



Universiteit
Leiden
The Netherlands

Wireless random-access networks and spectra of random graphs

Sfragara, M.

Citation

Sfragara, M. (2020, October 28). *Wireless random-access networks and spectra of random graphs*. Retrieved from <https://hdl.handle.net/1887/137987>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/137987>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/137987> holds various files of this Leiden University dissertation.

Author: Sfragara, M.

Title: Wireless random-access networks and spectra of random-graphs

Issue Date: 2020-10-28



CHAPTER 3

Arbitrary bipartite interference graphs

This chapter is based on:

S.C. Borst, F. den Hollander, F.R. Nardi, M. Sfragara. *Wireless random-access networks with bipartite interference graphs*. [arXiv:2001.02841], 2020.

Abstract

We consider random-access networks where each node represents a server with a queue. Each node can be either active or inactive. A node deactivates at unit rate, and activates at a rate that depends on its queue length, provided none of its neighbors is active. We consider arbitrary bipartite graphs in the limit as the queues become large, and we identify the transition time between the two states where one half of the network is active and the other half is inactive. We decompose the transition into a succession of transitions on complete bipartite subgraphs, and formulate a greedy algorithm that takes the graph as input and gives as output the set of transition paths the system is most likely to follow. Along each path we determine the mean transition time and its law on the scale of its mean. Depending on the activation rate functions, we identify three regimes of behavior.

§3.1 Introduction

This chapter is a continuation of Chapter 2. We turn our attention to general bipartite interference graphs: the node set can be partitioned into two nonempty sets U and V , but not necessarily all nodes in U interfere with all nodes in V .

In Section 3.1.1 we describe the setting and the mathematical model of interest in this chapter. In Section 3.1.2 we introduce the key idea behind our main results and give an outline of the remainder of the chapter.

§3.1.1 Setting

We refer to Sections 1.1.5 and 2.1.1 for a general introduction to the mathematical model. In this section we refine it with some extra notions we will need in the chapter.

Consider the bipartite graph $G = ((U, V), E)$, where $U \sqcup V$ is the set of nodes and E is the set of (undirected) edges that connect a node in U to a node in V , and vice versa (see Figure 3.1 for examples). Through the chapter we assume $|V| = N$.




Figure 3.1: Examples of bipartite graphs: cyclic ladder (left), hypercube (center), even torus (right).

We study the internal model with activation rates depending on the queue lengths as in (1.5). We assume the activation rates to satisfy Definition 1.1.4 and we focus on the following.

Definition 3.1.1 (Assumptions on the activation rates).

We assume *polynomial activation functions* of the form

$$\begin{aligned} g_U(x) &\sim Bx^\beta, & x \rightarrow \infty, \\ g_V(x) &\sim B'x^{\beta'}, & x \rightarrow \infty, \end{aligned} \tag{3.1}$$

with $B, B', \beta, \beta' \in (0, \infty)$. We assume that nodes in V are *much more aggressive* than nodes in U , namely,

$$\beta' > \beta + 1. \tag{3.2}$$

As we will see later, this ensures that the transition from u to v can be decomposed into a succession of transitions on complete bipartite subgraphs.

We begin by recalling the results for complete bipartite graphs from Chapter 2 (Theorem 2.1.6). Note that they are strongly related to the initial queue lengths at the nodes in U , which are assumed in (1.11) to be $Q_U(0) = \gamma_U r$.

Theorem (Theorem 2.1.6).

Let G be a complete bipartite graph.

(I) $\beta \in (0, \frac{1}{|U|-1})$: *subcritical regime. The transition time satisfies*

$$\mathbb{E}_u[\mathcal{T}_G] = F_{\text{sub}} Q_U(0)^{\beta(|U|-1)} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.3)$$

with $F_{\text{sub}} = \frac{1}{|U|B^{-(|U|-1)}}$, and

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_G}{\mathbb{E}_u[\mathcal{T}_G]} > x \right) = \int_x^\infty \mathcal{P}_{\text{sub}}(y) dy = e^{-x}, \quad x \in [0, \infty) \quad (3.4)$$

with

$$\mathcal{P}_{\text{sub}}(z) = e^{-z}, \quad z \in [0, \infty). \quad (3.5)$$

(II) $\beta = \frac{1}{|U|-1}$: *critical regime. The transition time satisfies*

$$\mathbb{E}_u[\mathcal{T}_G] = F_{\text{cr}} Q_U(0) [1 + o(1)], \quad r \rightarrow \infty, \quad (3.6)$$

with $F_{\text{cr}} = \frac{1}{|U|B^{-(|U|-1)} + (c - \rho_U)}$, and

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_G}{\mathbb{E}_u[\mathcal{T}_G]} > x \right) &= \int_x^\infty \mathcal{P}_{\text{cr}}(y) dy \\ &= \begin{cases} (1 - Cx)^{\frac{1-C}{C}}, & \text{if } x \in [0, \frac{1}{C}), \\ 0, & \text{if } x \in [\frac{1}{C}, \infty), \end{cases} \end{aligned} \quad (3.7)$$

with

$$\mathcal{P}_{\text{cr}}(z) = \begin{cases} (1 - C)(1 - Cz)^{\frac{1}{C}-2}, & \text{if } z \in [0, \frac{1}{C}), \\ 0, & \text{if } z \in [\frac{1}{C}, \infty), \end{cases} \quad (3.8)$$

and $C = F_{\text{cr}}(c - \rho_U) \in (0, 1)$.

(III) $\beta \in (\frac{1}{|U|-1}, \infty)$: *supercritical regime. The transition time satisfies*

$$\mathbb{E}_u[\mathcal{T}_G] = F_{\text{sup}} Q_U(0) [1 + o(1)], \quad r \rightarrow \infty, \quad (3.9)$$

with $F_{\text{sup}} = \frac{1}{c - \rho_U}$, and

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\tau_{1_V}}{\mathbb{E}_u[\mathcal{T}_G]} > x \right) = \int_x^\infty \mathcal{P}_{\text{sup}}(y) dy = \begin{cases} 1, & \text{if } x \in [0, 1), \\ 0, & \text{if } x \in [1, \infty), \end{cases} \quad (3.10)$$

with

$$\mathcal{P}_{\text{sup}}(z) = \delta_1(z), \quad z \in [0, \infty), \quad (3.11)$$

where $\delta_1(z)$ is the Dirac function at 1.

Theorem 2.1.6 shows that there is a trichotomy: depending on the value of β the transition exhibits a subcritical regime, a critical regime or a supercritical regime. Our goal is to extend Theorem 2.1.6 to arbitrary bipartite graphs. Note how the mean transition time depends on the actual value of the initial queue lengths at nodes in U : for complete bipartite graphs, those are fixed and equal to $\gamma_U r$; for arbitrary bipartite graphs, we will see how the mean transition time depends on the way the queue lengths are changing when nodes in V activate.

Next, we define some key notions that we will need in the chapter.

Definition 3.1.2 (Fork).

For a node $v \in V$, we define the set of neighbors of v as $N(v) = \{u \in U : uv \in E\}$ and the degree of v as $d(v) = |N(v)|$. Given a node $v \in V$, we refer to *fork* of v as the complete bipartite subgraph of G containing only node v , its neighbors $N(v) \subseteq U$ and the edges between them. We talk about a d -fork when $d(v) = d$ with $d \in \mathbb{N}$.

Definition 3.1.3 (Updated queue lengths).

Let $Q_U = \{Q_{U,i}\}_{i=1}^{|U|}$ be the sequence of queues associated with the nodes in U , and $Q_V = \{Q_{V,j}\}_{j=1}^{|V|}$ the sequence of queues associated with the nodes in V . Put $Q = (Q_U, Q_V)$, and let $Q^k = (Q_U^k, Q_V^k)$ be the pair of sequences representing the *updated queue lengths* after k nodes in V activated (see Definition 3.2.9 later for more details).

We denote by \mathcal{T}_G^Q the transition time of the graph G when the initial queue lengths are $Q = (Q_U, Q_V)$. It represents the time it takes to reach v starting from u . Below, we define the nucleation time in order to distinguish between the full transition of G and the successive transitions (nucleations) of the subgraphs of G related to each node activating in V .

Definition 3.1.4 (Nucleation time).

We call *nucleation time* of the fork of v the time it takes for the nodes $N(v)$ to deactivate and for v to activate. We denote this time by $\mathcal{T}_v^Q = \mathcal{T}_{N(v),v}^Q$, where v represents the activating node and Q represents the initial queue lengths. It can be seen as the transition time of the complete bipartite subgraph of G represented by the fork of v . Note that, for $v, w \in V$, \mathcal{T}_v^Q and \mathcal{T}_w^Q are dependent random variables when $N(v) \cap N(w) \neq \emptyset$.

§3.1.2 Key idea and outline

The key idea behind this chapter is to define an algorithm that allows us to identify the set of paths \mathcal{A} the network is mostly likely to follow while nodes in U deactivate and nodes in V activate. We label the nodes in V based on their first activation and we denote by a^* the path that the network follows. More precisely, $a^* = (v_1^*, \dots, v_N^*)$ with v_1^*, \dots, v_N^* all distinct, where v_1^* is the first node that activates and v_N^* the last one. Let $\mathcal{E}(a^*)$ denote the event that any of the paths in \mathcal{A} occurs. We will prove that

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{E}(a^*)) = 1. \quad (3.12)$$

In particular, we will show that if we condition on the event

$$A_a = \{\text{the network follows path } a \in \mathcal{A}\}, \quad (3.13)$$

then we are able to identify how the mean transition time $\mathbb{E}_u[\mathcal{T}_G^Q | A_a]$ depends on the sequence of nucleation times of the forks of the nodes in V , ordered as in the path a (Theorem 3.3.2). We derive the asymptotics of the mean transition time as $r \rightarrow \infty$ (Theorem 3.3.3) and identify the law of the transition time divided by its mean (Theorem 3.3.5). To do so, we determine how the queue lengths change along the given path (Theorem 3.4.8). Similarly as for the complete bipartite graph in Theorem 2.1.6, we distinguish between three regimes for the value of β (subcritical, critical and supercritical), in which the queues behave differently and, consequently, so does the transition time.

Outline of the chapter. The remainder of this chapter is organized as follows. In Section 3.2 we introduce the algorithm, show that it has two important properties, *greediness* and *consistency*, and give an example of how it works. In Section 3.3 we state our main theorems. In particular, we show how both the mean transition time and its law on the scale of its mean can be determined according to the path that the algorithm chooses. In Section 3.4 we show how the nucleation times depend on the graph structure and we analyze how the queue lengths at the nodes change along each path that the algorithm chooses. In Section 3.5 we provide the proof of the two algorithm properties mentioned above and we discuss the algorithm complexity. In Section 3.6 we prove our main theorems. In Appendix C, we show some technical computations for the mean nucleation time in the special setting of disjoint forks competing for activation.

§3.2 The algorithm

In this section we introduce the algorithm that describes, step by step, how the network behaves while nodes in U deactivate and nodes in V activate. The presentation is organised into a series of definitions and lemmas. In Section 3.2.1 we define how the algorithm works iteratively. In Section 3.2.2 we show that the algorithm is greedy and consistent (Propositions 3.2.6–3.2.7). In Section 3.2.3 we explain how the algorithm is used to capture the nucleation of the forks. An example of a bipartite graph and how the algorithm acts on it are given in Section 3.2.4.

§3.2.1 Definition of the algorithm

Let $N = |V|$ be the number of nodes in V . The algorithm takes as input the bipartite graph $G = ((U, V), E)$ and gives as output a sequence of triples that is needed to characterise the transition time, namely,

$$G \rightarrow (Y_k, \bar{d}_k, n_k)_{k=1}^N, \quad (3.14)$$

where Y_k is a random variable with values in $\{1, \dots, N\}$ describing the index of the node selected at step k , $\bar{d}_k \in \mathbb{N}$ is the degree of the selected node and $n_k \in \mathbb{N}$ is

a parameter that counts how many possibilities there are at step k to choose the next node in V (uniformly at random) from the remaining nodes with least degree. Sometimes we will write v_k^* instead of v_{Y_k} to emphasise that the network is following a specific order of activation for the nodes in V .

Definition 3.2.1 (Algorithm).

Set $G = G_1 = ((U_1, V_1), E_1)$. Given the graph $G_k = ((U_k, V_k), E_k)$, find the graph $G_{k+1} = ((U_{k+1}, V_{k+1}), E_{k+1})$ by iterating the following procedure until V_{k+1} is empty.

- Start from the graph G_k .
- Look at the nodes in V_k and at the *minimum degree* \bar{d}_k in G_k .
- Pick a node uniformly at random from the ones with minimum degree in G_k .
- Denote the chosen node by v_k^* and the *number of choices* by n_k .
- Eliminate the node v_k^* and all its neighbors in U_k , together with all their adjacent edges. Denote the resulting bipartite graph by G_{k+1} .

The idea of eliminating step by step the nodes in U that deactivated comes from the fact that when a node in V activates, it “blocks” all its neighbors in U , which, with high probability as $r \rightarrow \infty$, will remain inactive for the rest of the time. This is due to the aggressiveness of the nodes in V compared to the nodes in U (recall Definition 3.1.1). The following lemma will be proved in Section 3.6.2.

Lemma 3.2.2 (Activation sticks).

Consider a node $u \in U$ and let $N(u) \subseteq V$ be the set of neighbors of u . Denote by t_u the first time a node $v \in N(u)$ activates. Then, with high probability as $r \rightarrow \infty$, u remains inactive after t_u , i.e., $X_u(t) = 0$ for all $t \geq t_u$.

Definition 3.2.3 (Mean nucleation time for the algorithm).

The algorithm generates a sequence v_1^*, \dots, v_N^* of successively activating nodes in V . Associated with step k of the algorithm is the nucleation time of the fork of node v_k^* (see Definition 3.1.2), which according to Theorem 2.1.6 satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q_U^{k-1}}] = F^k (\mathbb{E}_u[Q_U^{k-1}])^{1 \wedge \beta(\bar{d}_k - 1)} [1 + o(1)], \quad r \rightarrow \infty. \quad (3.15)$$

Here F^k is a pre-factor that depends on the degree \bar{d}_k , which plays the role of $|U|$ in Theorem 2.1.6, and on its relation with β . The term $\mathbb{E}_u[Q_U^{k-1}]$ represents the mean updated queue lengths at the nodes in U_k in the subgraph G_{k-1} (see Definition 3.1.3), and plays the role of the initial queue lengths in Theorem 2.1.6. Note that Q_U^0 is fixed, while $Q_U^1, Q_U^2, \dots, Q_U^{N-1}$ are random.

Intuitively, the sum of the mean nucleation times associated with the path generated by the algorithm gives the mean transition time along that path. We will see in Section 3.4.2 that the pre-factors F^k actually need to be adjusted by certain weights that depend on the graph structure.

§3.2.2 Properties of the algorithm

Definition 3.2.4 (Maximum least degree).

Given the sequence $(\bar{d}_k)_{k=1}^N$ generated by the algorithm, let

$$d^* = \max_{1 \leq k \leq N} \bar{d}_k \quad (3.16)$$

be the *maximum least degree* of the path associated with $(\bar{d}_k)_{k=1}^N$.

Each time we run the algorithm, it may generate a different sequence, since it decides uniformly at random which node in V with minimum degree to pick next. We know that the set of paths \mathcal{A} generated by the algorithm is the set of most likely paths the network follows. The order of the nodes in a path is given by their successive activation in V .

The following lemma and two propositions will be proved in Section 3.5.2.

Lemma 3.2.5 (Comparing maximum least degrees of different paths).

Consider two different paths a, b such that $a \in \mathcal{A}$ is generated by the algorithm. For $k = 1, \dots, N$, denote by $\bar{d}_{k,a}$ and $\bar{d}_{k,b}$ the minimum degrees at step k in paths a and b . Let $d_a^* = \max_{1 \leq k \leq N} \bar{d}_{k,a}$ and $d_b^* = \max_{1 \leq k \leq N} \bar{d}_{k,b}$. Then $d_a^* \leq d_b^*$.

In other words, given any path b , its maximum least degree cannot be smaller than the maximum least degree of a path a generated by the algorithm. We will see how the maximum least degree d^* determines the order of the mean transition time. Depending on how β is related to d^* , we distinguish between the following three different regimes.

- (I) *Sucritical regime*, if $\beta \in (0, \frac{1}{d^*-1})$.
- (II) *Critical regime*, if $\beta = \frac{1}{d^*-1}$.
- (III) *Supercritical regime*, if $\beta \in (\frac{1}{d^*-1}, \infty)$.

The algorithm is greedy, in the sense that it always chooses the node that adds the least to the total transition time along the path, simply because this node is likely to be the first to activate. The greedy way in which the algorithm picks the nodes ensures that the transition time along the chosen path is the shortest possible.

Proposition 3.2.6 (Greediness).

The mean transition time along a path generated by the algorithm is the shortest possible.

The algorithm is consistent, in the sense that d^* is unique. Different paths generated by the algorithm lead to the same order of the mean transition time.

Proposition 3.2.7 (Consistency).

All the paths generated by the algorithm lead to the same order of the mean transition time.

§3.2.3 Structure of the algorithm

A node in V activates because it is the one whose complete bipartite fork has the fastest nucleation, and occurs because of the randomness in the activation and deactivation Poisson clocks and the randomness of the queue length processes that appear as the arguments of the activation rates.

Definition 3.2.8 (Next nucleation time).

Given that $k - 1$ nodes in V already activated, we define by *next nucleation time* $\bar{\tau}_k$ the time it subsequently takes for the k -th node in V to activate, i.e.,

$$\bar{\tau}_k = \min_{v \in V_k} \mathcal{T}_v^{Q^{k-1}}. \quad (3.17)$$

By keeping track of which nodes have been picked, we can compute the updated queue lengths for the successive mean nucleation times.

Definition 3.2.9 (Updated queue lengths).

For $k = 1, \dots, N$, define the *updated queue lengths* Q^{k-1} by

$$Q^{k-1} = (Q_U^{k-1}, Q_V^{k-1}) = \left(Q_U \left(\sum_{l=1}^{k-1} \bar{\tau}_l \right), Q_V \left(\sum_{l=1}^{k-1} \bar{\tau}_l \right) \right). \quad (3.18)$$

When a node in V activates, its fork can be of three different types depending on how its degree is related to β .

Definition 3.2.10 (Subcritical, critical and supercritical nodes).

Given that $k - 1$ nodes in V already activated, consider the k -th activating node and its fork of degree \bar{d}_k . If $\beta \in (0, \frac{1}{\bar{d}_k - 1})$, then the node (or its fork) is subcritical. If $\beta = \frac{1}{\bar{d}_k - 1}$, then it is critical. If $\beta \in (\frac{1}{\bar{d}_k - 1}, \infty)$, then it is supercritical.

In the subcritical and critical regimes, with high probability as $r \rightarrow \infty$, the next nucleation time $\bar{\tau}_k$ is given by the minimum over the nodes with least degree in V_k . Indeed, with high probability as $r \rightarrow \infty$, nodes with least degree activate first. The following lemma will be proved in Section 3.6.2.

Lemma 3.2.11 (Activation selects low degree).

For $k = 1, \dots, N$, consider two nodes $v, w \in V_k$ such that $d_k(w) > d_k(v) = \bar{d}_k$. Suppose that $\beta \in (0, \frac{1}{\bar{d}_k - 1}]$. Then the probability of w activating before v satisfies

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_w^{Q^{k-1}} < \mathcal{T}_v^{Q^{k-1}}) = 0. \quad (3.19)$$

In the supercritical regime the situation is more delicate. If at step k the least degree fork has degree \bar{d}_k such that $\beta \in (\frac{1}{\bar{d}_k - 1}, \infty)$, then the mean nucleation time of the next activating fork is the same for all the remaining forks in the graph. The network does not distinguish between the nodes according to their degree anymore, since all possibilities contribute equally to the total mean transition time. Indeed, the mean nucleation time is given by the expected time it takes for the queue lengths at nodes in U to hit zero. Hence, after the nucleation of the first supercritical fork,

all the queues in U are $o(r)$ and the transition occurs very fast (see Section 3.4.3 for more details).

In Section 3.3 we will see how the transition time can be computed given the set of possible paths generated by the algorithm. Moreover, for each fixed path we will identify the mean transition time and its law on the scale of its mean. Given a path, we know in which order the nodes activate. In Section 3.6 we will see how we can identify the nucleation time of a node given in Definition 3.2.8 with the nucleation time of the complete bipartite fork of the activating node, as written in (3.15). The sum of all the nucleation times gives us the transition time of the graph. Not all the terms in the sum contribute significantly as $r \rightarrow \infty$. We will need to identify which are the leading order terms. The answer depends on the sequence of degrees $(\bar{d}_k)_{k=1}^N$ generated by the algorithm and on how the queue lengths change along the path.

§3.2.4 Example

Consider the bipartite graph $G = ((U, V), E)$ with $|U| = 6$ and $|V| = 4$ in Figure 3.2. This graph serves as a simple example of how the algorithm works.




Figure 3.2: The initial bipartite graph $G = G_1 = ((U_1, V_1), E_1)$.

Step $k = 1$. We start with $G = G_1 = ((U_1, V_1), E_1)$. There are two nodes v_2, v_4 with minimum degree $\bar{d}_1 = 2$, so $n_1 = 2$. Pick uniformly at random one of them (with probability $\frac{1}{n_1} = \frac{1}{2}$), say $Y_1 = 2$. Eliminate node v_2 , all its neighbors u_2, u_3 , and all their edges $u_2v_1, u_2v_2, u_2v_3, u_3v_1, u_3v_2, u_3v_3$. Denote the new bipartite graph by $G_2 = ((U_2, V_2), E_2)$. The nucleation time associated with this node satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_1}}^{Q^0}] = \mathbb{E}_u[\mathcal{T}_{v_2}^{Q^0}] = F^1(Q_U^0)^{1 \wedge \beta} [1 + o(1)], \quad r \rightarrow \infty. \quad (3.20)$$

Step $k = 2$. Node v_1 has the minimum degree $\bar{d}_2 = 1$, so $Y_2 = 1$. Eliminate node v_1 , all its neighbors, and all their edges. Denote the new graph by $G_3 = ((U_3, V_3), E_3)$. The nucleation time associated with this node satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_2}}^{Q^0}] = \mathbb{E}_u[\mathcal{T}_{v_1}^{Q^1}] = F^2(\mathbb{E}_u[Q_U^1])^0 [1 + o(1)] = o(1), \quad r \rightarrow \infty. \quad (3.21)$$

Step $k = 3$. Node v_4 has the minimum degree $\bar{d}_3 = 2$, so $Y_3 = 4$. Eliminate node v_4 , all its neighbors, and all their edges. Denote the new graph by $G_4 = ((U_4, V_4), E_4)$.

The nucleation time associated with this node satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_3}}^{Q^0}] = \mathbb{E}_u[\mathcal{T}_{v_4}^{Q^2}] = F^3 (\mathbb{E}_u[Q_U^2])^{1 \wedge \beta} [1 + o(1)], \quad r \rightarrow \infty. \quad (3.22)$$

Step $k = 4$. Node v_3 is the only node left, with degree $\bar{d}_4 = 1$, so $Y_4 = 3$. Eliminate node v_3 , all its neighbors, and all their edges, after which the empty graph is left. The nucleation time associated with this node satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_4}}^{Q^0}] = \mathbb{E}_u[\mathcal{T}_{v_3}^{Q^3}] = F^4 (\mathbb{E}_u[Q_U^3])^0 [1 + o(1)] = o(1), \quad r \rightarrow \infty. \quad (3.23)$$

The above scenario forms a path that is described by nodes in V activating in the order v_2, v_1, v_4, v_3 (see Figure 3.3).




Figure 3.3: The sequence of bipartite graphs $G_2 = ((U_2, V_2), E_2)$, $G_3 = ((U_3, V_3), E_3)$ and $G_4 = ((U_4, V_4), E_4)$ generated by the algorithm.

Note that the algorithm may pick node v_4 at the first step by setting $Y_1 = 4$, since the choice of the node with minimum degree is uniformly at random. If so, then the algorithm would follow a different path. At the first step we would get $Y_1 = 4$ and $\mathbb{E}_u[\mathcal{T}_{v_4}^{Q^0}] = F^1 (Q_U^0)^{1 \wedge \beta} [1 + o(1)]$. At the second step, $Y_2 = 2$ and $\mathbb{E}_u[\mathcal{T}_{v_2}^{Q^1}] = F^2 (\mathbb{E}_u[Q_U^1])^{1 \wedge \beta} [1 + o(1)]$. At the third step, $Y_3 = 1$ and $\mathbb{E}_u[\mathcal{T}_{v_1}^{Q^2}] = o(r)$. At the fourth step, $Y_4 = 3$ and $\mathbb{E}_u[\mathcal{T}_{v_3}^{Q^3}] = o(r)$. This choice leads to a different path, where the nodes in V activate in the order v_4, v_2, v_1, v_3 .

Each possible scenario is identified with a path in the algorithm, described by the nodes in V according to the order of their first activation. The total mean transition time along a path can be thought as a sum of the mean nucleation times associated with each activating node in the path (see Theorem 3.3.2). We will prove in Section 3.5.2 that all the paths generated by the algorithm lead to the same order of the mean transition time.

§3.3 Main results

In this section we present our main theorems regarding the transition time. In Section 3.3.1 we show that \mathcal{E} , the event that the network follows the algorithm, occurs with high probability as $r \rightarrow \infty$ (Theorem 3.3.2(i)). We analyze the contributions along a given path, noting that not all the nucleation times are significant for the total mean transition time (Theorem 3.3.2(ii)). In Section 3.3.2 we compute the asymptotics of the mean transition time, including the pre-factor, focusing on the significant

terms only (Theorem 3.3.3). In Section 3.3.3 we identify the law of the transition time divided by its mean, which turns out to be a convolution of the laws found for the complete bipartite graph in Theorem 2.1.6 (Theorem 3.3.5). There is again a trichotomy, depending on the value of β . Proofs will be given in Section 3.6.

§3.3.1 Most likely paths

Let Ω be the set of all possible orderings (permutations) of nodes in V . Denote by $\mathcal{A} \subseteq \Omega$ the subset of orderings generated by the algorithm, and denote by \mathcal{A}_{sc} the subset of orderings generated by the algorithm truncated at the first supercritical node (if there is any). Recall that, according to Definition 3.2.10, a supercritical node is a node that activates through a supercritical fork. If $a = (v_1, \dots, v_N)$ is an element of \mathcal{A} , then $a_{sc} = (v_1, \dots, v_{sc})$ is an element of \mathcal{A}_{sc} , where v_{sc} denotes the last node of each truncated ordering. We allow this node to be any of the remaining supercritical nodes not already present in the sequence.

Definition 3.3.1 (The network follows the algorithm).

Denote by $a^* = (v_1^*, \dots, v_N^*)$ the ordering of the nodes in V along the path a^* followed by the network. For fixed a^* , let

$$\mathcal{E}(a^*) = \{\exists a \in \mathcal{A}: a = a^*\} \cup \left\{ \exists a_{sc} = (v_1, \dots, v_{sc}) \in \mathcal{A}_{sc}: v_1 = v_1^*, \dots, v_{sc} = v_{sc}^* \right\} \quad (3.24)$$

be the event that the network follows any of the paths generated by the algorithm up to the first supercritical node (if there is any).

Our first main theorem shows how the algorithm helps us to find the mean transition time. The first statement holds for all three regimes. The second and third statements hold in the subcritical and critical regimes only (for which the network follows the algorithm until the last activating node). The idea is that the mean transition time can be seen as a weighted sum of the mean nucleation times associated with each activation and of negligible terms representing the time it takes after each activation to bring the network back in the state with all the remaining nodes in U active. In the supercritical regime we do not need any statement, because the mean transition time is known to be the expected time it takes for the queue lengths to hit zero.

Theorem 3.3.2 (Most likely paths).

Consider the bipartite graph G with initial queue lengths Q^0 .

- (i) *With high probability as $r \rightarrow \infty$, the network follows the algorithm, i.e.,*

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{E}(a^*)) = 1. \quad (3.25)$$

Consider $\beta \in (0, \frac{1}{d^-1}]$: subcritical or critical regime.*

(ii) With high probability as $r \rightarrow \infty$, the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] = \sum_{k=1}^N \sum_{\substack{i_1, \dots, i_k: \\ (v_{i_1}, \dots, v_{i_k}) \in V_1 \times \dots \times V_k}} \left(\prod_{l=1}^k \frac{1}{n_l} \right) f_k \mathbb{E}_u[\mathcal{T}_{v_{i_k}}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)}] [1 + o(1)],$$

$$r \rightarrow \infty, \quad (3.26)$$

where $n_k \in \mathbb{N}$ is the number of possible nodes that the algorithm can pick at step k , while the factor $f_k \in (0, 1)$ (to be identified in Theorem 3.3.3) comes from the fact that the node activating at step k is the one that activates first among the n_k nodes with the same least degree. Both n_k and f_k depend on the sequence of nodes that activated before step k .

(iii) Conditional on the path $a = (v_1, \dots, v_N) \in \mathcal{A}$ and the event

$$A_a = \{a^* = a\} = \{v_1 = v_1^*, \dots, v_N = v_N^*\}, \quad (3.27)$$

with high probability as $r \rightarrow \infty$, the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] = \sum_{k=1}^N f_k \mathbb{E}_u[\mathcal{T}_{v_k}^{Q^{k-1}}] [1 + o(1)], \quad r \rightarrow \infty. \quad (3.28)$$

Theorem 3.3.2 will be proved in Section 3.6.3. Note that the mean transition time can be split as

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] + \mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)^c}]. \quad (3.29)$$

The second term in the right-hand side represents the mean transition time when the network does not follow the algorithm, and equals

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)^c}] = \mathbb{E}_u[\mathcal{T}_G^{Q^0} | \mathcal{E}(a^*)^c] \mathbb{P}_u(\mathcal{E}(a^*)^c). \quad (3.30)$$

Even though we know from Theorem 3.3.2(i) that $\mathbb{P}_u[\mathcal{E}(a^*)^c]$ tends to zero as $r \rightarrow \infty$, a priori this term may still affect the total mean transition time, since the conditional expectation may be substantial. In what follows we focus on the first term in the right-hand side, since this captures the typical behavior of the network.

We will see in Theorem 3.3.3 below that, in the supercritical regime, the mean transition time is the expected time it takes for the queues in U to hit zero, independently of which path the network took before the activation of the first supercritical node. Theorem 3.3.2(ii) gives us a way, in the subcritical and critical regimes, to split the total mean transition time into a sum of mean nucleation times of successive forks, by taking into account all possible paths that the algorithm may follow, each with its own probability. Theorem 3.3.2(iii) shows that we can also think of the total mean transition time as a sum over all possible paths, each with its own probability and mean transition time, namely,

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] = \sum_{a \in \mathcal{A}} \mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{A_a}] = \sum_{a \in \mathcal{A}} \mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] \mathbb{P}_u(A_a). \quad (3.31)$$

The above expression allows us to compute the mean transition time along a single path. For every $a \in \mathcal{A}$,

$$\mathbb{P}_u(A_a) = \prod_{k=1}^N \frac{1}{n_k}. \quad (3.32)$$

We already saw in Proposition 3.2.7 that the order of the mean transition time does not depend on which path the algorithm generates.

§3.3.2 Mean of the transition time

Consider a path $a \in \mathcal{A}$ generated by the algorithm and the event A_a that the network follows this path. Recall that $d^* = \max_{1 \leq k \leq N} \bar{d}_k$ is the maximum degree among the sequence of minimum degrees $(\bar{d}_k)_{k=1}^N$. Let v_k^* be the k -th activating node in path a . According to Definition 3.2.3, the mean nucleation time $\mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}]$ is given by

$$\mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}] = \begin{cases} F_{\text{sub}}^k (\mathbb{E}_u[Q_U^{k-1}])^{\beta(\bar{d}_k-1)} [1 + o(1)], & \text{if } \beta \in (0, \frac{1}{\bar{d}_k-1}), \\ F_{\text{cr}}^k \mathbb{E}_u[Q_U^{k-1}] [1 + o(1)], & \text{if } \beta = \frac{1}{\bar{d}_k-1}, \\ F_{\text{sup}}^k \mathbb{E}_u[Q_U^{k-1}] [1 + o(1)], & \text{if } \beta = (\frac{1}{\bar{d}_k-1}, \infty), \end{cases} \quad r \rightarrow \infty, \quad (3.33)$$

with

$$F_{\text{sub}}^k = \frac{1}{\bar{d}_k B^{-(\bar{d}_k-1)}}, \quad F_{\text{cr}}^k = \frac{1}{\bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}, \quad F_{\text{sup}}^k = \frac{1}{c - \rho_U}, \quad (3.34)$$

are constants depending on \bar{d}_k, B, c, ρ_U . Note that F_{sub}^k really depends on k , while $F_{\text{cr}}^k = \frac{1}{\bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}$ is the same for every critical node, and $F_{\text{sup}}^k = F_{\text{sup}}$ is independent of k . Moreover, note that the first mean nucleation time depends on the initial queue lengths Q_U^0 at the nodes in U , but in general the mean nucleation time associated with a fork depends on the mean queue lengths at the nodes in U at the moment the fork starts the nucleation.

Our second main theorem identifies the mean transition time along a given path.

Theorem 3.3.3 (Mean transition time).

Consider the bipartite graph G with initial queue lengths Q^0 .

(I) $\beta \in (0, \frac{1}{d^*-1})$: subcritical regime. The transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] = \sum_{\substack{1 \leq k \leq N \\ k: \bar{d}_k = d^*}} f_k \frac{\gamma_U^{\beta(d^*-1)}}{d^* B^{-(d^*-1)}} r^{\beta(d^*-1)} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.35)$$

with

$$f_k = \frac{1}{n_k}. \quad (3.36)$$

- (II) $\beta = \frac{1}{d^*-1}$: *critical regime*. Denote by $h_k \in \mathbb{N}_0$ the number of nodes in V at step k that already activated through a fork of degree d^* . Then the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] = \sum_{\substack{1 \leq k \leq N \\ k: \bar{d}_k = d^*}} f_k \frac{\gamma_U^{(h_k)}}{d^* B^{-(d^*-1)} + (c - \rho_U)} r [1 + o(1)], \quad r \rightarrow \infty, \quad (3.37)$$

with

$$f_k = \frac{\bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}{n_k \bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)} \quad (3.38)$$

and

$$\gamma_U^{(h_k)} = \gamma_U - (c - \rho_U) \sum_{\substack{1 \leq i \leq k \\ i: \bar{d}_i = d^*}} f'_i, \quad (3.39)$$

where for a critical node v_i the coefficient f'_i is defined in a recursive way as

$$f'_i = \frac{1}{n_i \bar{d}_i B^{-(\bar{d}_i-1)} + (c - \rho_U)} \left(\gamma_U - (c - \rho_U) \sum_{\substack{1 \leq j \leq i-1 \\ j: \bar{d}_j = d^*}} f'_j \right) > 0. \quad (3.40)$$

- (III) $\beta \in (\frac{1}{d^*-1}, \infty)$: *supercritical regime*. The transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{\gamma_U}{c - \rho_U} r [1 + o(1)], \quad r \rightarrow \infty. \quad (3.41)$$

Theorem 3.3.3 will be proved in Section 3.6.4. Both in the subcritical and supercritical regimes, Theorem 3.3.3 provides explicit formulas for the mean transition time in terms of the parameters c, γ_U, ρ_U and B, β in our model and the sequence of numbers $(\bar{d}_k, n_k)_{k=1}^N$ that are produced by the algorithm, with $d^* = \max_{1 \leq k \leq N} \bar{d}_k$. In the critical regime, however, the formula is more delicate, since the pre-factor depends on how long the critical nucleations take. Indeed, $\gamma_U^{(h_k)}$ in (3.39) represent the mean updated queue lengths at step k after h_k nodes in V activate through critical forks (see Section 3.4.3 for more details). Recall from Chapter 2 that the queue lengths all have a good behavior, in the sense that, with high probability as $r \rightarrow \infty$, they are always close to their mean (see Remark 3.4.7). Note that the mean transition time in the subcritical and critical regimes depends on the path, while in the supercritical regime it does not.

§3.3.3 Law of the transition time

Theorem 3.3.2 shows how the mean transition time along a path is a sum of terms related to the successive mean nucleation times of complete bipartite subgraphs of G . Theorem 3.3.3 tells us that, depending on the value of β , this sum reduces to a smaller sum of only a few significant terms. It also tells us how to compute the pre-factors of these terms.

Definition 3.3.4 (Multiplicity of d^*).

Consider a path $a \in \mathcal{A}$ generated by the algorithm and its associated degree sequence $(\bar{d}_k)_{k=1}^N$. Write m_{sub}^a and m_{cr}^a to denote the multiplicity of d^* in the path a in the subcritical and critical regimes, i.e.,

$$m_{\text{sub}}^a = |\{k: \bar{d}_k = d^* < \beta^{-1} + 1\}|, \quad (3.42)$$

$$m_{\text{cr}}^a = |\{k: \bar{d}_k = d^* = \beta^{-1} + 1\}|. \quad (3.43)$$

Our third main theorem identifies the law of $\mathcal{T}_G^{Q^0}/\mathbb{E}_u[\mathcal{T}_G^{Q^0}]$. Recall the laws $\mathcal{P}_{\text{sub}}, \mathcal{P}_{\text{cr}}, \mathcal{P}_{\text{sup}}$ arising from Theorem 2.1.6. Write \otimes to denote convolution.

Theorem 3.3.5 (Law of the transition time).

Consider the bipartite graph G with initial queue lengths Q^0 .

- (I) $\beta \in (0, \frac{1}{d^*-1})$: *subcritical regime*. With f_k as in (3.36) and m_{sub}^a as in (3.42), the transition time satisfies

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} > x \mid A_a \right) = \int_x^\infty \left(\otimes_{k=1}^{m_{\text{sub}}^a} \mathcal{P}_{\text{sub}}^{f_k, S_{m_{\text{sub}}^a}} \right)(y) dy, \quad x \in [0, \infty), \quad (3.44)$$

with

$$\mathcal{P}_{\text{sub}}^{f_k, S_{m_{\text{sub}}^a}}(z) = \frac{S_{m_{\text{sub}}^a}}{f_k} \exp \left(-\frac{S_{m_{\text{sub}}^a}}{f_k} z \right), \quad z \in [0, \infty), \quad (3.45)$$

and with $S_{m_{\text{sub}}^a} = \sum_{i: \bar{d}_i = d^*} f_i$.

- (III) $\beta \in (\frac{1}{d^*-1}, \infty)$: *supercritical regime*. The transition time satisfies

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0}]} > x \right) = \int_x^\infty \mathcal{P}_{\text{sup}}(y) dy = \begin{cases} 1, & \text{if } x \in [0, 1), \\ 0, & \text{if } x \in [1, \infty), \end{cases} \quad (3.46)$$

with

$$\mathcal{P}_{\text{sup}}(z) = \delta_1(z), \quad z \in [0, \infty), \quad (3.47)$$

where $\delta_1(z)$ is the Dirac function at 1.

Theorem 3.3.5 will be proved in Section 3.6.5. There we will also see why there is no statement for the critical regime (II).

§3.3.4 Discussion

Intuition. Analyzing the transition time for arbitrary bipartite graphs is much harder than for complete bipartite graphs. The key idea is to view the transition time as a sum of subsequent nucleation times for complete bipartite subgraphs. The order in which nodes activate in V is random, because it depends on the fluctuations of the activation rates via the queue lengths. However, with high probability as $r \rightarrow \infty$, the nodes with the least number of active neighbors in U activate first. After each activation, the underlying bipartite graph changes according to which node activates

and which nodes deactivate. Hence the subsequent activations in V depend on how the graph changes, as well as on the evolution of the network, since the queue lengths (and hence the activation rates) change over time as well. Recall that the queue lengths all have a good behavior (see Remark 3.4.7).

Theorems. To keep track of this evolution, we defined a greedy algorithm in Section 3.2. If we run the algorithm once, then it generates a specific path of activating nodes in V . This is enough to determine the leading order of the transition time as $r \rightarrow \infty$, since it only depends on the maximum least degree d^* , which is the same for all the paths that can be generated. Moreover, given d^* , we can immediately determine whether we are in the subcritical, critical or supercritical regime. If we are interested in the pre-factor of the mean transition time and in its law, then we need to generate all possible paths. Theorem 3.3.2 shows that we can split the mean transition time into a weighted sum over all possible paths of the mean nucleation times associated with each activation in the path. Theorem 3.3.3 gives the mean transition time conditional on the path and shows that the outcome is non-trivial both in the subcritical and critical regimes. Theorem 3.3.5 gives the law conditional on the path, but fails to capture the critical regime. The reason is that there are intricate dependencies between the subsequent nucleation times along the path.

§3.4 Nucleation times and queue lengths

In Section 3.4.1 we introduce the concept of asymptotic independence of forks and we show that in the subcritical and critical regimes competing forks can be treated as if they were disjoint, in the limit as $r \rightarrow \infty$ (Proposition 3.4.1). In Section 3.4.2 we study the mean and the law of the next nucleation time by using techniques from metastability and results from Section 3.4.1 (Propositions 3.4.3 and 3.4.6). In Section 3.4.3 we show how the mean queue lengths change according to which node activates in V (Theorem 3.4.8).

§3.4.1 Asymptotic independence of forks

In this section we show that, as $r \rightarrow \infty$, forks can be treated as being independent of each other even when they share some nodes. We introduce the concept of *asymptotic independence* of forks, which holds only as $r \rightarrow \infty$ and which allows us to treat overlapping forks as if they were disjoint. We show that the nucleation time of a fork is not influenced by the behavior of other forks sharing nodes with it.

In Chapter 2 it is shown that, as soon as all the nodes in U of a complete bipartite graph become simultaneously inactive, the first node in V (and subsequently all the others nodes) activates in a very short time interval, negligible compared to the time it takes for all the nodes in U to deactivate. Hence, the time it takes for the nodes in U to become all simultaneously inactive is the same as the time it takes for the first node in V to activate, up to an error term that is negligible as $r \rightarrow \infty$. In our setting, to study the nucleation times of forks it is enough to study the time it takes for all their respective nodes in U to deactivate, without considering the set V .

Proposition 3.4.1 (Asymptotic independence).

Consider the graph G_k and the \bar{d}_k -fork W , where \bar{d}_k is the minimum degree of the nodes in V_k . Denote by T_W the time it takes for fork W to nucleate for the first time. Consider the event

$$\mathcal{E} = \left\{ \exists \{s_1, \dots, s_\alpha\} \subset \{u_1, \dots, u_{\bar{d}_k}\} = W \cap U_k : \exists 0 \leq t < \bar{\tau}_k \text{ s.t.} \right. \\ \left. X_{s_i} \left(\sum_{j=1}^{k-1} \bar{\tau}_j + t \right) = 0 \ \forall i = 1, \dots, \alpha \right\} \quad (3.48)$$

of having a subset of α nodes in U_k belonging to fork W that are simultaneously inactive at a time t after the last nucleation. The following statements hold.

(i) The mean nucleation time of W satisfies

$$\mathbb{E}_u[T_W | \mathcal{E}] = \mathbb{E}_u[T_W] [1 + o(1)], \quad r \rightarrow \infty. \quad (3.49)$$

(ii) The law of the nucleation time of W satisfies

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(T_W > x | \mathcal{E}) = \lim_{r \rightarrow \infty} \mathbb{P}_u(T_W > x), \quad x \in [0, \infty). \quad (3.50)$$

Proof. We prove the two statements separately.

(i) We denote by \mathcal{S} the event that after time t all the nodes of W that are still active become simultaneously inactive before any of the inactive nodes in $\{s_1, \dots, s_\alpha\}$ activates again. We know that the time it takes for $\bar{d}_k - \alpha$ nodes to become simultaneously inactive is an exponential random variable $T_{\mathcal{S}}$ with mean of order $r^{\beta(\bar{d}_k - \alpha - 1)}$, while the time it takes for one of the α inactive nodes to activate is an exponential random variable with mean of order $1/r^\beta$. Hence the probability of \mathcal{S} is of order $r^{-\beta(\bar{d}_k - \alpha)} = o(1)$. If \mathcal{S} occurs, then W nucleates in time $T_W = t + T_{\mathcal{S}}$. Note that t must be of order $r^{\beta(\alpha - 1)}$, hence

$$\mathbb{E}_u[T_W | \mathcal{E} \cap \mathcal{S}] = O(r^{\beta(\alpha - 1)}) + O(r^{\beta(\bar{d}_k - \alpha - 1)}) = o(r^{\beta(\bar{d}_k - 1)}), \quad r \rightarrow \infty. \quad (3.51)$$

On the other hand, if the complementary event \mathcal{S}^C occurs, then, with high probability as $r \rightarrow \infty$, in a negligible time $o(1)$ the network reaches the state with all the nodes $u_1, \dots, u_{\bar{d}_k}$ active, and from there it takes time $\mathbb{E}_u[T_W]$ for W to nucleate. Hence

$$\mathbb{E}_u[T_W | \mathcal{E} \cap \mathcal{S}^C] = o(1) + \mathbb{E}_u[T_W], \quad r \rightarrow \infty. \quad (3.52)$$

Putting the two complementary events together, we obtain that

$$\begin{aligned} \mathbb{E}_u[T_W | \mathcal{E}] &= \mathbb{E}_u[T_W | \mathcal{E} \cap \mathcal{S}] \mathbb{P}_u(\mathcal{S}) + \mathbb{E}_u[T_W | \mathcal{E} \cap \mathcal{S}^C] \mathbb{P}_u(\mathcal{S}^C) \\ &= o(r^{\beta(\bar{d}_k - 1)}) o(1) + (o(1) + \mathbb{E}_u[T_W]) (1 - o(1)) \\ &= \mathbb{E}_u[T_W] [1 + o(1)], \quad r \rightarrow \infty. \end{aligned} \quad (3.53)$$

(ii) Using the complementary events \mathcal{S} and \mathcal{S}^C , we can write for all $x \geq 0$

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}_u(T_W > x \mid \mathcal{E}) &= \lim_{r \rightarrow \infty} \mathbb{P}_u(\{T_W > x\} \cap \mathcal{S} \mid \mathcal{E}) \mathbb{P}_u(\mathcal{S}) \\ &\quad + \lim_{r \rightarrow \infty} \mathbb{P}_u(\{T_W > x\} \cap \mathcal{S}^C \mid \mathcal{E}) \mathbb{P}_u(\mathcal{S}^C) \\ &= \lim_{r \rightarrow \infty} \mathbb{P}_u(T_W > x), \end{aligned} \quad (3.54)$$

since $\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{S}) = 0$ and, when conditioning on \mathcal{S}^C , with high probability as $r \rightarrow \infty$, the network reaches the initial state in a negligible time after t , hence it behaves as if at time t all nodes in U were active.

□

The above proposition shows that, as $r \rightarrow \infty$, the mean nucleation time of a fork and its law are not influenced by the fact that some of its nodes might be simultaneously inactive at some time. The intuition is that, as $r \rightarrow \infty$, the nucleation of a fork is so hard to achieve and takes so long that sharing some nodes with other forks does not help to make the nucleation happen appreciably faster. The network tends to quickly reach the initial state with all the remaining nodes in U active, and hence the nucleation time of a fork can be seen as the time it takes for its nodes in U to deactivate starting from all of them being active. In particular, in case of overlapping forks, the nucleation time of a fork is not influenced by the behavior of other forks sharing nodes with it.

§3.4.2 Next nucleation time

Given the graph G_k , consider the next nucleation time $\bar{\tau}_k$ from Definition 3.2.8. The next node that activates is the one that completes the fastest nucleation among the n_k nodes with least degree. We want to find an expression for $\mathbb{E}_u[\bar{\tau}_k]$.

In Appendix A we show the computations for the mean next nucleation time in the case when the competing forks are disjoint, hence described by i.i.d. random variables. Recall that in the subcritical regime we are considering a minimum of nucleation times that are exponential random variables, while in the critical regime we are considering a minimum of nucleation times that follow a truncated polynomial law (see Theorem 2.1.6). By using Proposition 3.4.1, we are also able to give explicit asymptotics for the mean next nucleation time without assuming the forks being independent.

Each nucleation of a fork can be seen as a successful escape from a metastable state, which is represented by the initial state where the nodes in U_k in the fork are active and the node in V_k in the fork is inactive. When considering multiple forks, we can view the network as an ergodic Markov process on a state space Ω representing the collection of all the feasible joint activity states of G_k . The first nucleation can be described by a regenerative process where the Markov process leaves a metastable state x_0 (with all the nodes in U_k active) and reaches a stable set S , which represents the set of states where at least one of the forks of minimum degree has all its nodes in U_k simultaneously inactive. The set S is rare for the Markov process, in the sense

that the probability of reaching S starting from x_0 is small. We denote by $T_{x_0 \rightarrow S}^k = \bar{\tau}_k$ the time it takes to go from x_0 to S .

Lemma 3.4.2 (Mean return time to metastable state).

For $k = 1, \dots, N$, suppose that $k - 1$ nodes in V already activated. Then, with high probability as $r \rightarrow \infty$, the time $R_{U_k}^x$ it takes for the network G_k to reach the state with all the nodes in U_k active (the metastable state x_0) starting from any other state x is negligible, i.e.,

$$\mathbb{E}_u[R_{U_k}^x] = o(1), \quad r \rightarrow \infty. \quad (3.55)$$

In particular, let $R_{U_k}^{k-1}$ be the time it takes for the network G_k to reach the state with all the nodes in U_k active starting from the moment the $(k-1)$ -th node in V activated. Then, with high probability as $r \rightarrow \infty$,

$$\mathbb{E}_u[R_{U_k}^{k-1}] = o(1), \quad r \rightarrow \infty. \quad (3.56)$$

Proof. At any time t , the activation and deactivation of a node $u \in U_k$ are described by i.i.d. exponential random variables with rates $g_U(Q_u(t))$ and 1, respectively. Hence, an active node takes on average one unit of time to deactivate, while an inactive node u takes on average $1/g_U(Q_u(t))$ time to activate. Since in the subcritical and critical regimes the queue lengths at any node at any moment are of order r (see Section 3.4.3 for more details), we can say that $1/g_U(Q_u(t)) = o(1)$ for each $u \in U_k$. Suppose that, at some time t , node $u_1 \in U_k$ is inactive and node $u_2 \in U_k$ is active, i.e., $X_{u_1}(t) = 0$ and $X_{u_2}(t) = 1$. Since

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(u_1 \text{ activates} < u_2 \text{ deactivates}) = 1, \quad (3.57)$$

and there is a finite number of nodes in U_k , with high probability as $r \rightarrow \infty$, starting from any state x all the nodes in U_k will be active on average in time $o(1)$. Hence, as $r \rightarrow \infty$, $\mathbb{E}_u[R_{U_k}^x] = o(1)$, and in particular $\mathbb{E}_u[R_{U_k}^{k-1}] = o(1)$. \square

We are now ready to state a result for the mean next nucleation time in the subcritical and critical regimes.

Proposition 3.4.3 (Mean next nucleation time).

Consider the graph G_k . Recall that \bar{d}_k is the minimum degree of a node in V_k , n_k is the number of forks of degree \bar{d}_k in G_k , and h_k is as in (3.74).

(I) $\beta \in (0, \frac{1}{\bar{d}_k - 1})$: subcritical regime. The mean next nucleation time satisfies

$$\mathbb{E}_u[\bar{\tau}_k] = f_k \mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}] = f_k F_{\text{sub}}^k \mathbb{E}_u[Q_U^{k-1}]^{\beta(\bar{d}_k - 1)} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.58)$$

with

$$f_k = \frac{1}{n_k}. \quad (3.59)$$

(II) $\beta = \frac{1}{\bar{d}_k - 1}$: critical regime. The mean next nucleation time satisfies

$$\mathbb{E}_u[\bar{\tau}_k] = f_k \mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}] = f_k F_{cr}^k \mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}] [1 + o(1)], \quad r \rightarrow \infty, \quad (3.60)$$

with

$$f_k = \frac{\bar{d}_k B^{-(\bar{d}_k - 1)} + (c - \rho_U)}{n_k \bar{d}_k B^{-(\bar{d}_k - 1)} + (c - \rho_U)}. \quad (3.61)$$

Proof. By Proposition 3.4.1, as $r \rightarrow \infty$ we may consider arbitrarily overlapping forks as if they were disjoint. Therefore the computations for the mean next nucleation time carried out in Appendix C for the case of disjoint forks can be used for the case of overlapping forks as well. For completeness, in the subcritical regime (I) we offer a proof that uses a different argument, which cannot be used in the critical regime (II) because the queues are changing on scale r over time.

Consider the stationary distribution π of the Markov process mentioned above. The probability of the set S is given by

$$\pi(S) = \sum_{j=1}^{n_k} \pi(S_j) [1 + o(1)] = n_k \left(\frac{1}{B (Q_U^{k-1})^\beta} \right)^{\bar{d}_k} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.62)$$

where S_j is the event that the j -th fork has all its nodes simultaneously inactive. The terms representing multiple forks with all their nodes simultaneously inactive contribute in a negligible way to $\pi(S)$. Moreover, we know that, for $j = 1, \dots, n_k$,

$$\begin{aligned} \pi(S_j) &= \frac{\mathbb{E}_u[\text{time spent in } S_j]}{\mathbb{E}_u[\text{time spent in } S_j] + \mathbb{E}_u[T_{x_0 \rightarrow S_j}^k]} \\ &= \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta} + F_{\text{sub}}^k (Q_U^{k-1})^{\beta(\bar{d}_k-1)}} [1 + o(1)] \\ &= \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{F_{\text{sub}}^k (Q_U^{k-1})^{\beta(\bar{d}_k-1)}} [1 + o(1)] = \left(\frac{1}{B (Q_U^{k-1})^\beta} \right)^{\bar{d}_k} [1 + o(1)], \quad r \rightarrow \infty. \end{aligned} \quad (3.63)$$

This proves (3.62).

Using the same type of argument, we can compute $\mathbb{E}_u[T_{x_0 \rightarrow S}]$. Indeed,

$$\begin{aligned} \pi(S) &= \frac{\mathbb{E}_u[\text{time spent in } S]}{\mathbb{E}_u[\text{time spent in } S] + \mathbb{E}_u[T_{x_0 \rightarrow S}^k]} \\ &= \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta} + \mathbb{E}_u[T_{x_0 \rightarrow S}^k]} [1 + o(1)] = \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{\mathbb{E}_u[T_{x_0 \rightarrow S}^k]} [1 + o(1)], \quad r \rightarrow \infty. \end{aligned} \quad (3.64)$$

After inverting, we get

$$\begin{aligned} \mathbb{E}_u[\bar{\tau}_k] &= \mathbb{E}_u[T_{x_0 \rightarrow S}^k] = \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{\pi(S)} [1 + o(1)] = \frac{\frac{1}{\bar{d}_k} \frac{1}{B (Q_U^{k-1})^\beta}}{n_k \left(\frac{1}{B (Q_U^{k-1})^\beta} \right)^{\bar{d}_k}} [1 + o(1)] \\ &= f_k F_{\text{sub}}^k (Q_U^{k-1})^{\beta(\bar{d}_k-1)} [1 + o(1)], \quad r \rightarrow \infty, \end{aligned} \quad (3.65)$$

with

$$f_k = \frac{1}{n_k}. \quad (3.66)$$

This completes the proof. \square

Corollary 3.4.4 (Pre-factor adjustment).

Given the graph G_k , conditional on the next activating node of degree \bar{d}_k ,

$$\mathbb{E}_u[\bar{\tau}_k | Y_k = i_k] = \mathbb{E}_u \left[\min_{v \in V_k} \mathcal{T}_v^{Q^{k-1}} \mid Y_k = i_k \right] = f_k \mathbb{E}_u[\mathcal{T}_{v_{i_k}}^{Q^{k-1}}], \quad r \rightarrow \infty, \quad (3.67)$$

where f_k is as in (3.59) or (3.61) when a subcritical node or a critical node activates, respectively.

Proof. The claim follows from Proposition 3.4.3. \square

In the subcritical regime (I), the queue lengths do not change on scale r and therefore the renewal theory developed in [49] applies, which is tailored to exponential behavior in metastable regimes. In the critical regime (II), however, the queue lengths do change on scale r and [49] does not apply. For details, see Section 3.4.3. Recall that Ω is the state space of the Markov process and that, in our notation, $\bar{\tau}_k = T_{x_0 \rightarrow S}$.

Definition 3.4.5 (Recurrence property).

Let $H > 0$ and $h \in (0, 1)$. We say that the pair (x_0, S) satisfies the property $\text{Rec}(H, h)$ if

$$\sup_{x \in \Omega} \mathbb{P}(T_{x \rightarrow \{x_0, S\}} > H) \leq h. \quad (3.68)$$

The following result is the equivalent of [49, Theorem 2.3].

Proposition 3.4.6 (Law of the next nucleation time).

Consider the pair (x_0, S) such that the property $\text{Rec}(H, h)$ holds for $0 < H < \mathbb{E}_u[\bar{\tau}_k]$, with $\epsilon = H/\mathbb{E}_u[\bar{\tau}_k]$ and h sufficiently small. Then there exist functions $C(\epsilon, h)$ and $\lambda(\epsilon, h)$, satisfying $C(\epsilon, h), \lambda(\epsilon, h) \rightarrow 0$ as $\epsilon, h \rightarrow 0$, such that, for any $t > 0$,

$$\left| \mathbb{P} \left(\frac{\bar{\tau}_k}{\mathbb{E}_u[\bar{\tau}_k]} > t \right) - e^{-t} \right| \leq C e^{-(1-\lambda)t}. \quad (3.69)$$

Proof. We choose H to be a constant, and without loss of generality set $H = 1$. We claim that the pair (x_0, S) satisfies the property $\text{Rec}(H, h)$ with h sufficiently small. Indeed, starting from any state $x \in \Omega$, the network reaches the set $\{x_0, S\}$ in a small time $o(1)$.

If the starting state x is one of the states S_j , $j = 1, \dots, n_k$, corresponding to the set S , then we are done. Otherwise, by Lemma 3.4.2, the metastable state x_0 attracts in time $o(1)$ every state x for which some forks have some nodes in U inactive. It is therefore immediate that, with high probability as $r \rightarrow \infty$, $T_{x \rightarrow \{x_0, S\}}$ is smaller than H , which is what we need in order to claim that (3.68) holds when h is sufficiently small. Note that we can let $h \rightarrow 0$ as $r \rightarrow \infty$.

We recover from Proposition 3.4.3 that the ratio between H and the mean next nucleation time is sufficiently small. Indeed, $\epsilon = H/\mathbb{E}_u[\bar{\tau}_k] \rightarrow 0$ as $r \rightarrow \infty$. Hence a straightforward application of [49, Theorem 2.3] allows us to conclude that the next nucleation time divided by its mean follows an exponential law with unit rate. \square

§3.4.3 Updated queue lengths

In this section we analyze in more detail how the mean queue lengths change over time and how they affect the mean nucleation times associated with each step of the algorithm.

Remark 3.4.7 (Good behavior).

Recall from Definition 2.3.1 and Lemma 2.3.2 in Chapter 2 that, with high probability as $r \rightarrow \infty$, the queue lengths all have a good behavior in the interval $[0, T_U(r)]$ with $T_U(r) = \frac{\gamma_U}{c - \rho_U} r [1 + o(1)]$ as $r \rightarrow \infty$, representing the expected time it takes for the queue lengths to hit zero. More precisely, for $\delta > 0$ small enough and for all $t \in [0, T_U(r)]$,

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\mathbb{E}_u[Q_U(t)] - \delta r \leq Q_U(t) \leq \mathbb{E}_u[Q_U(t)] + \delta r \right) = 1, \quad (3.70)$$

which means that the queue lengths are always close to their mean for all times smaller than $T_U(r)$.

Recall that we start with initial queue lengths $Q^0 = (Q_U^0, Q_V^0) = (\gamma_U r, \gamma_V r)$. We are interested in studying how the queue lengths change along a fixed path, depending on which types of forks we encounter at each activation. Fix a path and consider the sequence of nodes activating in V .

Similarly to (3.33), the next nucleation time $\bar{\tau}_k = \min_{v \in V_k} \mathcal{T}_v^{Q^{k-1}}$ (recall Definition 3.2.8) satisfies

$$\mathbb{E}_u[\bar{\tau}_k] = f'_k r^{1 \wedge \beta(\bar{d}_k - 1)} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.71)$$

where f'_k depends on f_k , on the constants $F_{\text{sub}}^k, F_{\text{cr}}^k, F_{\text{sup}}^k$ (for the three regimes, respectively), and on the mean updated queue lengths. The following theorem shows how the mean queue lengths change according to which type of node activates in V .

Theorem 3.4.8 (Mean updated queue lengths).

Let $(\bar{d}_k)_{k=1}^N$ be the sequence of degrees in a fixed path and $d^* = \max_{1 \leq k \leq N} \bar{d}_k$.

(I) $\beta \in (0, \frac{1}{d^* - 1})$: *subcritical regime*. After step k , the mean queue length at a node in U is

$$\mathbb{E}_u[Q_U^k] = \gamma_U r [1 + o(1)], \quad r \rightarrow \infty. \quad (3.72)$$

(II) $\beta = \frac{1}{d^* - 1}$: *critical regime*. After step k , the mean queue length at a node in U , after h_k critical nodes in V activated, is

$$\mathbb{E}_u[Q_U^k] = \gamma_U^{(h_k)} r [1 + o(1)], \quad r \rightarrow \infty, \quad (3.73)$$

with

$$\gamma_U^{(h_k)} = \gamma_U - (c - \rho_U) \sum_{\substack{1 \leq i \leq k \\ i: \bar{d}_i = d^*}} f'_i > 0, \quad (3.74)$$

where for a critical node v_i the coefficient f'_i is defined in a recursive way as

$$f'_i = \frac{1}{n_i \bar{d}_i B^{-(\bar{d}_i - 1)} + (c - \rho_U)} \left(\gamma_U - (c - \rho_U) \sum_{\substack{1 \leq j \leq i-1 \\ j: \bar{d}_j = d^*}} f'_j \right) > 0. \quad (3.75)$$

- (III) $\beta \in (\frac{1}{d^*-1}, \infty)$: *supercritical regime*. After step k , the mean queue length at a node in U , if any supercritical node in V activated, is

$$\mathbb{E}_u[Q_U^k] = o(r), \quad r \rightarrow \infty. \quad (3.76)$$

Proof. We treat the three regimes separately.

- (I) $\beta \in (0, \frac{1}{d^*-1})$. All the nodes in V are subcritical, in particular the first node $v_1 \in V$. Then $\mathbb{E}_u[\bar{\tau}_1] = o(r)$ as $r \rightarrow \infty$. The mean queue lengths at nodes in U after node v_1 activates are

$$\begin{aligned} \mathbb{E}_u[Q_U(\bar{\tau}_1)] &= \mathbb{E}_u[\gamma_U r - (c - \rho_U)\bar{\tau}_1] = \gamma_U r - (c - \rho_U)\mathbb{E}_u[\bar{\tau}_1] \\ &= \gamma_U r [1 + o(1)], \quad r \rightarrow \infty, \end{aligned} \quad (3.77)$$

which means that after the first activation the mean queue lengths are the same as before, up to an error term $o(1)$. Iterating this reasoning, we conclude that the mean queue lengths remain approximately the same as long as subcritical nodes in V activate.

- (II) $\beta = \frac{1}{d^*-1}$. If the first node $v_1 \in V$ is subcritical, then the time it takes to nucleate its fork does not influence the mean queue lengths by much, as seen in (I). Without loss of generality, we may therefore assume that v_1 is critical. Then $\mathbb{E}_u[\bar{\tau}_1] = f'_1 r$ is of order r . The mean queue lengths at nodes in U after node v_1 activates are

$$\begin{aligned} \mathbb{E}_u[Q_U(\bar{\tau}_1)] &= \mathbb{E}_u[\gamma_U r - (c - \rho_U)\bar{\tau}_1] = \gamma_U r - (c - \rho_U)\mathbb{E}_u[\bar{\tau}_1] \\ &= (\gamma_U - (c - \rho_U)f'_1)r [1 + o(1)] = \gamma_U^{(1)} r [1 + o(1)], \quad r \rightarrow \infty, \end{aligned} \quad (3.78)$$

where $\gamma_U^{(1)} = \gamma_U - (c - \rho_U)f'_1 > 0$.

If the second node $v_2 \in V$ is subcritical, then again the time it takes to nucleate its fork does not influence the mean queue lengths by much. Assume therefore that v_2 is critical. Then the fork requires a nucleation time of order r , namely, $\mathbb{E}_u[\bar{\tau}_2] = f'_2 r$. The mean queue lengths at nodes in U after node $v_2 \in V$ activates are

$$\begin{aligned} \mathbb{E}_u[Q_U(\bar{\tau}_1 + \bar{\tau}_2)] &= \mathbb{E}_u[\gamma_U r - (c - \rho_U)(\bar{\tau}_1 + \bar{\tau}_2)] \\ &= \gamma_U r - (c - \rho_U)(\mathbb{E}_u[\bar{\tau}_1] + \mathbb{E}_u[\bar{\tau}_2]) \\ &= (\gamma_U - (c - \rho_U)(f'_1 + f'_2))r [1 + o(1)] \\ &= \gamma_U^{(2)} r [1 + o(1)], \quad r \rightarrow \infty, \end{aligned} \quad (3.79)$$

where $\gamma_U^{(2)} = \gamma_U - (c - \rho_U)(f'_1 + f'_2) > 0$.

More generally, assume that h_k critical nodes activated in the first k steps. Then

$$\mathbb{E}_u[Q_U^k] = \gamma_U^{(h_k)} r [1 + o(1)], \quad r \rightarrow \infty, \quad (3.80)$$

with

$$\gamma_U^{(h_k)} = \gamma_U - (c - \rho_U) \sum_{\substack{1 \leq i \leq k \\ i: \bar{d}_i = d^*}} f'_i > 0, \quad (3.81)$$

where the last sum is over all the h_k critical nodes. Each of them contributes with a positive coefficient f'_i which is given by the recursive relation

$$\begin{aligned} f'_i &= f_i F_{\text{cr}}^i \gamma_U^{(h_{i-1})} \\ &= \frac{\bar{d}_i B^{-(\bar{d}_i-1)} + (c - \rho_U)}{n_i \bar{d}_i B^{-(\bar{d}_i-1)} + (c - \rho_U)} \frac{1}{\bar{d}_i B^{-(\bar{d}_i-1)} + (c - \rho_U)} \gamma_U^{(h_{i-1})} \\ &= \frac{1}{n_i \bar{d}_i B^{-(\bar{d}_i-1)} + (c - \rho_U)} \left(\gamma_U - (c - \rho_U) \sum_{\substack{1 \leq j \leq i-1 \\ j: \bar{d}_j = d^*}} f'_j \right). \end{aligned} \quad (3.82)$$

Note that the coefficients f'_k introduced in (3.71) are well defined for every $k = 1, \dots, N$, but in the above computations we are only interested in the ones associated with the critical nodes. For example,

$$f'_1 = \begin{cases} \frac{1}{n_1} \frac{1}{\bar{d}_1 B^{-(\bar{d}_1-1)}} \gamma_U, & \text{if } \bar{d}_1 < d^*, \\ \frac{1}{n_1 \bar{d}_1 B^{-(\bar{d}_1-1)} + (c - \rho_U)} \gamma_U, & \text{if } \bar{d}_1 = d^*. \end{cases} \quad (3.83)$$

- (III) $\beta \in [\frac{1}{d^*-1}, \infty)$. If the first node $v_1 \in V$ is subcritical, then its nucleation time does not influence the mean queue lengths by much, as seen in (I). If v_1 is critical, then the mean queue lengths decrease but remain of order r , as seen in (II). We therefore assume that v_1 is supercritical. Then $\mathbb{E}_u[\bar{\tau}_1] = \frac{\gamma_U}{c - \rho_U} r [1 + o(1)]$, as $r \rightarrow \infty$. Indeed, from Theorem 2.1.6 we know that the mean nucleation time of a supercritical fork is given by the expected time it takes for the queue length to hit zero. This holds for every supercritical node in V and therefore it is true also for $\mathbb{E}_u[\bar{\tau}_1]$. Hence, the mean queue lengths at nodes in U after node $v_1 \in V$ activates are

$$\mathbb{E}_u[Q_U(\bar{\tau}_1)] = \mathbb{E}_u[\gamma_U r - (c - \rho_U) \bar{\tau}_1] = \gamma_U r - (c - \rho_U) \mathbb{E}_u[\bar{\tau}_1] = o(r), \quad r \rightarrow \infty. \quad (3.84)$$

More generally, the mean queue lengths become $o(r)$ as soon as the first supercritical node activates, independently of which nodes activated before. Thus, after any step k the mean queue length at a node in U , if any supercritical node activated, is

$$\mathbb{E}_u[Q_U^k] = o(r), \quad r \rightarrow \infty. \quad (3.85)$$

□

In summary, we have shown that if a subcritical node activates, then we do not change the mean queue lengths at nodes in U by much: they only decrease by a factor $o(1)$. On the other hand, if a critical node activates, then the mean queue lengths drop significantly, but still remain of order r . Finally, if a supercritical node activates,

then the mean queue lengths become $o(r)$, and remain so during all the successive nucleations. With the help of (3.71) we know how to relate the mean next nucleation times of the forks to the mean updated queue lengths after each activation. Hence we know that, once a node that contributes order r to the total mean transition time activates, we can ignore the contribution of all the previous and all the subsequent subcritical nodes. Once a supercritical node activates, we can ignore the contribution of all the subsequent nodes, since their queue lengths are $o(r)$.

§3.5 Analysis of the algorithm

In Section 3.5.1 we describe how the algorithm acts on an arbitrary bipartite graph. (In Section 3.2.4 we already illustrated this via an example.) In Section 3.5.2 we prove the greediness and the consistency of the algorithm.

§3.5.1 Recursion

Consider the graph $G = G_1 = ((U_1, V_1), E_1)$. The first node activating in V_1 is the one with the least degree, since this requires the least number of nodes in U_1 to become simultaneously inactive. Since the expected time for m nodes in U_1 to become simultaneously inactive is of order $r^{1 \wedge \beta(m-1)}$, with high probability as $r \rightarrow \infty$, the first node that activates in V_1 is v_{Y_1} such that $d(v_{Y_1}) = \bar{d}_1 = \min_{v \in V_1} d(v)$, where $d(v)$ denotes the degree of node v in the graph G_1 . We make the algorithm pick as first node a node v_{Y_1} with least degree in V_1 . If there are multiple nodes with the same least degree, then the algorithm chooses one of them uniformly at random. If the least degree \bar{d}_1 is such that $\beta(\bar{d}_1 - 1) > 1$, then the algorithm chooses a node uniformly at random among all nodes in V_1 . Let $G'_1(U'_1, V'_1)$ be the complete bipartite subgraph of G_1 with $U'_1 = \{u \in U_1 : uv_{Y_1} \text{ is an edge of } G_1\}$ and $V'_1 = \{v_{Y_1}\}$. According to Theorem 2.1.6, the associated nucleation time $\mathcal{T}_{v_{Y_1}}^{Q^0}$ satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_1}}^{Q^0}] = F^1(Q_U^0)^{1 \wedge \beta(\bar{d}_1 - 1)} [1 + o(1)], \quad r \rightarrow \infty. \quad (3.86)$$

Reasoning as above, we see that the algorithm picks as second node a node v_{Y_2} with the least number of active neighbors left in G . Consider the bipartite graph $G_2(U_2, V_2)$ with $U_2 = U_1 \setminus U'_1$ and $V_2 = V_1 \setminus V'_1 = V_1 \setminus \{v_{Y_1}\}$. If we denote by $d_2(v)$ the degree of a node $v \in V_2$ in G_2 , then v_{Y_2} is such that $d_2(v_{Y_2}) = \bar{d}_2 = \min_{v \in V_2} d_2(v)$. If there are multiple nodes with the same least degree, then the algorithm again chooses one uniformly at random. If the least degree \bar{d}_2 is such that $\beta(\bar{d}_2 - 1) > 1$, then we choose a node uniformly at random among all nodes in V_2 . Let $G'_2(U'_2, V'_2)$ be the complete bipartite subgraph of G with $U'_2 = \{u \in U_2 : uv_{Y_2} \text{ is an edge of } G_2\}$ and $V'_2 = \{v_{Y_2}\}$. The associated nucleation time $\mathcal{T}_{v_{Y_2}}^{Q^1}$ satisfies

$$\mathbb{E}_u[\mathcal{T}_{v_{Y_2}}^{Q^1}] = F^2(Q_U^1)^{1 \wedge \beta(\bar{d}_2 - 1)} [1 + o(1)], \quad r \rightarrow \infty, \quad (3.87)$$

Iterating this procedure until all the nodes in V_1 are active, we find one of the paths that the algorithm follows in terms of successive activation of the nodes in V_1 .

Note that, depending on the choice the algorithm makes at each step, there may be different paths for the activation.

§3.5.2 Greediness and consistency

We first prove Lemma 3.2.5. After that we prove Propositions 3.2.6–3.2.7.

Proof of Lemma 3.2.5. The proof is by contradiction. Suppose that $d_a^* > d_b^*$. Denote by $d_{k,a}(v)$ and $d_{k,b}(v)$ the degrees of node $v \in V_k$ at step $k = 1, \dots, N$ in paths a and b , respectively.

We start by considering node $w_1 \in V$ such that, at some step k_1^a in path a , $d_{k_1^a,a}(w_1) = \bar{d}_{k_1^a,a} = d_a^*$. Then $d(w_1) \geq d_a^*$ in G . On the other hand, in path b , when w_1 activates at some step k_1^b , it has degree $d_{k_1^b,b}(w_1) \leq d_b^*$. This implies that some of the edges of w_1 (at least $d_a^* - d_b^*$ edges) have already been processed via previous forks in path b . At least one of these forks must have nucleated before the fork of w_1 , in path b but not in path a , say, the fork of w_2 . Hence there exists a node $w_2 \in V$ such that, at some step $k_2^b < k_1^b$ in path b , $d_{k_2^b,b}(w_2) \leq d_b^*$. This node has not yet activated by step k_1^a in path a , so it must be that $d_{k_1^a,a}(w_2) \geq d_a^*$, otherwise the algorithm would choose node w_2 before node w_1 . Say that node w_2 will activate at step $k_2^a > k_1^a$ in path a . Then, $d(w_2) \geq d_a^*$ in G . As before, this implies that some of its edges have already been processed with previous forks in path b . Again, at least one of these forks must have nucleated before the fork of w_2 , in path b but not in path a , say, the fork of w_3 . Hence there exists a node $w_3 \in V$ such that, at some step $k_3^b < k_2^b$ in path b , $d_{k_3^b,b}(w_3) \leq d_b^*$. This node has not yet activated by step k_2^a in path a , nor by step k_1^a , so $d_{k_1^a,a}(w_3) \geq d_{k_1^a,a}(w_1) \geq d_a^*$, otherwise the algorithm would choose node w_3 before node w_1 . Hence $d(w_3) \geq d_a^*$ in G .

We can iterate this argument. Since there are only N nodes in V , we get a contradiction after we have considered all the nodes. \square

We are now able to prove the greediness and the consistency of the algorithm.

Proof of Proposition 3.2.6. By Lemma 3.2.5, we know that the maximum least degree of a path generated by the algorithm is the smallest possible. We know that the order of the mean transition time along a path is related to d^* and depends on the value of β . Hence, Lemma 3.2.5 implies that the mean transition time along a path generated by the algorithm is the shortest possible, in the sense that it has the smallest order of r possible. \square

Proof of Proposition 3.2.7. Lemma 3.2.5 proves equality for any two paths generated by the algorithm. This leads to the same order of the mean transition time. \square

Despite the fact that d^* does not depend on which path the algorithm generates, its multiplicity does. Figure 3.4 shows a graph on which the algorithm can generate two different paths with the same maximum least degree but with different multiplicity.

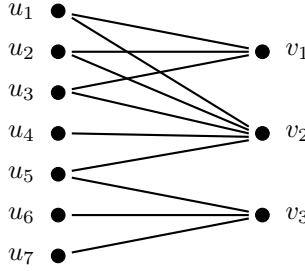


Figure 3.4: The algorithm may generate the path v_1, v_2, v_3 or the path v_3, v_1, v_2 with different multiplicity of d^* .

§3.5.3 Algorithm complexity

The algorithm we constructed can be implemented in different ways according to what we want to compute.

- In order to know the leading order of the mean transition time, it is enough to recover the maximum least degree d^* from the graph. By Proposition 3.2.7 we know that d^* is the same for all paths the algorithm can generate. Hence it is enough to run it once and comparing the value of d^* with the value of β we can immediately determine whether we are in the subcritical, critical or supercritical regime.

In this case the computational complexity of the algorithm is polynomial in the number of nodes in V , and so the leading order of the mean transition time is quickly determined. More precisely, the algorithm has a complexity of $\mathcal{O}(|U||V|^2)$.

- If we are interested in the precise asymptotics of the mean transition time and in its law as $r \rightarrow \infty$, then we need to compute the pre-factor of the leading order term. To do so, we need to run the algorithm multiple times, until all possible paths are generated, in order to recover all the possible sequences $(\bar{d}_k)_{k=1}^N$ and (n_k) . A proper approach is to let a (deterministic) depth-first search algorithm run through all possible paths and enumerate them. Theorem 3.3.2 shows that if we know the total mean transition time along each path, then we can recover the mean transition time of the graph.

In this case the computational complexity of the algorithm is factorial in the number of nodes in V , since it depends in a delicate manner on the architecture of the graph. More precisely, the algorithm has a complexity of $\mathcal{O}(|U||V|^2|V|!)$.

See [99] for a deeper analysis of the algorithm complexity.

§3.6 Proofs of the main results

The aim of this section is to prove the theorems in Section 3.3. In Section 3.6.1 we introduce some further definitions. In Section 3.6.2 we prove Lemmas 3.2.2 and 3.2.11. In the last three sections we prove the three main theorems, respectively: in Section 3.6.3 we prove Theorem 3.3.2, in Section 3.6.4 we prove Theorem 3.3.3, and in Section 3.6.5 we prove Theorem 3.3.5.

§3.6.1 Preparatory results

Consider an arbitrary bipartite graph $G = ((U, V), E)$ with $|V| = N$ and let v_1, \dots, v_N be the nodes in V . The activation path that the network follows is denoted by v_1^*, \dots, v_N^* , while the indices of the nodes that the algorithm picks are denoted by Y_1, \dots, Y_N (as in Definition 3.2.1). We want to study the transition time when the network follows a path generated by the algorithm. When conditioning the network on a specific activation order, we can write

$$\{Y_k = i\} = \{v_k^* = v_i\}, \quad (3.88)$$

in the sense that saying that the k -th index Y_k chosen by the algorithm equals i is equivalent to saying that the k -th activating node v_k^* equals v_i .

Definition 3.6.1 (Iteration graph).

For $k = 1, \dots, N$, suppose that $k - 1$ nodes in V already activated. Denote by $G_k = ((V_k, U_k), E_k)$ the subgraph of $G = ((U, V), E)$ consisting of:

- $V_k = V \setminus \{\{v_{Y_i}\}_{0 < i < k}\} \subseteq V$, the set of nodes in V that have not yet activated;
- $U_k = U \setminus \bigcup_{0 < i < k} N(v_{Y_i}) \subseteq U$, the set of nodes in U that are not neighbors of some of the nodes in V that already activated;
- $E_k = \{uv : u \in U_k, v \in V_k\} \subseteq E$, the set of edges between U_k and V_k .

Let \bar{d}_k be the minimum degree of the nodes in V_k and n_k be the number of least degree forks in G_k .

Definition 3.6.2 (Minimum degree subset).

Define the set of nodes with minimum degree in V as

$$M(V) = \{v' \in V : d(v') = \min_{v \in V} d(v)\}. \quad (3.89)$$

Lemma 3.6.3 (Probability of choosing the next node).

Given the graph G_k , in the subcritical and critical regimes, the probability that the next node activating in V_k is node v_i is

$$\mathbb{P}(Y_k = i) = \begin{cases} \frac{1}{n_k}, & \text{if } \beta \in (0, \frac{1}{\bar{d}_k - 1}], v_i \in M(V_k), \\ 0, & \text{if } \beta \in (0, \frac{1}{\bar{d}_k - 1}], v_i \in V_k \setminus M(V_k), \end{cases} \quad (3.90)$$

which depends on the sequence of nodes already active in V .

Proof. By construction, the algorithm picks nodes in $M(V_k)$ before before it picks nodes in $V_k \setminus M(V_k)$. It is therefore enough to count the number of forks of minimum degree at step k , which is n_k . \square

§3.6.2 Proof: activation sticks and selects low degrees

We next prove two lemmas from Section 3.2 that will be needed to prove Theorem 3.3.2 in Section 3.6.3.

Proof of Lemma 3.2.2. Recall Definition 3.1.1. We claim that if a node $u \in U$ deactivates and one of its neighbors in V activates at time t_u , then, with high probability as $r \rightarrow \infty$, it will not activate anymore after time t_u . The moments when u could possibly activate again are the moments when all its neighbors in V are simultaneously inactive. We consider the worst case scenario when u has only one active neighbor $v \in V$. Denote by t_v the first moment when v deactivates after t_u . This happens many times, since the activity period of a node is described by an exponential variable Z with rate 1. Instead, the inactivity periods are very short, since the nodes in V are very aggressive and the activation rates grow with the queue lengths, which tend to infinity as $r \rightarrow \infty$. We consider a time period of length equal to the total transition time, and we assume the transition time to be the longest possible (of order r). Then on average we have a number of possibilities for u to activate that is equal to

$$\frac{\mathbb{E}[\mathcal{T}_G^{Q^0}]}{\mathbb{E}[Z]} = \mathbb{E}[\mathcal{T}_G^{Q^0}] = Cr[1 + o(1)], \quad r \rightarrow \infty, \quad (3.91)$$

with C a positive constant. At each of these times, nodes u and v are both inactive and are competing with each other to activate again. Denote by Z_u and Z_v the lengths of the inactivity periods of u and v , respectively. Then $Z_u \simeq \text{Exp}(g_U(Q_u(t_v)))$ and $Z_v \simeq \text{Exp}(g_V(Q_v(t_v)))$ and so, with high probability as $r \rightarrow \infty$, node v activates before node u , i.e.,

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}(Z_v < Z_u) &= \lim_{r \rightarrow \infty} \frac{g_V(Q_v(t_v))}{g_U(Q_u(t_v)) + g_V(Q_v(t_v))} = \lim_{r \rightarrow \infty} \frac{K'r^{\beta'}}{Kr^{\beta} + K'r^{\beta'}} \\ &= \lim_{r \rightarrow \infty} \frac{1}{1 + (K/K')r^{-(\beta' - \beta)}} = 1, \end{aligned} \quad (3.92)$$

where we use that $\beta' > \beta$, and K, K' are positive constants.

Note that the queue lengths in U are always of order r , except when we are in the supercritical regime. In this regime we are not interested in the competition between u and v anymore, since we know how long the transition takes. The queue lengths in V start being of order r , increase while u is active and decrease when v is active, but remain of order r . Indeed, if there are other nodes in U that take long enough to activate so that the queue length of v becomes $o(r)$, then we must be in the supercritical regime. In the worst case scenario, nodes u and v compete with each other for the duration of the transition, i.e., order r times. The probability of v winning every competition is

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}(Z_v < Z_u)^r &= \lim_{r \rightarrow \infty} \left(\frac{1}{1 + (K/K')r^{-(\beta' - \beta)}} \right)^{Cr} = \lim_{r \rightarrow \infty} \left(e^{-(K/K')r^{-(\beta' - \beta)}} \right)^{Cr} \\ &= \lim_{r \rightarrow \infty} e^{-C(K/K')r^{-(\beta' - \beta - 1)}} = 1, \end{aligned} \quad (3.93)$$

where we use that $\beta' > \beta + 1$. Hence, with \mathbb{P}_u -probability tending to 1 as $r \rightarrow \infty$, node u will never win any competition against node v , and hence will remain blocked for the duration of the transition. \square

Proof of Lemma 3.2.11. We distinguish between node v being subcritical or critical.

- (I) $\beta \in (0, \frac{1}{d_k-1})$. From Theorem 2.1.6 we know the law of the nucleation time for the fork of v , namely,

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_v^{Q^{k-1}}}{\mathbb{E}_u[\mathcal{T}_v^{Q^{k-1}}]} > x \right) = \mathcal{P}_1(x) = e^{-x}, \quad x \in [0, \infty). \quad (3.94)$$

From the same equations we also know the law of the nucleation time for the fork of w , which depends on how β and $d(w)$ are related to each other. It is enough to verify that

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_w^{Q^{k-1}}}{\mathbb{E}_u[\mathcal{T}_w^{Q^{k-1}}]} > x \right) = \mathcal{P}(x), \quad x \in [0, \infty), \quad (3.95)$$

with $\mathcal{P}(x) \uparrow 1$ when $x \downarrow 0$ and $\mathcal{P}(x) \downarrow 0$ when $x \uparrow \infty$. To that end, assume that $\mathcal{T}_v^{Q^{k-1}}$ and $\mathcal{T}_w^{Q^{k-1}}$ deviate from their mean such that $\mathcal{T}_v^{Q^{k-1}} > \mathcal{T}_w^{Q^{k-1}}$. This happens with probability tending to 0 as $r \rightarrow \infty$, since the deviations must be of order r . Thus,

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_v^{Q^{k-1}} > \mathcal{T}_w^{Q^{k-1}}) = \lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{X}_{v,v}^{k-1} > \mathcal{X}_{v,w}^{k-1}), \quad (3.96)$$

where we abbreviate $\mathcal{X}_{v,w}^{k-1} = \mathcal{T}_w^{Q^{k-1}} / \mathbb{E}_u[\mathcal{T}_v^{Q^{k-1}}]$. For fixed $M < \infty$, we can split the right-hand side of (3.96) without the limit $r \rightarrow \infty$ as

$$\begin{aligned} \mathbb{P}_u(\mathcal{X}_{v,v}^{k-1} > \mathcal{X}_{v,w}^{k-1}) &= \mathbb{P}_u(\mathcal{X}_{v,v}^{k-1} > \mathcal{X}_{v,w}^{k-1} \mid \mathcal{X}_{v,w}^{k-1} > M) \mathbb{P}_u(\mathcal{X}_{v,w}^{k-1} > M) \\ &\quad + \mathbb{P}_u(\mathcal{X}_{v,v}^{k-1} > \mathcal{X}_{v,w}^{k-1} \mid \mathcal{X}_{v,w}^{k-1} \leq M) \mathbb{P}_u(\mathcal{X}_{v,w}^{k-1} \leq M) \\ &\leq \mathcal{P}_1(M) \mathbb{P}_u(\mathcal{X}_{v,w}^{k-1} > M) \\ &\quad + \mathbb{P}_u(\mathcal{X}_{v,v}^{k-1} > \mathcal{X}_{v,w}^{k-1} \mid \mathcal{X}_{v,w}^{k-1} \leq M) \mathbb{P}_u(\mathcal{X}_{v,w}^{k-1} \leq M). \end{aligned} \quad (3.97)$$

Pick $\epsilon > 0$ so small that, for $r > r_0(\epsilon)$,

$$\mathbb{P}_u(\mathcal{X}_{v,w}^{k-1} \leq M) = \mathbb{P}_u \left(\mathcal{X}_{w,w}^{k-1} \leq M \frac{\mathbb{E}_u[\mathcal{T}_v^{Q^{k-1}}]}{\mathbb{E}_u[\mathcal{T}_w^{Q^{k-1}}]} \right) \leq \mathbb{P}_u(\mathcal{X}_{w,w}^{k-1} \leq M\epsilon). \quad (3.98)$$

Letting $r \rightarrow \infty$ followed by $\epsilon \downarrow 0$, we get

$$\lim_{\epsilon \downarrow 0} \lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{X}_{w,w}^{k-1} \leq M\epsilon) = \lim_{\epsilon \downarrow 0} [1 - \mathcal{P}(M\epsilon)] = 0. \quad (3.99)$$

We can now let $M \rightarrow \infty$ and use (3.96)–(3.97) to arrive at

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_v^{Q^{k-1}} > \mathcal{T}_w^{Q^{k-1}}) = 0. \quad (3.100)$$

- (II) $\beta = \frac{1}{\bar{d}_k - 1}$. As before, we know the law of the nucleation time for the fork of v and w . As shown in Chapter 2, with high probability as $r \rightarrow \infty$, $\mathcal{T}_w^{Q^{k-1}} / \mathbb{E}_u[\mathcal{T}_w^{Q^{k-1}}]$ tends to 1. Moreover, with high probability as $r \rightarrow \infty$, any nucleation time of a complete bipartite graph in the critical regime (including the fork of v) is smaller than the transition time of the same graph in the supercritical regime. \square

§3.6.3 Proof: most likely paths

Proof of Theorem 3.3.2. We prove the three statements separately.

- (i) Assuming that the network does not follow the algorithm is equivalent to assuming that at some step k with $\beta \in (0, \bar{d}_k - 1]$ a node w that does not have a minimum degree is chosen instead of a node v with degree \bar{d}_k . The probability of a group of $d > \bar{d}_k$ nodes becoming simultaneously inactive before a group of \bar{d}_k nodes is equivalent to the probability of w activating before v , which satisfies

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_w^{Q^{k-1}} < \mathcal{T}_v^{Q^{k-1}}) = 0 \quad (3.101)$$

by Lemma 3.2.11. Hence, with high probability as $r \rightarrow \infty$, nodes in V activate in a greedy way, as described by the algorithm. By Lemma 3.2.2, we also know that the nodes in U that deactivated remain inactive for the duration of the transition process. Consequently, they do not influence any future activation attempt of the nodes in V , whose activation therefore follows the algorithm. In the supercritical regime, we are only interested in the order of activation of the nodes until the first supercritical node, for which the above reasoning still holds.

- (ii) Note that the queues Q^k depend on the sequence of indices (Y_1, \dots, Y_{k-1}) describing the order of the activating nodes in V . Indeed, we have seen in Section 3.4.3 that the queues change according to which nodes already activated. Moreover, for $k > 1$, also the probabilities $\frac{1}{n_k}$ depend on the sequence (Y_1, \dots, Y_{k-1}) . The reader should keep this in mind while going through the proof. The proof evolves in three steps.

1. Denote the graph $G = ((U, V), E)$ by $G_1 = ((U_1, V_1), E_1)$. Write

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] = \mathbb{E}_u[\mathcal{T}_{G_1}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] = \sum_{i_1: v_{i_1} \in V_1} \mathbb{E}_u[\mathcal{T}_{G_1}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_1 = i_1] \mathbb{P}(Y_1 = i_1). \quad (3.102)$$

By Lemma 3.6.3, when $\beta(\bar{d}_1 - 1) \leq 1$ not all the terms in the above sum have positive probability, only the ones corresponding to forks of minimum degree \bar{d}_1 (which all have the same probability). Recall that this probability is $\frac{1}{n_1}$. We can write the random variable $\mathcal{T}_{G_1}^{Q^0}$ as sum of three random variables

$$\mathcal{T}_{G_1}^{Q^0} = \bar{\tau}_1 + R_{U_2}^1 + \mathcal{T}_{G_2}^{Q^1}, \quad (3.103)$$

where $G_2 = ((U_2, V_2), E_2)$ is the subgraph with $U_2 = U_1 \setminus N(v_{Y_1})$, $V_2 = V_1 \setminus \{v_{Y_1}\}$ and $E_2 = E_1 \setminus \{(u, v) : u \in N(v_{Y_1})\}$, while $Q^1 = Q(\bar{\tau}_1)$ represents the updated queue lengths. The first variable represents the time it takes for the first node to activate, the second variable represents the time it takes (after the activation of the first node) to reach the state with all the nodes in U_2 active (see Lemma 3.4.2), while the third variable represents the transition time of the remaining graph when we take the first activating node out. Note that, by Corollary 3.4.4, if we condition the network to follow a path generated by the algorithm with a specific first activating node, then we get

$$\mathbb{E}_u[\bar{\tau}_1 \mid Y_1 = i_1] = f_1 \mathbb{E}_u[\mathcal{T}_{v_{i_1}}^{Q^0}], \quad r \rightarrow \infty, \quad (3.104)$$

where f_1 is the factor that arises from considering the minimum of random variables. Also the variable $\mathcal{T}_{G_2}^{Q^1}$ changes accordingly, but with an abuse of notation we may write it in the same way. Thus, with high probability as $r \rightarrow \infty$, by Lemma 3.4.2,

$$\begin{aligned} \mathbb{E}_u[\mathcal{T}_{G_1}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_1 = i_1] &= \mathbb{E}_u[(\bar{\tau}_1 + R_{U_2}^1 + \mathcal{T}_{G_2}^{Q^1}) \mid \mathcal{E}(a^*) \cap \{Y_1 = i_1\}] \\ &= \mathbb{E}_u[\bar{\tau}_1 \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_1 = i_1] + o(1) \\ &\quad + \mathbb{E}_u[\mathcal{T}_{G_2}^{Q^1} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_1 = i_1] \\ &= f_1 \mathbb{E}_u[\mathcal{T}_{v_{i_1}}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] + o(1) \\ &\quad + \mathbb{E}_u[\mathcal{T}_{G_2}^{Q^1} \mathbb{1}_{\mathcal{E}(a^*)}], \quad r \rightarrow \infty. \end{aligned} \quad (3.105)$$

We want to analyze the latter in a recursive way. The k -th iteration gives

$$\mathbb{E}_u[\mathcal{T}_{G_k}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)}] = \sum_{i_k: v_{i_k} \in V_k} \mathbb{E}_u[\mathcal{T}_{G_k}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_k = i_k] \mathbb{P}(Y_k = i_k). \quad (3.106)$$

2. We can again write the random variable $\mathcal{T}_{G_k}^{Q^{k-1}}$ as sum of three random variables

$$\mathcal{T}_{G_k}^{Q^{k-1}} = \bar{\tau}_k + R_{U_{k+1}}^k + \mathcal{T}_{G_{k+1}}^k, \quad (3.107)$$

where $G_{k+1} = ((U_{k+1}, V_{k+1}), E_{k+1})$, $U_{k+1} = U_k \setminus g(v_{Y_k})$, $V_{k+1} = V_k \setminus \{v_{Y_k}\}$, $E_{k+1} = E_k \setminus \{(u, v) : u \in g(v_{Y_k})\}$, while $Q^k = Q(\sum_{j=1}^{k-1} \mathcal{T}_{v_{i_j}}^{Q^{j-1}})$. By Corollary 3.4.4, we again have that

$$\mathbb{E}_u[\bar{\tau}_k \mid Y_k = i_k] = f_k \mathbb{E}_u[\mathcal{T}_{v_{i_k}}^{Q^{k-1}}], \quad r \rightarrow \infty, \quad (3.108)$$

and also the variable $\mathcal{T}_{G_{k+1}}^k$ changes accordingly when it is conditioned (again, with an abuse of notation we write it in the same way). With high probability

as $r \rightarrow \infty$, the conditional expectation in (3.106) can be written as

$$\begin{aligned}
 \mathbb{E}_u[\mathcal{T}_{G_k}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_k = i_k] &= \mathbb{E}_u[(\mathcal{T}_{v_{Y_k}}^{Q^{k-1}} + R_{U_{k+1}}^k + \mathcal{T}_{G_{k+1}}^{Q^k}) \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_k = i_k] \\
 &= \mathbb{E}_u[\mathcal{T}_{v_{Y_k}}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_k = i_k] + o(1) \\
 &\quad + \mathbb{E}_u[\mathcal{T}_{G_{k+1}}^{Q^k} \mathbb{1}_{\mathcal{E}(a^*)} \mid Y_k = i_k] \\
 &= f_k \mathbb{E}_u[\mathcal{T}_{v_{i_k}}^{Q^{k-1}} \mathbb{1}_{\mathcal{E}(a^*)}] + o(1) \\
 &\quad + \mathbb{E}_u[\mathcal{T}_{G_{k+1}}^{Q^k} \mathbb{1}_{\mathcal{E}(a^*)}], \quad r \rightarrow \infty.
 \end{aligned} \tag{3.109}$$

At each iteration the conditional expectation reduces to a sum of three terms: the first term represents the expected time it takes to switch the following node on (adjusted by a factor that keeps track of the fact that the node activates before the other nodes), the second term represents the expected time it takes (after the previous node activation) to reach the state with all the nodes remaining in U active, while the third term represents the mean transition time of the remaining network when we take the following activating node out.

3. Note that, for each $k = 1, \dots, N$, the graph G_{k+1} depends on the sequence of indices (Y_1, \dots, Y_k) . Moreover, we know that also the queue lengths Q^k depend on the indices (Y_1, \dots, Y_{k-1}) . Thus, all the conditional expectations depend on the sequence of indices of activated nodes. By Lemma 3.6.3, the first iteration comes with a probability $\frac{1}{n_1}$ of choosing the first node activating, while each iteration with $k > 1$ comes with a probability $\frac{1}{n_k}$, also depending on the sequence (Y_1, \dots, Y_{k-1}) . After $k = 2$ steps, using (3.106) and (3.109), with high probability as $r \rightarrow \infty$,

$$\begin{aligned}
 &\mathbb{E}_u[\mathcal{T}_{G_1}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] \\
 &= \sum_{i_1: v_{i_1} \in V_1} \frac{1}{n_1} \left(f_1 \mathbb{E}_u[\mathcal{T}_{v_{i_1}}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] + o(1) + \mathbb{E}_u[\mathcal{T}_{G_2}^{Q^1} \mathbb{1}_{\mathcal{E}(a^*)}] \right) \\
 &= \sum_{i_1: v_{i_1} \in V_1} \frac{1}{n_1} \left(f_1 \mathbb{E}_u[\mathcal{T}_{v_{i_1}}^{Q^0} \mathbb{1}_{\mathcal{E}(a^*)}] + o(1) \right. \\
 &\quad \left. + \sum_{i_2: v_{i_2} \in V_2} \frac{1}{n_2} \left(f_2 \mathbb{E}_u[\mathcal{T}_{v_{i_2}}^{Q^1} \mathbb{1}_{\mathcal{E}(a^*)}] + o(1) + \mathbb{E}_u[\mathcal{T}_{G_3}^{Q^2} \mathbb{1}_{\mathcal{E}(a^*)}] \right) \right), \quad r \rightarrow \infty.
 \end{aligned} \tag{3.110}$$

After N steps, the last node in V activates and the conditional expectation

becomes

$$\begin{aligned}
 \mathbb{E}_u[\mathcal{T}_{G_N}^{Q^{N-1}} \mathbb{1}_{\mathcal{E}(a^*)}] &= \sum_{i_N: v_{i_N} \in V_N} \frac{1}{n_N} \left(f_N \mathbb{E}_u[\mathcal{T}_{v_{i_N}}^{Q^{N-1}} \mathbb{1}_{\mathcal{E}(a^*)}] + \mathbb{E}_u[R_{U_{N+1}}^N] \right. \\
 &\quad \left. + \mathbb{E}_u[\mathcal{T}_{G_{N+1}}^{Q^N} \mathbb{1}_{\mathcal{E}(a^*)}] \right) \\
 &= \sum_{i_N: v_{i_N} \in V_N} \frac{1}{n_N} f_N \mathbb{E}_u[\mathcal{T}_{v_{i_N}}^{Q^{N-1}} \mathbb{1}_{\mathcal{E}(a^*)}], \quad r \rightarrow \infty.
 \end{aligned}
 \tag{3.111}$$

Indeed, as soon as the last node in V activates, we are actually done and we are not interested in what happens after. We can set $R_{U_{N+1}}^N = 0$ and we have $V_{N+1} = \emptyset$, which implies $\mathbb{E}_u[\mathcal{T}_{G_{N+1}}^{Q^N}] = 0$. Thus, we have arrived at (3.26).

- (iii) The claim follows from steps analogous to the ones in (ii), given any path $a \in \mathcal{A}$ that the algorithm generates.

□

§3.6.4 Proof: mean of the transition time

Proof of Theorem 3.3.3. Recall that, in the subcritical and critical regimes, we are computing the mean transition time conditioned on the event that the nucleation follows a fixed path $a = (v_1, \dots, v_N) \in \mathcal{A}$. We again distinguish between the three regimes.

- (I) $\beta \in (0, \frac{1}{d^*-1})$: subcritical regime. Every term in the sum is $o(r)$, which means that the significant terms are the ones with $\bar{d}_k = d^*$ only. The pre-factors of these terms are given by subcritical forks, and so

$$\begin{aligned}
 \mathbb{E}_u[\mathcal{T}_G^{Q^0}] &= \sum_{k: \bar{d}_k = d^*} f_k \mathbb{E}_u[\mathcal{T}_{v_k}^{Q^{k-1}}] = \sum_{k: \bar{d}_k = d^*} f_k \frac{\mathbb{E}_u[Q_U^{k-1}]^{\beta(d^*-1)}}{d^* B^{-(d^*-1)}} [1 + o(1)] \\
 &= \sum_{k: \bar{d}_k = d^*} f_k \frac{\gamma_U^{\beta(d^*-1)}}{d^* B^{-(d^*-1)}} r^{\beta(d^*-1)} [1 + o(1)], \quad r \rightarrow \infty,
 \end{aligned}
 \tag{3.112}$$

with $f_k = \frac{1}{n_k}$. The last equality comes from (3.72) in Theorem 3.4.8.

- (II) $\beta = \frac{1}{d^*-1}$: critical regime. Every term in the sum is $o(r)$, except the terms with $\bar{d}_k = d^*$, which is of order r . The significant terms are the ones with $\bar{d}_k = d^*$

only. The pre-factors of these terms are given by critical forks, and so

$$\begin{aligned}\mathbb{E}_u[\mathcal{T}_G^{Q^0}] &= \sum_{k: \bar{d}_k = d^*} f_k \mathbb{E}_u[\mathcal{T}_{v_k}^{Q^{k-1}}] = \sum_{k: \bar{d}_k = d^*} f_k \frac{\mathbb{E}_u[Q_U^{k-1}]}{d^* B^{-(d^*-1)} + (c - \rho_U)} [1 + o(1)] \\ &= \sum_{k: \bar{d}_k = d^*} f_k \frac{\gamma_U^{(k-1)}}{d^* B^{-(d^*-1)} + (c - \rho_U)} r [1 + o(1)], \quad r \rightarrow \infty,\end{aligned}\tag{3.113}$$

with $\gamma_U^{(k-1)}$ defined in (3.74) in Theorem 3.4.8 and

$$f_k = \frac{\bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}{n_k \bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}.\tag{3.114}$$

- (III) $\beta \in (\frac{1}{d^*-1}, \infty)$: supercritical regime. Denote by v_{sc} the first supercritical node. We know from (3.76) in Theorem 3.4.8 that, after v_{sc} activates, the queue lengths become negligible (order $o(r)$), and the mean transition time is given by the expected time it takes for them to hit zero, i.e.,

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{\gamma_U}{c - \rho_U} r [1 + o(1)], \quad r \rightarrow \infty.\tag{3.115}$$

□

§3.6.5 Proof: law of the transition time

Proof of Theorem 3.3.5. We again distinguish between the three regimes.

- (I) $\beta \in (0, \frac{1}{d^*-1})$: subcritical regime. Recall that the significant terms in the sum for the mean transition time are those coming from nodes with degree $\bar{d}_k = d^*$ with $d^* < \frac{1}{\beta} + 1$. There are m_{sub}^a such terms, where m_{sub}^a depends on the path $a \in \mathcal{A}$, and each term comes with a multiplicative factor f_k . We can write the transition time along path a divided by its mean as

$$\begin{aligned}\frac{\mathcal{T}_G^{Q^0} | A_a}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} &= \frac{\sum_{k=1}^N \bar{\tau}_k + \sum_{k=2}^N R_{U_k}^{k-1}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} \\ &= \frac{\sum_{k': \bar{d}_{k'} = d^*} \bar{\tau}_{k'} + \sum_{k'': \bar{d}_{k''} < d^*} \bar{\tau}_{k''} + \sum_{k=2}^N R_{U_k}^{k-1}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]}.\end{aligned}\tag{3.116}$$

We know that the law of a sum of independent random variables has a density given by the convolution of their densities. Here the nucleation times and the return times can be considered as independent, since they only depend on the queue lengths, which remain close to the initial value in the subcritical regime.

There are three types of sums in the numerator of the last line of (3.116). The first type of sum is of the form $\bar{\tau}_{k'}/\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]$, with k' such that $\bar{d}_{k'} = d^*$. As

$r \rightarrow \infty$, these are the significant terms in the sum, since they are of the same order as the mean transition time. For each of them, i.e., for each k' , we have

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\bar{\tau}_{k'}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} > x \right) &= \lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\bar{\tau}_{k'}}{\mathbb{E}_u[\bar{\tau}_{k'}]} > \frac{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]}{\mathbb{E}_u[\bar{\tau}_{k'}]} x \right) \\ &= \exp \left(- \frac{\sum_{i: \bar{d}_i = d^*} \mathbb{E}_u[\bar{\tau}_i]}{\mathbb{E}_u[\bar{\tau}_{k'}]} x \right) \\ &= \exp \left(- \frac{\sum_{i: \bar{d}_i = d^*} f_i}{f_{k'}} x \right), \quad x \in [0, \infty), \end{aligned} \quad (3.117)$$

where in the second step we use Proposition 3.4.6. We write the density as

$$\mathcal{P}_{\text{sub}}^{f_{k'}, S_{m_{\text{sub}}}^a}(x) = \frac{S_{m_{\text{sub}}}^a}{f_{k'}} \exp \left(- \frac{S_{m_{\text{sub}}}^a}{f_{k'}} x \right), \quad x \in [0, \infty), \quad (3.118)$$

with

$$S_{m_{\text{sub}}}^a = \sum_{i: \bar{d}_i = d^*} f_i. \quad (3.119)$$

The second type of sum is of the form $\bar{\tau}_{k''}/\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]$, with k'' such that $\bar{d}_{k''} < d^*$. As $r \rightarrow \infty$, these are negligible, since they are of smaller order than the mean transition time. For each of them, i.e., for each k'' , we have

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\bar{\tau}_{k''}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} > x \right) = \lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\bar{\tau}_{k''}}{\mathbb{E}_u[\bar{\tau}_{k''}]} > \frac{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]}{\mathbb{E}_u[\bar{\tau}_{k''}]} x \right), \quad x \in [0, \infty), \quad (3.120)$$

and the density is δ_0 , the Dirac function at 0. The third type of sum is of the form $R_{U_k}^{k-1}/\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]$, with $k = 2, \dots, N$. As $r \rightarrow \infty$, these are also negligible, since they are $o(1)$ by Lemma 3.4.2, and hence their density is also δ_0 .

The density of $\mathcal{T}_G^{Q^0} | A_a / \mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]$ is given by the convolution of the densities of the three types of terms. Since δ_0 gives the identity for the convolution, we can write

$$\lim_{r \rightarrow \infty} \mathbb{P}_u \left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a]} > x \mid A_a \right) = \int_x^\infty \left(\bigotimes_{k'=1}^{m_{\text{sub}}^a} \mathcal{P}_{\text{sub}}^{f_{k'}, S_{m_{\text{sub}}}^a} \right)(y) dy, \quad x \in [0, \infty), \quad (3.121)$$

and we can rename the index k' by k .

- (II) $\beta = \frac{1}{d^*-1}$: critical regime. For two reasons we do not know how to handle this regime: (a) We do not know the law of the next nucleation times because Proposition 3.4.6 only holds in the subcritical regime. (b) The next nucleation times are dependent random variables, and so convolution is no longer relevant.
- (III) $\beta \in (\frac{1}{d^*-1}, \infty)$: supercritical regime. Recall that $T_U = \frac{\gamma_U}{c-\rho_U} r [1 + o(1)]$ as $r \rightarrow \infty$. The law of the transition time is given by $\mathcal{P}_3(x)$ from Theorem 2.1.6.

Indeed, the mean transition time is the expected time it takes for the queue lengths in U to hit zero and. With high probability as $r \rightarrow \infty$, the transition does not occur before or after its mean. Since

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_G^{Q^0} > \mathbb{E}_u[\mathcal{T}_G^{Q^0}]) = \lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_G^{Q^0} > T_U) = 0, \quad (3.122)$$

we can write

$$\lim_{r \rightarrow \infty} \mathbb{P}_u\left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0}]} > x\right) = 0, \quad x \in [1, \infty). \quad (3.123)$$

We also have

$$\lim_{r \rightarrow \infty} \mathbb{P}_u\left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}_u[\mathcal{T}_G^{Q^0}]} > x\right) = 1, \quad x \in [0, 1). \quad (3.124)$$

Hence the density is the Dirac function at 1. □

§C Appendix: minimum of independent forks

In this appendix we compute the mean next nucleation time in the situation where the forks competing for nucleation are disjoint, i.e., they have no nodes in common. Recall that, in the subcritical regime, the nucleation time of a fork is given by an exponential random variable, while in the critical regime it is given by a “polynomial” random variable, in the sense that its law is truncated polynomial.

§C.1 Subcritical regime: exponential random variables

Let X_1, \dots, X_n be i.i.d. exponential random variables with rate λ and let their minimum be $Z = \min\{X_1, \dots, X_n\}$. Then

$$\mathbb{P}(Z > t) = \mathbb{P}(X_1 > t, \dots, X_n > t) = \mathbb{P}(X_1 > t)^n = e^{-n\lambda t}. \quad (3.125)$$

Hence, Z is an exponential random variable with rate $n\lambda$, and we have

$$\mathbb{E}[Z] = \frac{1}{n\lambda} = \frac{1}{n} \mathbb{E}[X_1]. \quad (3.126)$$

If we consider X_1, \dots, X_{n_k} to be the nucleation times of disjoint forks of degree \bar{d}_k , and Z to be the next nucleation time at step k , then we get

$$\mathbb{E}_u[\bar{\tau}_k] = f_k^{\text{iid}} \mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}], \quad r \rightarrow \infty, \quad (3.127)$$

with $f_k^{\text{iid}} = \frac{1}{n_k}$.

§C.2 Critical regime: polynomial random variables

Let X_1, \dots, X_n be i.i.d. polynomial random variables such that

$$\mathbb{P}\left(\frac{X_i}{\mathbb{E}[X_i]} > x\right) = \begin{cases} (1 - Cx)^{\frac{1-C}{C}}, & \text{if } x \in [0, \frac{1}{C}), \\ 0, & \text{if } x \in [\frac{1}{C}, \infty), \end{cases} \quad i = 1, \dots, n, \quad (3.128)$$

with

$$C = \frac{c - \rho_U}{\bar{d}_k B^{-(\bar{d}_k - 1)} + (c - \rho_U)}. \quad (3.129)$$

Let $Z = \min\{X_1, \dots, X_n\}$. Then, for $t = x \mathbb{E}[X_i]$,

$$\mathbb{P}(X_i > t) = \begin{cases} (1 - \frac{C}{\mathbb{E}[X_i]} t)^{\frac{1-C}{C}}, & \text{if } t \in [0, \frac{\mathbb{E}[X_i]}{C}), \\ 0, & \text{if } t \in [\frac{\mathbb{E}[X_i]}{C}, \infty), \end{cases} \quad i = 1, \dots, n. \quad (3.130)$$

Abbreviate $C = \frac{C_1}{C_1 + C_2}$, where $C_1 = c - \rho_U$ and $C_2 = \bar{d}_k B^{-(\bar{d}_k - 1)}$. Then the exponent $\frac{1-C}{C}$ becomes $\frac{C_2}{C_1}$. We have

$$\begin{aligned} \mathbb{P}(Z > t) &= \mathbb{P}(X_1 > t, \dots, X_n > t) \\ &= \mathbb{P}(X_1 > t)^n = \begin{cases} (1 - \frac{C}{\mathbb{E}[X_1]} t)^{n \frac{C_2}{C_1}}, & \text{if } t \in [0, \frac{\mathbb{E}[X_1]}{C}), \\ 0, & \text{if } t \in [\frac{\mathbb{E}[X_1]}{C}, \infty). \end{cases} \end{aligned} \quad (3.131)$$

The density function of Z is

$$f_z(t) = \frac{d}{dt} [1 - \mathbb{P}(Z > t)] = \begin{cases} \frac{C}{\mathbb{E}[X_1]} n \frac{C_2}{C_1} (1 - \frac{C}{\mathbb{E}[X_1]} t)^{n \frac{C_2}{C_1} - 1}, & \text{if } t \in [0, \frac{\mathbb{E}[X_1]}{C}), \\ 0, & \text{if } t \in [\frac{\mathbb{E}[X_1]}{C}, \infty). \end{cases} \quad (3.132)$$

Hence

$$\mathbb{E}[Z] = \int_0^{\frac{\mathbb{E}[X_1]}{C}} f_z(t) t dt = \frac{C}{\mathbb{E}[X_1]} n \frac{C_2}{C_1} \int_0^{\frac{\mathbb{E}[X_1]}{C}} \left(1 - \frac{C}{\mathbb{E}[X_1]} t\right)^{n \frac{C_2}{C_1} - 1} t dt. \quad (3.133)$$

Substituting $u = 1 - \frac{C}{\mathbb{E}[X_1]} t$, we get

$$\begin{aligned} \mathbb{E}[Z] &= \frac{\mathbb{E}[X_1]}{C} n \frac{C_2}{C_1} \int_0^1 u^{n \frac{C_2}{C_1} - 1} (1 - u) du \\ &= \frac{\mathbb{E}[X_1]}{C} n \frac{C_2}{C_1} \left[\int_0^1 u^{n \frac{C_2}{C_1} - 1} du - \int_0^1 u^{n \frac{C_2}{C_1}} du \right] = \frac{\mathbb{E}[X_1]}{C} n \frac{C_2}{C_1} \left[\frac{1}{n \frac{C_2}{C_1}} - \frac{1}{n \frac{C_2}{C_1} + 1} \right] \\ &= \frac{\mathbb{E}[X_1]}{C} n \frac{C_2}{C_1} \left[\frac{1}{n \frac{C_2}{C_1} (n \frac{C_2}{C_1} + 1)} \right] = \frac{\mathbb{E}[X_1]}{C} \left[\frac{1}{n \frac{C_2}{C_1} + 1} \right] = \mathbb{E}[X_1] \frac{C_1 + C_2}{n C_2 + C_1} \\ &= \frac{\bar{d}_k B^{-(\bar{d}_k - 1)} + (c - \rho_U)}{n \bar{d}_k B^{-(\bar{d}_k - 1)} + (c - \rho_U)} \mathbb{E}[X_1]. \end{aligned} \quad (3.134)$$

If we consider X_1, \dots, X_{n_k} to be the nucleation times of disjoint forks of degree \bar{d}_k , and Z to be the next nucleation time at step k , then we get

$$\mathbb{E}_u[\bar{\tau}_k] = f_k^{\text{iid}} \mathbb{E}_u[\mathcal{T}_{v_k^*}^{Q^{k-1}}], \quad r \rightarrow \infty, \quad (3.135)$$

with

$$f_k^{\text{iid}} = \frac{\bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}{n_k \bar{d}_k B^{-(\bar{d}_k-1)} + (c - \rho_U)}. \quad (3.136)$$