



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



Dynamic bipartite interference graphs

This chapter is based on:

M. Sfragara. *Adding edge dynamics to wireless random-access networks*. Preprint, 2020.

Abstract

We consider random-access networks with nodes representing servers with queues. The nodes can be either active or inactive: a node deactivates at unit rate, while it activates at a rate that depends on its queue length, provided none of its neighbors is active. In order to model the effects of user mobility in wireless networks, we analyze dynamic interference graphs where the edges are allowed to appear and disappear over time. We focus on bipartite graphs, and study the transition time between the two states where one half of the network is active and the other half is inactive, in the limit as the queues become large. Depending on the speed of the dynamics, we are able to obtain a rough classification of the effects of the dynamics on the transition time.

§4.1 Introduction and main results

This chapter is a continuation of Chapters 2–3. We introduce an edge dynamics on the bipartite interference graph, by allowing edges to appear and disappear over time. This represents a natural basic model to capture the effects of user mobility in wireless networks.

In Section 4.1.1 we motivate our interest in adding edge dynamics to random-access network models. In Section 4.1.2 we describe the setting and the mathematical model of interest in this chapter by specifying edge dynamics. In Section 4.1.3 we state out main results for the mean transition time with dynamics. In Section 4.1.4 we explain the main idea behind our analysis and give an outline of the remainder of the chapter.

§4.1.1 Motivation and background

User mobility is one of the major features in wireless networks. Different mobility patterns can be distinguished (pedestrians, vehicles, aerial, dynamic medium, robot, and outer space motion) and mathematical models can be developed in order to generalize such patterns and analyze their characteristics. Understanding the effects of user mobility in wireless networks is crucial in order to design efficient protocols and improve the performance of the network.

For example, consider radio communication protocols, for which central radio stations are used as base-stations for transmitting radio signals. The radio landscape is partitioned into cells and in each cell a station serves the users in its vicinity. In such cellular networks the users may be either stationary or mobile. User mobility leads to problems of handover: when a user moves from one cell to another, the transmitting signal has to be handed over from one station to another in order to ensure continuity of service and seamless mobility. If not enough capacity is available in the adjacent cell, then the transmission might be interrupted. Imagine that a node transmitting to particular station moves away from its cell and reaches a cell where another station serves for transmissions. Although initially the node interferes with a specific group of nodes sharing the same initial station, after the node has moved it interferes with the nodes in the new cell sharing the new station. In a similar fashion, imagine a network where nodes represent transmitter-receiver pairs. The signal of a node interferes with the signals of the nodes in its vicinity. Hence the protocol allows only one of the interfering nodes to transmit alternately. When allowing node mobility, we get new groups of nodes interfering with each other depending on their vicinity. We are therefore dealing with a network whose interference graph changes over time.

To the best of our knowledge, random-access models with user mobility in the context of interference graphs have so far not been considered in the literature. All the studies we are aware of that have examined the impact of user mobility in wireless networks are concerned with handover mechanisms (see [89], [98]), so-called opportunistic scheduling algorithms (see [11], [101]), capacity issues in ad hoc and cellular networks (see [15], [56]), and flow-level performance (see [7], [8], [16], [96]).

In this chapter we investigate a dynamic version of the random-access protocols in order to try to capture some features of user mobility in wireless networks. A natural paradigm for constructing *dynamic interference graphs* would be to use geometric graphs, such as unit-disk graphs, with node mobility, where each node follows a random trajectory and experiences interference from all nodes within a certain distance. A feasible state of the interference graph would then be generated by a specific instance of the geometric graph. We follow a different approach and, with an explorative intention, we consider a model where edges are allowed to appear and disappear from the graph according to i.i.d. Poisson clocks placed on each edge. We find different results for the mean transition time depending on the speed of the dynamics. The evolution of the network is captured by a continuous-time Markov process that keeps track of how the state, the queue length and the number of active neighbors change for each node.

We focus on queue-based activation rates, in line with the models and the results in Chapters 2–3. This leads to two level of complexity, driven by the queue dependencies of the activation rates and by the edge dynamics. In the appendix we briefly consider a simplified version of the model where the activation rates are fixed.

§4.1.2 Setting

We refer to Section 1.1.5 for a general introduction to the mathematical model. In this chapter we add an extra dynamics to the model.

We are interested in analyzing the behavior of the network when we allow the interference graph to change over time. We mainly focus on the model with queue-based activation rates and assume these rates to satisfy Definitions 1.1.4 and 3.1.1.

Definition 4.1.1 (Dynamic interference graphs).

We say that the interference graph is *dynamic* when the edges appear and disappear according to a continuous-time flip process. Consider the dynamic bipartite interference graph $G(\cdot) = (U \sqcup V, E(\cdot))$, where $U \sqcup V$ is the set of nodes, with $|U| = M$ and $|V| = N$, and $E(t)$ is the set of edges that are present between nodes in U and nodes in V at time t . The number of edges $|E(\cdot)|$ changes over time and can vary from a minimum of 0 to a maximum of MN . We set $G(0) = G$, where G is the initial bipartite graph. We denote by $G_{MN} = (U \sqcup V, E_{MN})$ the complete bipartite graph associated to (U, V) and, for every edge $e \in E_{MN}$, at time t we define the Bernoulli random variable $Y_e(t)$ as

$$Y_e(t) = \begin{cases} 0, & \text{if } e \notin E(t), \\ 1, & \text{if } e \in E(t). \end{cases} \quad (4.1)$$

In other words, $Y_e(t) = 0$ if edge e is not present in the graph at time t , while $Y_e(t) = 1$ if it is present. The *joint edge activity state* at time t is denoted by

$$Y(t) = \{Y_e(t)\}_{e \in E_{MN}} \quad (4.2)$$

and is an element of the state space

$$\mathcal{Y} = \{Y \in \{0, 1\}^{U \times V}\}. \quad (4.3)$$

The degree of node v at time t is denoted by $d_v(t)$.

We model the dynamics of the graph in the following way. If an edge is not present, then it *appears* according to a Poisson clock with rate λ , independently of the other edges. If an edge is present, then it *disappears* according to a Poisson clock with rate λ , independently of the other edges. This is equivalent to having a system of i.i.d. Poisson clocks with rate λ on the edges and letting an edge change its state every time its clock ticks. In order to study how the edge dynamics affects the transition time, we consider Poisson clocks with rates $\lambda = \lambda(r)$ depending on the parameter r .

Throughout the chapter we use the notation \prec, \succ to describe the asymptotic behavior in the limit $r \rightarrow \infty$. More precisely, $f(r) \prec g(r)$ means that $f(r) = o(g(r))$ as $r \rightarrow \infty$, and $f(r) \succ g(r)$ means that $g(r) = o(f(r))$ as $r \rightarrow \infty$.

Remark 4.1.2 (Rates on the edges).

We may allow different rates for the edges to change their state. Denote by $\lambda^+(r)$ and $\lambda^-(r)$ the rates at which edges appear and disappear, respectively. If these are of the same order, then we are in a situation similar to them being equal to $\lambda(r)$. If $\lambda^+(r) \rightarrow \infty$ and $\lambda^-(r) \prec \lambda^+(r)$, then, with high probability as $r \rightarrow \infty$, in time $o(1)$ the dynamics turns the initial graph into the complete bipartite graph with all the edges present. Analogously, if $\lambda^-(r) \rightarrow \infty$ and $\lambda^-(r) \succ \lambda^+(r)$, then, with high probability as $r \rightarrow \infty$, in time $o(1)$ the dynamics turns the initial graph into the empty graph with all the edges absent. Both these assumptions do not lead to interesting models. When $\lambda^+(r)$ and $\lambda^-(r)$ are of different order and do not tend to infinity, we have an intermediate situation where at any time t an edge is either present with high probability as $r \rightarrow \infty$ or absent with high probability as $r \rightarrow \infty$, but the total amount of time the edge has been absent or present, respectively, up to time t is not always negligible.

Remark 4.1.3 (Appearing edge).

When an edge disappears from the graph, the states of the nodes do not change. On the other hand, when an edge appears in the graph, it might appear between two active nodes. In this case, we assume that the active node in U deactivates, since the model does not allow two connected nodes to be simultaneously active. We could study alternative models, where the active node in V deactivates or where the deactivating node is chosen uniformly at random (or with certain probabilities). It is obvious that these alternative models slow down the transition and lead to a possible multiple counting of the time it takes for some nodes in V to activate.

§4.1.3 Main theorem

In Chapter 3 we analyzed the mean transition time for arbitrary bipartite graphs. We introduced a randomized algorithm that takes as input the graph and gives as output all the possible activation paths for nodes in V . We showed that, depending on the value of β , the transition exhibits a *subcritical regime*, a *critical regime* and a *supercritical regime*. Given a graph, the algorithm uniquely identifies the value

$$d^* = \max_{1 \leq k \leq N} \bar{d}_k, \tag{4.4}$$

which determines the leading order of the mean transition time and together with β determines the regime we are in. We were also able to identify the law of the transition time divided by its mean (in the subcritical and supercritical regime). The latter is beyond the scope of the present chapter, since understanding the effect of the edge dynamics is rather challenging. Our goal is to extend the results of Theorem 3.3.3 to dynamic bipartite graphs. We will distinguish between different types of dynamics and we will see how they affect the mean transition time. We denote by $\mathcal{T}_{G(\cdot)}^{Q^0}$ the transition time of the dynamic graph $G(\cdot)$ with initial queue lengths Q^0 .

Theorem 4.1.4 (Mean transition time for dynamic bipartite graphs).

Consider the dynamic bipartite graph $G(\cdot) = ((U, V), E(\cdot))$ with the edge dynamics governed by $\lambda(r)$ and initial queue lengths Q^0 .

(FD) If $\lambda(r) \rightarrow \infty$, then the dynamics is fast and, with high probability as $r \rightarrow \infty$, the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp \lambda(r)^{-1} = o(1), \quad r \rightarrow \infty. \quad (4.5)$$

(RD) If $\lambda(r) = C \in (0, \infty)$, then the dynamics is regular and, with high probability as $r \rightarrow \infty$, the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp \lambda(r)^{-1} = \mathcal{O}(1), \quad r \rightarrow \infty. \quad (4.6)$$

(SD) If $\lambda(r) \rightarrow 0$, then the dynamics is slow and the following cases occur.

(SDc) If $\lambda(r) \succeq r^{-(1 \wedge \beta(d^* - 1))}$, then the dynamics is competitive and, with high probability as $r \rightarrow \infty$, the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp \lambda(r)^{-1}, \quad r \rightarrow \infty. \quad (4.7)$$

More precisely, let $\lambda(r) = r^{-\alpha}$ with $0 < \alpha \leq 1 \wedge \beta(d^* - 1)$, and let $T_U(r)$ be the average time it takes for the queue lengths at nodes in U to hit zero.

(I) $\beta \in (0, \frac{1}{d^* - 1})$: subcritical regime. With high probability as $r \rightarrow \infty$,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp r^\alpha [1 + o(1)], \quad r \rightarrow \infty. \quad (4.8)$$

(II) $\beta = \frac{1}{d^* - 1}$: critical regime. With high probability as $r \rightarrow \infty$,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp r^\alpha [1 + o(1)], \quad r \rightarrow \infty. \quad (4.9)$$

In particular, when $\alpha = 1$, with positive probability,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.10)$$

(III) $\beta \in (\frac{1}{d^* - 1}, \infty)$: supercritical regime. When $0 < \alpha \leq 1$, with high probability as $r \rightarrow \infty$,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp r^\alpha [1 + o(1)], \quad r \rightarrow \infty. \quad (4.11)$$

In particular, when $\alpha = 1$, with positive probability,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.12)$$

When $\alpha > 1$, with high probability as $r \rightarrow \infty$,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.13)$$

(SDnc) If $\lambda(r) \prec r^{-(1 \wedge \beta(d^* - 1))}$, then the dynamics is non-competitive and, with high probability as $r \rightarrow \infty$, the transition time satisfies Theorem 3.3.3.

Note that the order of the mean transition time depends on the speed of the dynamics. When the dynamics is fast (FD), the edges quickly appear and disappear, reaching in time $o(1)$ the state where nodes in V have no edges connecting them to U . Since nodes in V are aggressive, they eventually activate in time $o(1)$. When the dynamics is regular (RD), the situation is similar, but it takes time $\mathcal{O}(1)$ to reach the state where all the edges are simultaneously absent. When the dynamics is slow (SD), a node in V can also activate through the nucleation of its fork (recall Definition 3.1.2). In the case of competitive dynamics (SDc), the relation between the speed of the dynamics and the aggressiveness of the nodes in U plays a key role, while in the case of non-competitive dynamics (SDnc), the network behaves as if the edges were fixed at the initial configuration and there were no dynamics. Note that, in the cases of fast, regular and competitive dynamics, the order of the mean transition time is given by the reciprocal of the rate $\lambda(r)$.

§4.1.4 Discussion and outline

Intuition. A node in V can activate for two reasons. It can activate when its neighbors are simultaneously inactive or when there are no edges connecting it to nodes in U . Interpolation between these two situations gives rise to different cases, which mainly depend on the speed of the dynamics. In the case of competitive dynamics, we are able to distinguish between different behaviors for the mean transition time by analyzing the subcritical, critical and supercritical regimes separately. To summarize, with high probability as $r \rightarrow \infty$, the order of activation of nodes in V follows one of the paths generated by the algorithm until the edge dynamics of rate $\lambda(r)$ becomes competitive. The competition begins on time scale $\lambda(r)^{-1}$, the time scale on which all the remaining nodes in V activate, if there are any, and the transition occurs.

Pre-factor. In order to give precise asymptotics, including the pre-factor for the mean transition time, we must analyze a more complicated Markov process describing how the states of the nodes, the queue lengths and the states of the edges change over time. This is beyond the scope of the present chapter, but in Section 4.4 we give an overview of the main challenges.

Outline of the chapter. The remainder of this chapter is organized as follows. In Section 4.2 we discuss the main effects of the dynamics on the mean transition time and we explain how it can slow down or speed up the activation of each node in V .

In Section 4.3 we prove Theorem 4.1.4 by discussing the different types of dynamics separately. In Section 4.4 we describe the graph evolution and discuss what needs to be considered in order to compute the pre-factor of the mean transition time. In Appendix D we consider a model where the activation rates are fixed and not queue-dependent. We adapt results from this chapter and the previous chapters in order to study how the dynamics affects the transition time.

§4.2 The edge dynamics

In this section we analyze the effects that different types of dynamics have on the mean transition time of the network.

§4.2.1 Disconnection time

Recall from Section 3.1.1 that the nucleation time \mathcal{T}_v^Q of the fork of a node $v \in V$ given the initial queue lengths Q is the time it takes for its neighbors to become simultaneously inactive, so that v can activate as soon as its clock ticks. Due to the dynamics, a node $v \in V$ does not necessarily activate through the nucleation of its fork, but it can also activate if at some point there are no edges connecting it to nodes in U . The dynamics, indeed, might sometimes bring the graph to a configuration where the degree of v is temporarily 0, so that v can activate as soon as its clock ticks in $o(1)$.

Definition 4.2.1 (Disconnection time).

Given $v \in V$, we call *disconnection time* of v the time it takes for v to be disconnected from U , i.e., to have all possible edges connecting it to U simultaneously absent. We denote the disconnection time of v by $D_v^{Q^0}$, where Q^0 indicates the initial queue lengths.

As introduced in Section 4.1.2, the dynamics affects the network by allowing the edges to appear and disappear according to a Poisson clock with rate $\lambda(r)$. The alternation between the states of each edge $e \in E_{MN}$ is described by an exponential random variable $S_e \simeq \text{Exp}(\lambda(r))$ with mean $\mu(r) = \lambda(r)^{-1}$. Note that, with high probability as $r \rightarrow \infty$, S_e takes values of the order of its mean, i.e., $S_e \asymp \mu(r)$. Indeed, if we pick $x \prec \mu(r)$, then

$$\lim_{r \rightarrow \infty} \mathbb{P}(S_e \leq x) = \lim_{r \rightarrow \infty} 1 - e^{-\lambda(r)x} = 0, \tag{4.14}$$

and the same holds for $x \succ \mu(r)$. In other words, if an edge is absent at time t , then, with high probability as $r \rightarrow \infty$, it will take an amount of time of order $\mu(r)$ for the Poisson clock to tick and for the edge to become present. Vice versa, if an edge is present at time t , then it will take an amount of time of order $\mu(r)$ for the edge to become absent.

The arbitrary bipartite initial configuration of the graph plays an important role in understanding the transition time. Consider a node in $v \in V$ of initial degree $d_v(0) = d > 0$. Since $|U| = M$, there are M possible total edges connecting v to U .

We construct a continuous-time Markov chain \mathcal{M} where each state k represents the set of configurations of the M edges in which k edges are present and $M - k$ edges are absent. State 0 corresponds to all edges being absent, state 1 corresponds to the M possible configurations with exactly one edge present, and so on (see Figure 4.1 below).

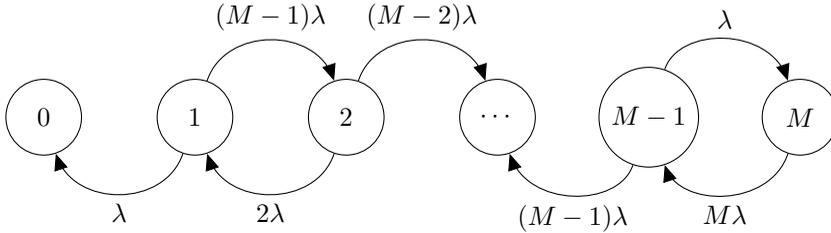


Figure 4.1: The Markov chain \mathcal{M} describing how the edge dynamics changes the degree of a node in V . It is a birth-death process with M transient states and one absorbing state.

We consider state 0 as an absorbing state, since we are interested in computing the hitting times to state 0 starting from any other state. From state M we can only jump to state $M - 1$, when one of the M present edges disappears, which happens with rate $M\lambda$. From each state $0 < k < M$ we jump to the neighboring states also with rate $M\lambda$. Indeed, as soon as the clock of one of the M possible edges ticks, we jump to the state $k + 1$ if the edge was absent and becomes present, while we jump to the state $k - 1$ if the edge was present and becomes absent. Hence, we jump from state k to state $k + 1$ with probability $\frac{M-k}{M}$, while we jump from state k to state $k - 1$ with probability $\frac{k}{M}$.

The transition rate matrix H of the Markov chain \mathcal{M} is given by

$$H = \begin{matrix} & \begin{matrix} 0 & 1 & 2 & \dots & M-1 & M \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \\ \vdots \\ M-1 \\ M \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda & -M\lambda & (M-1)\lambda & 0 & 0 & 0 \\ 0 & 2\lambda & -M\lambda & \dots & 0 & 0 \\ \vdots & 0 & \dots & \dots & \dots & 0 \\ 0 & 0 & 0 & \dots & -M\lambda & \lambda \\ 0 & 0 & 0 & 0 & M\lambda & -M\lambda \end{pmatrix} \end{matrix} \quad (4.15)$$

and can be written as

$$H = \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{S}^0 & S \end{pmatrix}, \quad (4.16)$$

where S is an $M \times M$ matrix and $\mathbf{S}^0 = -S\mathbf{1}_M$, where $\mathbf{1}_M$ represents the M -dimensional column vector with every element being 1. Let

$$(a_0, \mathbf{a}) = (a_0, a_1, \dots, a_M) \quad (4.17)$$

be the $M + 1$ dimensional row vector describing the probability of starting in one of the $M + 1$ states. Since $d_v(0) = d$, we have that the d -th entry of \mathbf{a} equals 1 and all

the other entries equal 0. Computing the disconnection time of a node with initial degree d is equivalent to computing the hitting time of the Markov chain \mathcal{M} to state 0 starting from state d .

Lemma 4.2.2 (Mean and law of the disconnection time).

Consider a node $v \in V$ of initial degree $d_v(0) = d > 0$, and let the edge dynamics be such that $\mathbb{E}_u[S_e] = \mu(r)$ for each $e \in E_{MN}$.

(i) The disconnection time $D_v^{Q^0}$ satisfies

$$\mathbb{E}_u[D_v^{Q^0}] = C_d \mu(r) [1 + o(1)], \quad r \rightarrow \infty, \quad (4.18)$$

where $(C_1 \mu(r), \dots, C_M \mu(r))$ is the solution of the linear system of equations

$$\begin{cases} x_1 &= \frac{1}{M} \frac{\mu(r)}{M} + \frac{M-1}{M} \left(\frac{\mu(r)}{M} + x_2 \right) \\ x_2 &= \frac{2}{M} \left(\frac{\mu(r)}{M} + x_1 \right) + \frac{M-2}{M} \left(\frac{\mu(r)}{M} + x_3 \right) \\ \dots &= \dots \\ x_{M-1} &= \frac{M-1}{M} \left(\frac{\mu(r)}{M} + x_{M-2} \right) + \frac{1}{M} \left(\frac{\mu(r)}{M} + x_M \right) \\ x_M &= \frac{\mu(r)}{M} + x_{M-1}. \end{cases} \quad (4.19)$$

(ii) The law of the disconnection time $D_v^{Q^0}$ follows a phase-type distribution $\text{PH}(\mathbf{a}, S)$ and is given by

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(D_v^{Q^0} > x) = \mathbf{a} \exp(Sx) \mathbf{1}, \quad x \in (0, \infty), \quad (4.20)$$

where \mathbf{a} and S are as in (4.17) and (4.16), respectively. In particular, the above probability equals the sum of the entries in d -th row of the matrix $\exp(Sx)$.

Proof. We prove the two statements separately.

(i) Consider the Markov chain \mathcal{M} described above. We know that from each state $k > 0$, we jump to a neighboring state with rate $M\lambda$. The jump occurs exactly when the first of the M possible edges changes its state. This corresponds to the minimum of M i.i.d. exponential random variables, which is known to follow an exponential distribution with mean $\frac{\mu(r)}{M}$. If v has initial degree d , then we start from state d . We denote by x_k the mean hitting times of state 0 starting from state k . The above system of equations allows us to compute the mean disconnection time of v .

Since, the system of equations is linear in $\mu(r)$ and in the variables x_i 's, its solution is linear in $\mu(r)$. Hence the mean disconnection time is of order $\mu(r)$.

(ii) The disconnection time of a node $v \in V$ of initial degree $d > 0$ is the hitting time of state 0 of the Markov chain \mathcal{M} starting from state d . The distribution of the hitting time to the unique absorbing state, starting from any of the other

finite transient states, is said to be phase-type and is denoted by $\text{PH}(\mathbf{a}, S)$, with \mathbf{a} and S as in (4.17) and (4.16), respectively.

The distribution function of $D_v^{Q^0}$ is given by

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(D_v^{Q^0} \leq x) = \int_0^x \mathcal{P}(y) dy = 1 - \mathbf{a} \exp(Sx) \mathbf{1}, \quad x \in (0, \infty), \quad (4.21)$$

where $\exp(\cdot)$ indicates the matrix exponential, and

$$\mathcal{P}(z) = \mathbf{a} \exp(Sz) \mathbf{S}^0, \quad z \in (0, \infty), \quad (4.22)$$

with \mathbf{S}^0 as in (4.16). Since the vector \mathbf{a} has its d -th entry equal to 1 and all the other entries equal to 0, we have that the product $\mathbf{a} \exp(Sx) \mathbf{1}$ equals the sum of the entries in the d -th row of the matrix $\exp(Sx)$. □

Note that the results of Lemma 4.2.2 hold even without letting $r \rightarrow \infty$.

§4.2.2 Nucleation vs. dynamics

Without loss of generality, we may consider interference graphs with no isolated nodes in V , since after time $o(1)$ we would be in such a scenario anyway.

Lemma 4.2.3 (Isolated nodes).

Nodes in V with initial degree 0 activate in time $o(1)$ as $r \rightarrow \infty$.

Proof. Consider the situation where $\lambda(r) < g_V(0)$, i.e., the dynamics is slower than the average time it takes for the activation clock of nodes in V to tick. Then a node $v \in V$ with initial degree 0 activate as soon as its clock ticks, hence in time $o(1)$. Next, consider the situation where the dynamics is very fast, $\lambda(r) \succ g_V(0)$. Then a node $v \in V$ with initial degree 0 might be blocked by some active neighbors in U by the time its activation clock ticks for the first time. Recall that $|U| = M$ and note that there are 2^M possible configurations of edges connecting v to U . Each time the activation clock of v ticks, the probability of being in each of the possible configurations tends to the uniform probability $1/2^M$ as $r \rightarrow \infty$. Therefore, after a finite number of attempts, v eventually activates. Since each tick of the activation clock of v takes time $o(1)$, v activates in time $o(1)$. Lastly, consider the situation where $\lambda(r) \asymp g_V(0)$. If the activation clock of a node $v \in V$ with initial degree 0 ticks before any of its potential edges appear, then v activates in time $o(1)$. Otherwise, each subsequent activation attempt will not be successful unless the edge configuration is such that v has no neighbors. In other words, v can activate only when the Markov chain describing how its degree changes over time is in state 0. In this case, v activates with a probability that at time t is given by $\frac{g_V(t)}{g_V(t) + M\lambda(r)} > 0$ as $r \rightarrow \infty$. Since $\lambda(r)^{-1} = o(1)$, by using similar arguments as in the proof of Lemma 4.2.2, the time it takes for the Markov chain to return to state 0 when starting from state 0 is $o(1)$. Hence, v has the chance to activate with positive probability every period of time $o(1)$. Therefore, after a finite number of attempts, v eventually activates in time $o(1)$. □

We call *activation time* of $v \in V$ the time it takes for v to activate. Depending on the dynamics, this can be given either by its nucleation time $\mathcal{T}_v^{Q^0}$ or by its disconnection time $D_v^{Q^0}$. When the dynamics is fast enough, nodes in V eventually activate because their clocks tick and no edges connect them to nodes in U . On the other hand, when the dynamics is particularly slow, it is more likely for nodes in V to activate through the nucleation of its fork, and the network tends to behave as if the edges were frozen at the initial configuration. In between these two scenarios the dynamics is more interesting and, depending on its speed, we distinguish between different behaviors. Proposition 4.2.4 below describes the competition between the nucleation and the dynamics.

Proposition 4.2.4 (Nucleation vs. dynamics).

Let $v \in V$ be the node of minimum degree at time $t = 0$, with $d_v(0) = d > 0$.

- (i) If $\lambda(r) \succ r^{-(1 \wedge \beta(d-1))}$, then, with high probability as $r \rightarrow \infty$, the activation time of v is given by its disconnection time, i.e.,

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(D_v^{Q^0} < \mathcal{T}_v^{Q^0}) = 1. \tag{4.23}$$

- (ii) If $\lambda(r) \asymp r^{-(1 \wedge \beta(d-1))}$, then the activation time of v is given either by its nucleation time with positive probability or by its disconnection time with positive probability.

- (iii) If $\lambda(r) \prec r^{-(1 \wedge \beta(d-1))}$, then, with high probability as $r \rightarrow \infty$, the activation time of v is given by its nucleation time, i.e.,

$$\lim_{r \rightarrow \infty} \mathbb{P}_u(\mathcal{T}_v^{Q^0} < D_v^{Q^0}) = 1. \tag{4.24}$$

Proof. Recall that $\mu(r) = \lambda(r)^{-1}$ and that the disconnection time $D_v^{Q^0}$ is given by a phase-type random variable with mean of order $\mu(r)$. Since phase-type random variables are constructed by convolutions of exponential random variables, we have that, with high probability as $r \rightarrow \infty$, $D_v^{Q^0}$ takes values of order $\mu(r)$. Recall also that, depending on the relation between β and d , the nucleation time $\mathcal{T}_v^{Q^0}$ is given by an exponential random variable with mean of order $r^{\beta(d-1)}$, by a polynomial random variable with mean of order r , or by $T_U(r)$, which is the average time it takes for the queue lengths at nodes in U to hit zero. Hence, with high probability as $r \rightarrow \infty$, $\mathcal{T}_v^{Q^0}$ takes values of order $r^{1 \wedge \beta(d-1)}$. It is therefore immediate to distinguish between the three cases.

- (i) Since $\mu(r) \prec r^{1 \wedge \beta(d-1)}$, with high probability as $r \rightarrow \infty$, v activates due to absence of edges.
- (ii) Since $\mu(r) \asymp r^{1 \wedge \beta(d-1)}$, there is a competition between the nucleation time $\mathcal{T}_v^{Q^0}$ and the phase-type random variable $D_v^{Q^0}$. Depending on their parameters, each of them can occur before the other with positive probability.
- (iii) Since $\mu(r) \succ r^{1 \wedge \beta(d-1)}$, with high probability as $r \rightarrow \infty$, v activates through the nucleation of its fork.

□

§4.3 Proofs of the main results

In this section we prove Theorem 4.1.4 by analyzing the different types of dynamics separately.

§4.3.1 Proof: fast dynamics

Consider the fast dynamics (FD) where $\lambda(r) \rightarrow \infty$ as $r \rightarrow \infty$.

Proof of Theorem 4.1.4 (FD). With high probability as $r \rightarrow \infty$, for each edge the random intervals between clock ticks are of order $\lambda(r)^{-1} = o(1)$. By Lemma 4.2.2, the mean disconnection time of a node in V is of order $\lambda(r)^{-1}$. Moreover, by Proposition 4.2.4, with high probability as $r \rightarrow \infty$, each node activates due to absence of edges and not through the nucleation of its fork, and hence it activates in a time of order $\lambda(r)^{-1}$. In conclusion, with high probability as $r \rightarrow \infty$, the transition time of $G(\cdot)$ with initial queue lengths Q^0 satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp \lambda(r)^{-1} = o(1), \quad r \rightarrow \infty, \quad (4.25)$$

hence the claim is settled. \square

§4.3.2 Proof: regular dynamics

Consider the regular dynamics (RD) where $\lambda(r) = C \in (0, \infty)$.

Proof of Theorem 4.1.4 (RD). With high probability as $r \rightarrow \infty$, for each edge the random intervals between clock ticks are of order $\lambda(r)^{-1} = \mathcal{O}(1)$. By Lemma 4.2.2, the mean disconnection time of a node in V is of order $\lambda(r)^{-1}$. Note that nodes in V of initial degree 1 can activate either because their only neighbor deactivates in $\mathcal{O}(1)$ or due to absence of edges with a mean disconnection time of order $\lambda(r)^{-1}$. Moreover, by Proposition 4.2.4, with high probability as $r \rightarrow \infty$, nodes in V of initial degree greater than 1 activate due to absence of edges in a time of order $\lambda(r)^{-1}$. In conclusion, with high probability as $r \rightarrow \infty$, the transition time of $G(\cdot)$ with initial queue lengths Q^0 satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp \lambda(r)^{-1} = \mathcal{O}(1), \quad r \rightarrow \infty, \quad (4.26)$$

hence the claim is settled. \square

§4.3.3 Proof: non-competitive dynamics

Consider the slow dynamics where $\lambda(r) \rightarrow 0$ as $r \rightarrow \infty$ with $\lambda(r) \prec r^{-(1 \wedge \beta(d^* - 1))}$, called the non-competitive dynamics (SDnc). In this case the dynamics is so slow that it has no effect on the transition.

Proof of Theorem 4.1.4 (SDnc). The mean disconnection time of any node in V is of order larger than $r^{1 \wedge \beta(d^* - 1)}$. Hence each node in V activates through the nucleation of its fork, which is at most of order $r^{1 \wedge \beta(d^* - 1)}$. The dynamics is very slow, almost frozen, and so it does not affect the nucleation of the forks. Hence, with high probability as $r \rightarrow \infty$, the transition time of $G(\cdot)$ with initial queue lengths Q^0 satisfies Theorem 3.3.3 and the network behaves as if there were no dynamics. Hence the claim is settled. \square

§4.3.4 Proof: competitive dynamics

Consider the slow dynamics (SD) where $\lambda(r) \rightarrow 0$ as $r \rightarrow \infty$ with $\lambda(r) = r^{-\alpha}$, with $0 < \alpha \leq 1 \wedge \beta(d^* - 1)$, called the competitive dynamics (SDc). This is the most interesting type of dynamics, since it competes with the fork nucleations. The activation of the nodes in V can occur both because of the absence of their edges and because of the nucleation of their forks. Recall the algorithm defined in Section 3.2 in Chapter 3.

Proof of Theorem 4.1.4 (SDc). Denote by \hat{d} the largest integer such that $\beta(\hat{d} - 1) < \alpha$. Let the algorithm generate all possible activation paths for nodes in V and denote this set by \mathcal{A} . Fix a path $a \in \mathcal{A}$. Consider the sequence of activating nodes along the path a up to the step in which the degree is larger than \hat{d} . Say that at step k we have $\bar{d}_k > \hat{d}$. Consider only the first $k - 1$ steps. We indicate by $A_a(\alpha)$ the event that the network follows the path $a \in \mathcal{A}$ until time scale r^α . On time scale r^α the dynamics starts competing with the nucleation, and the order of activation of the remaining nodes described by the algorithm is not preserved anymore. In other words, the order of activation of nodes in V follows the order of activation of the path a only for the first $k - 1$ nodes. With each of these $k - 1$ nodes is associated a nucleation time of order less than or equal to $r^{1 \wedge \beta(\hat{d} - 1)}$. Hence, by Proposition 4.2.4, with high probability as $r \rightarrow \infty$, the activation time of these nodes is given by their nucleation time. We apply Proposition 4.2.4 to each iteration of the graph, each time by considering a node with minimum degree \bar{d}_j for $j = 1, \dots, k - 1$. Indeed, we know from Lemma 4.2.2, that the mean disconnection time of a node is of order r^α . We treat the subcritical, critical and supercritical regimes separately.

- (I) $\beta \in (0, \frac{1}{d^* - 1})$: subcritical regime. We have $0 < \alpha \leq \beta(d^* - 1) < 1$. The activation time of the next activating node is of order r^α . It cannot be of smaller order since at step k we have $\bar{d}_k > \hat{d}$ by construction. It cannot be of higher order either since the disconnection time of any of the remaining nodes is of order r^α . After this activation, there might be nodes whose degree has decreased and whose nucleation time is of smaller order. When we sum the mean activation times of the nodes in V to compute the mean transition time, we see that these nodes will not contribute significantly as $r \rightarrow \infty$. All the remaining nodes are likely to activate in any possible order, but none of them will have an activation time of order larger than r^α . To know how many nodes contribute to the transition time with an activation time of order r^α , we need to have more control on how the degrees of the nodes evolve over time. To

conclude, the order of activation of nodes in V follows the path a as long as the nucleation times associated to the nodes are of order smaller than r^α . After that, the remaining nodes can activate with positive probability in any order with an activation time of order at most r^α . Hence, the transition time conditional on the event $A_a(\alpha)$ satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0} | A_a(\alpha)] \asymp r^\alpha [1 + o(1)], \quad r \rightarrow \infty, \quad (4.27)$$

and we get

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp r^\alpha [1 + o(1)], \quad r \rightarrow \infty. \quad (4.28)$$

- (II) $\beta = \frac{1}{d^* - 1}$: critical regime. For $0 < \alpha < 1$, the situation is the same as in the subcritical regime described above. For $\alpha = 1$, the activation time of the next activating node is of order r . After this activation, all the remaining nodes are likely to activate in any possible order, but none of them will have an activation time of order larger than r . The order of activation of nodes in V follows the path a as long as the nucleation times associated to the nodes are of order smaller than r . After that, the remaining nodes can activate with positive probability in any order with an activation time of order at most r . Hence, the transition time conditional on the event $A_a(\alpha)$ satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0} | A_a(\alpha)] \asymp r [1 + o(1)], \quad r \rightarrow \infty, \quad (4.29)$$

and we get

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] \asymp r [1 + o(1)], \quad r \rightarrow \infty. \quad (4.30)$$

Note that if any of the nodes has an activation time of order r but larger than $T_U(r)$, then the transition time conditional on the event $A_a(\alpha)$ is the time it takes for the queue lengths at nodes in U to hit zero, which satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0} | A_a(\alpha)] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.31)$$

Hence,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.32)$$

- (III) $\beta \in (\frac{1}{d^* - 1}, \infty)$: supercritical regime. For $0 < \alpha < 1$, the situation is the same as in the subcritical regime described above. For $\alpha = 1$, the situation is the same as in the critical regime described above. For $\alpha > 1$, the transition time conditional on the event $A_a(\alpha)$ is the time it takes for the queue lengths at nodes in U to hit zero, which satisfies

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0} | A_a(\alpha)] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.33)$$

Hence,

$$\mathbb{E}_u[\mathcal{T}_{G(\cdot)}^{Q^0}] = T_U(r) [1 + o(1)], \quad r \rightarrow \infty. \quad (4.34)$$

□

Note that the order of the transition time does not depend on the path along which we compute it. The algorithm generates all possible activation paths of the nodes nucleating before time scale $\lambda(r)^{-1} = r^\alpha$. The remaining nodes can activate in any order depending on the dynamics. To compute the pre-factor of the mean transition time along these paths, we need to analyze in detail the Markov process describing the graph evolution, in particular, the degrees of the nodes changing over time. Our methods do not capture this detail and we are only able to state a result for the leading order term.

§4.4 The graph evolution

In this section we discuss the Markov process describing the graph evolution under the dynamics. Control on this process is the key to obtaining a more precise asymptotics for the mean transition time of the network.

§4.4.1 The graph evolution process

Consider a dynamics with rate $\lambda(r) = r^{-\alpha}$. We have seen in Proposition 4.2.4 that each node in V whose nucleation time is of smaller order than r^α activates through the nucleation of its fork. On time scale r^α the dynamics starts competing with the nucleation and the order of activation of the remaining nodes described by the algorithm is not preserved anymore. Note that the algorithm updates the graph at each iteration in order to keep track of the degree of the remaining nodes after each activation. When introducing the dynamics on the edges, we need information about the states of the nodes and the edges in the graph. We assume that the algorithm does not update the graph at each iteration anymore, but we focus on the number of active neighbors each node has.

Definition 4.4.1 (Active degree).

We define the *active degree* of a node as the number of its active neighbors. For $u \in U$, the active degree at time t is given by

$$\tilde{d}_u(t) = |\{v \in V : uv \in E(t), X_v(t) = 1\}|. \quad (4.35)$$

Analogously, for $v \in V$, the active degree at time t is given by

$$\tilde{d}_v(t) = |\{u \in U : uv \in E(t), X_u(t) = 1\}|. \quad (4.36)$$

Note that for a node to activate, its active degree must be 0. It is immediate to see that the active degree of a node cannot exceed its degree, i.e., for any $u \in U$ and $v \in V$,

$$\tilde{d}_u(t) \leq d_u(t), \quad \tilde{d}_v(t) \leq d_v(t). \quad (4.37)$$

The main challenge in describing the graph evolution is that any of the remaining nodes could activate next with positive probability. The activation of a node due to absence of edges is captured by the scenario in which its active degree hits 0. The activation of a node through the nucleation of its fork depends on the aggressiveness of

the activation rates and on the number of active neighbors. Both types of activation are determined by the degree evolution. Assume, for example, that an edge between two active nodes appears. By our model assumptions (see Remark 4.1.3), the node in U deactivates, implying that the active degrees of its neighbors in V decrease by 1. If the mean nucleation time of the new fork of one of the neighbors is of order less than or equal to r^α , then this neighbor will be more likely to activate through the nucleation of its fork. The degree evolution induced by the dynamics affects both the disconnection and the nucleation times of the nodes.

The node activity process $(X(t), Q(t))_{t \geq 0}$ and the edge activity process $(Y(t))_{t \geq 0}$ form a continuous-time Markov process on

$$\mathcal{X} \times \mathbb{R}_{\geq 0} \times \mathcal{Y} \tag{4.38}$$

that describes the evolution of the graph under the effect of the dynamics. We refer to this process as the *graph evolution process*. Note that if we know which nodes are active and which edges are present, then we can recover the degree and the active degree of each node in the graph. Hence, understanding the graph evolution process is crucial to describe how the degrees of the nodes change over time and how nodes activate.

§4.4.2 Transitions

Consider a feasible state where some nodes are active and some edges are present. By feasible we mean that it respects the constraints given by the edges, for which two connected nodes cannot be active simultaneously. Recall that $|U| = M$, $|V| = N$ and $|E_{MN}| = MN$. Hence, an arbitrary feasible state at time t has h active nodes in U with $h = 0, \dots, M$, k active nodes in V with $k = 0, \dots, N$, and l present edges with $l = 0, \dots, MN$. Consequently, there are $M - h$ inactive nodes in U , $N - k$ inactive nodes in V , and $MN - l$ absent edges. Note that the initial state u is described by $h = M$, $k = 0$ and $l = |E(0)|$, while the transition occurs as soon as state v is reached, for which $k = N$.

Clock ticks. The graph evolution is governed by different Poisson clocks ticking at various rates: the activation clocks, the deactivation clocks and the edge clocks. We analyze how the network evolves each time one of these clock ticks. Moreover, note that the queue lengths, hence the input process (recall Definition 1.1.3), also play a role, since the activation rates depend on them.

- The activation clock of a node $u \in U$ ticks at rate $g_U(Q_u(t))$ at time t . The probability of this clock being the first one to tick is given by

$$\frac{g_U(Q_u(t))}{Z}, \tag{4.39}$$

with

$$Z = \sum_{i=1}^{M-h} g_U(Q_i(t)) + \sum_{j=1}^{N-k} g_V(Q_j(t)) + h + k + MN\lambda(r). \tag{4.40}$$

The tick has two possible effects on the network. If the neighbors of u are all inactive, then u activates and the active degrees of all its neighbors increase by 1. If there is at least one active neighbor of u , then the activation attempt fails and nothing happens.

- The deactivation clock of a node $u \in U$ ticks at rate 1. The probability of this clock being the first one to tick is given by

$$\frac{1}{\bar{Z}}. \quad (4.41)$$

Node u deactivates and the active degrees of all its neighbors decrease by 1.

- The activation clock of a node $v \in V$ ticks at rate $g_V(Q_v(t))$ at time t . The probability of this clock being the first one to tick is given by

$$\frac{g_V(Q_v(t))}{Z}. \quad (4.42)$$

The tick has two possible effects on the network. If the neighbors of v are all inactive, then v activates and the active degrees of all its neighbors increase by 1. If there is at least one active neighbor of v , then the activation attempt fails and nothing happens.

- The deactivation clock of a node $v \in V$ ticks at rate 1. The probability of this clock being the first one to tick is given by

$$\frac{1}{\bar{Z}}. \quad (4.43)$$

Node v deactivates and the active degrees of all its neighbors decrease by 1.

- The activation clock of an edge $e \in E_{MN}$ ticks at rate $\lambda(r)$. The probability of this clock being the first one to tick is given by

$$\frac{\lambda(r)}{Z}. \quad (4.44)$$

Depending on which edge appears or disappears and on the nodes involved, the tick has different effects on the network, which are described below.

Edge appearing and disappearing. If we know the number of active nodes in U and V , then we can compute the probabilities of each of the following scenarios with simple combinatorial arguments. There are four possible scenarios in which an edge can appear.

- ($\circ\circ$) When an edge between two inactive nodes appears, their degrees increase by 1.
- ($\circ\bullet$) When an edge between an inactive node in U and an active node in V appears, the active degree of the node in U increases by 1 and the degree of the node in V increases by 1.

- (•◦) When an edge between an active node in U and an inactive node in V appears, the degree of the node in U increases by 1 and the active degree of the node in V increases by 1.
- (••) When an edge between two active nodes appears, the node in U deactivates, its active degree increases by 1, the active degrees of all its neighbors in V decrease by 1 and the degree of the node in V increases by 1.

In a similar fashion, there are three possible scenarios in which an edge can disappear. Recall that there cannot be an edge between two active nodes.

- (◦◦) When an edge between two inactive nodes disappears, their degrees decrease by 1.
- (◦•) When an edge between an inactive node in U and an active node in V disappears, the active degree of the node in U decreases by 1 and the degree of the node in V decreases by 1.
- (•◦) When an edge between an active node in U and an inactive node in V disappears, the degree of the node in U decreases by 1 and the active degree of the node in V decreases by 1.

The transition time is related to the graph evolution process, since the activation times of the nodes in V depend on the activation rates, the speed of the dynamics and the degree evolution. The complicated nature of the process prevents us from deriving an explicit formula for the pre-factor of the mean transition time, which would require a better control on the precise asymptotics of each activation.

§D Appendix: a model with fixed activation rates

We have seen how the dynamics influences the mean transition time of wireless random-access models where the activation rates depend on the current queue lengths at the nodes. The model is quite challenging and deals with two levels of complexity, namely, the queue-based activation rates and the edge dynamics. Not much is known in the literature for random-access protocols with dynamic interference graph, even for models with *fixed activation rates*. In this section we adapt the theory built up in Chapters 2–3 to study the effect of the dynamics on such type of models. Assume that the activation rates are of the form

$$r_i(t) = \begin{cases} r^\beta, & \text{if } i \in U, \\ r^{\beta'}, & \text{if } i \in V, \end{cases} \quad (4.45)$$

with $\beta, \beta' \in (0, \infty)$ and $\beta' > \beta + 1$. We recall that we are interested in the transition time asymptotics as $r \rightarrow \infty$.

We start by adapting the results for complete bipartite graphs in Chapter 2 to the model with fixed activation rates. The following theorem is consistent with [59, Example 4.1].

Theorem D.1 (Complete bipartite graphs with fixed activation rates).

Consider the complete bipartite graph $G = ((U, V), E)$ with initial queue lengths Q^0 as in (1.11). Suppose that (4.45) holds.

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

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{1}{|U|} r^{\beta(|U|-1)} [1 + o(1)], \quad r \rightarrow \infty. \quad (4.46)$$

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

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{1}{|U|} r [1 + o(1)], \quad r \rightarrow \infty. \quad (4.47)$$

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

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{\gamma U}{c - \rho U} r [1 + o(1)], \quad r \rightarrow \infty. \quad (4.48)$$

Proof. The claims follow from Sections 2.4.1–2.4.2 in Chapter 2. We compute the critical time scale and the mean transition time using fixed activation rates instead of time depending ones. In both the critical and subcritical regimes, the pre-factor turns out to be $\frac{1}{|U|}$ and the law is exponential. In the critical regime, we know that the queue lengths decrease significantly after a time of order r . However, this does not affect the transition time, since now the activation rates do not depend on the queue lengths. In the supercritical regime, we still have the same behavior as in the model with queue-dependent activation rates. Indeed, when the queue lengths at nodes in U hit zero, the nodes in U deactivate by assumption and the transition occurs. \square

Next, we state a result for arbitrary bipartite graphs with fixed activation rates (analogue of Theorem 3.3.3 in Chapter 3). Note that the algorithm still plays a crucial role in determining the mean transition time.

Theorem D.2 (Arbitrary bipartite graphs with fixed activation rates).

Consider the bipartite graph $G = ((U, V), E)$ with initial queue lengths Q^0 as in (1.11). Suppose that (4.45) holds. Let A_a be the event that the network follows the path $a \in \mathcal{A}$, among the paths generated by the algorithm.

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

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] = \sum_{\substack{1 \leq k \leq N \\ k: \bar{d}_k = d^*}} \frac{1}{n_k d^*} r^{\beta(d^*-1)} [1 + o(1)], \quad r \rightarrow \infty. \quad (4.49)$$

(II) $\beta = \frac{1}{d^*-1}$: critical regime. Then the transition time satisfies

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0} | A_a] = \sum_{\substack{1 \leq k \leq N \\ k: \bar{d}_k = d^*}} \frac{1}{n_k d^*} r [1 + o(1)], \quad r \rightarrow \infty. \quad (4.50)$$

The above result holds as long as the pre-factor is below the value $\frac{\gamma_U}{c-\rho_U}$, which corresponds to the time it takes for the queue lengths at nodes in U to hit zero. Otherwise, the supercritical regime applies.

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

$$\mathbb{E}_u[\mathcal{T}_G^{Q^0}] = \frac{\gamma_U}{c-\rho_U} r [1 + o(1)], \quad r \rightarrow \infty. \quad (4.51)$$

Proof. The claims follow from Theorem D.1 and the analysis of the algorithm and the next nucleation times in Sections 3.2 and 3.4.2 in Chapter 3. We derive the mean transition time along the paths generated by the algorithm by computing the next nucleation times at each step. In the subcritical regime, the nucleation times of nodes in V are all exponentially distributed and independent of each other. Indeed, the activation rates are the same, independently of the queue lengths decreasing over time. At each step k , the next nucleation time is the minimum of n_k i.i.d. exponential random variables, and hence its mean exhibits the term $f_k = \frac{1}{n_k}$ in the pre-factor. In the critical regime, the pre-factor of the mean transition time along each path must be below the value $\frac{\gamma_U}{c-\rho_U}$, otherwise the supercritical regime applies and the transition occurs because the queue lengths at nodes in U hit 0. If we assume that $\frac{\gamma_U}{c-\rho_U} > 1$, then the nucleation of a fork occurs before the queue lengths at nodes in U hit zero. We are able to derive the law of the transition time along each path for both the subcritical and critical regimes. Both are described by convolutions of the exponential laws of the next nucleation times of the activating nodes in V . In the supercritical regime, we have the same behavior as in the model with queue-dependent activation rates. \square

Finally, we show that the results from Theorem 4.1.4 also hold when we consider a dynamic bipartite graph with fixed activation rates. We are able to compute the order of the mean transition time, while the pre-factor still depends on the graph evolution described in Section 4.4.

Theorem D.3 (Dynamic bipartite graphs with fixed activation rates).

Consider the dynamic bipartite graph $G(\cdot) = ((U, V), E(\cdot))$ with the edge dynamics governed by $\lambda(r)$ and initial queue lengths Q^0 . Suppose that (4.45) holds. Then the results of Theorem 4.1.4 hold.

Proof. The claim follows from Theorem D.2 and the intuition behind Proposition 4.2.4. The order of the mean transition time in the model with fixed activation rates is the same as in the model with queue-dependent activation rates. The dynamics competes with the nucleations of the nodes in the same way, depending on its speed. The different type of dynamics (fast, regular and slow) lead to the same results as in Theorem 4.1.4. \square

