Combinatorics, Probability and Computing (2005) 14, 697–721. © 2005 Cambridge University Press
doi:10.1017/S0963548305006930 Printed in the United Kingdom

Degree Distribution of Competition-Induced Preferential Attachment Graphs

N. BERGER[1], C. BORGS[1], J. T. CHAYES[1],
R. M. D'SOUZA[1] and R. D. KLEINBERG[2†]

[1]Microsoft Research, One Microsoft Way, Redmond WA 98052, USA
(e-mail: berger@its.caltech.edu, borgs@microsoft.com,
jchayes@microsoft.com, raissa@alum.mit.edu)

[2]MIT CSAIL, 77 Massachusetts Ave, Cambridge MA 02139, USA
(e-mail: rdk@math.mit.edu)

Received 25 May 2004; revised 8 February 2005

For Béla Bollobás on his 60th birthday

We introduce a family of one-dimensional geometric growth models, constructed iteratively by locally optimizing the trade-offs between two competing metrics, and show that this family is equivalent to a family of preferential attachment random graph models with upper cut-offs. This is the first explanation of how preferential attachment can arise from a more basic underlying mechanism of local competition. We rigorously determine the degree distribution for the family of random graph models, showing that it obeys a power law up to a finite threshold and decays exponentially above this threshold.

We also rigorously analyse a generalized version of our graph process, with two natural parameters, one corresponding to the cut-off and the other a 'fertility' parameter. We prove that the general model has a power-law degree distribution up to a cut-off, and establish monotonicity of the power as a function of the two parameters. Limiting cases of the general model include the standard preferential attachment model without cut-off and the uniform attachment model.

1. Introduction

1.1. Network growth models

This paper is dedicated, with great affection and admiration, to Béla Bollobás on the occasion of his 60th birthday. Two of us (C. B. and J. T. C.) are privileged to count Béla among our dearest friends. And all of us have been inspired by his pioneering work on graph processes in general, and scale-free graphs in particular. We use the opportunity of this birthday volume to provide complete proofs of results on a new graph model, first announced in [6].

† Supported by a Fannie and John Hertz Foundation Fellowship.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

There is currently tremendous interest in understanding the mathematical structure of networks – especially as we discover the pervasiveness of network structures in natural and engineered systems. Much recent theoretical work has been motivated by measurements of real-world networks, indicating that they have certain 'scale-free' properties, such as a power-law distribution of degrees. For the Internet graph, in particular, both the graph of routers and the graph of autonomous systems (AS) seem to obey power laws [15, 16]. However, these observed power laws hold only for a limited range of degrees, presumably due to physical constraints and the finite size of the Internet.

Many random network growth models have been proposed which give rise to power-law degree distributions. Most of these models rely on a small number of basic mechanisms, mainly preferential attachment[1] [20, 4] or copying [18], extending ideas known for many years [13, 21, 23, 22] to a network context. Variants of the basic preferential attachment mechanism have also been proposed, and some of these lead to changes in the values of the exponents in the resulting power laws. For extensive reviews of work in this area, see Albert and Barabási [2], Dorogovtsev and Mendes [12], and Newman [19]; for a survey of the relatively limited amount of mathematical work see [7]. Most of this work concerns network models without reference to an underlying geometric space. Nor do most of these models allow for heterogeneity of nodes, or address physical constraints on the capacity of the nodes. Thus, while such models may be quite appropriate for geometry-free networks, such as the web graph, they do not seem to be ideally suited to the description of other observed networks, *e.g.*, the Internet graph.

In this paper, instead of assuming preferential attachment, we show that it can arise from a more basic underlying process, namely competition between opposing forces. The idea that power laws can arise from competing effects, modelled as the solution of optimization problems with complex objectives, was proposed originally by Carlson and Doyle [10]. Their 'highly optimized tolerance' (HOT) framework has reliable design as a primary objective. Fabrikant, Koutsoupias and Papadimitriou (FKP) [14] introduce an elegant network growth model with such a mechanism, which they called 'heuristically optimized trade-offs'. As in many growth models, the FKP network is grown one node at a time, with each new node choosing a previous node to which it connects. However, in contrast to the standard preferential attachment types of models, a key feature of the FKP model is the underlying geometry. The nodes are points chosen uniformly at random from some region, for example a unit square in the plane. The trade-off is between the geometric consideration that it is desirable to connect to a nearby point, and a networking consideration, that it is desirable to connect to a node that is 'central' in the network as a graph. Centrality is measured by using, for example, the graph distance to the initial node. The model has a tunable, but fixed, parameter, which determines the relative weights given to the geometric distance and the graph distance.

The suggestion that competition between two metrics could be an alternative to preferential attachment for generating power-law degree distributions represents an important paradigm shift. Though FKP introduced this paradigm for network growth,

[1] As Aldous [3] points out, proportional attachment may be a more appropriate name, stressing the linear dependence of the attractiveness on the degree.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

and FKP networks have many interesting properties, the resulting distribution is not a power law in the standard sense [5]. Instead the overwhelming majority of the nodes are leaves (degree one), and a second substantial fraction heavily connected 'stars' (hubs), producing a node degree distribution which has clear bimodal features.[2]

Here, instead of directly producing power laws as a consequence of competition between metrics, we show that such competition can give rise to a preferential attachment mechanism, which in turn gives rise to power laws. Moreover, the power laws we generate have an upper cut-off, which is more realistic in the context of many applications.

1.2. Overview of competition-induced preferential attachment

We begin by formulating a general competition model for network growth. Let $x_0, x_1, \ldots, x_t$ be a sequence of random variables with values in some space Λ. We think of the points $x_0, x_1, \ldots, x_t$ arriving one at a time according to some stochastic process. For example, we typically take Λ to be a compact subset of $\mathbb{R}^d$, x_0 to be a given point, say the origin, and $x_1, \ldots, x_t$ to be i.i.d. uniform on Λ. The network at time t will be represented by a graph, $G(t)$, on $t+1$ vertices, labelled $0, 1, \ldots, t$, and at each time step, the new node attaches to one or several nodes in the existing network. For simplicity, here we assume that each new node connects to a single node, resulting in $G(t)$ being a tree.

Given $G(t-1)$, the new node, labelled t, attaches to that node j in the existing network that minimizes a certain cost function representing the trade-off of two competing effects, namely connection or startup cost, and routing or performance cost. The connection cost is represented by a metric, $g_{ij}(t)$, on $\{0, \ldots, t\}$ which depends on $x_0, \ldots, x_t$, but not on the current graph $G(t-1)$, while the routing cost is represented by a function, $h_j(t-1)$, on the nodes which depends on the current graph, but not on the physical locations $x_0, \ldots, x_t$ of the nodes $0, \ldots, t$. This leads to the cost function

$$c_t = \min_j [\alpha g_{tj}(t) + h_j(t-1)], \tag{1.1}$$

where α is a constant which determines the relative weighting between connection and routing costs. We think of the function $h_j(t-1)$ as measuring the centrality of the node j; for simplicity, we take it to be the hop distance along the graph $G(t-1)$ from j to the root 0.

To simplify the analysis of the random graph process, we will assume that nodes always choose to connect to a point which is closer to the root, *i.e.*, they minimize the cost function

$$\tilde{c}_t = \min_{j:\|x_j\|<\|x_t\|} [\alpha g_{tj}(t) + h_j(t-1)], \tag{1.2}$$

where $\|\cdot\|$ is an appropriate norm.

In the original FKP model, Λ is a compact subset of $\mathbb{R}^2$, say the unit square, and the points x_i are independently uniformly distributed on Λ. The cost function is of the form

[2] In simulations of the FKP model, this can be clearly discerned by examining the probability distribution function (pdf); for the system sizes amenable to simulations, it is less prominent in the cumulative distribution function (cdf).

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

(1.1), with $g_{ij} = d_{ij}$, the Euclidean metric (modelling the cost of building the physical transmission line), and $h_j(t)$ is the hop distance along the existing network $G(t)$ from j to the root. A rigorous analysis of the degree distribution of this two-dimensional model was given in [5], and the analogous one-dimensional problem was treated in [17].

Our model is defined as follows.

Definition 1. (Border toll optimization process) Let $x_0 = 0$, and let $x_1, x_2, \ldots$ be i.i.d., uniformly at random in the unit interval $\Lambda = [0, 1]$, and let $G(t)$ be the following process. At $t = 0$, $G(t)$ consists of a single vertex 0, the root. Let $h_j(t)$ be the hop distance to 0 along $G(t)$, and let $g_{ij}(t) = n_{ij}(t)$ be the number of existing nodes between x_i and x_j at time t, which we refer to as the *jump cost* of i connecting to j. Given $G(t-1)$ at time $t-1$, a new vertex, labelled t, attaches to the node j which minimizes the cost function (1.2). Furthermore, if there are several nodes j that minimize this cost function and satisfy the constraint, we choose the one whose position x_j is nearest to x_t. The process so defined is called the *border toll optimization process* (BTOP).

As in the FKP model, the routing cost is just the hop distance to the root along the existing network. However, in our model the connection cost metric measures the number of 'borders' between two nodes: hence the name BTOP. Note the correspondence to the Internet, where the principal connection cost is related to the number of AS domains crossed – representing, *e.g.*, the overhead associated with BGP, monetary costs of peering agreements, *etc*. In order to facilitate a rigorous analysis of our model, we took the simpler cost function (1.2), so that the new node always attaches to a node to its left.

It is interesting to note that the ratio of the BTOP connection cost metric to that of the one-dimensional FKP model is just the local density of nodes: $n_{ij}/d_{ij} = \rho_{ij}$. Thus the transformation between the two models is equivalent to replacing the constant parameter α in the FKP model with a variable parameter $\alpha_{ij} = \alpha\rho_{ij}$ which changes as the network evolves in time. That α_{ij} is proportional to the local density of nodes in the network reflects a model with an increase in cost for local resources that are scarce or in high demand. Alternatively, it can be thought of as reflecting the economic advantages of being first to market.

Somewhat surprisingly, the BTOP is equivalent to a special case of the following process, which closely parallels the preferential attachment model and makes no reference to any underlying geometry.

Definition 2. (Generalized preferential attachment with fertility and aging) Let A_1, A_2 be two positive integer-valued parameters. Let $G(t)$ be the following Markov process, whose states are finite rooted trees in which each node is labelled either *fertile* or *infertile*. At time $t = 0$, $G(t)$ consists of a single fertile vertex. Given the graph at time t, the new graph is formed in two steps: first, a new vertex, labelled $t+1$ and initialized as infertile, connects to an old vertex j with probability zero if j is infertile, and with probability

$$Pr(t+1 \to j) = \frac{\min\{d_j(t), A_2\}}{W(t)} \tag{1.3}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

if j is fertile. Here, $d_j(t)$ is equal to 1 plus the out-degree of j, and $W(t) = \sum_j' \min\{d_j(t), A_2\}$ with the sum running over fertile vertices only. We refer to vertex $t+1$ as a child of j. If after the first step, j has more than $A_1 - 1$ infertile children, one of them, chosen uniformly at random, becomes fertile. The process so defined is called a *generalized preferential attachment process with fertility threshold A_1 and aging threshold A_2*.

The special case $A_1 = A_2$ is called the *competition-induced preferential attachment process* with parameter A_1.

The last definition is motivated by the following theorem, to be proved in Section 2. To state it, we define a graph process as a random sequence of graphs $G(0), G(1), G(2), \ldots$ on the vertex sets $\{0\}, \{0,1\}, \{0,1,2\}, \ldots$, respectively.

Theorem 1.1. *As a graph process, the border toll optimization process has the same distribution as the competition-induced preferential attachment process with parameter $A = \lceil \alpha^{-1} \rceil$.*

Certain other limiting cases of the generalized preferential attachment process are worth noting. If $A_1 = 1$ and $A_2 = \infty$, we recover the standard model of preferential attachment as considered in [20, 4]. If $A_1 = 1$ and A_2 is finite, the model is equivalent to the standard model of preferential attachment with a cut-off. On the other hand, if $A_1 = A_2 = 1$, we get a uniform attachment model.

The degree distribution of our random trees is characterized by the following theorem, which asserts that almost surely (a.s.) the fraction of vertices having degree k converges to a specified limit q_k, and moreover that this limit obeys a power law for $k < A_2$, and decays exponentially above A_2.

Theorem 1.2. *Let A_1, A_2 be positive integers. Consider the generalized preferential attachment process with fertility parameter A_1 and aging parameter A_2. Let $N_0(t)$ be the number of infertile vertices at time t, and let $N_k(t)$ be the number of fertile vertices with $k-1$ children at time t, $k \geqslant 1$. Then we have the following.*

(1) *There are numbers $q_k \in [0,1]$ such that, for all $k \geqslant 0$,*

$$\frac{N_k(t)}{t+1} \to q_k \quad \text{a.s., as } t \to \infty. \tag{1.4}$$

(2) *There exists a number $w = w(A_1, A_2) \in [0,2]$ such that the q_k are determined by the following equations:*

$$q_i = \left(\prod_{k=2}^{i} \frac{k-1}{k+w} \right) q_1 \quad \text{if} \quad 1 \leqslant i \leqslant A_2, \tag{1.5}$$

$$q_i = \left(\frac{A_2}{A_2 + w} \right)^{i - A_2} q_{A_2} \quad \text{if} \quad i > A_2, \tag{1.6}$$

$$1 = \sum_{i=0}^{\infty} q_i, \quad \text{and} \quad q_0 = \sum_{i=1}^{\infty} q_i \min\{i-1, A_1 - 1\}.$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

(3) *There are positive constants* c_1 *and* C_1*, independent of* A_1 *and* A_2*, such that*

$$c_1 k^{-(w+1)} < q_k/q_1 < C_1 k^{-(w+1)} \tag{1.7}$$

for $1 \leqslant k \leqslant A_2$.

(4) *If* $A_1 = A_2$*, the parameter* w *is equal to* 1*, and for general* A_1 *and* A_2*,* w *decreases with increasing* A_1*, and increases with increasing* A_2.

Equation (1.7) clearly defines a power-law degree distribution with exponent $\gamma = w + 1$ for $k \leqslant A_2$. Note that for measurements of the Internet the value of the exponent for the power law is $\gamma \approx 2$. In our border toll optimization model, where $A_1 = A_2$, we recover $\gamma = 2$.

The convergence claim of Theorem 1.2 is proved using a novel method which we believe is one of the main technical contributions of this work. For preferential attachment models which have been analysed in the past [1, 8, 9, 11], the convergence was established using the Azuma–Hoeffding martingale inequality. To establish the bounded-differences hypothesis required by that inequality, those proofs employed a clever coupling of the random decisions made by the various edges, such that the decisions made by an edge e only influence the decisions of subsequent edges which choose to imitate e's choices. A consequence of this coupling is that if e made a different decision, it would alter the degrees of only finitely many vertices. This in turn allows the required bounded-differences hypothesis to be established. No such approach is available for our models, because the coupling fails. The random decisions made by an edge e may influence the time at which some node v crosses the fertility or aging threshold, which thereby exerts a subtle influence on the decisions of *every* future edge, not only those which choose to imitate e.

Instead we introduce a new approach based on the second-moment method. The argument establishing the requisite second-moment upper bound is quite subtle; it depends on a computation involving the eigenvalues of a matrix describing the evolution of the degree sequence in a continuous-time version of the model.

2. Equivalence of the two models

2.1. Basic properties of the border toll optimization process

In this section we will turn to the BTOP defined in the introduction, establishing some basic properties which will enable us to prove that it is equivalent to the competition-induced preferential attachment model. In order to avoid complications we exclude the case that some of the x_is are identical, an event that has probability zero. We say that $j \in \{0, 1 \ldots, t\}$ lies to the right of $i \in \{0, 1 \ldots, t\}$ if $x_i < x_j$, and we say that j lies directly to the right of i if $x_i < x_j$ but there is no $k \in \{1, \ldots, t\}$ such that $x_i < x_k < x_j$. In a similar way, we say that j is the first vertex with a certain property to the right of i if j has that property and there exists no $k \in \{1, \ldots, t\}$ such that $x_i < x_k < x_j$ and k has the property in question. Similar notions apply with 'left' in place of 'right'.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

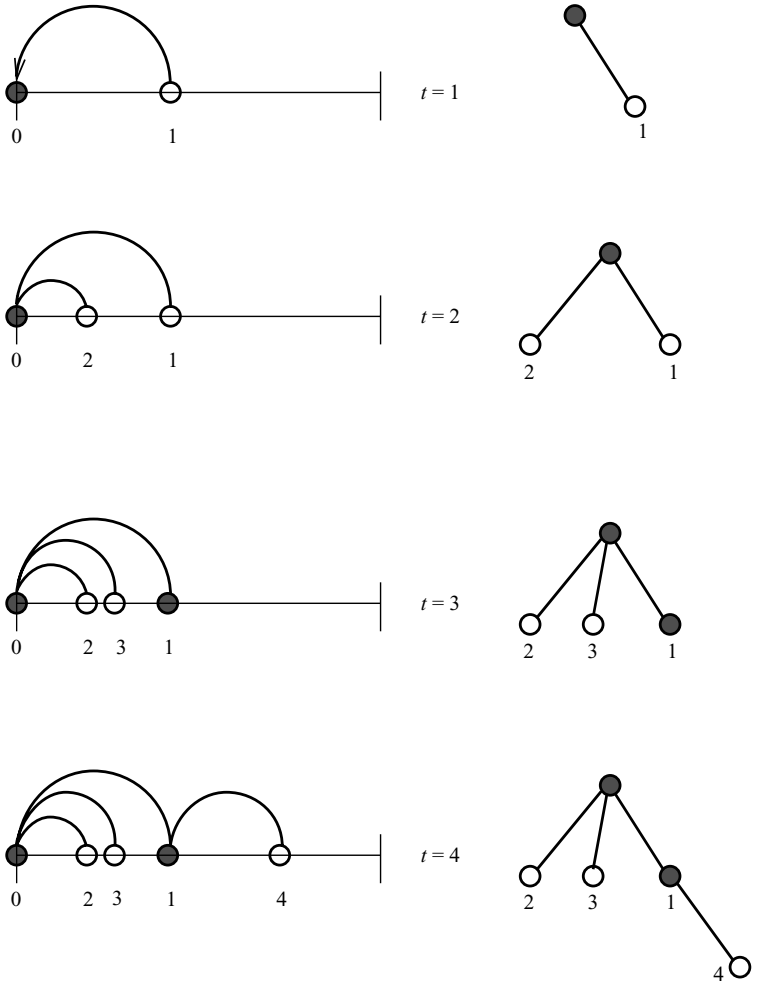


Figure 1. A sample instance of BTOP for $\alpha = 1/3, A = 3$, showing the process on the unit interval (on the left), and the resulting tree (on the right). Fertile vertices are shaded, infertile ones are not. Note that vertex 1 became fertile at $t = 3$.

Definition 3. A vertex i is called *fertile at time* t if a hypothetical new point arriving at time $t+1$ and landing directly to the right of x_i would attach itself to the node i. Otherwise i is called *infertile* at time t.

This definition is illustrated in Figure 1.

Lemma 2.1. *Let* $0 < \alpha < \infty$, *let* $A = \lceil \alpha^{-1} \rceil$, *and let* $0 < t < \infty$. *Then we have the following.*

(i) *The node* 0 *is fertile at time* t.

(ii) *Let* i *be fertile at time* t. *If* i *is the rightmost fertile vertex at time* t (*case 1*), *let* ℓ *be the number of infertile vertices to the right of* i. *Otherwise* (*case 2*), *let* j *be the next*

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

fertile vertex to the right of i, and let $\ell = n_{ij}(t)$. *Then* $0 \leqslant \ell \leqslant A-1$, *and the* ℓ *infertile vertices located directly to the right of i are children of i. In case 2, if* $h_j > h_i$, *then j is a fertile child of i and* $\ell = A-1$. *As a consequence, the hop count between two consecutive fertile vertices never increases by more than* 1 *as we move to the right, and if it increases by* 1, *there are* $A-1$ *infertile vertices between the two fertile ones.*

(iii) *Assume that the new vertex at time* $t+1$ *lands between two consecutive fertile vertices i and j, and let* $\ell = n_{ij}(t)$. *Then* $t+1$ *becomes a child of i. If* $\ell + 1 < A$, *the new vertex is infertile at time* $t+1$, *and the fertility of all old vertices is unchanged. If* $\ell + 1 = A$ *and the new vertex lies directly to the left of j, the new vertex is fertile at time* $t+1$ *and the fertility of the old vertices is unchanged. If* $\ell + 1 = A$ *and the new vertex does not lie directly to the left of j, the new vertex is infertile at time* $t+1$, *the vertex directly to the left of j becomes fertile, and the fertility of all other vertices is unchanged.*

(iv) *If* $t+1$ *lands to the right of the rightmost fertile vertex at time t, the statements in* (iii) *hold with j replaced by the right endpoint of the interval* $[0,1]$, *and* $n_{ij}(t)$ *replaced by the number of vertices to the right of i.*

(v) *If i is fertile at time t, it is still fertile at time* $t+1$.

(vi) *If i has k children at time t, the* $\ell = \min\{A-1, k\}$ *leftmost of them are infertile at time t, and any others are fertile.*

Proof. The proof is straightforward but lengthy. We include the details of the argument here for completeness.

Statement (i) is trivial, statement (v) follows immediately from (iii) and (iv), and (vi) follows immediately from (ii). So we are left with (ii)–(iv). We proceed by induction on t. If (ii) holds at time t, and (iii) and (iv) hold for a new vertex arriving at time $t+1$, (ii) clearly also holds at time $t+1$. We therefore only have to prove that (ii) at time t implies (iii) and (iv) for a new vertex arriving at time $t+1$.

Assume thus that (ii) holds at time t. At time $t+1$, a new vertex arrives, and falls directly to the right of some vertex k. Let i be the nearest vertex to the left of k that was fertile at time t (if k is fertile at time t, we set $i = k$) and let j be the nearest vertex to the right of i that was fertile at time t (we assume for the moment that i is not the rightmost fertile vertex at time t), let ℓ be the number of vertices between i and j at time t.

Let us first prove that the vertex $t+1$ connects to i. If $i = k$, this is obvious, since i is fertile at time t. We may therefore assume that $k \neq i$. For the new vertex $t+1$, the cost of connecting to the vertex i is then equal to $\alpha(n_{ik}(t)+1)$. Let us first compare this cost to the cost of connecting to a fertile vertex i' to the left of i. Let $i_0 = i'$, let $i_s = i$, and let $i_1, \ldots, i_{s-1}$ be the fertile vertices between i' and i, ordered from left to right. If $h_{i_{m-1}} < h_{i_m}$, we use the inductive assumption (ii) to conclude that the number of infertile vertices between i_{m-1} and i_m is equal to $A-1$, and $h_{i_{m-1}} = h_{i_m} - 1$. A decrease of q in the hop cost is therefore accompanied by an increase in the jump cost of at least $\alpha A q \geqslant q$. As a consequence, it never pays to connect to a fertile vertex i' to the left of i. The cost of connecting to an infertile vertex to the left of i is even higher, since the hop count of an infertile vertex is at best equal to the hop count of the next fertile vertex to the right. We therefore only have to consider the connection cost to some of the infertile children of i. But again, the hop count is worse by 1 when compared to the hop count of i, and

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

the jump cost is at best reduced by $(A-1)\alpha < 1$, proving that the cost of connecting to i is minimal.

To discuss the fertility of the vertices in the graph $G(t+1)$, we need to consider the arrival of a second vertex, labelled $t+2$. If $t+2$ falls to the left of $t+1$, it will face an optimization problem that has not been changed by the arrival of the vertex $t+1$, implying that the fertility of the vertices to the left of $t+1$ is unchanged. If $t+2$ falls to the right of j, the cost of connecting to j or one of the vertices to the right of j is the same as before, and the cost of connecting to a vertex to the left of j is at best equal (the cost of connecting to any vertex to the left of $t+1$ is in fact higher, because of the additional cost of jumping over the vertex $t+1$). Therefore, the vertex $t+2$ will still prefer to connect to either j or one of the vertices to the right of j, implying that the fertility of the vertices to the right of j has not changed at all. We are therefore left with analysing the case where $t+2$ falls between $t+1$ and j. Again, the vertex $t+2$ will prefer i over any vertex to the left of i (the cost analysis is the same as the one used for $t+1$ above), so we just have to compare the costs of connecting to the different vertices between i and j. If $\ell + 1 < A$, this will again imply that $t+2$ connects to i; but if $\ell + 1 = A$, the vertex $t+2$ will only connect to i if it does not fall to the right of the rightmost of the now $\ell + 1$ vertices between i and j. If it falls to the right of this vertex, it will be as expensive to connect to the rightmost of the now $\ell + 1$ vertices between i and j as it is to connect to i. Recalling our convention of connecting to the nearest vertex to the left if there is a tie in costs, this proves that now $t+2$ connects to the rightmost vertex between i and j, implying that this vertex is fertile.

The above considerations prove the fertility statements in (iii), and thus complete the proof of (iii). The case where i is the rightmost fertile vertex at time t is similar (in fact, it is slightly easier since it involves fewer cases), and leads to the proof of (iv). This completes the proof of Lemma 2.1. □

2.2. Proof of Theorem 1.1

In the BTOP, note that our cost function

$$\min_j [\alpha n_{tj}(t) + h_j(t-1)], \tag{2.1}$$

and hence the graph $G(t)$, only depends on the order of the vertices $x_0, \ldots, x_t$, and not on their actual positions in the interval $[0,1]$. Let $\vec{\pi}(t)$ be the permutation of $\{0,1,\ldots,t\}$ which orders the vertices $x_0, \ldots, x_t$ from left to right, so that

$$x_0 = x_{\pi_0(t)} < x_{\pi_1(t)} < \cdots < x_{\pi_t(t)}. \tag{2.2}$$

(Recall that the vertices $x_0, x_1, \ldots, x_t$ are pairwise distinct with probability one.) Note that $\vec{\pi}(t)$ and $\vec{\pi}(t+1)$ are related as follows: there exists $i_0 \in \{1,2,\ldots,t+1\}$ such that

$$\pi_i(t+1) = \begin{cases} \pi_i(t) & \text{if } i < i_0, \\ t+1 & \text{if } i = i_0, \\ \pi_{i-1}(t) & \text{if } i > i_0. \end{cases} \tag{2.3}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

Informally, the permutation $\vec{\pi}(t+1)$ is obtained by inserting the new element $t+1$ into the permutation $\vec{\pi}(t)$ in a random position i_0, where $x_{\pi_{i_0}(t)}$ is the left endpoint of the subinterval of $(0,1)$ into which x_{t+1} falls. The distribution of the random variable i_0 may be deduced as follows. Since $x_0 = 0$ and $x_1, x_2, \ldots, x_t$ are i.i.d., we know that, for all t, the permutation $\vec{\pi}(t)$ is uniformly distributed among permutations of $\{0, 1, \ldots, t\}$ which fix the element 0. This means that, conditioned on a given such permutation $\vec{\pi}(t)$, the permutation $\vec{\pi}(t+1)$ is uniformly distributed among all permutations related to $\vec{\pi}(t)$ by the transformation (2.3). In other words, i_0 is uniformly distributed in the set $\{1, 2, \ldots, t+1\}$.

With the help of Lemma 2.1, we now easily derive a description of the graph $G(t)$ which does not involve any optimization problem. To this end, let us consider a vertex i with ℓ infertile children at time t. If a new vertex falls into the interval directly to the right of i, or into one of the intervals directly to the right of an infertile child of i, it will connect to the vertex i. Since there is a total of $t+1$ intervals at time t, the probability that a vertex i with ℓ infertile children grows an offspring is $(\ell+1)/(t+1)$. By Lemma 2.1(vi), this number is equal to $\min\{A, k_i\}/(t+1)$, where $k_i - 1$ is the number of children of i. Note that fertile children do not contribute to this probability, since vertices falling into an interval directly to the right of a fertile child will connect to the child, not the parent.

Assume now that i did get a new offspring, and that it had $A-1$ infertile children at time t. Then the new vertex is either born fertile, or makes one of its infertile siblings fertile. Using the principle of deferred decisions, we may assume that with probability $1/A$ the new vertex becomes fertile, and with probability $(A-1)/A$ an old one, chosen uniformly at random among the $A-1$ candidates, becomes fertile.

We thus have shown that the solution $G(t)$ of the optimization problem (2.1) can alternatively be described by the competition-induced preferential attachment model with parameter A.

3. Convergence of the degree distribution

3.1. Overview

To characterize the behaviour of the degree distribution, we will derive a recursion which governs the evolution of the vector $\vec{N}(t)$, whose components are the number of vertices of each degree, at the time when there are t nodes in the network. The conditional expectation of $\vec{N}(t+1)$ is given by an evolution equation of the form

$$\mathbb{E}\big(\vec{N}(t+1) - \vec{N}(t) \mid \vec{N}(t)\big) = M(t)\vec{N}(t),$$

where $M(t)$ depends on t through the random variable $W(t)$ introduced in Definition 2. Owing to the randomness of the coefficient matrix $M(t)$, the analysis of this evolution equation is not straightforward. We avoid this problem by introducing a continuous-time process, with time parameter τ, which is equivalent to the original discrete-time process up to a (random) reparametrization of the time coordinate. The evolution equation for the conditional expectations in the continuous-time process involves a coefficient matrix M that is not random and does not depend on τ. We will first prove that the *expected* degree distribution in the continuous-time model converges to a scalar multiple of the eigenvector $\hat{p}$ of M associated with the largest eigenvalue w. This is followed by the much

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

more difficult proof that the *empirical* degree distribution converges a.s. to the same limit. Finally, we translate this continuous-time result into a rigorous convergence result for the original discrete-time system.

3.2. Notation

Let A be any integer greater than or equal to $\max(A_1, A_2)$. Let $N_0(t)$ be the number of infertile vertices at (discrete) time t, and, for $k \geqslant 1$, let $N_k(t)$ be the number of fertile vertices with $k-1$ children at time t. Let $\widetilde{N}_A(t) = N_{\geqslant A}(t) = \sum_{k \geqslant A} N_k(t)$, and $\widetilde{N}_k(t) = N_k(t)$ if $k < A$. The combined attractiveness of all vertices is denoted by $W(t) = \sum_{k=1}^{A} \min\{k, A_2\} \widetilde{N}_k(t)$. Let $n_k(t) = \frac{1}{t+1} N_k(t)$ and $\tilde{n}_k(t) = \frac{1}{t+1} \widetilde{N}_k(t)$. Finally, the vectors $(\widetilde{N}_k(t))_{k=1}^{A}$ and $(\tilde{n}_k(t))_{k=1}^{A}$ are denoted by $\widetilde{N}(t)$ and $\tilde{n}(t)$ respectively. Note that the index k runs from 1 to A, not 0 to A.

3.3. Evolution of the expected value

From the definition of the generalized preferential attachment model, it is easy to derive the probabilities for the various alternatives which may happen upon the arrival of the $(t+1)$st node.

- With probability $A_2 \widetilde{N}_A(t)/W(t)$, it attaches to a node of degree $\geqslant A$. This increments $\widetilde{N}_1$, and leaves $\widetilde{N}_A$ and all $\widetilde{N}_j$ with $1 < j < A$ unchanged.
- With probability $\min(A_2, k)\widetilde{N}_k(t)/W(t)$, it attaches to a node of degree k, where $1 \leqslant k < A$. This increments $\widetilde{N}_{k+1}$, decrements $\widetilde{N}_k$, increments $\widetilde{N}_0$ or $\widetilde{N}_1$ depending on whether $k < A_1$ or $k \geqslant A_1$, and leaves all other $\widetilde{N}_j$ with $j < A$ unchanged.

It follows that the discrete-time process $(\widetilde{N}_k(t))_{k=0}^{A}$ at time t is equivalent to the state of the following continuous-time stochastic process $(\widehat{N}_k(\tau))_{k=0}^{A}$ at the random stopping time $\tau = \tau_t$ of the t th event.

- With rate $A_2 \widehat{N}_A(\tau)$, $\widehat{N}_1$ increases by 1.
- For every $0 < k < A$, with rate $\widehat{N}_k(\tau) \min(k, A_2)$, the following happens:

$$\widehat{N}_k \to \widehat{N}_k - 1, \quad \widehat{N}_{k+1} \to \widehat{N}_{k+1} + 1, \quad \widehat{N}_{g(k)} \to \widehat{N}_{g(k)} + 1$$

  where $g(k) = 0$ for $k < A_1$ and $g(k) = 1$ otherwise.

Note that the above rules need to be modified if $A_1 = 1$. Here the birth of a child of a degree-one vertex does not change the net number of fertile degree-one vertices, N_1.

Let M be the following $A \times A$ matrix:

$$M_{i,j} = \begin{cases} -1 & \text{if } i = j = 1 < A_1, \\ -\min(j, A_2) & \text{if } 2 \leqslant i = j \leqslant A-1, \\ \min(j, A_2) & \text{if } 2 \leqslant i = j+1 \leqslant A, \\ \min(j, A_2) & \text{if } i = 1 \text{ and } j \geqslant \max(A_1, 2), \\ 0 & \text{otherwise.} \end{cases} \tag{3.1}$$

Then, for every $\tau > \sigma$, the conditional expectation of the vector $\widehat{N}(\tau) = (\widehat{N}_k(\tau))_{k=1}^{A}$ is given by

$$\mathbb{E}\big(\widehat{N}(\tau) \mid \widehat{N}(\sigma)\big) = e^{(\tau-\sigma)M} \widehat{N}(\sigma). \tag{3.2}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

It is easy to see that the matrix e^M has all positive entries, and therefore (by the Perron–Frobenius theorem) M has a unique eigenvector $\hat{p}$ of ℓ_1-norm 1 having all positive entries. Let w be the eigenvalue corresponding to $\hat{p}$. Then w is real, it has multiplicity 1, and it exceeds the real part of every other eigenvalue. Therefore, for every nonzero vector y with nonnegative entries,

$$\lim_{\tau\to\infty} e^{-\tau w} e^{\tau M} y = \langle \hat{a}, y\rangle \hat{p},$$

where $\hat{a}$ is the eigenvector of M^{T} corresponding to w, normalized so that $\langle \hat{a}, \hat{p}\rangle = 1$. Note that $\langle \hat{a}, y\rangle > 0$ because y is nonzero and nonnegative, and $\hat{a}$ is positive, again by Perron–Frobenius. Therefore, the vector $\mathbb{E}\big(e^{-\tau w}\widehat{N}(\tau)\big)$ converges to a positive scalar multiple of $\hat{p}$, say $\lambda\hat{p}$, as $\tau \to \infty$.

In order to prove concentration for the continuous-time model, we will prove that the difference $\widehat{N}_k(\tau)/q_k - \widehat{N}_j(\tau)/q_j$ has an exponential growth rate which is at most the real part of the second eigenvalue of M, which is strictly less than w, the growth rate of the individual terms $\widehat{N}_k(\tau)/q_k$ and $\widehat{N}_j(\tau)/q_j$. From this, we will conclude that the ratio $\widehat{N}_k(\tau)/\widehat{N}_j(\tau)$ converges almost surely to q_k/q_j, for all k and j, which in turn implies the convergence of the normalized degree sequence to the vector $(q_i)_{i=0}^{\infty}$.

In order to prove bounds on the growth rate of the differences $\widehat{N}_k(\tau)/q_k - \widehat{N}_j(\tau)/q_j$, we will need some auxiliary bounds involving the well-known standard birth process, to be defined below.

3.4. Standard birth process

We start with the definition of the standard birth process with rate ρ. The standard birth process was first introduced by Yule in 1924 [22], and is a special case of the well-known Yule Process, defined in that paper.

Definition 4. Let $\rho > 0$ and let $\{o_n\}_{n=1}^{\infty}$ be independent exponential random variables so that $\mathbb{E}(o_n) = \frac{1}{\rho} n^{-1}$. For $\tau \in [0,\infty)$, let $X_\tau = \min\{n \geqslant 1 : \sum_{k=1}^{n} o_k > \tau\}$. Then X is called the *standard birth process with rate* ρ.[3]

The standard birth process is connected to our discussion through the following easy claim.

Claim 3.1. *Let* $\|\widehat{N}(\tau)\| = \sum_{k=1}^{A} \widehat{N}_k(\tau)$. *Let* $T \geqslant 0$, *let* $x \geqslant y$, *and let* X *be a standard birth process with rate* 2. *Then* $\{\{X_\tau\}_{\tau\geqslant T} \mid X_T = x\}$ *stochastically dominates* $\{\{\|\widehat{N}(\tau)\|\}_{\tau\geqslant T} \mid \|\widehat{N}(T)\| = y\}$.

Proof. Let us start with the observation that $\sum_{k=1}^{n} o_k$ is the first time τ for which $X_\tau = n+1$. Let $\{r_n\}_{n=0}^{\infty}$ be i.i.d. exponential random variables with mean 1. Then $\sum_{k=1}^{n} o_k$ has the same distribution as $\sum_{k=0}^{n-1} r_k/(2k+2)$. The time τ_n at which the node n is born has

[3] The name 'standard birth process' is due to the fact that X_τ is equivalent to the following process. Start with one cell at time 0. At each time, every cell divides into two cells with rate ρ. Then X_τ is the number of cells at time τ.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

the same distribution as $\sum_{k=0}^{n-1} r_k / W(k)$, where $W(k)$ denotes the combined attractiveness of all nodes at the random time τ_k. The claim now follows from the observation that $W(k) \leqslant 2k+1 \leqslant 2k+2$. □

The main purpose of this section is the proof of the following claims.

Claim 3.2. *Let X be a standard birth process with rate ρ. Then X_τ is almost surely finite for every τ. Furthermore, there exists a constant $C_s = C_s(\rho)$ such that, for every $\tau_2 > \tau_1$, $x \geqslant 1$, and $k \geqslant 1$,*

$$\mathbf{P}\big(X_{\tau_2} > kxe^{\rho(\tau_2-\tau_1)} \mid X_{\tau_1} = x\big) < \frac{C_s}{x(k-1)^2}. \tag{3.3}$$

If, in addition, $\tau_2 - \tau_1 < 1$, then

$$\mathbf{P}\big(X_{\tau_2} - X_{\tau_1} > kx[e^{\rho(\tau_2-\tau_1)} - 1] \mid X_{\tau_1} = x\big) < \frac{C_s}{x(\tau_2-\tau_1)(k-1)^2}. \tag{3.4}$$

To see the finiteness of X_τ, we need to show that $\sum_{n=1}^{\infty} o_n = \infty$ a.s. This follows from the following simple argument. For every k, let

$$U_k = \sum_{j=2^k+1}^{2^{k+1}} o_j.$$

For $j \in [2^k+1, 2^{k+1}]$, with probability greater than $\frac{1}{2}$, $o_j > \frac{1}{\rho} 2^{-k-2}$. Therefore, $\mathbf{P}(U_k > \frac{1}{4\rho}) > \frac{1}{2}$. The random variables $\{U_k\}_{k=1}^{\infty}$ are independent, and therefore $\sum_{n=1}^{\infty} o_n \geqslant \sum_{k=1}^{\infty} U_k = \infty$ almost surely.

To see (3.3) and (3.4), we use the following lemma, which is proved in [22, Section II]. Since the proof is short and simple, we choose to include it for the sake of making the exposition more self-contained.

Lemma 3.3. (Yule, 1924) *For every $\tau > 0$ and every positive integer k, $\mathbb{E}(X_\tau^k) < \infty$. Furthermore,*

$$\mathbb{E}(X_\tau) = \exp(\rho\tau), \tag{3.5}$$

and

$$\operatorname{var}(X_\tau) = \exp(2\rho\tau) - \exp(\rho\tau). \tag{3.6}$$

In particular,

$$\operatorname{var}(X_\tau) = O(\exp(2\rho\tau)), \tag{3.7}$$

and for $\tau < 1$ there exists a constant $C_v = C_v(\rho)$ so that

$$\operatorname{var}(X_\tau) \leqslant C_v \tau. \tag{3.8}$$

Proof. An equivalent description of the standard birth process is the following. Let α be an exponential variable with expected value ρ, and let $\{G_t\}$ be a Poisson point process

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

with rate $\alpha e^{\rho t}$. Then $X_\tau = 1 + G_\tau$ has the same distribution as the standard birth process. To see this, all we need is to show that for every τ and n, the rate of the process $\{G_t\}$ at time τ conditioned on $X_\tau = n$ is ρn. Indeed,

$$\begin{aligned}\text{rate}(\tau|X_\tau = n) &= \frac{\int_0^\infty \alpha e^{\rho\tau}\mathbf{P}(X_\tau = n|\alpha)\frac{1}{\rho}e^{-\alpha/\rho}d\alpha}{\int_0^\infty \mathbf{P}(X_\tau = n|\alpha)\frac{1}{\rho}e^{-\alpha/\rho}d\alpha} \\ &= \frac{\int_0^\infty \alpha e^{\rho\tau}e^{-\frac{\alpha}{\rho}(\exp(\rho\tau)-1)}\left(\frac{\alpha}{\rho}(\exp(\rho\tau)-1)\right)^{(n-1)}((n-1)!)^{-1}\frac{1}{\rho}e^{-\alpha/\rho}d\alpha}{\int_0^\infty e^{-\frac{\alpha}{\rho}(\exp(\rho\tau)-1)}\left(\frac{\alpha}{\rho}(\exp(\rho\tau)-1)\right)^{(n-1)}((n-1)!)^{-1}\frac{1}{\rho}e^{-\alpha/\rho}d\alpha} \\ &= \rho n.\end{aligned}$$

Here the second equality follows from the fact that $X_\tau - 1$ is a Poisson variable with rate $\frac{\alpha}{\rho}(e^{\rho\tau} - 1)$, and the last equality follows by integration by parts.

From this we get that the distribution of X_τ is geometric with expected value $\exp(\rho\tau)$. To see this, we again use the fact that $X_\tau - 1$ is a Poisson variable with rate $\frac{\alpha}{\rho}(e^{\rho\tau} - 1)$, where α is an exponential variable with expectation ρ. Therefore, for every n,

$$\begin{aligned}\mathbf{P}(X_\tau = n+1) &= \int_0^\infty \mathbf{P}(X_\tau - 1 = n|\alpha)\frac{1}{\rho}e^{-\alpha/\rho}d\alpha \\ &= (n!)^{-1}\int_0^\infty e^{-\frac{\alpha}{\rho}(e^{\rho\tau}-1)}\left(\frac{\alpha}{\rho}(e^{\rho\tau}-1)\right)^n\frac{1}{\rho}e^{-\alpha/\rho}d\alpha \\ &= \left(1 - e^{-\rho\tau}\right)\mathbf{P}(X_\tau = n)\end{aligned}$$

where, again, the last step follows from integration by parts.

The relations (3.5) and (3.6) follow immediately, and (3.7) and (3.8) follow from (3.6). □

Proof of (3.3) and (3.4) in Claim 3.2. Equations (3.3) and (3.4) will follow from Chebyshev's inequality if we show that

$$\mathbb{E}(X_{\tau_2}|X_{\tau_1}) = X_{\tau_1}e^{\rho(\tau_2-\tau_1)} \tag{3.9}$$

and

$$\text{var}(X_{\tau_2}|X_{\tau_1}) = X_{\tau_1}O\left(e^{2\rho(\tau_2-\tau_1)}\right) \tag{3.10}$$

for $\tau_2 > \tau_1$, and

$$\text{var}(X_{\tau+\tau_1}|X_{\tau_1}) = O(\tau)\cdot X_{\tau_1} \tag{3.11}$$

for $\tau < 1$.

Equations (3.9), (3.10) and (3.11) follow from (respectively) (3.5), (3.7) and (3.8) and the fact that, conditioned on X_{τ_1}, the process $X_{\tau+\tau_1}$ is the sum of X_{τ_1} independent copies of X_τ. □

Remark. From now on we will always assume that $\rho = 2$. In particular, whenever we use the term 'standard birth process', it should be understood as 'standard birth process with rate 2'.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

3.5. Concentration of the continuous-time process

In order to show concentration of the degree distribution for the continuous-time process, we will first prove the following lemma. To state it, we observe for any b with $b^{\mathsf{T}}\hat{p} = 0$,

$$\|b^{\mathsf{T}} e^{(T-\tau)M}\|_\infty \leqslant \|b\|_\infty e^{(T-\tau)v'} \tag{3.12}$$

for some $v' < w$. Without loss of generality, we may assume that $v' > w/2$. Also, for a general vector b,

$$\|b^{\mathsf{T}} e^{(T-\tau)M}\|_\infty \leqslant \|b\|_\infty e^{(T-\tau)w}. \tag{3.13}$$

Lemma 3.4. *Let b be a vector in $\mathbb{R}^A$ with $\|b\|_\infty \leqslant 1$. Then there exists a constant $C < \infty$, such that, for all $T > 0$,*

$$\operatorname{var}\big(b^{\mathsf{T}}\widehat{N}(T)\big) < C\exp(2uT) \tag{3.14}$$

where $u = w$ if $b^{\mathsf{T}}\hat{p} \neq 0$, and $u = v'$ if $b^{\mathsf{T}}\hat{p} = 0$.

Proof. We use a martingale to bound the variance. Fix T, and let

$$L_\tau = \mathbb{E}\big(b^{\mathsf{T}}\widehat{N}(T) \mid \widehat{N}(\tau)\big).$$

Clearly, L_τ is a (continuous-time) martingale. By (3.2), we know that $L_\tau = b^{\mathsf{T}} e^{(T-\tau)M}\widehat{N}(\tau)$. Let $0 < \epsilon < \exp(-10T)$ be such that $K = T/\epsilon$ is an integer number. Then, $\{U_k = L_{k\epsilon}\}_{k=0}^K$ is a martingale and

$$\operatorname{var}\big(b^{\mathsf{T}}\widehat{N}(T)\big) = \sum_{k=0}^{K-1} \operatorname{var}(U_{k+1} - U_k).$$

We want to estimate the variance of $U_{k+1} - U_k$. Let $v_k = \widehat{N}((k+1)\epsilon) - \widehat{N}(k\epsilon)$. For two vectors $\widehat{N}_1$ and $\widehat{N}_2$,

$$\big(b^{\mathsf{T}} e^{(T-(k+1)\epsilon)M}\widehat{N}_1 - b^{\mathsf{T}} e^{(T-(k+1)\epsilon)M}\widehat{N}_2\big)^2 \leqslant \|\widehat{N}_1 - \widehat{N}_2\|^2 e^{2u(T-(k+1)\epsilon)},$$

where the norm $\|\cdot\|$ refers to the L^1-norm here and throughout this section, unless otherwise noted. Choose $\widehat{N}(k\epsilon)$ according to its distribution, and let $\widehat{N}_1$ and $\widehat{N}_2$ be chosen independently, according to the distribution of $\widehat{N}((k+1)\epsilon)$ conditioned on $\widehat{N}(k\epsilon)$. Then

$$\begin{aligned}\operatorname{var}(U_{k+1} - U_k) &= \frac{1}{2}\mathbb{E}\big[\big(b^{\mathsf{T}} e^{(T-(k+1)\epsilon)M}\widehat{N}_1 - b^{\mathsf{T}} e^{(T-(k+1)\epsilon)M}\widehat{N}_2\big)^2\big] \\ &\leqslant \frac{1}{2}\mathbb{E}(\|\widehat{N}_1 - \widehat{N}_2\|^2) e^{2u(T-(k+1)\epsilon)}.\end{aligned}$$

On the other hand, using the fact that, for every vector x in $\mathbb{R}^d$,

$$\left(\sum_{i=1}^d x_i\right)^2 \leqslant d\sum_{i=1}^d x_i^2,$$

we get

$$\sum_{j=1}^A \operatorname{var}(v_k(j)) \geqslant \frac{1}{2A}\mathbb{E}(\|\widehat{N}_1 - \widehat{N}_2\|^2)$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

where $v_k(j)$ is the jth component of v_k. Therefore,

$$\operatorname{var}(U_{k+1} - U_k) \leqslant A \sum_{j=1}^{A} \exp[2u(T - (k+1)\epsilon)] \operatorname{var}(v_k(j)). \tag{3.15}$$

By Claim 3.1, (3.9), and (3.11), for every $j = 1, 2, \ldots, A$,

$$\operatorname{var}\big(v_k(j) \mid \widehat{N}(k\epsilon)\big) \leqslant \mathbb{E}\big[(v_k(j))^2 \mid \widehat{N}(k\epsilon)\big] \leqslant C_v \epsilon \|\widehat{N}(k\epsilon)\|,$$
$$|\mathbb{E}\big(v_k(j) \mid \widehat{N}(k\epsilon)\big)| \leqslant (e^{2\epsilon} - 1)\|\widehat{N}(k\epsilon)\| \leqslant 4\epsilon \|\widehat{N}(k\epsilon)\|$$

and

$$\mathbb{E}\big[(\|\widehat{N}(k\epsilon)\|)^2\big] < e^{4k\epsilon}.$$

Therefore,

$$\begin{aligned}\operatorname{var}(v_k(j)) &= \mathbb{E}\big(\operatorname{var}\big(v_k(j) \mid \widehat{N}(k\epsilon)\big)\big) + \operatorname{var}\big(\mathbb{E}\big(v_k(j) \mid \widehat{N}(k\epsilon)\big)\big) \\ &\leqslant C_v \epsilon \exp(wk\epsilon) + 16\epsilon^2 \exp(4k\epsilon) \\ &< C_0 \epsilon \exp(wk\epsilon)\end{aligned}$$

for $C_0 = C_v + 1$, by the choice of ϵ. Therefore,

$$\begin{aligned}\operatorname{var}\big(b^{\mathsf{T}} \widehat{N}_k(T)\big) &< A^2 C_0 \epsilon \sum_{k=0}^{K-1} \exp\big(wk\epsilon + 2u(T - (k+1)\epsilon)\big) \\ &\leqslant A^2 C_0 e^{2uT} \int_0^T e^{(w-2u)\tau} d\tau < C_u \exp(2uT)\end{aligned}$$

for

$$C_u = A^2 C_0 \int_0^\infty e^{(w-2u)\tau} d\tau < \infty. \qquad \square$$

In addition, note that by (3.12) and (3.13),

$$|\mathbb{E}\big(b^{\mathsf{T}} \widehat{N}(T)\big)| \leqslant e^{uT}$$

and therefore there exists C so that

$$\mathbb{E}\big[\big(b^{\mathsf{T}} \widehat{N}(T)\big)^2\big] \leqslant C e^{2uT}. \tag{3.16}$$

We are now ready to state and prove the two main lemmas used to prove concentration, as follows.

Lemma 3.5. *For every $w' < w$ and every $1 \leqslant k \leqslant A$, a.s. for every τ large enough,*

$$\widehat{N}_k(\tau) > e^{w'\tau}. \tag{3.17}$$

Lemma 3.6. *There exists $v < w$ s.t. for every $1 \leqslant k < j \leqslant A$ a.s. for every τ large enough,*

$$p_j \widehat{N}_k(\tau) - p_k \widehat{N}_j(\tau) < e^{v\tau},$$

where $p_i, i = 1, \ldots, A$ are the components of the vector $\widehat{p}$.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

The following corollary is an immediate consequence of Claim 3.2, Claim 3.1 and Lemma 3.5.

Corollary 3.7. $w \leqslant 2$.

Proof of Lemma 3.6. Choose some v strictly between v' and w in a way that $w - v < 0.25\min(0.1, v - v', w/10)$ and let $\delta = \min(0.1, v - v', w/10)$. The vector

$$b_i = \begin{cases} p_j & \text{if } i = k, \\ -p_k & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases}$$

satisfies $b^{\mathsf{T}}\hat{p} = 0$, and therefore, using (3.16) and Markov's inequality,

$$\mathbf{P}\left(p_j\widehat{N}_k(T) - p_k\widehat{N}_j(T) = b^{\mathsf{T}}\widehat{N}(T) > \frac{1}{3}e^{vT}\right) \leqslant 9Ce^{-2\delta T}. \tag{3.18}$$

Let $\{T_i\}_{i=1,2,\ldots}$ be such that $e^{2\delta T_i} = i^2$. By Borel–Cantelli, almost surely there exists i_0 such that for all $i > i_0$,

$$p_j\widehat{N}_k(T_i) - p_k\widehat{N}_j(T_i) < \frac{1}{2}e^{vT_i}. \tag{3.19}$$

Note that

$$T_i = \frac{\log i}{\delta}$$

and therefore

$$T_{i+1} - T_i = \Theta(i^{-1}). \tag{3.20}$$

We want to show that almost surely, for all T large enough,

$$p_j\widehat{N}_k(T) - p_k\widehat{N}_j(T) < e^{vT}. \tag{3.21}$$

Section 3.3 tells us that $\mathbb{E}(\|\widehat{N}(T_i)\|) = O(\exp(wT_i))$, and Lemma 3.4 tells us that $\mathrm{var}(\|\widehat{N}(T_i)\|) = O(\exp(2wT_i))$. Therefore

$$\mathbf{P}(\|\widehat{N}(T_i)\| > e^{(w+0.6\delta)T_i}) < C_l e^{-1.2\delta T_i} = C_l i^{-1.2}$$

for some constant C_l, so that, if $m(i)$ is the number of vertices arriving between T_i and T_{i+1}, then

$$\begin{aligned} &\mathbf{P}\left(m(i) > \frac{1}{2}e^{vT_i}\right) \\ &\quad\leqslant \mathbf{P}\left(\|\widehat{N}(T_i)\| > e^{(w+0.6\delta)T_i}\right) + \mathbf{P}\left(m(i) > \frac{1}{2}e^{vT_i} \mid \|\widehat{N}(T_i)\| \leqslant e^{(w+0.6\delta)T_i}\right) \\ &\quad\leqslant \mathbf{P}\left(\|\widehat{N}(T_i)\| > e^{(w+0.6\delta)T_i}\right) + \mathbf{P}\left(m(i) > \frac{1}{2}e^{vT_i} \mid \|\widehat{N}(T_i)\| = e^{(w+0.6\delta)T_i}\right) \\ &\quad\leqslant C_l i^{-1.2} + C_s e^{-(w+0.6\delta)T_i}(T_{i+1} - T_i)^{-1}, \end{aligned} \tag{3.22}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

where the last inequality uses (3.4) in Claim 3.2 and the fact that

$$\frac{1}{2}e^{vT_i} > 2e^{(w+0.6\delta)T_i}(\exp(2(T_{i+1}-T_i))-1)$$

for i large enough.

Clearly, the first part of the right side of (3.22) is a convergent sum. We need to show that so is the second part. Remember the choice $\delta \leqslant w/10$. Then, using (3.20),

$$\begin{aligned} C_s e^{-(w+0.6\delta)T_i}(T_{i+1}-T_i)^{-1} &= \Theta\big(i\cdot e^{-(w+0.6\delta)T_i}\big) = \Theta\big(e^{\delta T_i}\cdot e^{-(w+0.6\delta)T_i}\big) \\ &= \Theta\big(e^{-(w-0.4\delta)T_i}\big) = O\big(e^{-9\delta T_i}\big) = O(i^{-9}). \end{aligned}$$

Using Borel–Cantelli, we conclude that, almost surely,

$$\sum_{k=1}^{A}|\widehat{N}_k(T)-\widehat{N}_k(T_i)| < \frac{1}{2}e^{vT_i} \tag{3.23}$$

for all k and all i large enough, and all T between T_i and T_{i+1}. Equation (3.21) follows from (3.23). □

Proof of Lemma 3.5. By Lemma 3.4, $\mathrm{var}(\widehat{N}_1(\tau)) < C_1 e^{2w\tau}$, while $\mathbb{E}(\widehat{N}_1(\tau)) > C_2 e^{w\tau}$ by Section 3.3. Therefore there exists $\rho > 0$ such that

$$\mathbf{P}\big(\widehat{N}_1(\tau) > \rho e^{w\tau}\big) > \rho. \tag{3.24}$$

Fix some large T, and let $\tau_i = iT$. For each vertex v which is a fertile leaf at time τ_{i-1}, let ℓ_v denote the number of descendants of v (including v itself) at time τ_i which are fertile leaves. The random variables $\{\ell_v\}$ are independent, their sum is $\widehat{N}_1(\tau_i)$, and the distribution of each of them is the same as the unconditional distribution of $\widehat{N}_1(T)$. Using this fact and (3.24), we get

$$\mathbf{P}\left(\widehat{N}_1(\tau_i) > \frac{\rho^2}{2}e^{wT}\widehat{N}_1(\tau_{i-1}) \mid \widehat{N}_1(\tau_{i-1})\right) \geqslant 1 - e^{-\frac{1}{16}\widehat{N}_1(\tau_{i-1})} \tag{3.25}$$

via Chernoff's bound. From (3.25), we get that almost surely there exists a constant $C_3 > 0$ such that, for all i large enough,

$$\widehat{N}_1(\tau_i) > C_3 \exp\left(i\left[wT + \log\left(\frac{\rho^2}{2}\right)\right]\right).$$

From Lemma 3.6, we may conclude that the same holds for $\widehat{N}_A(\tau_i)$, *i.e.*, for any constant $C_4 < C_3$,

$$\widehat{N}_A(\tau_i) > C_4 \exp\left(i\left[wT + \log\left(\frac{\rho^2}{2}\right)\right]\right).$$

$\widehat{N}_A(\tau)$ is monotone increasing, and therefore there exists $C_5 > 0$ such that

$$\widehat{N}_A(\tau) > C_5 \exp\left(\tau\left[w + \frac{1}{T}\log\left(\frac{\rho^2}{2}\right)\right]\right) \tag{3.26}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

for all τ large enough. Using Lemma 3.6 again, we conclude that there exists $C_6 > 0$ such that

$$\widehat{N}_k(\tau) > C_6 \exp\left(\tau\left[w + \frac{1}{T}\log\left(\frac{\rho^2}{2}\right)\right]\right)$$

for all k and large enough τ. We get (3.17) by taking T so large that

$$w + \frac{1}{T}\log\left(\frac{\rho^2}{2}\right) > w'. \qquad \square$$

Proposition 3.8. *For every k and j, almost surely*

$$\lim_{t\to\infty} \frac{\widehat{N}_k(\tau)}{\widehat{N}_j(\tau)} = \frac{p_k}{p_j}. \tag{3.27}$$

Proof. This follows immediately from Lemmas 3.5 and 3.6. $\square$

3.6. Back to discrete time

Proposition 3.9. *For the discrete-time process, and $A > \max\{A_1, A_2\}$ there exists a vector $\hat{q}$ such that, for $k \leqslant A$, we have*

$$\lim_{t\to\infty} \frac{\widetilde{N}_k(t)}{t+1} = q_k. \tag{3.28}$$

Proof. The number of infertile vertices increases at step t with probability

$$\frac{\sum_{k=1}^{A_1-1} \widetilde{N}_k(t)}{\sum_{k=1}^{A} \widetilde{N}_k(t)}$$

(their number cannot decrease). However, by (3.27), this expression tends to a limit, and therefore, using the law of large numbers,

$$\lim_{t\to\infty} \frac{N_0(t)}{t+1} = q_0 = \frac{\sum_{k=1}^{A_1-1} p_k}{\sum_{k=1}^{A} p_k}. \tag{3.29}$$

Using (3.27) once more, the proposition now follows for $k \geqslant 1$ with $q_k = (1 - q_0)p_k$. $\square$

Note that the above proposition implies that q_k and hence p_k is independent of A if $A > k$, since the left-hand side of (3.28) does not depend on A if $A > k$. So, in particular, p_1 does not depend on A.

4. Power law with a cut-off

In the previous section, we saw that for every $A > \max\{A_1, A_2\}$, the limiting proportions up to $A - 1$ are $\lambda\hat{p}$ where $\hat{p}$ is the eigenvector corresponding to the highest eigenvalue w of the A-by-A matrix M defined in (3.1). Therefore, the components $p_1, p_2, \ldots, p_A$ of the vector $\hat{p}$ satisfy the equation

$$wp_i = -\min(i, A_2)p_i + \min(i-1, A_2)p_{i-1} \quad i \geqslant 2, \tag{4.1}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

where the normalization is determined by $\sum_{i=1}^{A} p_i = 1$. From (4.1) we get that, for $i \leqslant A_2$,

$$p_i = \left(\prod_{k=2}^{i} \frac{k-1}{k+w} \right) p_1, \tag{4.2}$$

and for $i > A_2$,

$$p_i = \left(\frac{A_2}{A_2 + w} \right)^{i-A_2} p_{A_2}. \tag{4.3}$$

Clearly, (4.3) is exponentially decaying. There are many ways to see that (4.2) behaves like a power law with degree $1 + w$. Indeed,

$$\begin{aligned} \frac{p_i}{p_1} &= \left(\prod_{k=2}^{i} \frac{k-1}{k+w} \right) = \exp\left(\sum_{k=2}^{i} \log\left(\frac{k-1}{k+w} \right) \right) \\ &= \exp\left(\sum_{k=2}^{i} \left(\frac{-1-w}{k+w} \right) + O(1) \right) = \exp\left((-1-w) \left(\sum_{k=2}^{i} (k+w)^{-1} \right) + O(1) \right) \\ &= \exp\left((-1-w) \left(\sum_{k=2}^{i} k^{-1} \right) + O(1) \right) = \exp\left((-1-w) \left(\sum_{k=2}^{i} \log\left(\frac{k+1}{k} \right) \right) + O(1) \right) \\ &= \exp\big((-1-w) \log(i/2) + O(1) \big) = O(1) i^{-1-w}. \end{aligned} \tag{4.4}$$

Note that the constants implicit in the $O(\cdot)$ symbols do not depend on A_1, A_2 or i, owing to the fact that $0 < w \leqslant 2$. Equation (4.4) can be stated in the following way.

Proposition 4.1. *There exist $0 < c < C < \infty$ such that, for every A_1, A_2 and $i \leqslant A_2$, if $w = w(A_1, A_2)$ is as in* (4.1)*, then*

$$c i^{-1-w} \leqslant \frac{p_i}{p_1} \leqslant C i^{-1-w}. \tag{4.5}$$

The vector $(q_1, q_2, \ldots, q_{A-1})$ is a scalar multiple of the vector $(p_1, p_2, \ldots, p_{A-1})$, so equations (1.5), (1.6), and (1.7) in Theorem 1.2 (and the comment immediately following it) are consequences of equations (4.2), (4.3), and (4.5) derived above. It remains to prove the normalization conditions

$$\sum_{i=0}^{\infty} q_i = 1 \quad \text{and} \quad q_0 = \sum_{i=1}^{\infty} q_i \min(i-1, A_1 - 1)$$

stated in Theorem 1.2. These follow from the equations

$$\sum_{i=0}^{\infty} N_i(t) = t + 1 \quad \text{and} \quad N_0(t) = \sum_{i=1}^{\infty} N_i(t) \min(i-1, A_1 - 1).$$

The first of these simply says that there are $t + 1$ vertices at time t; the second equation is proved by counting the number of infertile children of each fertile node.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

5. Monotonicity properties of w

In this section we will prove that the exponent $1 + w$ of the power law in Proposition 4.1 is monotonically decreasing in A_1 and monotonically increasing in A_2. For this purpose, it will be useful to define a family of matrices, parametrized by two vectors $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$, which generalizes the matrix M appearing in (3.1), whose top eigenvalue is w.

Given vectors $\mathbf{y} = (y_1, y_2, \ldots, y_n)$, $\mathbf{z} = (z_1, z_2, \ldots, z_n) \in \mathbb{R}^n$, let $M(\mathbf{y}, \mathbf{z})$ denote the n-by-n matrix whose (ij)th entry is:

$$M_{i,j}(\mathbf{y}, \mathbf{z}) = \begin{cases} z_1 - y_1 & \text{if } 1 = i = j, \\ -y_j & \text{if } 2 \leqslant i = j \leqslant n, \\ y_j & \text{if } 2 \leqslant i = j + 1 \leqslant n, \\ z_j & \text{if } i = 1 \text{ and } j \geqslant 2, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, for instance, the matrix M defined in (3.1) is $M(\mathbf{y}, \mathbf{z})$, where $n = A$ and

$$y_j = \begin{cases} \min(j, A_2) & \text{if } 1 \leqslant j < A, \\ 0 & \text{if } j = A, \end{cases}$$

$$z_j = \begin{cases} 0 & \text{if } 1 \leqslant j < A_1, \\ \min(j, A_2) & \text{if } A_1 \leqslant j \leqslant A. \end{cases}$$

For the remainder of this section, we will assume:

$$y_i > 0 \quad \text{for } 1 \leqslant i < n, \tag{5.1}$$

$$z_i \geqslant 0 \quad \text{for } 1 \leqslant i < n, \tag{5.2}$$

$$y_n = 0, \quad z_n > 0. \tag{5.3}$$

All of these criteria will be satisfied by the matrices $M(\mathbf{y}, \mathbf{z})$ which arise in proving the desired monotonicity claim. It follows from (5.1), (5.2), and (5.3) that if we add a suitably large scalar multiple of the identity matrix to $M(\mathbf{y}, \mathbf{z})$, we obtain an irreducible matrix $M(\mathbf{y}, \mathbf{z}) + BI$ with nonnegative entries. The Perron–Frobenius theorem guarantees that $M(\mathbf{y}, \mathbf{z}) + BI$ has a positive real eigenvalue R of multiplicity 1, such that all other complex eigenvalues have modulus $\leqslant R$; consequently $M(\mathbf{y}, \mathbf{z})$ has a real eigenvalue $w = R - B$, of multiplicity 1, such that the real part of every other eigenvalue is strictly less than w.

We will study how w varies under perturbations of the parameters $\mathbf{y}, \mathbf{z}$. Let $P(\lambda, \mathbf{y}, \mathbf{z})$ be the characteristic polynomial of $M(\mathbf{y}, \mathbf{z})$, *i.e.*,

$$P(\lambda, \mathbf{y}, \mathbf{z}) = \det(\lambda I - M(\mathbf{y}, \mathbf{z})).$$

This is a polynomial of degree n in λ (with coefficients depending smoothly on $\mathbf{y}, \mathbf{z}$), whose largest real root $w(\mathbf{y}, \mathbf{z})$ exists and has multiplicity 1, provided $(\mathbf{y}, \mathbf{z})$ belongs to the region $V \subset \mathbb{R}^n \times \mathbb{R}^n$ determined by (5.1), (5.2), and (5.3). It follows from the implicit function theorem that $w(\mathbf{y}, \mathbf{z})$ is a smooth function of $(\mathbf{y}, \mathbf{z})$ in V, satisfying

$$\left(\frac{\partial P}{\partial y_i} + \frac{\partial w}{\partial y_i} \cdot \frac{\partial P}{\partial \lambda} \right) \Bigg|_{(w, \mathbf{y}, \mathbf{z})} = 0, \qquad \left(\frac{\partial P}{\partial z_i} + \frac{\partial w}{\partial z_i} \cdot \frac{\partial P}{\partial \lambda} \right) \Bigg|_{(w, \mathbf{y}, \mathbf{z})} = 0. \tag{5.4}$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

If $\mathbf{x}$ is any vector in $\mathbb{R}^n \times \mathbb{R}^n$, and $\partial_{\mathbf{x}}$ is the corresponding directional derivative operator, we have from (5.4):

$$\partial_{\mathbf{x}} w(\mathbf{y}, \mathbf{z}) = -\frac{\partial_{\mathbf{x}} P(w, \mathbf{y}, \mathbf{z})}{(\partial P/\partial\lambda)|_{(w,\mathbf{y},\mathbf{z})}}. \tag{5.5}$$

We know that $(\partial P/\partial\lambda)|_{(w,\mathbf{y},\mathbf{z})} > 0$ because P is a polynomial with positive leading coefficient, w is its largest real root, and w has multiplicity 1. Thus we have established the following.

Claim 5.1. *For any vector $\mathbf{x} \in \mathbb{R}^n \times \mathbb{R}^n$, and any $(\mathbf{y}, \mathbf{z}) \in V$, put $w = w(\mathbf{y}, \mathbf{z})$. Then the directional derivatives $\partial_{\mathbf{x}} w(\mathbf{y}, \mathbf{z})$ and $\partial_{\mathbf{x}} P(w, \mathbf{y}, \mathbf{z})$ have opposite signs.*

This allows monotonicity properties of w to be deduced from calculations involving directional derivatives of P. Given the definition of $M(\mathbf{y}, \mathbf{z})$, it is straightforward to compute that

$$P(\lambda, \mathbf{y}, \mathbf{z}) = \det(\lambda I - M(\mathbf{y}, \mathbf{z})) = P_1(\lambda, y, z) - \sum_{j=2}^{n} P_j(\lambda, \mathbf{y}, \mathbf{z}), \tag{5.6}$$

where

$$P_1(\lambda, y, z) = (\lambda + y_1 - z_1) \prod_{i=2}^{n} (\lambda + y_i), \tag{5.7}$$

$$P_j(\lambda, \mathbf{y}, \mathbf{z}) = \left(\prod_{i=1}^{j-1} y_i \right) z_j \left(\prod_{i=j+1}^{n} (\lambda + y_j) \right). \tag{5.8}$$

As an easy consequence of this formula, w is strictly positive.

Lemma 5.2. *w is strictly positive.*

Proof. From (5.6)–(5.8) and the fact that $y_n = 0$, we have $P(0, \mathbf{y}, \mathbf{z}) = -P_n(0, \mathbf{y}, \mathbf{z}) = -\left(\prod_{i=1}^{n-1} y_i\right) z_n$, and this is strictly negative by (5.1) and (5.3). For sufficiently large positive λ, we know that $P(\lambda, \mathbf{y}, \mathbf{z}) > 0$ because P is a polynomial whose leading coefficient in λ is positive. By the intermediate value theorem, $P(\lambda, \mathbf{y}, \mathbf{z})$ has a strictly positive real root. □

The following three lemmas encapsulate the requisite directional derivative estimates for P.

Lemma 5.3. $(\partial P/\partial z_k)|_{(w,\mathbf{y},\mathbf{z})} < 0$ *for* $(\mathbf{y}, \mathbf{z}) \in V$.

Proof. For $k > 1$,

$$\partial P/\partial z_k = -\partial P_k/\partial z_k = -\left(\prod_{i=1}^{k-1} y_i \right) \left(\prod_{i=k+1}^{n} (w + y_i) \right) < 0.$$

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

For $k = 1$,

$$\partial P/\partial z_1 = \partial P_1/\partial z_1 = -\prod_{i=2}^{n}(w + y_i) < 0.$$

Corollary 5.4. *w is monotonically decreasing in A_1.*

Proof. Increasing A_1 from k to $k+1$ has no effect on $\mathbf{y}$, and its only effect on $\mathbf{z}$ is to decrease z_k from $\min(k, A_2)$ to 0. As we move in the $-z_k$ direction, the directional derivative of P is positive, so the directional derivative of w is negative by Claim 5.1. Thus w decreases as we increase A_1 from k to $k+1$. □

Lemma 5.5. *For $1 < k < n$, $(\partial P/\partial y_k)|_{(w,\mathbf{y},\mathbf{z})} < 0$ if $(\mathbf{y}, \mathbf{z}) \in V$ and $z_k = 0$.*

Proof.

$$\begin{aligned}
\frac{\partial P}{\partial y_k} &= \frac{\partial P_1}{\partial y_k} - \sum_{j=2}^{n} \frac{\partial P_j}{\partial y_k} \\
&= \frac{1}{w + y_k} P_1 - \frac{1}{w + y_k} \sum_{j=2}^{k-1} P_j - \frac{1}{y_k} \sum_{j=k+1}^{n} P_j \\
&< \frac{1}{w + y_k} P_1 - \frac{1}{w + y_k} \sum_{j=2}^{k-1} P_j - \frac{1}{w + y_k} \sum_{j=k+1}^{n} P_j \\
&= \frac{P(w, \mathbf{y}, \mathbf{z})}{w + y_k} \\
&= 0.
\end{aligned}$$

Lemma 5.6. *For $k > 1$, $(\partial P/\partial y_k + \partial P/\partial z_k)|_{(w,\mathbf{y},\mathbf{z})} < 0$ if $(\mathbf{y}, \mathbf{z}) \in V$ and $y_k = z_k$.*

Proof.

$$\begin{aligned}
\frac{\partial P}{\partial y_k} + \frac{\partial P}{\partial z_k} &= \frac{\partial P_1}{\partial y_k} - \sum_{j=2}^{n} \frac{\partial P_j}{\partial y_k} - \frac{\partial P_k}{\partial z_k} \\
&= \frac{1}{w + y_k} P_1 - \frac{1}{w + y_k} \sum_{j=2}^{k-1} P_j - \frac{1}{y_k} \sum_{j=k+1}^{n} P_j - \frac{1}{z_k} P_k \\
&< \frac{1}{w + y_k} P_1 - \frac{1}{w + y_k} \sum_{j=2}^{k-1} P_j - \frac{1}{w + y_k} \sum_{j=k+1}^{n} P_j - \frac{1}{w + y_k} P_k \\
&= \frac{P(w, \mathbf{y}, \mathbf{z})}{w + y_k} \\
&= 0.
\end{aligned}$$

Corollary 5.7. *w is monotonically increasing in A_2.*

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

Proof. If we change A_2 from k to $k+1$, this changes $\mathbf{y}$ into a new vector $\mathbf{y}'$ satisfying

$$y'_j - y_j = \begin{cases} 1 & \text{if } k < j < n, \\ 0 & \text{otherwise.} \end{cases}$$

It changes $\mathbf{z}$ into a new vector $\mathbf{z}'$ satisfying

$$z'_j - z_j = \begin{cases} 1 & \text{if } \max(A_1, k+1) \leqslant j \leqslant n, \\ 0 & \text{otherwise.} \end{cases}$$

Letting $\mathbf{e}_j^{(y)}$ denote a unit vector in the $+y_j$ direction, and $\mathbf{e}_j^{(z)}$ a unit vector in the $+z_j$ direction, the direction of change is expressed by the vector

$$\mathbf{x} = (\mathbf{y}', \mathbf{z}') - (\mathbf{y}, \mathbf{z}) = \left[\sum_{k+1 \leqslant j < A_1} \mathbf{e}_j^{(y)} \right] + \left[\sum_{\max(k+1, A_1) \leqslant j < n} (\mathbf{e}_j^{(y)} + \mathbf{e}_j^{(z)}) \right] + \mathbf{e}_n^{(z)},$$

and $\partial_{\mathbf{x}} P$ is negative, by the preceding three lemmas. By Claim 5.1, this means w increases monotonically as we move along this path. □

Acknowledgement

We wish to thank the referee for many good suggestions that helped us improve the paper.

References

[1] Aiello, W., Chung, F. and Lu, L. (2002) Random evolution of massive graphs. In *Handbook of Massive Data Sets*, Kluwer, pp. 97–122.
[2] Albert, R. and Barabási, A.-L. (2002) Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74** 47–97.
[3] Aldous, D. J. (2003) A stochastic complex network model. *Electron. Res. Announc. Amer. Math. Soc.* **9** 152–161.
[4] Barabási, A.-L. and Albert, R. (1999) Emergence of scaling in random networks. *Science* **286** 509–512.
[5] Berger, N., Bollobás, B., Borgs, C., Chayes, J. T. and Riordan, O. (2003) Degree distribution of the FKP network model. In *Proc. 30th International Colloquium on Automata, Languages and Programming*, Vol. 2719 of *Lecture Notes in Computer Science*, Springer, pp. 725–738.
[6] Berger, N., Borgs, C., Chayes, J. T., D'Souza, R. M. and Kleinberg, R. D. (2004) Competition-induced preferential attachment. In *Proc. 31st International Colloquium on Automata, Languages and Programming*. Vol. 3142 of *Lecture Notes in Computer Science*, Springer, pp. 208–221.
[7] Bollobás, B. and Riordan, O. (2003) Mathematical results on scale-free random graphs. In *Handbook of Graphs and Networks*, Wiley-VCH, Weinheim, pp. 1–34.
[8] Bollobás, B., Borgs, C., Chayes, J. T. and Riordan, O. (2003) Directed scale-free graphs. In *Proc. 14th ACM–SIAM Symposium on Discrete Algorithms*, pp. 132–139.
[9] Bollobás, B., Riordan, O., Spencer, J. and Tusnady, G. E. (2001) The degree sequence of a scale-free random graph process. *Random Struct. Alg.* **18** 279–290.
[10] Carlson, J. M. and Doyle, J. (1999) Highly optimized tolerance: a mechanism for power laws in designed systems. *Phys. Rev. E* **60** 1412.
[11] Cooper, C. and Frieze, A. M. (2001) A general model of web graphs. In *Proc. 9th European Symposium on Algorithms*, pp. 500–511.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press

[12] Dorogovtsev, S. N. and Mendes, J. F. F. (2002) Evolution of networks. *Adv. Phys.* **51** 1079.
[13] Eggenberger, F. and Pólya, G. (1923) Über die statistik verketteter Vorgänge. *Z. Agnew. Math. Mech.* **3** 279–289.
[14] Fabrikant, A., Koutsoupias, E. and Papadimitriou, C. H. (2002) Heuristically optimized trade-offs: a new paradigm for power laws in the internet. In *Proc. 29th International Colloquium on Automata, Languages and Programming*, Vol. 2380 of *Lecture Notes in Computer Science*, Springer, pp. 110–122.
[15] Faloutsos, M., Faloutsos, P. and Faloutsos, C. (1999) On the power-law relationships of the Internet topology. *Comput. Commun. Rev.* **29** 251.
[16] Govindan, R. and Tangmunarunkit, H. (2000) Heuristics for Internet map discovery. In *Proc. INFOCOM*, pp. 1371–1380.
[17] Kenyon, C. and Schabanel, N. Personal communication.
[18] Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A. and Upfal, E. (2000) Stochastic models for the web graph. In *Proc. 41st IEEE Symp. on Foundations of Computer Science*, pp. 57–65.
[19] Newman, M. E. J. (2003) The structure and function of complex networks. *SIAM Review* **45** 167–256.
[20] Price, D. J. de S. (1976) A general theory of bibliometric and other cumulative advantage processes. *J. Amer. Soc. Inform. Sci.* **27** 292–306.
[21] Simon, H. A. (1955) On a class of skew distribution functions. *Biometrika* **42** 425–440.
[22] Yule, G. U. (1924) A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis. *Philos. Trans. Roy. Soc. London, Ser. B* **213** 21–87.
[23] Zipf, G. K. (1949) *Human Behavior and the Principle of Least Effort*, Addison-Wesley, Cambridge, MA.

https://doi.org/10.1017/S0963548305006930 Published online by Cambridge University Press