both bin M and M' have the same load at time t_0, and apply Azuma's inequality on the sequence of random variables, X_t, starting at time t_0 and ending at the end of phase B, in order to conclude that the fractional load of bin M and bin M' remain close at then end of phase B (whp). Thus, Claim 5 follows.

4 Linked Decompositions with Exceptions

The proof in the previous section that $\mathcal{PA}(m)$ graphs have linked decompositions requires m to be large for small ϵ. However, the Internet has average degree about four, and thus, the $\mathcal{PA}(m)$ model is not considered to be a realistic model of the Internet unless m is very small, since the average degree of a node in a $\mathcal{PA}(m)$ graph is $2m$. Our next result holds when $m \geq 2$, and achieves a slightly better $a = \sqrt{n \log n}$, but it achieves a weaker form of linked decomposition. In particular, we define a *linked decomposition with exceptions* to be a decomposition satisfying the same 4 requirements as before, with parameters a, b, c, d, except that each component is allowed to contain one "promiscuous" node, which may belong to more than b components. It is not hard to see that the routing properties of the decomposition (see Section 5) are preserved in the face of a single exception per component (the remaining nodes essentially "route around it"). We can show:

Theorem 4 *A graph G_n generated by the $\mathcal{PA}(2)$ model has, with high probability, a linked decomposition with exceptions and $a = \Theta(\sqrt{n \log n})$, $b = 2$, $c = \Theta(\sqrt{n/\log n})$, and $d = O(\log n)$.*

Proof. We start by running $\mathcal{PA}(2)$ for $T_0 = n/\log n$ steps, and then we decompose the resulting graph G_{T_0} into $c = \sqrt{n/\log n}$ components of small diameter and each with total degree (the sum of the degrees of its nodes) between s and $3s$, where $s = \sqrt{n \log n}$. These components will be completely disjoint, except for at most one exceptional node in each.

Due to space limitations, we only provide a rough sketch regarding how to obtain this starting decomposition. We start by computing a breadth-first search tree of G_{T_0} starting from the first node, thus achieving $\log n/\log\log n$ diameter with high probability [12]. From here, it is not hard to show that one can start from the bottom of the tree and iteratively find connected subtrees (components) of total degree between s and $2s$, until one subtree (component) remains of total degree at most $3s$. These subtrees cover all nodes, and each subtree only intersects other subtrees at one of its nodes (the one node closest to the root).

Once we have this initial decomposition, with total degrees balanced within a constant, we continue the $\mathcal{PA}(m)$ process for the remaining $n - T_0$ steps. To assign these last $n - T_0$ nodes to components, we follow the same idea as before, and assign each new node i to two components each as follows: we look at where the first two edges of node i point, say to nodes j and j', and assign node i to a component to which j belongs, and a component to which j' belongs. If j or j' belong to two or more components, then we choose a random component of j or j' for node i (not necessarily uniformly at random). We omit formal details, but the random component of j or j' is chosen in a manner such that node i gets assigned to two random components, which are as close as possible to two uniformly random components.

Now, since the total degree of each component is roughly balanced to start with one can use similar techniques (essentially Azuma's inequality) as Lemma 3 in the Appendix to prove that the total degree of each component remains roughly balanced throughout the remainder of the process (whp). Furthermore, this implies that each new node i is assigned to two approximately uniformly random components, and in particular, one can show that each new node has probability at least $\Omega(1/c^2)$ of achieving each of the $O(c^2)$ intersections. Lastly, we can then apply a coupon collector argument to show that after the final $n - T_0 > \Theta(c^2 \log c)$ nodes arrive, all components intersect with high probability, and thus, we have a proper linked decomposition. □

Another important model of Internet-like graphs is called the *degree sequence model* [6, 19], where we start with a degree sequence $d_1 \geq d_2 \geq \cdots \geq d_n$, such that $d_i = d_1 \cdot i^{-\alpha}$ for some α between 1/2 and 1. Given the degrees, d_i, we then add an edge between each pair of nodes $\{i, j\}$ independently at random with probability proportional to $d_i \cdot d_j$. This model ensures that in expectation the degrees of the nodes are proportional to the d_i's. By a similar argument and construction, which we omit, we can show the following:

Theorem 5 *Linked decompositions with exceptions and $a = \Theta(\sqrt{n \log n})$ can be obtained in the degree sequence model, as well as for $G_{n,p}$, with $p = \Theta(\frac{1}{n})$.*

Finally, we note that by another simple argument, random graphs in the $G_{n,p}$ model have linked decompositions (without exceptions), provided that p is above $\log n/n$.

5 The Internet, Routing, and Experiments

In recent years, we have seen a surge of research activity aiming at a theoretical and foundational understanding of the Internet. The motivation for such a research agenda is twofold: first, the Internet is a novel, fascinating, and intellectually challenging computational artifact of central importance to computing technology and society in general, and hence it is naturally an attractive subject for theoreticians. Second, even though the Internet has emerged without much deliberate design, such design may become necessary in the future; foundational understanding and theoretical insights would be handy at such a juncture. Indeed, as challenges of scale accumulate, there are several serious efforts underway to "redesign" the Internet [3, 4, 2, 13]. It is within this framework that we see the concept of linked decompositions, and the theoretical and experimental evidence presented here is that it is feasible in the context of the Internet.

It is not hard to see that a linked decomposition with parameters a, b, c, d enables a novel form of Internet routing with routing tables of size $\tilde{O}(ab + n/a)$, delay $O(d)$, and $O(d)$ packets sent per message routed. To show how this can be done, first note that each node v only needs $\tilde{O}(ab)$ space to store enough information to route to any node u which belongs to one of v's components. Furthermore, in order to route to any node u, which does not belong to one of v's components, v just needs to find an intermediate node w, which belongs both to one of u's components and one of v's components. Note that such a node w must exist due to property (4) of our linked decomposition. Once we have found our node w, it is easy to route our message from v to u through w.

Conceptually, storing an intermediate node w for each destination node u appears to require a table of size $O(n)$, but fortunately node v does not have to store the entire table, as it is not hard to distribute this table among the nodes in v's components, such that each node only has to store $\tilde{O}(n/a)$ information. In particular, we can use a hash function to distribute the information in a balanced manner among the nodes, such that u can find the required intermediate node w for any destination node v, while only storing $\tilde{O}(n/a)$ information per node. Thus, given a linked decomposition with parameters a, b, c, and d, we have a routing protocol that uses $\tilde{O}(ab + n/a)$ storage and routes with delay $O(d)$, and in particular, if we have a linked decomposition with $a \approx \sqrt{n}$, $b = 2$, and $d = O(\log n)$, then we have a simple routing scheme with routing tables of size about $\sqrt{n}$ and delay $O(\log n)$.

For a concrete example, see Figure 1 below. To route a packet from vertex 1 to vertex 2 in the given example, vertex 1 knows how to reach any node in component S_1 and can route the packet entirely within component S_1 by following the links in S_1. To route a packet from vertex 1 to vertex 3, which is not in any component of vertex 1, we a send lookup message within component S_1 to determine the intermediate vertex to use, in this case vertex 4. The packet is then routed to vertex 4, and vertex 4 uses its routing table to complete the route to vertex 3.

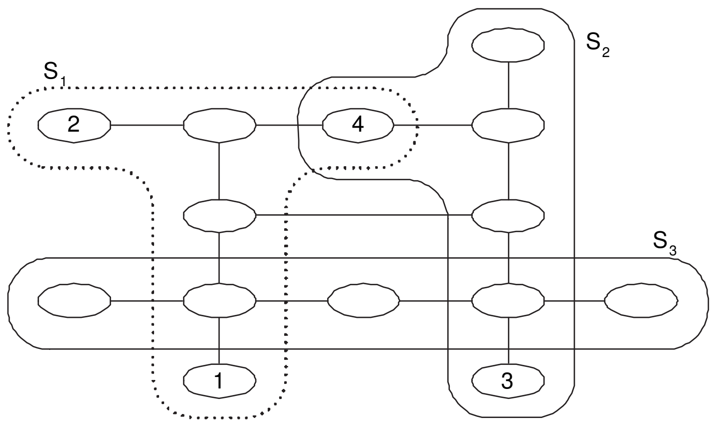

Figure 1: In this example, we show a linked decomposition with $a = 6$, $b = 2$, $c = 3$, and $d = 4$.

Note that the above routing scheme can be implemented even if the addresses of the vertices are totally amorphous strings (called "flat labels" in the Internet redesign literature [13]), as opposed to today's hierarchical and geographically specific IP addresses — a key feature of today's Internet, which is also the source of some of the most challenging problems of scale and evolution.

By giving up the "flat labels" goal, we can generate a routing scheme with even less information stored, about $n^{1/3}$ per node, by relaxing the fourth requirement of linked decomposition as follows: instead of requiring that all components intersect, we only require that any two components *either intersect with one another, or both intersect at some other component.* That is, the intersection graph of the components is no longer a clique, but has diameter two. It turns out that we can achieve this slightly weaker form of linked decomposition in $\mathcal{PA}(m)$ graphs. In particular, we can apply a combination of coupon collector and birthday paradox arguments (details omitted), in order to show that this type of linked decomposition can be generated in $\mathcal{PA}(m)$ graphs by decomposing the graph into $c \approx n^{2/3}$ components of size $a \approx n^{1/3}$ each. Given this decomposition, one can then show that if the addresses of each node are assigned so that they also indicate a component to which the node belongs, then we only need routing tables of size about $a \approx n^{1/3}$ (details omitted).

There has been much work done on network routing, although the one that provided direct inspiration is the *routing on flat labels* proposal [13], a novel routing architecture, which utilizes methodologies from peer-to-peer networks. We came up with the concept of linked decompositions while trying to identify the limits of their approach. The routing on flat labels work is validated experimentally, and does not provide nontrivial performance guarantees. Another related body of work is that on *compact routing* [7, 5, 17, 18] seeking routing algorithms with small routing tables and small *stretch* (worst-case ratio between routing delay and distance in the network). The strongest known such result achieves, for any graph, $\tilde{O}(\sqrt{n})$ tables and stretch 3 [5], and it is in fact known empirically [17] that this algorithm runs better on the real Internet.

Although the existing compact routing work dominates our results, our approach is conceptually simpler, and does not require the use of landmark nodes. Avoiding the use of landmark nodes, may provide benefits in terms of congestion and robustness against failures, since as many as $\Omega(n^{3/2})$ pairs of nodes may route packets through a single landmark node in the existing $\tilde{O}(\sqrt{n})$ routing scheme. As we describe in the open problems section, we hope to explore the potential benefits which linked decompositions may provide in terms of balanced congestion and robustness. Furthermore, our decomposition may prove useful in maintaining various control and management functions locally within each decomposition component. In addition, in our experiments, the stretch of our routing algorithm is very rarely over 3. Finally, our work also provides a rigorous explanation for the surprising empirical finding (see Appendix L) that the actual Internet can be decomposed in such a demanding way.

6 Open Problems

Our proofs are in some sense existential and non-constructive: Even though we give a decomposition algorithm, this algorithm needs to "see" the actual running of the $\mathcal{PA}(m)$ process in order to work with high probability; what if we are given ex post a graph that has been generated by $\mathcal{PA}(m)$, but with its nodes permuted? Or if we are actually given the graph of the Internet? In our experiments with permuted $\mathcal{PA}(m)$ graphs and Internet graphs, we run our decomposition algorithms on the graph *with nodes ordered in decreasing degree.* The intuition is that this is our best guess for the creation order. This works well in practice (see Appendix L), and we would love to prove that it works with high probability on $\mathcal{PA}(m)$ graphs.

There are also interesting questions regarding the existence of linked decompositions with parameters $a \approx \sqrt{n}$, $b = 2$, $c \approx \sqrt{n}$, $d \approx \text{diameter}(G)$. We have found that graphs generated by the $G_{n,p}$ model, the degree sequence model, and the $\mathcal{PA}(m)$ model all have linked decompositions (whp).

Yet if the graph G is a tree, a linked decomposition with the described parameters cannot exist, essentially because the required component intersections cannot be created without forcing at least one node to be included in too many components. This begs the question: what are necessary and sufficient conditions for a graph G to have a linked decomposition with the parameters described? Can all expander graphs be decomposed in this way?

In addition, some other open questions we would like to answer are:

- Does the Polya urn process with the power of *two* choices also balance after $O(n^{2+\epsilon})$ steps? In the case of interest, when each bin starts with fractional load at least $\frac{1}{2n}$, we can prove the power of two choices balances loads after $n^{3+\epsilon}$ steps (whp), but is an exponent arbitrarily close to two possible?

- Does our routing algorithm yield better worst-case congestion than the existing $\tilde{O}(\sqrt{n})$ routing scheme? We have removed the use of landmark nodes, but is it enough to balance congestion?

- Are there efficient ways to update the routing tables with the addition of new nodes and edges, or more importantly, node and edge deletions?

- Can autonomous systems in a network be appropriately incentivized to organize themselves in linked decompositions?

Acknowledgment: We want to thank Kevin Fall, Robert Kleinberg, Dima Krioukov, Satish Rao, Luca Trevisan, Grant Schoenbeck, Scott Shenker, Ion Stoica, and Mikkel Thorup for valuable discussions. In particular, Satish Rao provided the distributed storage idea in our routing scheme.

References

[1] Cooperative Association for Internet Data Analysis. http://www.caida.org/home/.

[2] NewArch Project: Future-Generation Internet Architecture. http://www.isi.edu/newarch/.

[3] FIND: Future Internet Network Design. http://find.isi.edu/, 2005.

[4] GENI: Global Environment for Network Innovations. http://www.geni.net/, 2005.

[5] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. *The Sixteenth ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 04)*, 2004.

[6] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. In *STOC*, 2000.

[7] Marta Arias, Lenore J. Cowen, Kofi A. Laing, Rajmohan Rajaraman, and Orjeta Taka. Compact routing with name independence. In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures (SPAA 03)*, pages 184–192, New York, NY, USA, 2003. ACM Press.

[8] Baruch Awerbuch, Andrew Goldberg, Michael Luby, and Serge Plotkin. Network decompositions and locality in distributed computation. 1989.

[9] Baruch Awerbuch and David Peleg. Sparse partitions. 1990.

[10] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations (extended abstract). In *STOC '94: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 593–602, New York, NY, USA, 1994. ACM Press.

[11] A. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[12] B. Bollobás and O. Riordan. Mathematical results on scale-free random graphs. In *Handbook of Graphs and Networks*. Wiley-VCH, Berlin, 2002.

[13] M. Caesar, T. Condie, J. Kumar Kannan, K. Lakshminarayanan, I. Stoica, and S. Shenker. ROFL: Routing on Flat Labels. *SIGCOMM*, 2006.

[14] F. Chung, S. Handjani, and D. Jungreis. Generalizations of polya's urn problem. *Annals of Combinatorics*, 7:141–153, 2003.

[15] A. Frieze, A. Flaxman, and T. Fenner. High degree vertices and eigenvalues in the preferential attachment graph. *Internet Mathematics*, 2:1–20, 2005.

[16] Richard M. Karp, Michael Luby, and Friedhelm Meyer auf der Heide. Efficient pram simulation on a distributed memory machine. In *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 318–326, New York, NY, USA, 1992. ACM Press.

[17] D. Krioukov, K. Fall, and X. Yang. Compact routing on internet-like graphs. *INFOCOM*, 2004.

[18] D. Krioukov and kc claffy. Toward compact interdomain routing. `arXiv:cs.NI/0508021`.

[19] Milena Mihail, Christos Papadimitriou, and Amin Saberi. On certain connectivity properties of the internet topology. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 28, Washington, DC, USA, 2003. IEEE Computer Society.

[20] Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.

[21] David Peleg and Alejandro A. Schaffer. Graph spanners. 1989.

A Proof Claim 2 for Theorem 3

The proof of claim 2 is long and divided into three subsections: the first and last subsections prove two key lemmas, and the main body of the proof is contained in Subsection A.2.

A.1 Proof of First Lemma Needed for Claim 2

To prove Claim 2, we make use of the following lemma, stating that the fractional load of a bin decreases rapidly, provided the probability of throwing a ball into the bin is less than its fractional load:

Lemma 1 *Consider a set of n bins with $\hat{N}_0$ balls total at time 0, and let i be any bin that starts with at least $\omega(\log n)$ balls. Let be p_t be the fraction of balls in bin i at time t, and consider any process that throws balls into bins such that the following* shrinkage condition *holds for bin i for some constant $c < (\frac{1}{1.01})$, and at any time t when the system contains less than or equal to $\beta \equiv \hat{N}_0 \delta^{(\frac{1}{1.01c-1})}$ balls:*

$$\Pr[\textit{Throwing a ball into bin } i \textit{ at time } t] \leq c p_t$$

where $\delta \in (0, 1)$. Then the fractional load of bin i becomes less than or equal to δp_0, at some time before the bins obtain strictly more than β balls, with probability at least $(1 - \frac{1}{n^\alpha})$, where $\alpha > 0$ can be any arbitrarily large constant, provided n is sufficiently large.

Proof. To prove the lemma, we start by showing the fractional load of bin i decreases with high probability, when $\epsilon \hat{N}_0$ balls are thrown, for some small constant $\epsilon > 0$. To show this first claim, we define the event $\mathcal{E}$ to be the event that bin i never obtains more than $(1+\epsilon)p_0$ fraction of balls while the first $\epsilon \hat{N}_0$ balls are thrown, and contains at most $p_0(e^{(1.01c-1)\epsilon})$ fraction of balls after $\epsilon \hat{N}_0$ balls are thrown. Note that if we can prove event $\mathcal{E}$ happens with high probability, then this implies the fractional load decreases with high probability, since $c < (\frac{1}{1.01})$. To prove event $\mathcal{E}$ happens with high probability, consider an alternative process that throws $\epsilon \hat{N}_0$ balls, such that each ball has exactly $c(1+\epsilon)p_0$ probability of landing in bin i. A simple coupling argument can show that the probability event $\mathcal{E}$ happens for this new process must be less than or equal to the probability that event $\mathcal{E}$ happens for our original balls and bins process. The coupling exists more or less because if both the shrinkage condition and $\mathcal{E}$ hold for the first $\epsilon \hat{N}_0$ balls thrown in our original process, then the probability of throwing each ball into bin i is upper bounded by $c(1+\epsilon)p_0$. Thus, if we can prove $\mathcal{E}$ happens with high probability for this alternative process, it must also happen with high probability for our original process.

Now to analyze our alternative process, note that for the alternative process, the expected number of balls added to bin i is $(c(1+\epsilon)p_0)\epsilon \hat{N}_0$. Moreover, since bin i starts with at least $p_0 \hat{N}_0 = \omega(\log n)$ balls, there exists a sufficiently large n such that the expected number of balls added is at least $\frac{4\alpha'}{\epsilon^2} \log n$, for any fixed $\epsilon > 0$ and $\alpha' > 0$. Thus, we can apply a standard Chernoff bound to show that less than or equal to $(1+\epsilon)(c(1+\epsilon)p_0)\epsilon \hat{N}_0$ balls are added to bin i, with probability at least $(1 - \frac{1}{n^{\alpha'}})$, for any arbitrarily large $\alpha' > 0$, provided n is sufficiently large. Therefore, after $\epsilon \hat{N}_0$ balls are thrown, bin i contains at most $p_0 \hat{N}_0 + (c(1+\epsilon)^2 p_0)\epsilon \hat{N}_0$ balls with probability at least $(1 - \frac{1}{n^{\alpha'}})$, and the fractional load of bin i does not exceed $p_0(1 + c(1+\epsilon)^2 \epsilon)$ with high probability while the $\epsilon \hat{N}_0$ balls are thrown. Thus, provided ϵ is sufficiently small, the fractional load of bin i does not exceed $p_0(1+\epsilon)$ with probability at least $(1-\frac{1}{n^{\alpha'}})$, and the first condition of event $\mathcal{E}$ holds with high probability. Furthermore, after $\epsilon\hat{N}_0$ balls have been thrown, the fractional load of bin i is at most $p_0(1+c(1+\epsilon)^2\epsilon)/(1+\epsilon) \leq p_0(e^{(1.01c-1)\epsilon})$ with probability at least $(1-\frac{1}{n^{\alpha'}})$, for sufficiently small $\epsilon > 0$. Therefore, we have shown that event $\mathcal{E}$ happens with probability at least $(1-\frac{1}{n^{\alpha'}})$.

We can now conclude that for our original process after $\epsilon\hat{N}_0$ balls have been thrown with the shrinkage condition holding, the fractional load of bin i decreases by a factor $e^{(1.01c-1)\epsilon}$ with high probability, while the total number of balls increases by a factor of $(1+\epsilon)$. Note that if the shrinkage condition holds while the system contains less than or equal to $\hat{N}_0(1+\epsilon)^r$ balls, then we can we repeat this analysis r times to conclude the fractional load decreases by a factor of $e^{(1.01c-1)\epsilon r}$ with probability at least $(1-\frac{r}{n^{\alpha'}})$, after the system obtains $\hat{N}_0(1+\epsilon)^r$ balls. Now take $r = (\frac{1}{\epsilon})(\frac{\log\delta}{1.01c-1})$, and we see that the fractional load of bin i decreases by a factor of δ with high probability, provided the shrinkage condition holds while the total number of balls is less than or equal to $\beta = \hat{N}_0\delta^{(\frac{1}{1.01c-1})}$. Finally, note that since we can make α' arbitrarily large, we can also make the previous statement hold with probability at least $(1-\frac{1}{n^{\alpha}})$, for any arbitrarily large $\alpha > 0$, provided n is sufficiently large. Therefore, we have proven our lemma. □

Remark 1 *In the previous analysis, for simplicity, we ignored the subtle point that $\epsilon\hat{N}_0$ and the number of rounds r might not be integer. To complete the analysis more precisely, one needs to run the analysis for r' rounds, where r' is an integer, and choose a potentially different ϵ for each of the r' rounds, such that the number of balls thrown (e.g. $\epsilon\hat{N}_0$) is always integer. For notational purposes, suppose ϵ_i is the value we choose for ϵ to ensure that the number of balls thrown in each round i is integer. For large n, it is easy to see that we can make each ϵ_i small, and roughly the same value as a single fixed ϵ. Furthermore, if we choose our r' and ϵ_i, such that $\sum_{i=1}^{r'}\epsilon_i$ is close to $(\frac{\log\delta}{1.01c-1})$, then we can apply roughly the same argument as before to obtain our result, where $r' \approx (\frac{1}{\epsilon})(\frac{\log\delta}{1.01c-1})$ is an integer. For simplicity, we also ignore the same rounding issue that occurs in Lemma 4, but it can be avoided in a manner similar to the one just described.*

A.2 Application of Lemma 1

To prove Claim 2, we would like to apply Lemma 1 to high bins in phase A to show that even if bin i is high and starts with fractional load close to 1, the shrinkage condition still holds for some constant c, and the fractional load eventually decreases to $(\frac{c_1}{n})$. Once a bin decreases to fractional load $(\frac{c_1}{n})$, we need to show the load stays below $1.01(\frac{c_1}{n})$ with high probability. We omit proving this second point, but it can be done with the same techniques used to prove Lemma 1. To prove the first point, our goal is to apply Lemma 1 with $\delta = (\frac{c_1}{n})$. To apply the lemma, note that if there are at least $\omega(n\log n)$ balls in the bins, then in phase A the definition of high bin implies that each high bin contains at least $\omega(\log n)$ balls. Therefore, we can apply our lemma to any high bin, provided we set $\hat{N}_0 = \omega(n\log n)$.

Now, if we can make sure the shrinkage condition holds for an arbitrarily small constant c, then our lemma shows that the fractional load of any high bin decreases to $(\frac{c_1}{n})$ with high probability sometime before $N_A = \Theta(n^{(2+\epsilon_0)})$ total balls have been thrown, where ϵ_0 can be arbitrarily small, provided c can be made arbitrarily small. In order to show that c can be made arbitrarily small, note that the probability of throwing a ball into a high bin at time t is equal to $(h_t)^m \cdot (\frac{p_t}{h_t}) = (h_t)^{m-1} \cdot p_t$, where h_t is the total fractional load of the high bins at time t and k is the number of choices allowed.

Thus, if we can show that with high probability h_t is strictly less than 1, then we can show that c can be made arbitrarily small with high probability, provided we allow a sufficient number of random choices m. Note that in order to apply our lemma, we need to show the previous statement is true with high probability for any time t in phase A after the bins contain at least $\hat{N}_0 = \omega(n \log n)$ balls.

To complete the proof, we now just need to prove that h_t is strictly less than 1 with high probability in phase A after the bins contain at least $\hat{N}_0 = \omega(n \log n)$ balls. To prove this, first note that at any time in phase A at most $(\frac{1}{c_1})$ fraction of the bins are high bins, and at least $(1 - \frac{1}{c_1})$ fraction of bins are low bins. Furthermore, the low bins begin with at least a constant fraction of the load when the process starts with $N_0 = O(n)$ balls, since at least $(1 - \frac{1}{c_1})$ fraction of bins are low bins, and since each bin contains at least 1 ball. To show that the low bins continue to have at least some constant fraction of the load, requires another lemma, Lemma 2, which we prove in the next subsection.

A.3 Proof of Second Lemma Needed for Claim 2

Lemma 2 *Let N_0 balls be distributed among n bins, such that every subset of cn bins contains at least $c'N_0$ balls, where $c, c' \in (0,1)$ are constants. Define $f(c') \equiv (c' + (1-c')\ln(1-c'))$ and let ϵ be any small constant in $(0, \hat{\epsilon}]$, where $\hat{\epsilon} > 0$ is a small constant defined in the proof. Then after $t \geq \frac{6N_0}{\epsilon^2(1-\epsilon)f(c')} = \Theta(N_0)$ balls have been thrown according to the standard Polya urns process, every subset of cn bins contains at least $(1-\epsilon)^3 f(c')$ balls with probability at least $(1 - \frac{1}{n^\alpha})$, where $\alpha > 0$ can be arbitrarily large, provided n is sufficiently large.*

Remark 2 *Note that over the interval $(0,1)$, $f(c')$ is strictly positive, increasing, and has range $(0,1)$. Furthermore, a simple coupling argument shows that the theorem also holds for any Polya process that sometimes uses the power of multiple choices.*

Proof. Although our original process starts with N_0 balls distributed among n bins, it will be easier to analyze a standard Polya process that starts with N_0 balls distributed evenly among $n' \equiv N_0$ bins such that there is one ball in each of the n' bins. Note that our new Polya process on n' bins can be used to simulate the random loads that occur in our original Polya process with n bins. For a bin i that starts with b_i balls in our original process, we represent bin i by using b_i distinct bins from the new Polya process on n' bins. Namely, we define the current load of bin i to be the sum of the current loads of the b_i bins that are used to represent bin i. It is not difficult to show that the loads we have defined, based on the new process on n' bins, are equivalent in distribution to the random loads generated by the original Polya process on n bins.

In order to lower bound the loads of any subset of cn bins in the original process on n bins, first note that any subset of cn bins is represented by at least $c'n'$ bins in our new process on n' bins. Furthermore, if we can lower bound the total fractional load on any subset of $c'n'$ bins for our new process with a sufficient exponentially high probability, then we can also lower bound the total fractional load on any subset of cn bins in our original process with high probability.

Following this idea, we now analyze the standard Polya process on n' bins starting with one ball each, and show that a fixed subset of $c'n'$ bins has exponentially high probability of containing at least $(1-\epsilon)^3 f(c')$ fractional load after $t \geq \frac{6n}{\epsilon^2(1-\epsilon)f(c')}$ balls have been thrown. From previous work on the Polya urn model [14], we know that when n' bins start with one ball each and t more balls are thrown, the random bin loads that are generated are equivalent in distribution to the loads generated by the following random process:

1. Pick $(n' - 1)$ points uniformly at random from the interval $[0, 1]$, and define x_i to be the position of the ith lowest point generated for $i \in \{1, ..., (n' - 1)\}$. For notational purposes, define $x_0 = 0$ and $x_{n'} = 1$.

2. Pick t points uniformly at random from the interval $[0, 1]$, and define the load of bin i to be number of points that fall in the interval (x_{i-1}, x_i) plus 1, for $i \in \{1, ..., n'\}$.

By analyzing this alternate process over $[0, 1]$, we can lower bound the fractional load of a fixed subset of $z \equiv c'n'$ bins with high probability. We start by defining Y to be the sum of the lengths of the smallest $c'n'$ intervals defined by the $(n' - 1)$ points generated by our alternate process. If we can lower bound Y with high probability, then we are not far from lower bounding the the fractional load of any fixed subset of $c'n'$ bins with high probability. To start, one can show that Y is equivalent in distribution to:

$$\frac{\sum_{i=0}^{z-1}((z-i)X_{n'-i})}{\sum_{i=0}^{n'-1}((n'-i)X_{n'-i})}$$

where each X_j variable denotes an independent exponential random variable with rate j.

To lower bound Y with high probability, first define $X = \sum_{i=0}^{z-1}(z-i)X_{n'-i}$ to represent the numerator, and $X' = \sum_{i=0}^{n'-1}(n'-i)X_{n'-i}$ to represent the denominator. Now, we can lower bound Y by lower bounding X and upper bounding X'. To lower bound X, we can use an argument similar to the one used to prove Cheroff bounds. First note that $\mu \equiv \mathbf{E}[X] = \sum_{i=0}^{z-1}\left(\frac{z-i}{n'-i}\right)$, and consider any arbitrarily small fixed $\delta > 0$. We can show for negative t sufficiently close to 0:

$$\begin{aligned}
\Pr[X \leq (1-\delta)\mu] &= \Pr[e^{tX} \geq e^{t(1-\delta)\mu}] \\
&\leq \frac{\mathbf{E}[e^{tX}]}{e^{t(1-\delta)\mu}} \\
&\leq \frac{e^{t(1-\frac{\delta}{2})\mu}}{e^{t(1-\delta)\mu}} \\
&\leq e^{t(\frac{\delta}{2})\mu}
\end{aligned}$$

where going from the second line to the third line follows by upper bounding $\mathbf{E}[e^{tX}]$:

$$
\begin{aligned}
\mathbf{E}[e^{tX}] &= \mathbf{E}[e^{t(\sum_{i=0}^{z-1}(z-i)X_{n'-i})}] \\
&= \prod_{i=0}^{z-1} \mathbf{E}[e^{t(z-i)X_{n'-i}}] \\
&= \prod_{i=0}^{z-1} \left(1 - t\left(\frac{z-i}{n'-i}\right)\right)^{-1} \\
&\leq \prod_{i=0}^{z-1} e^{t(1-\frac{\delta}{2})(\frac{z-i}{n'-i})} \\
&\leq e^{t(1-\frac{\delta}{2})\mu}
\end{aligned}
$$

To go from the third line to the fourth line in the calculation above, we assume that t is negative and sufficiently close to 0.

Lastly, observe that $\sum_{i=0}^{z-1}\left(\frac{z-i}{n'-i}\right) + \sum_{i=0}^{z-1}\left(\frac{n'-z}{n'-i}\right) = z$, so that we can write μ in a more convenient form:

$$
\begin{aligned}
\mu &= \sum_{i=0}^{z-1}\left(\frac{z-i}{n'-i}\right) \\
&\approx z - (n'-z)(\ln n' - \ln(n'-z)) \\
&\approx c'n' - n'(1-c')\ln\left(\frac{1}{1-c'}\right) \\
&\approx n'(c' + (1-c')\ln(1-c'))
\end{aligned}
$$

By observing that $\mu = \Omega(n')$, we have therefore shown that $X > (1-\delta)\mu$ with exponentially high probability. It can be similarly shown that $X' < (1+\delta)n'$ with exponentially high probability, which means $Y > (1-\epsilon)f(c')$ with exponentially high probability, where $\epsilon \equiv 1 - (\frac{1-\delta}{1+\delta})$ can be made arbitrarily small by making $\delta > 0$ arbitrarily small. Now, recall that since Y is the sum of the smallest $c'n'$ intervals, we have actually shown that *every* subset of $c'n'$ bins is represented by intervals of total length at least $(1-\epsilon)f(c')$ with exponentially high probability. Furthermore, for any fixed set of $c'n'$ bins, if we throw balls according to our alternate process on the real interval $[0,1]$, then the probability that any ball thrown lands in one of the $c'n'$ bins is at least $(1-\epsilon)f(c')$.

Finally, consider exactly $t = \frac{6n'}{\epsilon^2(1-\epsilon)f(c')}$ balls being thrown. By applying a standard Chernoff bound, one can show that at least $(1-\epsilon)^2 f(c')$ fraction of the t balls thrown fall into the $c'n'$ bins with probability that at least $1-e^{-3n'}$. Thus after $t = \frac{6n'}{\epsilon^2(1-\epsilon)f(c')}$ balls have been thrown, there are a total of $n'(1+\frac{6}{\epsilon^2(1-\epsilon)f(c')})$ balls, and our $c'n'$ bins contain at least $n'(\frac{6(1-\epsilon)}{\epsilon^2})$ balls with probability at least $1-e^{-3n'}$. Therefore, the fractional load of the $c'n'$ bins is at least $\frac{6(1-\epsilon)^2 f(c')}{\epsilon^2(1-\epsilon)f(c')+6} \geq (1-\epsilon)^3 f(c')$ with exponentially high probability, after $t = \frac{6n'}{\epsilon^2(1-\epsilon)f(c')}$ balls have been thrown. Note, the last inequality only follows provided ϵ is sufficiently small, which is the reason we require $\epsilon \leq \hat{\epsilon}$ for some small $\hat{\epsilon} > 0$. Additionally, note that although we proved the previous statement for $t = \frac{6n'}{\epsilon^2(1-\epsilon)f(c')}$, it is not hard to argue that the lower bound also holds when $t \geq \frac{6n'}{\epsilon^2(1-\epsilon)f(c')}$, by showing that the lower bound value only increases for higher values of t.

Lastly, note that there are only $2^{n'}$ subsets of n' bins, so we can apply a naive union bound to conclude that every subset of more than $c'n'$ bins contains at least $(1-\epsilon)^3 f(c')$ fraction of balls with probability at least $1-e^{n'}$. Thus, it follows that after $t \geq \frac{6N_0}{\epsilon^2(1-\epsilon)f(c')}$ balls have been thrown in our original process on n bins, every subset of cn bins contains at least $(1-\epsilon)^3 f(c')$ fractional load with probability at least $(1-\frac{1}{n^\alpha})$, where $\alpha > 0$ can be arbitrarily large, provided n is sufficiently large. □

B Proof of Claim 1 for Theorem 3

To prove Claim 1, we follow the same steps used to prove Lemma 2. For conciseness, we only sketch the main details. We start by lower bounding the number of balls in the least loaded bin when balls are thrown according the standard Polya process starting with $n' = N_0$ bins each with one ball each. By lower bounding the load of the least loaded bin for the standard Polya process on n' bins, a simple coupling argument shows that this also lower bounds the load of the least loaded bin for our original Polya process with the power choice on n bins. To lower bound the load of the least loaded bin for the standard Polya process on n' bins, we analyze the same alternative random process over the interval $[0,1]$ as described in Lemma 2, used to generate the random bin loads. If we define the random variable X to be size of the smallest interval generated with $n'-1$ points are thrown uniformly over the interval $[0,1]$, then it can be shown that X is equivalent in distribution to $X_{n'}/\sum_{i=0}^{n'-1}((n'-i)X_{n'-i})$, where each X_j is an exponential random variable with rate j. Then we can show:

$$\begin{aligned}
\Pr\left[X \leq \frac{1.01}{(n')^{-(2+\epsilon_0/2)}}\right] &\leq \Pr\left[\sum_{i=0}^{n'-1}((n'-i)X_{n'-i}) \geq 1.01n'\right] + \Pr\left[X_{n'} \leq \frac{1}{(n')^{-(1+\epsilon_0/2)}}\right] \\
&\leq 0.01n^{(\epsilon_0/2)} + n^{(\epsilon_0/2)} \leq 1.01n^{(\epsilon_0/2)}
\end{aligned}$$

Thus all bins are represented by intervals of size at least $\Omega((n')^{-(2+\epsilon_0/2)})$ with high probability (although the probability exponent is low). Furthermore, since $\Theta((n')^{2+\epsilon_0})$ balls are thrown, it is easy to apply a standard Chernoff bound to show that each bin obtains at least $\Theta((n')^{\epsilon_0/2})$ balls with high probability. Then our previous coupling argument implies that each of the n bins in our original process contains at least $\Omega(n^{\epsilon_0/2}) = \omega(\log n)$ balls with high probability, proving Claim 1.

C Proof of Claim 5 for Theorem 3

To prove Claim 5, we start by proving the following lemma:

Lemma 3 *Let A and B be two bins each with $a_0 = b_0$ balls starting at time 0. For any time $t \geq 0$, let a_t and b_t represent the number of balls at time t in bin A and bin B, respectively, and consider a random process that throws a new ball at time t into bin A with probability $(\frac{a_t}{a_t+b_t})$, and bin B with probability $(\frac{b_t}{a_t+b_t})$. Then for any $T \geq 0$,* $\Pr[((\frac{1}{2}-\epsilon)/(\frac{1}{2}+\epsilon)) \cdot b_T \leq a_T \leq ((\frac{1}{2}+\epsilon)/(\frac{1}{2}-\epsilon)) \cdot b_T] \geq 1 - 2e^{-\epsilon^2(a_0+b_0-1)/2}$.

Proof. To prove the lemma, we define a sequence of random variables $X_t = \left(\frac{a_t}{a_t+b_t}\right)$, which is a martingale since:

$$\begin{aligned}\mathbf{E}[X_{t+1}|\ X_t] &= \mathbf{E}[X_{t+1}|\ a_t, b_t] \\ &= \left(\frac{a_t}{a_t+b_t}\right)\left(\frac{a_t+1}{a_t+b_t+1}\right)+\left(\frac{b_t}{a_t+b_t}\right)\left(\frac{a_t}{a_t+b_t+1}\right) \\ &= \left(\frac{a_t}{a_t+b_t}\right) = X_t\end{aligned}$$

Furthermore, note that at time t moving one ball from bin A to bin B or vice versa, changes X_t by $(\frac{1}{a_t+b_t})$. Thus, it is easy to see that $|X_{t+1} - X_t| \leq (\frac{1}{a_t+b_t}) = (\frac{1}{a_0+b_0+t})$. Now we can apply Azuma's inequality obtain $\Pr[X_T \geq \frac{1}{2}+\epsilon] \leq e^{-\epsilon^2/(2\sum_{i=0}^{T-1}(a_0+b_0+i)^{-2})} \leq e^{-\epsilon^2(a_0+b_0-1)/2}$, where the last inequality follows because $\sum_{i=0}^{T-1}(a_0+b_0+i)^{-2} \leq \int_{x=(a_0+b_0)}^{\infty}(x-1)^{-2} = (a_0+b_0-1)^{-1}$. Rearranging, we get $\Pr[a_T \geq ((\frac{1}{2}-\epsilon)/(\frac{1}{2}+\epsilon)) \cdot b_T] \leq e^{-\epsilon^2(a_0+b_0-1)/2}$. Similarly, we can also use Azuma's inequality to show $\Pr[a_T \leq ((\frac{1}{2}+\epsilon)/(\frac{1}{2}-\epsilon)) \cdot b_T] \leq e^{-\epsilon^2(a_0+b_0-1)/2}$, and the lemma follows. □

Remark 3 *It is not hard to show the lemma also holds for any process that throws a ball at time t into bin A with probability $\gamma_t(\frac{a_t}{a_t+b_t})$, bin B with probability $\gamma_t(\frac{b_t}{a_t+b_t})$, and neither bin with probability $(1-\gamma_t)$, where $\gamma_t \in [0,1]$ can be a random variable dependent on past events.*

Now to complete the proof of Claim 5, we just need to apply Lemma 3 to bin M and any bin i that obtains the same load as bin M, starting at the time they obtain the same load. Note that we can apply Lemma 3 to our two bins due to the remark. Furthermore, since bin M contains at least $\omega(\log n)$ balls by Claim 1, with high probability the load of that bin i is at most $(1+\epsilon_1)$ times the load of bin M, and at least $(1-\epsilon_1)$ times the load of bin M at the end of phase B, where ϵ_1 can be made arbitrarily small, provided we set ϵ small enough and n is large enough. Finally, note that for sufficiently large n the previous statement happens with high enough probability, so that we can take a naive union bound over all $n-1$ possible bins i to conclude that the previous statement happens for all bins i that obtain the same load as bin M. Therefore, Claim 5 holds with high probability.

D Proof of Claim 3 for Theorem 3

To prove Claim 3, we start by following the steps used to prove Claim 2. For conciseness, we only sketch the main details. By following the same steps as Claim 2, one can show that the low and middle bins always contain a constant (dependent on c_2) fraction of the total load with high probability. Consequently, one can show that the shrinkage condition of Lemma 1 holds for some fixed constant dependent on c_2, for any high bin. From here, one can apply Lemma 1 to the high bins in phase B, with δ equal to a constant, to show that if any high bins remain after γN_A balls have been thrown, then they all contain at most $(\frac{1}{2})(\frac{1}{n})$ fractional load with high probability, where δ is a constant dependent on c_1, and γ is a constant dependent on δ and c_2.

The final step needed to prove Claim 3 is to show that with high probability it cannot be the case that there remains some set of high bins all with fractional load $(\frac{1}{2})(\frac{1}{n})$ with high probability. Claim 3 then follows because this implies that all high bins must have become middle bins at some point before $\Theta(N_A)$ balls have been thrown. To prove that with high probability it is not possible for some high bins to remain all with fractional load at most $(\frac{1}{2})(\frac{1}{n})$, we use a proof by contradiction. Note that if some high bins remain with fractional load at most $(\frac{1}{2})(\frac{1}{n})$, then bin M must have load at most $(\frac{1}{2})(\frac{1}{n})$. Furthermore, by applying the same analysis used in Claim 5, one can show that all the low and middle bins have fractional load at most $1.01(\frac{1}{2})(\frac{1}{n})$ with high probability. However, this implies the total fraction load of all bins is at most $1.01(\frac{1}{2}) < 1$ with high probability, which is a contradiction. Therefore, all high bins must have become middle bins sometime before γN_A balls have been thrown with high probability.

E Proof of Claim 4 for Theorem 3

Before we can prove Claim 4, we need to prove the following lemma:

Lemma 4 *Consider a set of n bins with $\hat{N}_0 = O(poly(n))$ balls total at time 0, and let i be any bin that starts with at least $\omega(\log n)$ balls. Let be p_t be the fraction of balls in bin i at time t, and consider any process that throws balls into bins such that the following* growth condition *holds for bin i, for some constant $c > \frac{1}{.99}$, and at any time t with less than or equal to N_f total balls:*

$$\Pr[\textit{Throwing a ball into bin } i \textit{ at time } t] \geq cp_t$$

where $N_f \geq \hat{N}_0$ is a random integer stopping time. Then N_f must be less than $\beta \equiv \hat{N}_0^{1+(\frac{1}{.99c-1})}$ with probability at least $(1 - \frac{1}{n^\alpha})$, where $\alpha > 0$ can be any arbitrarily large constant, provided n is sufficiently large.

Proof. To prove the statement, we prove that if the growth condition does hold until β balls are in the system, then the fractional load of bin i becomes strictly greater than 1, with probability at least $(1 - \frac{1}{n^\alpha})$. Furthermore, since it is impossible for a bin to have fractional load more than 1, it must be the case $N_f \geq \beta$ happens with probability at most $\frac{1}{n^\alpha}$. Therefore, we can conclude N_f must be less than β with probability at least $(1 - \frac{1}{n^\alpha})$. To complete the proof, we now just need to show the fractional load of bin i becomes strictly greater than 1, with probability greater than $(1 - \frac{1}{n^\alpha})$, if the growth condition holds whenever there are less than or equal to β balls are in the system.

To prove this last statement, we start by showing that when $\epsilon\hat{N}_0$ balls are thrown, the fractional load of bin i increases to at least $p_0(e^{(.99c-1)\epsilon})$ with high probability, where $\epsilon > 0$ is some small constant. Note that even if no new balls are added to bin i, p_t is still lower bounded by $(\frac{p_0}{1+\epsilon})$ over this time period. In order to analyze our process, we first consider a simpler process that throws a ball into bin i with probability exactly $c(\frac{p_0}{1+\epsilon})$. Note that this new process can be coupled with our original process, so that the new process always adds fewer balls to bin i, while the $\epsilon\hat{N}_0$ balls are thrown. Thus, if we can say that for this new process, the fractional load of bin i increases to at least $p_0(e^{(.99c-1)\epsilon})$ with high probability, then the fractional load also increases to at least $p_0(e^{(.99c-1)\epsilon})$ with high probability for our original process.

Now, if we look at this new process, the expected number of new balls added to bin i is $(\frac{cp_0}{1+\epsilon})\epsilon\hat{N}_0$, when $\epsilon\hat{N}_0$ balls are thrown. Furthermore, for any fixed $\epsilon > 0$ and constant $\alpha' > 0$, there exists a sufficiently large n such that the expected number of new balls added is at least $\frac{2\alpha'}{\epsilon^2} \log n$, since bin i starts with $p_0 \hat{N}_0 = \omega(\log n)$ balls. Moreover, by applying a standard Chernoff bound, we can conclude the number of balls added to bin i is at least $(1-\epsilon)(\frac{cp_0}{1+\epsilon})\epsilon\hat{N}_0$ with probability at least $(1-\frac{1}{n^{\alpha'}})$. So the total number of balls in bin i is at least $p_0\hat{N}_0 + (1-\epsilon)(\frac{cp_0}{1+\epsilon})\epsilon\hat{N}_0$ with high probability, and the fractional load of bin i is at least $p_0(1+c\epsilon\frac{1-\epsilon}{1+\epsilon})/(1+\epsilon)$ with high probability. For simplicity, we can lower bound the previous expression by $p_0(e^{(.99c-1)\epsilon})$, for sufficiently small ϵ. Therefore, by the previous coupling argument, we can conclude that for our original process, the fractional load increases by at least $p_0 e^{(.99c-1)\epsilon}$ with probability at least $(1-\frac{1}{n^{\alpha'}})$.

Thus, after $\epsilon\hat{N}_0$ balls have been thrown, we have increased the fractional load of bin i by a multiplicative factor of $e^{(.99c-1)\epsilon}$ with high probability, and we have increased the total number of balls in the system by a multiplicative factor of $(1+\epsilon)$. Furthermore, we can repeat the same analysis r times, increasing the total number of balls to $\hat{N}_0(1+\epsilon)^r$, and increasing the fractional load of bin i by a factor of $e^{(.99c-1)\epsilon r}$ with probability at least $(1-\frac{r}{n^{\alpha'}})$, provided the growth condition holds for the new balls thrown over this time. Now, assuming that the growth condition holds until there are at least β balls in the system, we know that the growth condition holds over $r = (\frac{1}{\epsilon})(\frac{\log \hat{N}_0}{.99c-1})$ rounds, since $\hat{N}_0(1+\epsilon)^r \le \hat{N}_0 \cdot \hat{N}_0^{1/(.99c-1)}$. Moreover, if the growth condition holds over $r = (\frac{1}{\epsilon})(\frac{\log \hat{N}_0}{.99c-1})$ rounds, then the fractional load of bin i increases by at least a factor of $\hat{N}_0$ with probability at least $(1-\frac{r}{n^{\alpha'}})$. Therefore, we can conclude the fractional load of bin i becomes strictly greater than 1 with high probability, since $p_0 > \frac{1}{\hat{N}_0}$. Lastly, note that since we can make α' arbitrarily large, we can also make the previous statement hold with probability at least $(1-\frac{1}{n^{\alpha}})$, for any arbitrarily large $\alpha > 0$, provided n is sufficiently large. Therefore, bin i obtains fractional load strictly greater than 1 with probability at least $(1-\frac{1}{n^{\alpha}})$, and by our previous reasoning, the lemma follows. □

To complete the proof of Claim 4, we would like to apply Lemma 4 to the low bins in phase B. Note that by starting in phase B, we know that each low bin contains at least $\omega(\log n)$ balls with high probability, by Claim 1. So by setting with $\hat{N}_0 = N_A$, we automatically meet the first condition of the lemma with high probability. To complete the proof, we need to show that for any low bin with fractional load p_t at time t in phase B, the probability of throwing a ball into the low bin is greater than or equal to cp_t, where $c > \frac{1}{.99}$ can be made arbitrarily large, provided c_2 can be arbitrarily large. By showing the growth condition holds for all low bins, we can apply Lemma 4 to show that with high probability the growth condition must fail to hold for the low bin before the bins obtain $\Theta(N_A^{(1+\epsilon_0)})$ balls. Claim 4 then follows immediately because the only way for the growth condition to stop holding is if the low bin obtains the same load as bin M and becomes a middle bin.

Therefore, we just need to show that the growth condition holds for low bins for arbitrarily large c, and then Lemma 4 implies Claim 4. To prove this, note that the probability of throwing a ball into the low bin at time t is equal to Pr[Throw a ball into a low bin]$(\frac{p_t}{l_t}) = (\frac{1-(1-l_t)^m}{l_t}) \cdot p_t$, where l_t is the fractional load of the low bins at time t. Thus, to complete the proof, we just need to show that by increasing c_2, l_t can be made arbitrarily small with high probability.

To complete the last step, we prove the total fractional load of the low bins is at most $\frac{1.01}{c_2-1}$ with high probability by showing that there are at most $(\frac{n}{c_2})$ low bins, and with high probability each low bin contains at most $1.01\frac{c_2}{n(c_2-1)}$ fractional load. The first part of the statement follows easily from the definition of low bin, and the second part of the statement follows by showing bin

M contains at most $1.01\frac{c_2}{n(c_2-1)}$ fractional load with high probability. We can bound the the load of bin M, by noting that there are at least $(1-\frac{1}{c_2})n$ middle and high bins, whose total fractional load cannot exceed 1. This implies the fractional load of the least loaded middle or high bin must be less than or equal to $\frac{c_2}{n(c_2-1)}$. Furthermore, note that the analysis used to prove Claim 5, also shows that the fractional load of bin M has at most 1.01 times the load of the lowest middle bin with high probability, and bin M can never have more load than a high bin. Therefore, the fractional load of bin M is at most $1.01\frac{c_2}{n(c_2-1)}$ with high probability, which completes the proof of the last step.

Therefore, we have shown that every low bin eventually obtains the same load as bin M and joins the middle group with high probability sometime before the bins obtain $\Theta(N_A^{(1+\epsilon_0)})$ balls, where ϵ_0 can be arbitrarily small, by taking c_2 large enough.

F Proof of the Theorem 2

We prove that our algorithm yields balanced component sizes (in terms of degree) with high probability when running on the preferential attachment model. To analyze the size of the components generated by our preferential attachment process, we analyze a random process on n bins, where each bin represents a component and the loads represent the size of the components. By looking at the preferential attachment process, we know that the random process starts with $2kn$ balls distributed arbitrarily among n bins, such that each bin contains at least k balls each. Furthermore, we also know that with high probability the process should start with at most $O(n^{1/2+\epsilon})$ balls each in bin [15], a fact which will be useful later on. At each step of the process, $2k$ new balls are thrown into the n bins, according to the following rule:

- Pick k bins i.i.d. at random, with probability proportional to bin load
  - (I) Throw 1 ball into each of the k random bins picked
  - (II) Throw k more balls into the least loaded of the k random bins picked.

Let's call this random process $\mathcal{P}$. As before, we'll analyze another random process $\mathcal{P}_0$ which does not always throw the k balls in step (II) into the least loaded bin, but sometimes throws the k balls into a heavier bin. We can show that the random process $\mathcal{P}_0$ achieves roughly balanced loads (whp), and we can show a coupling exists where $\mathcal{P}$ always majorizes $\mathcal{P}_0$. Therefore, $\mathcal{P}$ also achieves roughly balanced loads (whp). We ignore the coupling argument and finish by defining $\mathcal{P}_0$ below and showing it becomes roughly balanced (whp).

Our random process $\mathcal{P}_0$, also throws balls in two phases, analogous to the two phases defined in section [cite]. For each phase, we throw the same number of balls as before, and we also choose where to throw the k balls in step (II) using the same rules as before, based on whether or not the bins selected are low, high, or middle. The only slight modification we employ is in phase B, where we say a low bin becomes a middle bin whenever it obtains load equal to *or greater than* the load of bin M, and a high bin becomes a middle bin whenever it has obtains load equal to *or less than* the load of bin M. To prove the process $\mathcal{P}_0$ yields balanced loads, we prove the same five claims as before. The proof of each claim is roughly the same as before, but the details are slightly different as the two processes are slightly different.

G Claim 1 for Theorem 2

To prove Claim 1, we analyze a new random process $\mathcal{P}_1$ on n bins, which starts with the same initial load distribution as the random process $\mathcal{P}_0$, but adds balls to bins in a slightly different manner. At each step of the process, $2k$ new balls are thrown into the n bins, according to the following rule:

- Pick k bins i.i.d. at random, with probability proportional to bin load
  - (I) Throw 1 ball into each of the random bins picked
  - (II) Throw k more balls into the last random bin picked

A simple coupling can be defined so that $\mathcal{P}_0$ always majorizes $\mathcal{P}_1$, so that we just need to lower bound the load of the least loaded bin when running the random process $\mathcal{P}_1$, in order to prove lower bounds on the load of the least loaded bin when running the random process $\mathcal{P}_0$. We omit the coupling argument for conciseness, but finish by proving that all bins contain at least $\Omega(\log n)$ balls with high probability after throwing $\Theta(n^{2+\epsilon})$ balls when running the random process $\mathcal{P}_1$, where $\epsilon > 0$ can be arbitrarily close to 0.

To prove the previous statement, we show that after throwing $\Theta(n^{2+\epsilon})$ balls, any fixed bin has $\Omega(\log n)$ balls with probability at least $1 - n^{-\alpha}$, where $\alpha > 0$ can be arbitrarily large, provided we can set k arbitrarily large. Our claim then follows by applying a simple union bound, which implies that after throwing $\Theta(n^{2+\epsilon})$ balls, all bins have load at least $\Omega(\log n)$ with probability at least $1 - n^{-\alpha+1}$.

To analyze the load of a single bin B, note that after t iterations of our process (i.e. after $2kt$ new balls have been thrown), bin B is picked as a random bin in the next iteration with probability at least $\frac{k}{2kn+2kt} = (\frac{1}{2})(\frac{1}{n+t})$. Now consider the random variable $Y \equiv \sum_{t=0}^{T} \sum_{j=1}^{k} X_{t,j}$, where $X_{t,j}$ are bernoulli random variables with success probability $(\frac{1}{2})(\frac{1}{n+t})$, for $t \in \{0, \ldots, T\}$, $j \in \{1, \ldots, k\}$. It is easy to that Y is stochastically dominated by the number of balls thrown by step (I) of our process, and furthermore can be used to lower bound the number of balls added to bin B. Note that for $T = n^{2+\epsilon}$, $\mathbf{E}[Y] = \frac{k-1}{2} \sum_{t=0}^{T} (\frac{1}{n+t}) \approx \frac{k-1}{2}(1+\epsilon) \log n$. Now, we can apply a standard Chernoff bound on Y to conclude the at least $\Omega(\log n)$ balls are added to bin B with probability at least $1 - n^{-\alpha}$, where α can be made arbitrarily large, provided k can be arbitrarily large. Therefore, we have shown that after throwing $\Theta(n^{2+\epsilon})$ balls, all bins have load at least $\Omega(\log n)$ with probability at least $1 - n^{-\alpha+1}$.

H Claim 2 for Theorem 2

H.1 Lemma 3 for Theorem 2

Consider running the random process $\mathcal{P}_1$ defined above in Claim 1, for $T_f = O(poly(n))$ number of steps. Assuming we have a constant $c \in [\frac{1}{2}, 1)$, which is sufficiently close to 1, we can prove the following lemma about the process above:

Lemma 5 *The lowest cn bins always contain at least $\frac{1}{8}$ fraction of the load, during all T_f iterations of the random process, with high probability.*

Remark 4 *A simple coupling argument implies that this statement also holds for our original random process $\mathcal{P}_0$.*

To prove the lemma, let's consider a fixed subset of cn bins. Let X_t denote the fraction of balls in cn bins after t iterations of our process (i.e. after $2kt$ new balls have been thrown). For notation, let $N_t \equiv 2kn + 2kt$ be the total number of balls in the n bins after iteration t, and let B_t represent the number of new balls added to our subset of cn bins in the tth iteration. Note that X_t is a martingale, since:

$$\begin{aligned}
\mathbf{E}[X_{t+1}|\ X_t] &= \frac{X_t \cdot N_t + \mathbf{E}[B_{t+1}|X_t]}{N_t + 2k} \\
&= \frac{X_t \cdot N_t + 2k \cdot X_t}{N_t + 2k} \\
&= X_t
\end{aligned}$$

Since X_t is a martingale, we make use of Azuma's inequality. Note that $c_t \equiv |X_{t+1} - X_t| \leq (\frac{2k}{N_t}) = (\frac{1}{n+t})$ and $\sum_{t=0}^{\infty} {c_t}^2 \leq \sum_{t=0}^{\infty}(n+t)^{-2} \leq \int_{x=n}^{\infty}(x-1)^{-2} = (n-1)^{-1}$. Furthermore, $X_0 \geq \frac{cnk}{2nk} = \frac{c}{2} \geq \frac{1}{4}$. Therefore, we can apply Azuma's inequality to obtain $\Pr[X_T \leq \frac{1}{8}] \leq e^{-(n-1)/128}$, for any fixed time T, which means that any fixed subset of cn bins has probability at most $e^{-(n-1)/128}$ of having less than $\frac{1}{8}$ load at time T. Provided c is sufficiently close to 1, we can then apply a union bound over all subsets of bins of size cn, and over all times $T \in \{0, \ldots, T_f\}$ to conclude that the lowest cn bins always have fractional load at least $\frac{1}{8}$ with high probability.

H.2 Using Lemma 3

Note that Lemma 3 (5) also lower bounds the fractional load of the lowest cn bins when running the $\mathcal{P}_0$ process, via a simple coupling/majorization argument. Lemma 3 is useful for proving Claim 2 because a lower bound the fractional load of the low bins in phase A of the $\mathcal{P}_0$ process can be used to upper bound the probability that balls get added to the high bins in phase A.

To start our proof, note that at most $\frac{1}{c_1}$ fraction of the bins have fractional load greater than or equal to $\frac{c_1}{n}$, which implies that at least $(1 - \frac{1}{c_1})$ fraction of bins are low bins. Thus, if we make c_1 sufficiently large and take $c = (1 - \frac{1}{c_1})$, then we can use Lemma 3 to conclude with high probability that the low bins in phase A always contain at least $\frac{1}{8}$ of the load and the high bins in phase A contain at most $\frac{7}{8}$ fractional load.

Now let p_t represent the fractional load of some high bin H at time t and let h_t represent the total fractional load of all high bins at time t. It is not hard to see the expected number of balls that get added to high bin H in phase A at time t equals $kp_t + k(h_t)^k(\frac{p_t}{h_t})$, which is upper bounded by $2k \cdot p_t(\frac{1}{2} + (\frac{7}{8})^{k-1})$ with high probability. Thus, we have shown with high probability that the fraction of $2k$ new balls that get added to bin H at time t is less than or equal to γp_t, in expectation, where $\gamma = (\frac{1}{2} + (\frac{7}{8})^{k-1}) \approx \frac{1}{2}$ for large k. In other words, in expectation bin H is only obtaining $\gamma p_t \approx (\frac{1}{2})p_t$ fraction of the $2k$ new balls thrown at each time step, which should cause the fractional load of high bin H to decrease rapidly.

So in essence, the high bin H satisfies a *shrinkage condition* analogous to the one described in Lemma 2, with shrinkage constant $\gamma \approx \frac{1}{2}$. Furthermore, note that for our modified Polya urns process $\mathcal{P}_0$, we can assume that each bin starts with at most $O(n^{\frac{1}{2}+\epsilon})$ balls (or in other words $O(n^{-1/2+\epsilon})$ fractional load) with high probability [15], where $\epsilon > 0$ can be arbitrarily small. Thus, the fractional load of bin H must only decrease by a factor of $\delta = \Theta(n^{-1/2-\epsilon})$ to ensure H has fractional load at most $\frac{c_1}{n}$. From here, we can essentially follow the same steps used in proving Lemma 2, to prove that bin H obtains load $\frac{c_1}{n}$ with high probability, sometime before at most $\beta = \hat{N}_0 \delta^{(\frac{1}{(1+\epsilon)\gamma-1})}$ balls are thrown by our process, where $\hat{N}_0$ is the total number of balls that start in our bins and $\epsilon > 0$ can be arbitrarily small. We omit the proof for conciseness, but it can also be proved using a series of Chernoff bounds as before.

Note that since γ can be arbitrarily close to $\frac{1}{2}$ for k sufficiently large and ϵ arbitrarily close to 0, then our new version of Lemma 2 implies that bin H obtains fractional load $\frac{c_1}{n}$ sometime before $\beta = O(\hat{N_0}^{2+\epsilon_0})$ balls are thrown. After bin H obtains fractional load $\frac{c_1}{n}$, it is not hard to show that with high probability bin H always maintains fractional load less than $1.01(\frac{c_1}{n})$ until the end of phase A, thus proving Claim 2.

I Claim 5 for Theorem 2

To prove Claim 5, we follow the same steps as the original proof of Claim 5. Note that if a bin B first becomes a middle bin at time T_0 (i.e. after $2kT_0$ new balls have been thrown), then $|b_{T_0} - m_{T_0}| \leq 2k$, where we use the notation b_t to represent the number of balls in bin B at time t, and we use m_t to represent the number of balls in the "yardstick" middle bin M at time t. Now if we define $X_t = \frac{b_t}{m_t + b_t}$, then we know $X_{T_0} \in [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$ (whp), where $\epsilon > 0$ can be arbitrarily small provided n is sufficiently large, since (whp) $m_{T_0} = \Omega(\log n)$ and $|b_{T_0} - m_{T_0}| \leq 2k$. Furthermore, it is not hard to show X_t is a martingale for $t \geq T_0$, and we can apply Azuma's inequality as before to prove that at the end phase B, $X_t \in [\frac{1}{2} - 2\epsilon, \frac{1}{2} + 2\epsilon]$, where $\epsilon > 0$ can be arbitrarily small. Therefore, we have proved Claim 5.

J Claim 4 for Theorem 2

To prove Claim 4, we follow the same steps as the original proof of Claim 4. We can show the total fractional load of the low bins is at most $\frac{1.01}{c_2 - 1}$ (whp), following the same steps detailed in the next to last paragraph of the original proof of Claim 4. Now, let p_t represent the fractional load a low bin L at time t of our process $\mathcal{P}_0$ (i.e. after $2kt$ new balls have been thrown in t iterations), and let l_t represent the total fractional load of all lows bins at time t. Then, the expected number of new balls added to bin L on the $t+1$th iteration of $\mathcal{P}_0$ is $kp_t + k(1 - (1 - l_t)^k)(\frac{p_t}{l_t})$, and (whp) this is greater than or equal to $k(\frac{1}{2})(\frac{p_t}{l_t}) \geq kp_t(\frac{c_2 - 1}{3})$, provided $c_2 \geq 4$.

Thus, we have shown with high probability that the expected fraction of the $2k$ new balls that get added to bin L at time t is greater than or equal to γp_t, where $\gamma = \frac{c_2 - 1}{6}$. In other words, (whp) the low bin L is satisfying a *growth condition* analogous to the one described in Lemma 5, where bin L obtains in expectation least $\gamma p_t = (\frac{c_2 - 1}{6}) p_t$ fraction of the $2k$ new balls thrown at each time step. The growth condition implies the fractional load of bin L increases rapidly, and we can follow the same steps as in Lemma 5 to prove that bin L becomes a middle bin within $\beta \equiv \hat{N_0}^{1+(\frac{1}{.99\gamma - 1})}$ iterations with high probability, where $\hat{N}_0 = N_A$ is the total number of balls in the bins at the start of phase B. Furthermore, note that γ can be made arbitrarily large by taking c_2 large enough, which implies that bin L becomes a middle bin sometime before $N_A^{1+\epsilon}$ balls are thrown, where ϵ

can be arbitrarily small provided by choosing c_2 large enough. Finally, we note bin L becomes a middle bin with sufficiently high probability, so that a union bound implies that all low bins must become middle bins with high probability, thus proving Claim 4.

K Claim 3 for Theorem 2

To prove Claim 3, we follow roughly the same steps as those used in the original proof of Claim 3. We start by proving that after $N_A^{1+\epsilon}$ new balls are thrown in phase B, the fractional load of the low and middle bins must be at least $\frac{1.01}{c_2-1}$ (whp). To prove this, we take all bins that start as low and middle bins in phase B, say bins $B_1, \ldots, B_i$, and imagine a fictional *super* bin $\mathcal{L}$, which contains bins $B_1, \ldots, B_i$ and whose load is the sum of the loads of bins $B_1, \ldots, B_i$. Note that if the low and middle bins contain less than $\frac{1.01}{c_2-1}$ fractional load, then super bin $\mathcal{L}$ satisfies the same growth condition as described previously, with the same growth constant $\gamma = \frac{c_2-1}{6}$. Furthermore, one can again use Lemma 5 to show that within $N_A^{1+\epsilon}$ steps, $\mathcal{L}$ must have fractional load at least $\frac{1.01}{c_2-1}$ (whp). Therefore, the low and middle bins must contain fractional load at least $\frac{1.01}{c_2-1}$ (whp) after $N_A^{1+\epsilon}$ balls have been thrown.

As a result, the high bins must contain fractional load at most $(1 - \frac{1.01}{c_2-1})$ after $N_A^{1+\epsilon}$ new balls have been thrown (whp), and therefore, we can show any high bin H, satisfies a shrinkage condition with shrinkage constant $\gamma = (\frac{1}{2} + (1 - \frac{1.01}{c_2-1})^{k-1})$, which can be arbitrarily close to $\frac{1}{2}$. At this point in our process, it is not hard to show that each high bin still contains at most $1.01\frac{c_1}{n}$ fractional load (whp), and must become a middle bin before $\hat{c}N_A^{1+\epsilon}$ more balls are thrown (whp), for sufficiently large $\hat{c}$. The latter statement must be true because if H does not become a middle bin, then H must satisfy the shrinkage condition for $\hat{c}N_A^{1+\epsilon}$ steps. One can then use the same analysis from Lemma 2 to conclude the fractional load of bin H must decrease to $\frac{1}{2n}$ (whp). However, following the same argument as in last paragraph of the original proof of Claim 3, we can show that this cannot happen (whp), and bin H must become a middle bin (whp). Finally, a simple union bound implies that all high bins must become middle bins (whp) sometime before $O(N_A^{1+\epsilon})$ new balls are thrown in phase B.

L Experiments

L.1 Preferential Attachment Graphs

In order to validate the linked decomposition idea, we apply the following simple algorithm to graphs generated according to $\mathcal{PA}(4)$: *Fix the number c of components (colors). For each new vertex, if it connects two components that are not yet connected, add it to both of them, otherwise to the smallest one. Continue until all components are connected.*

The figure below shows the dependence of the final size of the graph on the number of colors; as expected, it is quadratic.

Below we plot the largest (smallest, average) component size as a function of the number of components; we note occasional large deviations from the expected linear growth of the largest component.

These experiments allow $b = 4$ (a node to belong to up to 4 components). If we decrease b to 2 (the standard value used in the rest of the paper) we notice that, for the same graph, the maximum

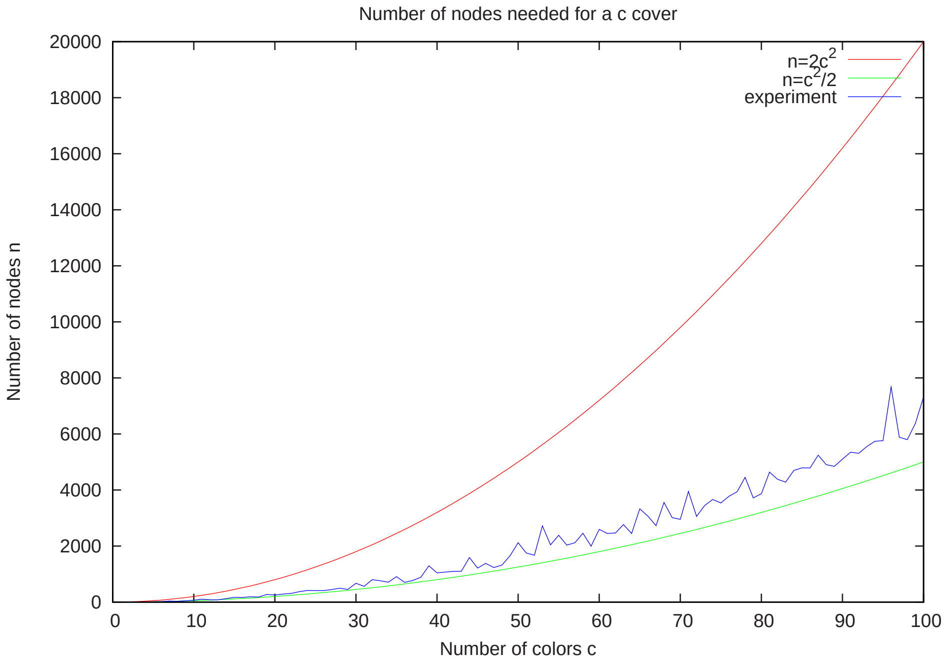

Figure 2: $\beta = 4, m = 4$

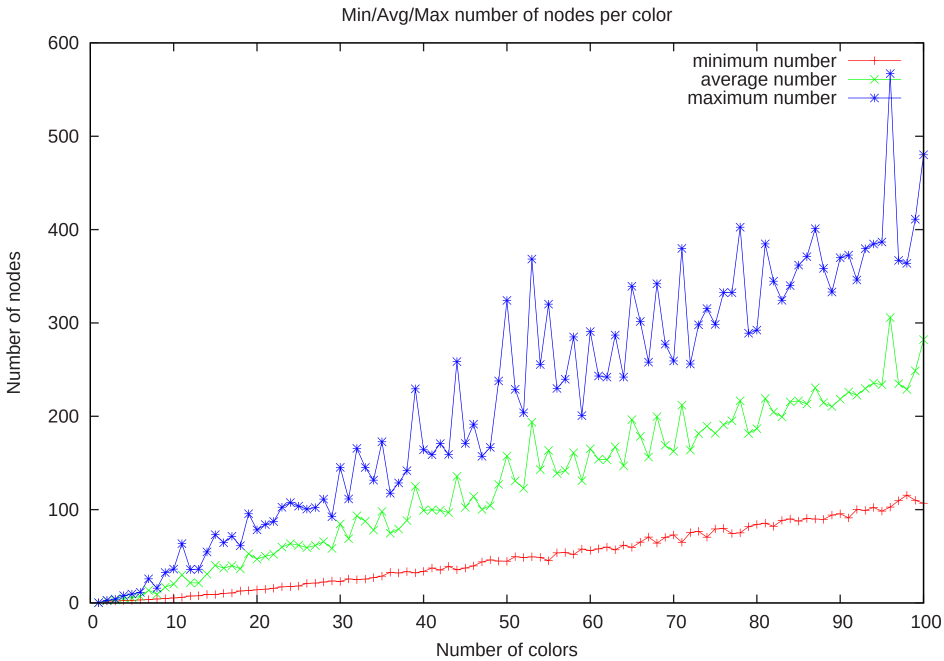

Figure 3:

number of components drops by a factor of approximately two, and the maximum number of nodes per component is increased also by a factor of about two.

L.2 Internet Graphs

We also applied our ideas to two actual Internet graphs, one showing connections between autonomous systems (the "BGP graph), and one the connections between Internet routers (the "router graph), both obtained from the CAIDA website [1].

To run our algorithms on these graphs, we first preprocess them by repeatedly removing degree-one nodes. We approximate the (unknown) presumed arrival order of $\mathcal{PA}(m)$ by the order of decreasing degrees. One problem arises: Sometimes the next node is not connected to previous nodes. For this reason we maintain a priority queue containing such nodes, and remove them when a neighbor has been processed.

The BGP graph has, after the removal of leaves, 12358 nodes. With $b = 4$ it can be decomposed into 56 linked components, and with $b = 2$ to 35. Similarly, the router graph has, after the removal of leaves, 141,509 nodes. With $b = 4$ it can be decomposed into 215 components, but with $b = 2$ to only 93.

Interestingly, if we apply the same algorithm (sort by degree and then process, using a priority queue) to $\mathcal{PA}(m)$ data, after the end of the generation process, we get slightly worse results for small graphs, but slightly better results for larger graphs.

L.3 Routing

We simulated routing on the decompositions of real graphs that we obtained, by selecting 10,000 random source-destination pairs and measuring the stretch and congestion (number of times each node was used). The stretch results for both BGP and router graphs (below) show that the stretch very rarely exceeds 3.

But these experiments also showed that our scheme has a problem with congestion: In both cases, as roughly 10% of the traffic was directed *through one particular node!* Remedies are discussed briefly in the main paper's last section.

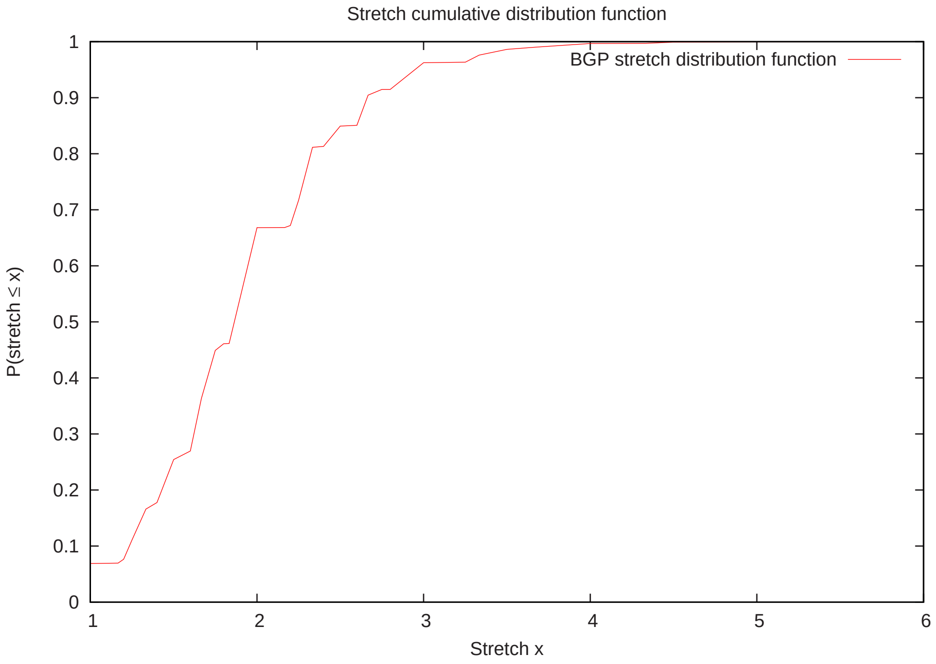

Figure 4: Routing experiment on BGP graph

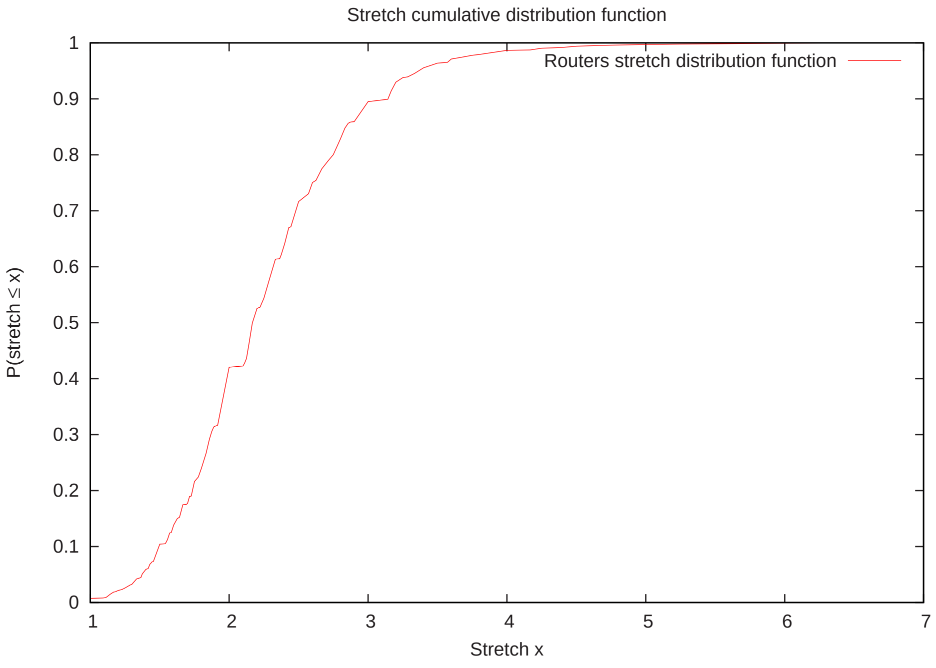

Figure 5: Routing experiment on routers graph