

Delay performance in random-access grid networks



A. Zocca^{a,*}, S.C. Borst^{a,b}, J.S.H. van Leeuwen^a, F.R. Nardi^a

^a Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

^b Bell Laboratories, Alcatel-Lucent, P.O. Box 636, Murray Hill, NJ 07974, USA

ARTICLE INFO

Article history:

Available online 29 August 2013

Keywords:

Wireless random-access networks

Grid topologies

Delay scaling

Hitting times

Conductance

Mixing times

ABSTRACT

We examine the impact of torpid mixing and meta-stability issues on the delay performance in wireless random-access networks. Focusing on regular meshes as prototypical scenarios, we show that the mean delays in an $L \times L$ toric grid with normalized load ρ are of the order $(\frac{1}{1-\rho})^L$. This superlinear delay scaling is to be contrasted with the usual linear growth of the order $\frac{1}{1-\rho}$ in conventional queueing networks. The intuitive explanation for the poor delay characteristics is that (i) high load requires a high activity factor, (ii) a high activity factor implies extremely slow transitions between dominant activity states, and (iii) slow transitions cause starvation and hence excessively long queues and delays.

Our proof method combines both renewal and conductance arguments. A critical ingredient in quantifying the long transition times is the derivation of the communication height of the uniformized Markov chain associated with the activity process. We also discuss connections with Glauber dynamics, conductance, and mixing times. Our proof framework can be applied to other topologies as well, and is also relevant for the hard-core model in statistical physics and sampling from independent sets using single-site update Markov chains.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Emerging wireless mesh networks typically lack any centralized control entity, and instead vitally rely on a distributed mechanism for regulating the access of the various nodes to the shared medium. A popular mechanism for distributed medium access control is provided by the so-called Carrier-Sense Multiple-Access (CSMA) protocol, various incarnations of which are implemented in IEEE 802.11 networks. In the CSMA protocol, each node attempts to access the medium after a random back-off time, but nodes that sense activity of interfering nodes freeze their back-off timer until the medium is sensed idle.

While the CSMA protocol is fairly easy to understand at a local level, the interaction among interfering nodes gives rise to quite intricate behavior and complex throughput characteristics on a macroscopic scale. In recent years, relatively parsimonious models have emerged that provide a useful tool in evaluating the throughput characteristics of CSMA-like networks [1–4]. These models essentially assume that the interference constraints can be represented by a general conflict graph, and that the various nodes activate asynchronously whenever none of their neighbors are presently active. Although such a representation of the IEEE 802.11 back-off mechanism is not as detailed as in the landmark work of Bianchi [5], the general conflict graph provides far greater versatility in describing a broad range of topologies. Experimental results in

* Corresponding author.

E-mail addresses: a.zocca@tue.nl (A. Zocca), s.c.borst@tue.nl (S.C. Borst), j.s.h.v.leeuwen@tue.nl (J.S.H. van Leeuwen), f.r.nardi@tue.nl (F.R. Nardi).

Liew et al. [6] demonstrate that these models, while idealized, yield throughput estimates that match remarkably well with measurements in actual IEEE 802.11 networks.

Besides long-term throughput estimates, the above-described models also provide valuable qualitative insights in transient characteristics. In particular, the models reveal that the activity process of the various nodes may exhibit slow mixing and meta-stable behavior due to slow transitions between dominant states [7,8]. As a result, individual nodes will experience long sequences of transmissions in rapid succession, interspersed with extended periods of starvation. Since the above-described models assume saturated buffers, they do not explicitly capture the effect of temporal starvation on queue lengths and delays though. In the present paper, we therefore augment the basic model with queueing dynamics, and analyze average steady-state delays in order to examine the performance repercussions of temporal starvation.

Only a few results for average steady-state delays in random-access networks have been obtained so far. An interesting paper by Shah et al. [9] showed that low-complexity schemes cannot achieve low delay in arbitrary topologies (unless P equals NP), since that would imply that certain NP-hard problems could be solved efficiently. However, the notion of delay in [9] is a transient one, and it is not exactly clear what the implications are for the average steady-state delay in specific networks, if any. Jiang et al. [10] derived *upper* bounds for the average steady-state delay. The bounds show that for low load the delay grows only polynomially with the number of nodes in bounded-degree graphs. Subramanian and Alanyali [11] presented similar upper bounds for bounded-degree graphs with low load based on analysis of neighbor sets and stochastic coupling arguments.

Focusing on regular meshes as prototypical scenarios, we establish that the average steady-state delays in an $L \times L$ toric grid with normalized load ρ scale as $(\frac{1}{1-\rho})^L$, in contrast to the typical linear growth of the order $\frac{1}{1-\rho}$ in conventional queueing networks. The intuitive explanation for the adverse delay scaling is that (i) high load requires a high activity factor, (ii) a high activity factor implies extremely slow transitions between dominant activity states, and (iii) slow transitions cause starvation and hence excessively long queues and delays.

Our novel proof method combines both renewal and conductance arguments. A critical role in determining the long transition times in (ii) is played by the derivation of the communication height of the uniformized Markov chain associated with the activity process. As a side-result, we gain valuable insight in the exponentially small conductance of Glauber dynamics on the independent sets in the $L \times L$ grid, and hence in the resulting exponentially large mixing time. Our proof methodology can be applied to other graphs as well, and is also relevant for the hard-core model in statistical physics [12] and sampling from independent sets using single-site update Markov chains in the theoretical computer science literature [13–15].

The remainder of the paper is organized as follows. In Section 2, we present a detailed model description. In Section 3, we give an overview of the main results for the delays, transition times, communication height, and mixing times, and provide heuristic interpretations and high-level sketches of the proof arguments. In Section 4, we establish the intimate connection between long transition times and long delays. Section 5 is devoted to the derivation of the communication height, which plays a pivotal role in characterizing the slow transitions between dominant activity states, proved in Section 6. In Section 7, we make some concluding remarks and sketch some directions for further research.

2. Model description

In this section, we present a detailed description of the model for random-access networks with queueing dynamics. In Section 2.1, we specify the network dynamics for arbitrary conflict graphs with node-specific arrival, transmission, and activation rates. To facilitate a more insightful analysis, we introduce in Section 2.2 several symmetry assumptions, amounting to toric grids with homogeneous arrival, transmission, and activation rates. This symmetric network will be the focus throughout the paper. In Section 2.3, we establish the stability conditions for the network.

2.1. Network dynamics

Consider a network of N nodes sharing a wireless medium. The network is represented by an undirected graph $G = (V, E)$, called the *conflict graph* or *interference graph*, where the set of vertices $V = \{1, \dots, N\}$ corresponds to the various nodes and the set of edges $E \subseteq V \times V$ indicates which pairs of nodes interfere. Nodes that are neighbors in the conflict graph are prevented from simultaneous activity, and thus the independent sets of G (sets of vertices not sharing any edge) correspond to the feasible joint activity states of the network. A node is said to be *blocked* whenever the node itself or any of its neighbors is active, and *unblocked* otherwise. Denote by $\Omega \subseteq \{0, 1\}^N$ the collection of all independent sets of G . Packets arrive at node i as a Poisson process of rate λ_i . The packet transmission times at node i are independent and exponentially distributed with mean $1/\mu_i$. Let $X(t) \in \Omega$ represent the joint activity state of the network at time t , with $X_i(t)$ indicating whether node i is active at time t or not. Let $\{Q(t)\}_{t \geq 0}$ be the joint queue-length process, with $Q_i(t)$ the number of packets waiting at node i at time t (excluding the packet that may be in the process of being transmitted).

The various nodes share the medium according to a random-access mechanism. When a node ends a packet transmission, it either starts a back-off period with probability p_i before the next packet transmission, or immediately starts the next packet transmission otherwise. The back-off periods of node i are independent and exponentially distributed with mean $1/v_i$.

The back-off period of a node is suspended whenever it becomes blocked by activity of any of its neighbors, and only resumes once the node becomes unblocked again.

Under this random-access mechanism, the process $\{(X(t), Q(t))\}_{t \geq 0}$ evolves as a continuous-time Markov process with state space $\Omega \times \mathbb{N}^N$. Transitions from a state (X, Q) to $(X, Q + e_i)$ due to arrivals occur at rate λ_i , transitions due to activations from a state (X, Q) with $X_i = 0$ and $X_j = 0$ for all neighbors j of node i to $(X + e_i, Q - e_i I_{\{Q_i > 0\}})$ occur at rate ν_i , transitions due to transmission completions followed by a back-off period from a state (X, Q) with $X_i = 1$ to $(X - e_i, Q)$ occur at rate $p_i \mu_i$, and transitions due to transmission completions immediately followed by another transmission (without an intermediate back-off period) from a state (X, Q) with $X_i = 1$ to $(X, Q - e_i I_{\{Q_i > 0\}})$ occur at rate $(1 - p_i) \mu_i$.

For any $x \in \Omega$, define $\pi(x) = \lim_{t \rightarrow \infty} \mathbb{P}(X(t) = x)$ as the stationary probability that the activity process resides in state x . The activity process $\{X(t)\}_{t \geq 0}$ does not depend on the process $\{Q(t)\}_{t \geq 0}$, and in fact constitutes a reversible Markov process with a product-form stationary distribution [1]:

$$\pi(x) = Z^{-1} \prod_{i=1}^N \sigma_i^{x_i}, \quad x \in \Omega,$$

where $\sigma_i = \nu_i / (p_i \mu_i)$ and Z is the normalization constant defined as

$$Z = \sum_{x \in \Omega} \prod_{i=1}^N \sigma_i^{x_i}.$$

In contrast, the queue-length process $\{Q(t)\}_{t \geq 0}$ does strongly depend on the activity process $\{X(t)\}_{t \geq 0}$, and is considerably harder to analyze. Since the evolution of $\{Q(t)\}_{t \geq 0}$ is modulated by that of $\{X(t)\}_{t \geq 0}$, the former process can be viewed as a queueing network in a random environment, and in principle the queue-length process at each individual node could be considered as a quasi-birth-and-death process, and analyzed using matrix-analytic techniques. However, this would only yield numerical solutions and not provide explicit insight into how the queue-length process and delay performance are affected by the temporal starvation issues.

Remark. We have implicitly assumed that a node may become active even when its buffer is empty, and then transmits dummy packets. It would be interesting to extend the analysis to the case where nodes with empty buffers refrain from transmission activity, but this would be challenging, since the behavior of the activity process then does depend on the queue-length process.

2.2. Symmetry assumptions

To obtain more transparent results, we henceforth focus on a symmetric scenario with homogeneous arrival, transmission, and activation rates, as well as equal back-off probabilities at all nodes:

$$\lambda_i \equiv \lambda, \quad \mu_i \equiv \mu, \quad \nu_i \equiv \nu, \quad p_i \equiv p, \quad i = 1, \dots, N.$$

The stationary distribution of the activity process then simplifies to

$$\pi(x) = Z^{-1} \sigma^{\sum_{i=1}^N x_i}, \quad x \in \Omega, \tag{1}$$

where $\sigma = \nu / (p\mu)$ and the normalization constant may be written as $Z = \sum_{x \in \Omega} \sigma^{\sum_{i=1}^N x_i}$. As (1) shows, when $\sigma > 1$, the stationary probability of an activity state $x \in \Omega$ increases with its cardinality in an exponential fashion.

As prototypical examples of a symmetric topology, we will focus on toric grids, i.e., those with a periodic (wrap-around) boundary to preserve symmetry. More precisely, we consider as conflict graph $G = (V, E)$ the even discrete torus of size L . The vertex set is $V = \{1, \dots, L\}^2$; hence $N = |V| = L^2$, and two vertices in V are adjacent if they differ by $1 \pmod L$ in only one coordinate. A vertex is called *even* (*odd*) if the sum of its two coordinates is even (odd). The vertex set V is thus partitioned into two classes: the set \mathcal{E} of even nodes and the set \mathcal{O} of odd nodes. Clearly, $|\mathcal{E}| = L^2/2 = |\mathcal{O}|$, and these are the two maximum independent sets of the graph G . In the case of the toric grid, (1) shows that, as σ grows large, the probability mass will concentrate on the maximum independent sets \mathcal{E} and \mathcal{O} , to which we will also refer as *dominant activity states*.

For the toric grid, the conflict graph G is a bipartite graph, and in the special case $L = 2$, it becomes even a *complete* bipartite graph with components \mathcal{E} and \mathcal{O} . This conflict graph has additional features which simplify its analysis [8], and in the remainder of the paper we will therefore assume that $L \geq 4$.

By virtue of the symmetry, all the nodes will be active the same fraction of time θ , which will approach $1/2$ as σ grows large, since the probability mass will then concentrate on the two dominant activity states \mathcal{E} and \mathcal{O} . However, the transitions between the two dominant activity states will occur at an increasingly slow rate, resembling meta-stability phenomena in statistical physics [16,17]. As a result, individual nodes will experience long sequences of transmissions in rapid succession, during which their buffers drain, interspersed with extended periods of starvation, during which large numbers of packets accumulate. Our objective in this paper is to examine the effect of the temporal starvation on queue lengths and delays. As it

turns out, the analysis will revolve around the transition times $T_{\mathcal{E} \rightarrow \mathcal{O}}$ and $T_{\mathcal{O} \rightarrow \mathcal{E}}$ between the two dominant activity states, defined as

$$T_{\mathcal{E} \rightarrow \mathcal{O}} := \inf\{t > 0 : X(t) = \mathcal{O} \mid X(0) = \mathcal{E}\} \quad \text{and} \quad T_{\mathcal{O} \rightarrow \mathcal{E}} := \inf\{t > 0 : X(t) = \mathcal{E} \mid X(0) = \mathcal{O}\}.$$

Note that these two random variables only depend on the activity process $\{X(t)\}_{t \geq 0}$, and are independent of the queue-length process $\{Q(t)\}_{t \geq 0}$.

2.3. Stability conditions

In view of the above observations, it is natural to define the normalized load as $\rho = \lambda/2\mu$ and to assume $\rho < 1$, since otherwise all the queues will be unstable. The latter condition is not only necessary for stability, but in fact also sufficient for the queue-length processes at each individual node to be positive recurrent for large enough values of σ , as stated in the next proposition.

Proposition 2.1. *Assume that $\rho < 1$. Then the queue-length process $\{Q_i(t)\}_{t \geq 0}$ at each individual node is positive recurrent for sufficiently large values of σ .*

The above proposition may be intuitively explained as follows. As mentioned earlier, the fraction of time θ that each individual node is active will approach $1/2$ arbitrarily close for sufficiently large values of σ . This suggests that $\lambda < \mu/2$, i.e., $\rho < 1$, is sufficient for stability for large enough values of σ .

Proof. We say that the activity process resides in an even period if the last dominant state it visited was \mathcal{E} , and likewise in an odd period if the last dominant state it visited was \mathcal{O} . We refer to a period consisting of an even and subsequent odd period as a cycle. Define $\{\hat{Q}_n\}_{n \in \mathbb{N}}$ to be the queue-length process of an odd node embedded at the epochs when even periods start (or equivalently when cycles start). Thanks to the strong Markov property, this latter process is an irreducible time-homogeneous Markov chain on \mathbb{N} . Denote by S_n the number of packets served during the n th odd period and by A_n the total number of packets which arrived during the n th cycle. By construction, these two random variables are independent of each other and independent from those of different cycles. We claim that the following inequality holds:

$$\hat{Q}_{n+1} \leq (\hat{Q}_n - S_n)^+ + A_n.$$

Indeed we underestimate the total number of packets served by allowing transmission only during the odd period (instead of continuously) and only the awaiting packets which arrived in the previous cycles.

Consider now another fictitious Markov chain $\{\bar{Q}_n\}_{n \in \mathbb{N}}$ defined by the initial condition $\bar{Q}(0) = \hat{Q}(0)$ and by the classical Lindley recursion $\bar{Q}_{n+1} = (\bar{Q}_n - S_n + A_n)^+$. Clearly, $\hat{Q}_n \leq \bar{Q}_n + A_n$ for every $n \in \mathbb{N}$. Let Y be a random variable distributed as $A - S$. It is well known that 0 is a positive recurrent state for $\{\bar{Q}_n\}_{n \in \mathbb{N}}$ if $\mathbb{E}Y < 0$. This is indeed the case, since $\mathbb{E}Y = \lambda \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}} + (\lambda - \mu) \mathbb{E}T_{\mathcal{O} \rightarrow \mathcal{E}} = (2\lambda - \mu) \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}} = \mu(\rho - 1) < 0$. Therefore, also for the chain $\{\hat{Q}_n\}_{n \in \mathbb{N}}$, the state 0 is positive recurrent and, since the cycle time has finite expectation (the Markov process $\{X(t)\}_{t \geq 0}$ is irreducible on a finite state space, and hence positive recurrent), it follows that the original continuous-time process $\{Q_i(t)\}_{t \geq 0}$ is positive recurrent as well. \square

The next proposition provides a lower bound for the value of σ that is required for stability in terms of the normalized load ρ .

Proposition 2.2. *In order for all queues to be stable, it is required that*

$$\sigma > \frac{\rho}{2(1 - \rho)}.$$

Proof. In order to be stable, each individual node should be active at least a fraction $\lambda/\mu = \rho/2$ of the time, and thus $\theta > \rho/2$. The fraction of time that a node is active may be expressed as $\theta = \sigma \mathbb{P}(\text{node is unblocked})$. Now consider an arbitrary node and any of its neighbors. Then the fraction of time that the former node is unblocked is bounded from above by the fraction of time that both nodes are inactive, so $\mathbb{P}(\text{node is unblocked}) \leq 1 - \rho$. Combining the above three inequalities and some rearranging yields the result. \square

3. Main results

In Section 3.1, we present the two main theorems, which lead to bounds for the long-term average delay in random-access grid networks. The first result follows from a third theorem on the so-called communication height, presented in Section 3.2, along with a sketch of the proof. In Section 3.3, we show how the result for the communication height can be leveraged to characterize the mixing time of the activity process, and we also discuss connections with other results in the areas of statistical physics and theoretical computer science.

3.1. Transition times and delays

As mentioned earlier, the bounds we derive revolve around three simple observations: (i) high load requires high activity factors for stability; (ii) high activity factors cause long mixing times, in particular *slow transitions between dominant activity states*; (iii) slow transitions between dominant states imply long starvation periods, and hence large queue lengths and delays.

The next two theorems formalize the last two observations. We henceforth assume μ to be fixed, and when we write $\sigma \rightarrow \infty$ we allow for either $\nu \rightarrow \infty, p \downarrow 0$, or both. For compactness, we only attach σ in brackets to various quantities to reflect the dependence on both ν and p in limit statements for $\sigma \rightarrow \infty$. Define $\alpha := \liminf_{\sigma \rightarrow \infty} -\frac{\log p}{\log \sigma} \geq 0$.

Theorem 3.1 (*Transition Time Between Dominant States*).

$$\liminf_{\sigma \rightarrow \infty} \frac{\log(\mu \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma))}{\log \sigma} \geq L + \alpha \geq L.$$

With minor abuse of notation, let $\mathbb{E}W(\sigma)$ be the long-term average delay (waiting time plus service time) of an arbitrary packet, where again the attached σ in brackets describes the dependence on both ν and p .

Theorem 3.2 (*Long-Term Average Delay*).

$$\liminf_{\sigma \rightarrow \infty} \frac{\mathbb{E}W(\sigma)}{\mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)} \geq \frac{1}{4} \frac{\mu}{\mu - \lambda} = \frac{1}{4 - 2\rho}.$$

In view of Proposition 2.2 and Theorem 3.2, an immediate corollary of the above theorem is that the long-term average delay must scale as $(\frac{1}{1-\rho})^L$ as $\rho \uparrow 1$. This is to be contrasted with the usual linear scaling in $\frac{1}{1-\rho}$ in conventional queueing networks, and further reveals that the exponential rate of growth increases with the size of the grid.

Theorem 3.2 may be understood as follows. First, by virtue of the symmetry, we may as well consider the average delay of an arbitrary packet at a particular node, say an odd one. In view of the nature of the transitions between the dominant states \mathcal{E} and \mathcal{O} , this node will hardly be active during the transition time $T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ for large values of σ . As a result, the queue at this node will roughly grow at rate λ , reaching an expected length of $\lambda \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ at the end of the transition period, and thus having an average size of approximately $\frac{1}{2} \lambda \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ over this period. During the subsequent transition time $T_{\mathcal{O} \rightarrow \mathcal{E}}(\sigma)$, the queue will decrease at a rate of at most $\lambda - \mu$, and thus requires a time period of at least $\frac{\lambda}{\mu - \lambda} \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ to empty, so the average size over that period is at least

$$\frac{1}{2} \frac{\lambda^2}{\mu - \lambda} \frac{\mathbb{E}((T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma))^2)}{\mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)}.$$

Noting that $\mathbb{E}((T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma))^2) \geq (\mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma))^2$ and $\mathbb{E}T_{\mathcal{O} \rightarrow \mathcal{E}}(\sigma) = \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$, we conclude that the average queue length over the course of the cycle consisting of the transition periods $T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ and $T_{\mathcal{O} \rightarrow \mathcal{E}}(\sigma)$ is at least

$$\frac{1}{2} \left(\frac{1}{2} \lambda + \frac{1}{2} \frac{\lambda^2}{\mu - \lambda} \right) \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma) = \frac{\lambda}{4} \frac{\mu}{\mu - \lambda} \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma) = \frac{\lambda}{4 - 2\rho} \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$$

for sufficiently large values of σ . Invoking Little's law, we then obtain the lower bound as stated in the theorem for the long-term average delay of an arbitrary packet.

Of course, the above arguments are heuristic in nature. We will provide a rigorous proof of Theorem 3.2 in Section 4. Specifically, we will establish that an odd node receives hardly any service during a transition time $T_{\mathcal{E} \rightarrow \mathcal{O}}(\sigma)$ for large values of σ , and justify the informal calculations based on drifts and expected values.

3.2. Communication height

The high-level idea of the proof of Theorem 3.1 is that, in order for a transition from the even dominant state \mathcal{E} to the odd dominant state \mathcal{O} to occur, the activity process $X(t)$ must follow a transition path through some highly unlikely configurations consisting of mixtures of even and odd active nodes, and the time to reach such configurations is correspondingly long. Theorem 3.3 below shows a key feature of all such transition paths.

In the rest of the paper, we will interchangeably refer to a feasible configuration $x \in \Omega$ or to the corresponding independent set I of the conflict graph G . Clearly, $\sum_{i=1}^N x_i = |I|$. For any independent set $I \in \Omega$, define $\Delta(I)$ to be the difference between the size of a maximum independent set and that of the set I , i.e., $\Delta(I) := L^2/2 - |I|$. The quantity $\Delta(I)$ describes how “inefficient” the configuration I is with respect to the maximum independent sets \mathcal{E} and \mathcal{O} , and hence how much less likely it is in view of (1). For this reason, we will refer to $\Delta(I)$ as the *efficiency gap* of configuration I . A path

is a sequence $\omega = (I_1, \dots, I_n)$, $n \in \mathbb{N}$, where $I_k \in \Omega$ for $k = 1, \dots, n$ and such that I_k and I_{k+1} differ by a single flip, for $k = 1, \dots, n-1$. We write $\omega : I \rightarrow I'$ to denote a path from I to I' . The communication height between two states $I, I' \in \Omega$ is

$$\phi(I, I') := \min_{\omega: I \rightarrow I'} \max_{J \in \omega} \Delta(J).$$

The quantity $\phi(I, I')$ is the minimum number of nodes which must be simultaneously inactive with respect to the maximum number of active nodes to make the transition $I \rightarrow I'$ possible. The cornerstone of the proof of Theorem 3.1 is the derivation of the communication height $\Gamma := \phi(\mathcal{E}, \mathcal{O})$ between the two dominant configurations \mathcal{E} and \mathcal{O} .

Theorem 3.3 (Communication Height).

$$\Gamma = L + 1.$$

The proof of Theorem 3.3, given in Section 5, is technically challenging, since all possible paths between the dominant states need to be taken into consideration. In order to do so, we use a similar approach as in [13–15] which associates with each independent set $I \in \Omega$ a collection of cutsets and contours, that effectively translates the problem into geometric properties of these contours.

Theorem 3.3 is a powerful result. In Section 6, we show that together with a general result for exit times for interacting particle systems in statistical physics [17] it readily leads to Theorem 3.1. Theorem 3.3 also gives a way to derive results for mixing times. Since the majority of related literature has been dealing with mixing times, instead of transition times, we present in Section 3.3 the mixing time result that can be obtained using Theorem 3.3.

3.3. Mixing times

Let $d(t)$ be the maximal distance over $x \in \Omega$, measured in terms of total variation, between the distribution at time t and the stationary distribution, i.e., $d(t) := \max_{x \in \Omega} \|\mathbb{P}(X(t) \in \cdot \mid X(0) = x) - \pi\|_{TV}$, where $\|p - q\|_{TV} = \frac{1}{2} \sum_{x \in \Omega} |p(x) - q(x)|$ is the total variation distance between two probability measures p, q on the finite state space Ω . The mixing time of the process $\{X(t)\}_{t \geq 0}$ is defined as

$$t_{\text{mix}}(\varepsilon, \sigma) = \inf\{t \geq 0 : d(t) \leq \varepsilon\},$$

and it expresses the time needed for a process to converge to stationarity.

An important observation is that, when σ is large, the distribution in (1) favors dense configurations with $\Delta(I)$ small and with extremes \mathcal{E} and \mathcal{O} , while configurations that are roughly half odd and half even are highly unlikely. This ‘bimodality’ causes slow convergence to equilibrium, because the Markov process is either in the even or odd sublattice, while it can hardly pass through the less likely set in the middle (the bottleneck) to go from the even to the odd sublattice or vice versa. Hence, a classical strategy for proving slow mixing is to separate the state space into three sets, two of which are exponentially larger (in terms of probability) than the third, such that moving from one large set to the other through local steps of the Markov process requires passing through the small bottleneck set in the middle. This strategy has been applied in [13–15] to prove slow mixing of local or single-site update Markov processes (like ours) on the grid. Traditionally, the two dominant sets are those configurations that lie predominantly on the even or odd sublattice, and the bottleneck set consists of those configurations that have a relatively balanced number of even and odd vertices. The challenge is then to show that this bottleneck or balanced set has exponentially small probability, in which case the conductance of the Markov process would be exponentially small, leading via a standard argument to exponentially large mixing times. This challenge essentially boils down to counting the number of balanced configurations, which is solved in [13,14] by powerful combinatorial enumeration methods. In [15], a different way of partitioning was introduced, which divides the state space according to a certain topological obstruction, rather than according to the relative number of even and odd vertices. In that case, again a counting problem needs to be solved, in order to establish that the set of configurations having this topological obstruction is exponentially small, but the counting problem is much simpler. Recently, the same approach was followed in [18], but with sharper estimates for the counting problem.

What is new about our approach is that we partition the state space in a fundamentally different way, such that there is no longer the need to solve a counting problem. The approach is best explained in terms of the set

$$\mathcal{S} := \{I \in \Omega \mid \phi(\mathcal{E}, I) \leq L\}. \tag{2}$$

From the perspective of the maximum configuration \mathcal{E} , the set \mathcal{S} contains all configurations that can be reached by the process by having at most L nodes simultaneously inactive with respect to the maximum number of active nodes. Moreover, Theorem 3.3 tells us that the Markov process starting in \mathcal{E} cannot reach the configuration \mathcal{O} without leaving the set \mathcal{S} . Hence, the two dominant sets are $\mathcal{S} = \{I \in \Omega \mid \phi(\mathcal{E}, I) \leq L\}$, and by symmetry $\mathcal{S}' := \{I \in \Omega \mid \phi(\mathcal{O}, I) \leq L\}$, while the remaining configurations together form the bottleneck. Now, since the distribution (1) readily provides an upper bound on the probability of this bottleneck set, there is no counting problem involved in proving exponentially small conductance and hence slow mixing; see Theorem 3.5 below.

Let us now demonstrate the above reasoning to get an upper bound on the conductance of the activity process and thus a lower bound for the mixing time. These two results are not necessary for deriving the delay bounds, but are of independent interest in view of Refs. [13–15]. Therefore, for simplicity, we assume that $p = 1$ and $\mu = 1$ in the rest of this subsection.

For $S \subseteq \Omega$, let $\pi(S) := \sum_{x \in S} \pi(x)$ be its stationary probability, and define the *conductance* of $S \subseteq \Omega$ as

$$\Phi(S) := \frac{\sum_{x \in S} \sum_{y \in S^c} \pi(x)q(x, y)}{\pi(S)},$$

where $q(x, y)$ is the transition rate from state x to state y . The conductance of a Markov process, also known in the literature as the *bottleneck ratio* or *Cheeger constant*, is effectively a measure of how well connected the state space is, taking into account possible geometric features, usually referred to as bottlenecks, that strongly influence the mixing time.

Let $\partial^- \mathcal{S}$ denote the inner boundary of the set \mathcal{S} , i.e., the states in \mathcal{S} that communicate with \mathcal{S}^c .

Lemma 3.4 (*Upper Bound on Conductance of \mathcal{S}*).

$$\Phi(\mathcal{S}) \leq \frac{|\partial^- \mathcal{S}|(L^2/2 - L)}{\sigma^L}.$$

Proof. From the definition of \mathcal{S} , it follows that all states $x \in \partial^- \mathcal{S}$ have the same stationary probability, namely $\pi(x) = Z^{-1} \sigma^{L^2/2-L}$. Moreover, the only way to exit from \mathcal{S} is by deactivating one node, and this happens at rate $q(x, y) = \mu = 1$, if $x \in \mathcal{S}, y \in \mathcal{S}^c$, and $\sum_{i=1}^N |x_i - y_i| = 1$. Hence

$$\sum_{x \in \mathcal{S}} \sum_{y \in \mathcal{S}^c} \pi(x)q(x, y) \leq Z^{-1} \sigma^{L^2/2-L} |\partial^- \mathcal{S}|(L^2/2 - L).$$

The last inequality comes from the fact that the exit from \mathcal{S} can occur only by deactivation of one node, and there are exactly $L^2/2 - L$ of them to be turned off in each configuration in $\partial^- \mathcal{S}$. Trivially, $\pi(\mathcal{S}) \geq \pi(\mathcal{E}) = Z^{-1} \sigma^{L^2/2}$, and thus

$$\Phi(\mathcal{S}) \leq \frac{\sigma^{L^2/2-L}}{\sigma^{L^2/2}} |\partial^- \mathcal{S}|(L^2/2 - L) = \frac{|\partial^- \mathcal{S}|(L^2/2 - L)}{\sigma^L},$$

which completes the proof. \square

Due to the relation (see [8])

$$t_{\text{mix}}(\epsilon) \geq \left(\frac{1}{2} - 2\epsilon\right) \max_{A \subseteq \Omega: \pi(A) \leq 1/2} \frac{1}{\Phi(A)},$$

the following result is then an immediate consequence of Theorem 3.3 and Lemma 3.4.

Theorem 3.5 (*Lower Bound on Mixing Time*). For any $\sigma > 1$, the mixing time of the process $\{X(t)\}_{t \geq 0}$ satisfies

$$t_{\text{mix}}(\epsilon, \sigma) \geq \left(\frac{1}{2} - 2\epsilon\right) \frac{\sigma^L}{|\partial^- \mathcal{S}|(L^2/2 - L)}.$$

Remark. At first sight, our novel approach might look deceptively simple, because no sophisticated enumeration method is needed. In fact, there is no counting problem to be solved. However, the technical challenge is shifted to establishing the communication height in Theorem 3.3. Indeed, this result is very powerful, and the proof relies on an approach similar to that in [13–15]. Once the communication height is established, Theorems 3.1 and 3.5 readily follow.

4. Proof of Theorem 3.2 (Long-Term Average Delay)

In this section, we present the proof of Theorem 3.2. Guided by the intuitive arguments in Section 3.1, we consider an alternating renewal process. As before, an even renewal period starts each time we observe a first entrance into the even dominant state \mathcal{E} after a visit to the odd dominant state \mathcal{O} , and likewise an odd renewal period starts each time the process enters the odd dominant state \mathcal{O} for the first time after a visit to the even dominant state \mathcal{E} . Thus the lengths of the even and odd renewal periods correspond to the transition times $T_{\mathcal{E} \rightarrow \mathcal{O}}$ and $T_{\mathcal{O} \rightarrow \mathcal{E}}$, respectively. We have defined a *cycle* as the period consisting of an even and subsequent odd renewal period. Let $S_{\mathcal{E}}$ and $U_{\mathcal{E}}$ be two random variables representing the amounts of time during an even renewal period that the activity process $\{X(t)\}_{t \geq 0}$ spends in the even dominant state \mathcal{E} and in other activity states, respectively. Since the random variable $S_{\mathcal{E}}$ tracks the total time the process spends in the state \mathcal{E} in an entire cycle, the renewal-reward theorem yields

$$\frac{\mathbb{E}S_{\mathcal{E}}}{\mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}} + \mathbb{E}T_{\mathcal{O} \rightarrow \mathcal{E}}} = \pi(\mathcal{E}) \uparrow \frac{1}{2}, \quad \sigma \rightarrow \infty,$$

where we suppress the implicit dependence on σ for notational compactness. By symmetry, $\mathbb{E}T_{\varepsilon \rightarrow \emptyset} = \mathbb{E}T_{\emptyset \rightarrow \varepsilon}$, so

$$\frac{\mathbb{E}S_{\varepsilon}}{\mathbb{E}T_{\varepsilon \rightarrow \emptyset}} \uparrow 1 \quad \text{and} \quad \frac{\mathbb{E}U_{\varepsilon}}{\mathbb{E}T_{\varepsilon \rightarrow \emptyset}} \downarrow 0, \quad \sigma \rightarrow \infty.$$

In particular, for any $\delta > 0$, there exists $\sigma_{\delta} > 0$ such that, for all $\sigma > \sigma_{\delta}$,

$$\frac{\mathbb{E}U_{\varepsilon}}{\mathbb{E}T_{\varepsilon \rightarrow \emptyset}} \leq \delta. \tag{3}$$

As mentioned earlier, by virtue of the symmetry, we may consider the long-term average delay of an arbitrary packet at a particular node, say an odd one. With minor abuse of notation, let

$$\mathbb{E}L = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T L(t) dt, \tag{4}$$

with $L(t)$ denoting the queue length at that particular node at time t , i.e., the total number of packets, including the packet that may possibly be in the process of being transmitted. Assuming that $\mathbb{E}W$ exists, Little’s law can then be invoked to conclude that $\mathbb{E}L = \lambda \mathbb{E}W$ exists as well, so it suffices to establish a lower bound for $\mathbb{E}L$.

By definition, $L(t)$ increases whenever a packet arrives at the node, as governed by a Poisson process of rate λ . Further observe that $L(t)$ cannot decrease when the activity process resides in the even dominant state ε , since that precludes activity of any odd node. When the activity process does not reside in the even dominant state, the node could potentially be active, and $L(t)$ decreases whenever a transmission is completed and there are packets in the node ($L(t) \neq 0$). Unfortunately, there are intricate dependences that arise between the transmission periods of the node under consideration and the dynamics of the activity process, and in particular the durations of time that it does not reside in the even dominant state. However, the number of transmissions that occur during such a period is a stopping time, and hence the expected number is bounded from above by the amount of time divided by the mean transmission time. Based on the above considerations, we conclude, conditioning on the durations of the even and odd renewal periods, that $\mathbb{E}_{\lambda, \mu} L(t)$ is bounded from below by $Z(t)$, where the subscripts λ and μ indicate expectation with respect to the arrival epochs and transmission periods, and $\{Z(t)\}_{t \geq 0}$ is a process with $Z(0) = L(0)$ that increases at rate λ when the activity process resides in the even dominant state and decreases at rate $\lambda - \mu$ at all other times, as long as $Z(t)$ is positive. Unconditioning, this implies in particular that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \mathbb{E}L(t) dt \geq \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \mathbb{E}Z(t) dt. \tag{5}$$

In order to evaluate the latter limit expression, let T_k and V_k be the durations of the k th even and k th odd renewal periods, respectively, let U_k be the amount of time that the activity process does not reside in the even dominant state during the k th even renewal period, and let $S_K = \sum_{k=1}^K (T_k + V_k)$ be the duration of the first K cycles. Note that T_k and V_k are i.i.d. copies of the random variables $T_{\varepsilon \rightarrow \emptyset}$ and $T_{\emptyset \rightarrow \varepsilon}$, respectively. Assume that an even renewal period starts at time 0 with $Z(0) = 0$, and define the random variable $M = \inf\{K \geq 1 : Z(S_K) = 0\}$. Observe that S_M is a regeneration epoch for the process $\{Z(t)\}_{t \geq 0}$, and hence the renewal-reward theorem implies that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T Z(t) dt = \frac{\mathbb{E} \int_{t=0}^{S_M} Z(t) dt}{\mathbb{E}S_M}. \tag{6}$$

Considering the denominator, it follows from Wald’s equation and symmetry considerations that $\mathbb{E}S_M = \mathbb{E}M(\mathbb{E}T_{\varepsilon \rightarrow \emptyset} + \mathbb{E}T_{\emptyset \rightarrow \varepsilon}) = 2\mathbb{E}M\mathbb{E}T_{\varepsilon \rightarrow \emptyset}$. Turning attention to the numerator, we first condition on the number of cycles M that have elapsed by the regeneration epoch, the durations T_1, T_2, \dots, T_M of the even renewal periods involved, and the amounts of time U_1, U_2, \dots, U_M that the activity process does not reside in the even dominant state.

When $U_k = 0$, inspection of Fig. 1 shows that the area of the triangle associated with the k th cycle is

$$\frac{1}{2} \lambda T_k^2 + \frac{1}{2} \frac{\lambda^2}{\mu - \lambda} T_k^2 = \frac{1}{2} \frac{\lambda \mu}{\mu - \lambda} T_k^2,$$

while the area of the parallelogram associated with the k th cycle is

$$Z_k \left(T_k + \frac{\lambda}{\mu - \lambda} T_k \right) = \frac{\mu}{\mu - \lambda} Z_k T_k,$$

with $Z_k = Z(S_k)$ the value of the process $Z(\cdot)$ at the start of the k th cycle.

In general, when U_k may not be zero, a similar geometric construction leads to the conclusion that the two areas are

$$\frac{1}{2} \lambda T_k \left(T_k - \frac{\mu}{\lambda} U_k \right) + \frac{1}{2} \frac{\lambda^2}{\mu - \lambda} \left(T_k - \frac{\mu}{\lambda} U_k \right)^2 \geq \frac{1}{2} \frac{\lambda \mu}{\mu - \lambda} \left(T_k - \frac{\mu}{\lambda} U_k \right)^2,$$




Fig. 1. Sample path of the process $Z(t)$, with $M = 4$.

and

$$Z_k \left(T_k + \frac{\lambda}{\mu - \lambda} \left(T_k - \frac{\mu}{\lambda} U_k \right) \right) = \frac{\mu}{\mu - \lambda} Z_k (T_k - U_k),$$

respectively. Unconditioning, i.e., taking expectations with respect to the number of cycles M and the random variables T_k and U_k , $k = 1, \dots, M$, we find

$$\mathbb{E} \int_{t=0}^{S_M} Z(t) dt \geq \frac{1}{2} \frac{\lambda \mu}{\mu - \lambda} \mathbb{E} \left(\sum_{k=1}^M \left(T_k - \frac{\mu}{\lambda} U_k \right)^2 \right) + \frac{\mu}{\mu - \lambda} \mathbb{E} \left(\sum_{k=1}^M Z_k (T_k - U_k) \right),$$

yielding

$$\begin{aligned} \frac{\mathbb{E} \int_{t=0}^{S_M} Z(t) dt}{\mathbb{E} S_M} &= \frac{1}{4} \frac{\lambda \mu}{(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \frac{\mathbb{E} \left(\sum_{k=1}^M \left(T_k - \frac{\mu}{\lambda} U_k \right)^2 \right)}{\mathbb{E} M} + \frac{\mu}{2(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \frac{\mathbb{E} \left(\sum_{k=1}^M Z_k (T_k - U_k) \right)}{\mathbb{E} M} \\ &= \frac{1}{4} \frac{\lambda \mu}{(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \mathbb{E} \left(T_{\varepsilon \rightarrow \emptyset} - \frac{\mu}{\lambda} U_{\varepsilon} \right)^2 + \frac{\mu}{2(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \mathbb{E} Z (T_{\varepsilon \rightarrow \emptyset} - U_{\varepsilon}), \end{aligned}$$

where the three random variables $T_{\varepsilon \rightarrow \emptyset}$, U_{ε} , and Z have the joint distribution of the duration of an even renewal period, the amount of time that the activity process does not reside in the even dominant state during that period, and the value of the process $Z(\cdot)$ at the start of that period.

Applying Jensen's inequality and noting that $T_{\varepsilon \rightarrow \emptyset}$ and U_{ε} are independent of Z , we obtain

$$\begin{aligned} \frac{\mathbb{E} \int_{t=0}^{T_M} Z(t) dt}{\mathbb{E} T_M} &\geq \frac{1}{4} \frac{\lambda \mu}{(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \left(\mathbb{E} T_{\varepsilon \rightarrow \emptyset} - \frac{\mu}{\lambda} \mathbb{E} U_{\varepsilon} \right)^2 + \frac{\mu}{2(\mu - \lambda) \mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \mathbb{E} Z (\mathbb{E} T_{\varepsilon \rightarrow \emptyset} - \mathbb{E} U_{\varepsilon}) \\ &\geq \frac{1}{4} \frac{\lambda \mu}{\mu - \lambda} \mathbb{E} T_{\varepsilon \rightarrow \emptyset} \left(1 - \frac{\mu}{\lambda} \frac{\mathbb{E} U_{\varepsilon}}{\mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \right)^2 + \frac{\mu}{2(\mu - \lambda)} \mathbb{E} Z \left(1 - \frac{\mathbb{E} U_{\varepsilon}}{\mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \right). \end{aligned}$$

In view of (3), we find that, for any $\delta > 0$, the liminf of the latter expression (for large values of σ) is bounded from below by

$$\frac{1}{4} \frac{\lambda \mu}{\mu - \lambda} \mathbb{E} T_{\varepsilon \rightarrow \emptyset} \left(1 - \delta \frac{\mu}{\lambda} \right)^2 + \frac{\mu}{2(\mu - \lambda)} \mathbb{E} Z (1 - \delta).$$

Since Z is non-negative and $\delta > 0$ is arbitrary, we deduce that

$$\liminf_{\sigma \rightarrow \infty} \frac{\mathbb{E} L}{\mathbb{E} T_{\varepsilon \rightarrow \emptyset}} \geq \frac{\lambda}{4} \frac{\mu}{\mu - \lambda},$$

where we also used the fact that the expectation signs in (5) can be moved in front of the integrals and then dropped in view of (4) and (6). The statement of the theorem then follows by applying Little's law.

5. Proof of Theorem 3.3 (Communication Height)

In this section, we obtain the communication height Γ . In Section 5.1, we introduce cutsets and contours, which we use in Section 5.2 to distinguish among feasible configurations. We then explore the geometric properties of the set \mathcal{A} in Section 5.3, and we present the proof of Theorem 3.3 in Section 5.4.

5.1. Cutsets and contours

We start by introducing some notation. Given $S \subseteq V$, define $S^c := V \setminus S$, the set $\partial^- S$ of vertices in S which are adjacent to a vertex in S^c , the set $\partial^+ S$ of vertices in S^c which are adjacent to a vertex in S (in particular, for $v \in V$, we write ∂v for $\partial^+ \{v\}$), and the sets $S^+ := S \cup \partial^+ S$, $S^\varepsilon := S \cap \varepsilon$ and $S^\theta := S \cap \theta$. Let $\nabla S := \{(v, w) \in E \mid v \in \partial^- S, w \in \partial^+ S\}$ denote the set of edges having one end in S and the other in S^c .

An independent set $I \in \Omega$ consists of vertices in I^ε and/or vertices in I^θ . In between these regions, there is always an unoccupied two-layer interface, which corresponds to a collection of edge cutsets. Recall that $\gamma \subseteq E$ is an *edge cutset* if the graph $(V, E \setminus \gamma)$ is disconnected, and it is said to be *minimal* if it is inclusion-wise minimal. We describe a way of associating with each independent set $I \in \Omega$ a collection of minimal edge cutsets, following the approach and some ideas of [13–15].

For a given $I \in \Omega$, the vertices in V can be partitioned into two sets, $\mathcal{R}_\theta(I) = I^\theta \cup (\varepsilon \setminus I^\varepsilon)$ and its complement $\mathcal{R}_\varepsilon(I) = I^\varepsilon \cup (\theta \setminus I^\theta)$. The collection $\mathcal{R}_\theta(I)$ ($\mathcal{R}_\varepsilon(I)$) consists of one or more connected components to which we will refer as *odd regions* (*even regions*). From this point onward, our focus will be on the odd regions only and, unless stated otherwise, all regions that will be mentioned will be odd.

For each region $R \subseteq \mathcal{R}_\theta(I)$, we consider the corresponding edge cutset $\gamma_R := \nabla(R)$, which can consist of several disjoint components. A well-known identity relates the size of a specific edge cutset to the number of even and odd vertices present in its interior; see [18,13,14]. This can be generalized and, in the case of a region $R \subseteq \mathcal{R}_\theta(I)$, it reads

$$|\gamma_R| = |\nabla(R)| = 4(|R^\varepsilon| - |R^\theta|). \tag{7}$$

One way to interpret (7) is the following. For a given region $R \subseteq \mathcal{R}_\theta(I)$, $|\gamma_R|$ is equal to 4 times the difference between the number of even inactive nodes in R , i.e., $|R^\varepsilon|$, and the number of odd active nodes in R , i.e., $|R^\theta|$. With this interpretation in mind, the idea behind the proof is easily explained. Indeed, one way to count all the edges of $\nabla(R)$, i.e., the “outgoing” edges of R , is counting first how many edges the subgraph induced by R^+ has and then subtracting the number of its “internal” edges. The total number of edges of R^+ is $4|R^\varepsilon|$, while the number of internal ones is $4|R^\theta|$, since each of them has precisely one endpoint in R^θ .

Consider the dual graph G' of the graph G , which is a discrete torus of the same size. For every region $R \subseteq \mathcal{R}_\theta(I)$, we associate with the edge cutset γ_R the edge set c_R on G' which consists of all the edges of G' orthogonal to edges in γ_R . The set c_R , to which we will refer as the *contour* of R , is a collection of one or more disjoint piecewise linear closed curves. By construction, $|c_R| = |\gamma_R|$, and $|c_R|$ has a natural interpretation as the total length of the contour of R . Define $\mathcal{L}(I) := \bigcup_{R \subseteq \mathcal{R}_\theta(I)} c_R$ to be the collection of contours of all the regions in $\mathcal{R}_\theta(I)$, and $l(I)$ to be their total length, i.e., $l(I) := \sum_{R \subseteq \mathcal{R}_\theta(I)} |c_R|$. The total contour length of a configuration $I \in \Omega$ is related to its efficiency gap $\Delta(I)$ as follows:

$$l(I) = 4 \cdot \Delta(I), \tag{8}$$

which follows from (7) and the definition of $\mathcal{R}_\theta(I)$. Indeed,

$$\begin{aligned} l(I) &= \sum_{R \subseteq \mathcal{R}_\theta(I)} |c_R| = \sum_{R \subseteq \mathcal{R}_\theta(I)} |\gamma_R| = \sum_{R \subseteq \mathcal{R}_\theta(I)} 4(|R^\varepsilon| - |R^\theta|) = 4 \left(\sum_{R \subseteq \mathcal{R}_\theta(I)} |R^\varepsilon| - \sum_{R \subseteq \mathcal{R}_\theta(I)} |R^\theta| \right) \\ &= 4(|\mathcal{R}_\theta(I)^\varepsilon| - |\mathcal{R}_\theta(I)^\theta|) = 4(|\varepsilon \setminus I^\varepsilon| - |I^\theta|) = 4(|\varepsilon| - |I|) = 4 \cdot \Delta(I). \end{aligned}$$

To help the reader, we will present later some examples of independent sets of the graph G , with the corresponding regions and contours. Displaying a configuration $I \in \Omega$, a vertex is drawn as a disk if it belongs to I , and otherwise as a circle. Moreover, we use the following color conventions:

vertex	odd	even
active	●	●
inactive	○	○

With these conventions, all regions in $\mathcal{R}_\theta(I)$ end up being colored in black (and those in $\mathcal{R}_\varepsilon(I)$ in gray).

5.2. Clusters, stripes, and crosses

In this subsection, we present a partition of the state space Ω exploiting geometric properties of the configurations. Let Λ_θ be the graph whose vertex set is θ and such that two odd vertices are connected if and only if they have distance 2 in G , and define Λ_ε analogously. A closed path or curve is said to be *contractible* in a graph H if it can be continuously deformed or contracted in H to a single vertex. Given an independent set $I \in \Omega$, a region $R \subseteq \mathcal{R}_\theta(I)$ is called a *stripe* if there exists a curve in c_R which is non-contractible in G' , a *cross* if all the curves in c_R are contractible in G' and there is a non-contractible closed path in $\Lambda_\theta \cap R^\theta$, and a *cluster* if all the curves in c_R are contractible in G' and all the closed paths are contractible in $\Lambda_\theta \cap R^\theta$. Notice that, if R is a stripe, by parity there must be two non-contractible closed curves in c_R . If instead R is a cross, then one can show that there exist two non-contractible paths in $\Lambda_\theta \cap R^\theta$ which are not homotopic and intersect each other (this motivates the name *cross*). Fig. 2 shows some examples of these types of region.

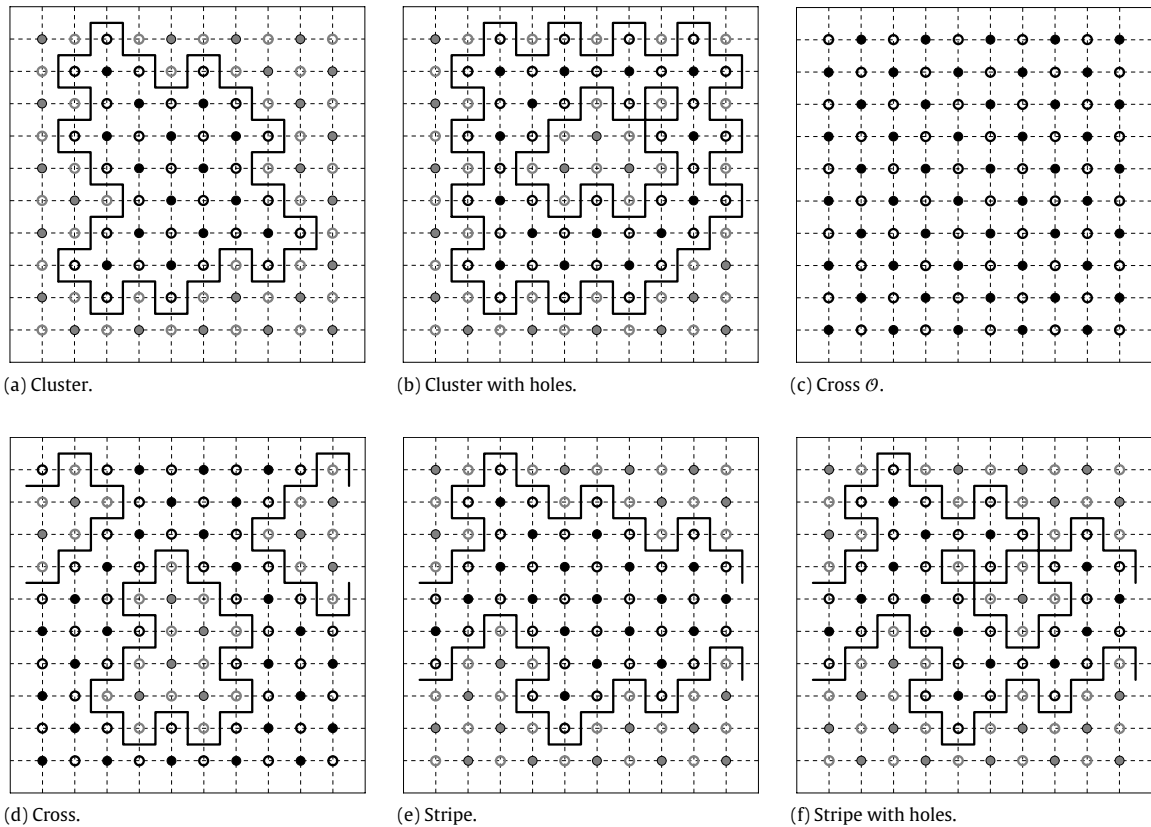


Fig. 2. Examples of configurations in Ω for $L = 10$, with the respective contours.

We partition the state space Ω into the following three subsets: Ω_s , the collection of independent sets which have at least one stripe; Ω_{cr} , the collection of independent sets which have a cross; Ω_{cl} , the collection of independent sets whose regions, if any, are clusters only. Subset Ω_{cl} is disjoint from the other two by definition. Then, to guarantee that $\Omega_s \cup \Omega_{cr} \cup \Omega_{cl}$ is indeed a partition of Ω , we need to prove that there are no configurations $I \in \Omega$ which display simultaneously a cross and a stripe, i.e., $\Omega_s \cap \Omega_{cr} = \emptyset$. This is indeed the case, since the non-contractible curve, which the contour of a stripe displays, is surrounded by an unoccupied two-layer interface that precludes the existence of one the two non-contractible paths on $\Lambda_{\mathcal{E}}$ which a cross must have.

Denote by $A \Delta B$ the symmetric difference between two sets A and B . Consider the collection of independent sets which exhibit a *critical cross*, defined as $\Omega_{cc} := \{I \in \Omega_{cr} \mid \exists I' \in \Omega_{cl} : |I \Delta I'| = 1\}$. In words, a critical cross is a cross which can be “broken” by deactivating just a single node directly into clusters, without becoming a stripe as an intermediate stage.

The next lemma gives some lower bounds on the length of the contour of certain types of region, and it will be crucial in understanding the geometrical properties of the set \mathcal{R} .

Lemma 5.1 (Minimal Contour Length for Stripes and Critical Crosses). *Let $R \subseteq \mathcal{R}_{\emptyset}(I)$ be a region, for some $I \in \Omega$. If R is a stripe, then $|c_R| \geq 4L$. If R is a critical cross, then $|c_R| \geq 8L - 12$.*

This lemma implies that, if I has a stripe, then $\Delta(I) \geq L$, while if I has a critical cross, then $\Delta(I) \geq 2L - 3$.

Proof. (a) Consider a stripe R . By definition, its contour c_R has at least one non-contractible curve, but, by parity, there are two disjoint non-contractible curves. A non-contractible curve has length at least $2L$, because it should be long enough in one direction to wind around the $L \times L$ toric grid. Moreover, every curve of the contour has by construction an equal number of vertical and horizontal edges, which means that both curves have length at least $2L$. Hence $|c_R| \geq 4L$, and the proof is concluded.

(b) Let $I \in \Omega_{cc}$ be an independent set which exhibits a critical cross R , and consider a configuration $I' \in \Omega_{cl}$ such that $|I \Delta I'| = 1$, i.e., $I' \in \partial^+ \Omega_{cc}$. A single cluster cannot evolve in a single step into a cross, so the configuration I' must have at least two clusters, say R_1 and R_2 , from which the critical cross R will be created. Moreover, such clusters must have specific features: R_1 and R_2 must miss the same odd vertex $v \notin I'$ to become (individually) a stripe; see Fig. 3. This means that R_1 and R_2 are clusters “stretched in different directions”. More precisely, one of them, say R_1 , has a projection of length $L - 1$ in the horizontal direction, and the other one, R_2 , a projection of length $L - 1$ in the vertical direction. This fact implies that

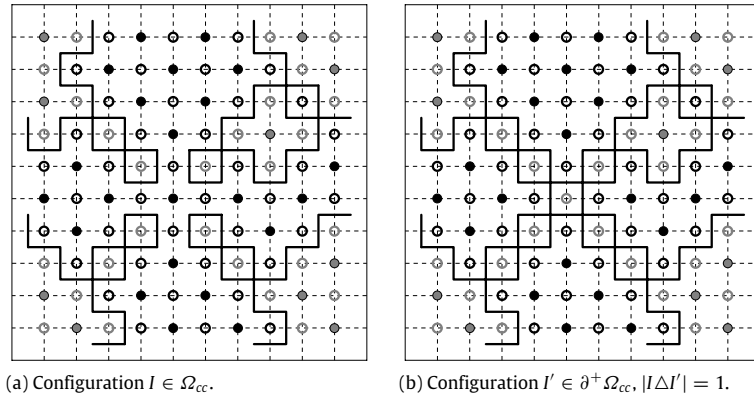


Fig. 3. Structure of configurations with critical crosses.

both R_1 and R_2 have a contour of length at least $4(L - 1)$. Indeed, to have a close curve, at each latitude and longitude there must be two edges, and a contour has the same number of vertical and horizontal edges. Thus $l(I') \geq |c_{R_1}| + |c_{R_2}| \geq 8(L - 1)$ and, by construction, $l(I) = |c_{R_1}| + |c_{R_2}| - 4 \geq 8(L - 1) - 4 = 8L - 12$. \square

5.3. The set \mathcal{S} and its properties

Consider two configurations $I, I' \in \Omega, I \neq I'$. By construction, if the Markov process $\{X(t)\}_{t \geq 0}$ can make the transition $I \rightarrow I'$ in a single step, then $|I \Delta I'| = 1$, and hence $\Delta(I') = \Delta(I) \pm 1$. In particular $\Delta(I') \neq \Delta(I)$.

The bottom $\mathcal{F}(S)$ of a non-empty set $S \subseteq \Omega$ is the set of the largest independent sets in S , i.e., $\mathcal{F}(S) = \{I \in S : \Delta(I) = \min_{I' \in S} \Delta(I')\}$. The independent sets in $\mathcal{F}(S)$ all have the same efficiency gap, which we will denote by $\Delta(\mathcal{F}(S))$. Thanks to (8) and the definition of the communication height, for any integer c ,

$$\phi(I, I') \leq c \Leftrightarrow \Delta(I) \leq 4 \cdot c, \quad \forall I \in \omega : I \rightarrow I'. \tag{9}$$

A non-trivial cycle is a connected set C such that $\max_{I \in C} \Delta(I) < \Delta(\mathcal{F}(\partial^+ C))$. Any singleton that is not a non-trivial cycle is called a trivial cycle. The depth of a cycle C is given by $D(C) := \Delta(\mathcal{F}(\partial^+ C)) - \Delta(\mathcal{F}(C))$.

Recall definition (2) of the set $\mathcal{S} = \{I \in \Omega \mid \phi(\mathcal{S}, I) \leq L\}$. Thanks to (9), the condition $\phi(\mathcal{S}, I) \leq L$ is equivalent to $l(I') \leq 4 \cdot L$ for all configurations $I' \in \omega : \mathcal{S} \rightarrow I$. Moreover, $\mathcal{S} \neq \Omega$, because for instance an independent set J consisting of a single vertex does not belong to \mathcal{S} . Indeed, $\Delta(J) = L^2/2 - 1 > L$ for $L > 2$. This fact and the irreducibility of the process on Ω imply that $\partial^+ \mathcal{S} \neq \emptyset$.

Lemma 5.2 (Stripes and Crosses Do Not Belong To \mathcal{S}). Consider an independent set $I \in \Omega$.

- (a) If I has at least one stripe, then any path $\omega : I \rightarrow \mathcal{S}$ satisfies $\max_{I' \in \omega} \Delta(I') > L$.
- (b) If I has a cross, then any path $\omega : I \rightarrow \mathcal{S}$ satisfies $\max_{I' \in \omega} \Delta(I') > L$.

Proof. (a) Consider a path $\omega : I \rightarrow \mathcal{S}, \omega = (I_1, \dots, I_n), n \in \mathbb{N}, I_1 = I, I_n = \mathcal{S}$. There exists $1 \leq m \leq n - 1$ such that I_m is the first configuration belonging to Ω_{cl} along the path ω . In the configuration I_{m-1} there is either a stripe S or a critical cross C .

Suppose that I_{m-1} has a stripe S . Its contour satisfies $|c_S| \geq 4L$, as proved above in Lemma 5.1. Due to our assumptions, the stripe S must “break” in a single step into clusters. Recall that, in a single transition, we can either add a vertex to or remove a vertex from the current independent set. The action of adding an (even or odd) vertex cannot break the stripe S . Therefore, the only way to break a stripe S is by removing at least one vertex (on its inside), obtaining a new independent set I' which satisfies $\Delta(I') = \Delta(I) + 1 \geq L + 1$.

Suppose instead that I_{m-1} has a cross C . By construction, C must be a critical cross, and then Lemma 5.1 implies that $\Delta(I_{m-1}) \geq 2L - 3$, which is greater than L , since $L \geq 4$.

(b) Consider a path $\omega' : I \rightarrow \mathcal{S}, \omega' = (I_1, \dots, I_{n'}), n' \in \mathbb{N}, I_1 = I, I_{n'} = \mathcal{S}$. There exists $1 \leq m' \leq n' - 1$ such that $I_{m'}$ is the first configuration belonging to Ω_{cl} along the path ω' . In the configuration $I_{m'-1}$ there is either a stripe S or a critical cross C , and then we can argue as in (a). \square

In particular, this latter lemma implies that $\mathcal{S} \subseteq \Omega_{cl}$. We remark that $\mathcal{S} \neq \Omega_{cl}$. Indeed there exist independent sets $I \in \Omega_{cl}$ which have a cluster with a contour larger than $4L$ and therefore do not belong to \mathcal{S} . The next proposition gives two crucial properties of the set \mathcal{S} , following the general method proposed in [16, p. 614].

Proposition 5.3. The set \mathcal{S} satisfies the following two properties.

- (a) \mathcal{S} is a connected set, $\mathcal{S} \in \mathcal{S}$, and $\emptyset \notin \mathcal{S}$.
- (b) There exists a path $\omega^* : \mathcal{S} \rightarrow \mathcal{O}$ such that the maximum of $\Delta(\cdot)$ is reached in $\mathcal{F}(\partial^+ \mathcal{S})$, namely

$$\arg \max_{\omega^*} \Delta \cap \mathcal{F}(\partial^+ \mathcal{S}) \neq \emptyset. \tag{10}$$

Proof. (a) Clearly, $\mathcal{E} \in \mathcal{S}$, and \mathcal{S} is connected by construction. We claim that every path $\omega : \mathcal{E} \rightarrow \mathcal{O}$ satisfies $\max_{I \in \omega} \Delta(I) > L$, which implies that $\mathcal{O} \notin \mathcal{S}$. Indeed, $\mathcal{E} \in \Omega_{cl}$, $\mathcal{O} \in \Omega_{cr} \setminus \Omega_{cc}$, and therefore along any path $\omega : \mathcal{E} \rightarrow \mathcal{O}$ there must be a configuration showing a stripe or a critical cross, and Lemma 5.2 proves our claim.

(b) The reference path ω^* can be constructed as follows. Starting from \mathcal{E} , we gradually create a linear cluster, until it contains $L/2 - 1$ odd active nodes. This configuration, say I , still belongs to \mathcal{S} , since its only cluster has a contour of length $4L - 4$, and hence $\Delta(I) = L - 1$. In order to add the odd vertex v which is missing to create a stripe, we first need to remove sequentially the two even vertices v_1, v_2 in $I \cap \partial v$. Let us call I_1 and I_2 the configurations corresponding to these two intermediate steps and I_3 the configuration where the stripe finally appears, adding the vertex v . By construction, $I_1 \in \mathcal{S}$, with $\Delta(I_1) = L$, while $\Delta(I_2) = L + 1$, and hence $I_2 \in \partial^+ \mathcal{S}$. From I_3 it is possible to reach the configuration \mathcal{O} visiting only configurations $J \in \Omega$ such that $\Delta(J) \leq L + 1$, just gradually expanding the stripe column by column. Therefore $I_2 \in \arg \max_{\omega^*} \Delta \cap \mathcal{F}(\partial^+ \mathcal{S})$. \square

5.4. Proof of Theorem 3.3

Proposition 5.3(a) implies that the process should exit from the set \mathcal{S} to reach the state \mathcal{O} , and so it must visit a configuration $I \in \partial^+ \mathcal{S}$. Therefore $\phi(\mathcal{E}, \mathcal{O}) \geq \Delta(\mathcal{F}(\partial^+ \mathcal{S}))$. Moreover, the existence of a path $\omega^* : \mathcal{E} \rightarrow \mathcal{O}$ such that the maximum of $\Delta(\cdot)$ is reached in $\mathcal{F}(\partial^+ \mathcal{S})$, guaranteed by Proposition 5.3(b), gives $\phi(\mathcal{E}, \mathcal{O}) \leq \max_{I \in \omega^*} \Delta(I)$, and by (10) the two bounds coincide and are equal to $\Delta(\mathcal{F}(\partial^+ \mathcal{S}))$. The definition of \mathcal{S} yields that any $I \in \mathcal{F}(\partial^+ \mathcal{S})$ is such that $\Delta(I) = L + 1$, and the conclusion follows.

6. Proof of Theorem 3.1 (Transition Time Between Dominant States)

In this section, we prove Theorem 3.1, exploiting the communication height Γ , together with a well-known result about exit times from cycles, namely Theorem 6.23 in [17]. In order to do so, we will consider the uniformized discrete-time version of the activity process. In more detail, we construct a discrete-time Markov chain $\{\tilde{X}(t)\}_{t \in \mathbb{N}}$ starting from the continuous-time Markov process $\{X(t)\}_{t \geq 0}$ by means of uniformization at rate $q_{\max} = L^2 \xi$, where $\xi = \max\{\rho\mu, \nu\}$. The transition probabilities of the new Markov chain are as follows:

$$p(I, J) = \begin{cases} \nu/q_{\max}, & \text{if } |I \Delta J| = 1, I \subset J, \\ \rho\mu/q_{\max}, & \text{if } |I \Delta J| = 1, J \subset I, \\ 1 - \sum_{I' \neq I} p(I, I'), & \text{if } I = J, \\ 0, & \text{otherwise.} \end{cases}$$

Since we are interested in the regime $\sigma \rightarrow \infty$, we will assume that $\sigma \geq 1$, so $\xi = \nu$, and the above transition probabilities may be written in the form

$$p(I, J) = \begin{cases} c(I, J)\sigma^{-[H(J)-H(I)]^+}, & \text{if } I \neq J, \\ 1 - \sum_{I' \neq I} p(I, I'), & \text{if } I = J, \end{cases}$$

where $H : \Omega \rightarrow \mathbb{R}$ is the Hamiltonian defined as $H(I) := -|I|$ and c is the connectivity function $c : \Omega^2 \setminus \{(I, I) : I \in \Omega\} \rightarrow [0, 1]$, defined as

$$c(I, J) = \begin{cases} 1/L^2, & \text{if } |I \Delta J| = 1, \\ 0, & \text{otherwise.} \end{cases}$$

The process $\{\tilde{X}(t)\}_{t \in \mathbb{N}}$ is a reversible Markov chain which, taking $\sigma = e^\beta$, satisfies condition M in [17, p. 336], and thus fits in that framework. Moreover, its stationary distribution is the same as that of the original continuous-time process $\{X(t)\}_{t \geq 0}$, see (1), and can be rewritten as

$$\pi(I) = \frac{e^{-\beta H(I)}}{\sum_{J \in \Omega} e^{-\beta H(J)}}, \quad I \in \Omega.$$

Let $(Y_n)_{n \in \mathbb{N}}$ be a sequence of i.i.d. exponential random variables with mean $1/q_{\max}$. If $\tilde{X}(0) = X(0)$ and $n(t) := \sup\{m : \sum_{n=1}^m Y_n \leq t\}$, then $\{X(t)\}_{t \geq 0} \stackrel{d}{=} \{\tilde{X}(n(t))\}_{t \geq 0}$.

For $I \in \Omega$ and $A \subseteq \Omega$, let $\tau_A^I(\sigma) := \inf\{t > 0 : \tilde{X}(t) \in A\}$ be the first hitting time of the set A for the process $\tilde{X}(t)$ starting in I at $t = 0$, and let $T_{I \rightarrow A}(\sigma) := \inf\{t > 0 : X(t) \in A\}$ be its continuous-time counterpart, namely the first hitting time of the set S for the process $X(t)$ starting in I at $t = 0$. These hitting times are closely related:

$$T_{I \rightarrow A} \stackrel{d}{=} \sum_{n=1}^{\tau_A^I} Y_n.$$

In particular,

$$\mathbb{E}T_{I \rightarrow A}(\sigma) = \frac{1}{q_{\max}} \mathbb{E}\tau_A^I(\sigma). \tag{11}$$

Let us now focus on the Markov chain $\{\tilde{X}(t)\}_{t \in \mathbb{N}}$. As shown in Proposition 5.3, this process must exit from the set \mathcal{I} in order to make the transition $\mathcal{E} \rightarrow \mathcal{O}$. In particular, it must hit the set $\partial^+ \mathcal{I}$, the external boundary of \mathcal{I} , before reaching \mathcal{E} . Therefore $\tau_{\mathcal{O}}^{\mathcal{E}} \geq_{st} \tau_{\partial^+ \mathcal{I}}^{\mathcal{E}}$, and hence

$$\mathbb{P}(\tau_{\mathcal{O}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}) \geq \mathbb{P}(\tau_{\partial^+ \mathcal{I}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}). \tag{12}$$

Theorem 6.23 in [17] says that, for any $\varepsilon > 0$, there exist $k(\varepsilon) > 0$ and $\beta_0 > 0$ such that, for all $\beta \geq \beta_0$, and for any $I \in \mathcal{I}$,

$$\mathbb{P}(\tau_{\partial^+ \mathcal{I}}^I < e^{\beta(D(\mathcal{I}) - \varepsilon)}) \leq e^{-k(\varepsilon)\beta}, \tag{13}$$

where $D(\mathcal{I})$ is the depth of the cycle \mathcal{I} , which is equal to Γ by construction. Eqs. (12) and (13) imply that for every $\varepsilon > 0$ there exist $k(\varepsilon) > 0$ and $\beta_0 > 0$ such that, for all $\beta \geq \beta_0$,

$$\mathbb{P}(\tau_{\mathcal{O}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}) \geq 1 - e^{-\beta k(\varepsilon)}.$$

In particular, $\exists \beta_1 > \beta_0$ such that $\mathbb{P}(\tau_{\mathcal{O}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}) \geq 1/2$. Thus

$$\mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}} \geq \mathbb{E}(\tau_{\mathcal{O}}^{\mathcal{E}} \mid \tau_{\mathcal{O}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}) \mathbb{P}(\tau_{\mathcal{O}}^{\mathcal{E}} \geq e^{\beta(\Gamma - \varepsilon)}) \geq \frac{1}{2} e^{\beta(\Gamma - \varepsilon)}.$$

Therefore, for $\beta > \beta_1$,

$$\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}} \geq \beta(\Gamma - \varepsilon) \log 2^{-1} = \beta(\Gamma - \varepsilon) - \log 2,$$

and hence

$$\frac{1}{\beta} \log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}} \geq \Gamma - \varepsilon - \frac{\log 2}{\beta},$$

or, equivalently, by replacing $\beta = \log \sigma$, for σ sufficiently large,

$$\frac{\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}}}{\log \sigma} \geq \Gamma - \varepsilon - \frac{\log 2}{\log \sigma}.$$

Taking the liminf for $\sigma \rightarrow \infty$ on both sides, we get

$$\liminf_{\sigma \rightarrow \infty} \frac{\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}}}{\log \sigma} \geq \Gamma - \varepsilon,$$

and, since ε is arbitrary, it follows that

$$\liminf_{\sigma \rightarrow \infty} \frac{\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}}}{\log \sigma} \geq \Gamma. \tag{14}$$

Using the fact that $q_{\max} = L^2 \nu = L^2 \sigma p \mu$, (11) implies that

$$\log(\mu \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}}) = \log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}} - \log(L^2) - \log \sigma - \log p,$$

and hence

$$\liminf_{\sigma \rightarrow \infty} \frac{\log(\mu \mathbb{E}T_{\mathcal{E} \rightarrow \mathcal{O}})}{\log \sigma} = \liminf_{\sigma \rightarrow \infty} \frac{\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}} - \log(L^2) - \log \sigma - \log p}{\log \sigma} = \liminf_{\sigma \rightarrow \infty} \frac{\log \mathbb{E}\tau_{\mathcal{O}}^{\mathcal{E}}}{\log \sigma} - 1 + \alpha \stackrel{(14)}{\geq} \Gamma - 1 + \alpha,$$

with $\alpha = \liminf_{\sigma \rightarrow \infty} \frac{-\log p}{\log \sigma}$. The conclusion then follows from Theorem 3.3, which states that $\Gamma = L + 1$.

7. Concluding remarks

We have obtained delay bounds for random-access grid networks, which show that the delays grow dramatically with both the load and the dimension of the network. For transparency, we have focused on symmetric grid networks, but the proof techniques and delay bounds are expected to extend to a far broader range of scenarios.

In the present paper, we have assumed the activation rate ν and back-off probability p to be static parameters. A natural question is whether dynamic activation rates and back-off probabilities can potentially achieve better delay performance.

Specifically, various algorithms have recently been proposed where the back-off probability at each node is dynamically adapted over time as a (decreasing) function $p(q)$ of the instantaneous queue length q [19–22]. Remarkably, for suitable choices of the function $p(\cdot)$, such algorithms are guaranteed to provide maximum stability in arbitrary topologies, without any explicit knowledge of the arrival rates. However, lower bounds in [23] demonstrate that for such choices of $p(\cdot)$ the delays are of the order $\exp\left(\frac{1}{2(1-\rho)}\right)$, for the $L \times L$ toric grid growing even faster with the load than the $(1 - \rho)^{-L}$ scaling we obtained.

In the specific case of an $L \times L$ grid, the fluid-limit results in [24] suggest that maximum stability would actually be maintained as long as the function $p(q)$ decays no faster than q^{-2/L^2} . The lower bounds in [23] then indicate that the delays grow as $(1 - \rho)^{-L^2/2}$, which is still faster than the bounds we obtained for fixed back-off probabilities. In order for the lower bounds for queue-based back-off probabilities in [23] to match the lower bounds in the present paper for fixed back-off probabilities, the function $p(q)$ should decay as $q^{-1/L}$, but it is not clear whether maximum stability would remain guaranteed in that case.

Acknowledgments

This work was financially supported by a TOP Grant from The Netherlands Organization for Scientific Research (NWO) and an ERC Starting Grant. The authors gratefully acknowledge helpful discussions with Patrick Thiran, Reint den Toonder, and Peter van de Ven during the early stages of this work.

References

- [1] R.R. Boorstyn, A. Kershenbaum, B. Maglaris, V. Sahin, Throughput analysis in multihop CSMA packet radio networks, *IEEE Transactions on Communications* 35 (1987) 267–274.
- [2] M. Durvy, O. Dousse, P. Thiran, Modeling the 802.11 protocol under different capture and sensing capabilities, in: Proc. IEEE INFOCOM Conf., 2007, pp. 2356–2360.
- [3] M. Durvy, P. Thiran, A packing approach to compare slotted and non-slotted medium access control, in: Proc. IEEE INFOCOM Conf., 2006, pp. 1–12.
- [4] X. Wang, K. Kar, Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks, in: Proc. IEEE INFOCOM Conf., Vol. 1, 2005, pp. 23–34.
- [5] G. Bianchi, Performance analysis of the IEEE 802.11 distributed coordination function, *IEEE Journal on Selected Areas in Communications* 18 (3) (2000) 535–547.
- [6] S.C. Liew, C.H. Kai, J. Leung, B. Wong, Back-of-the-envelope computation of throughput distributions in CSMA wireless networks, *IEEE Transactions on Mobile Computing* 9 (9) (2010) 1319–1331.
- [7] C.H. Kai, S.C. Liew, Temporal starvation in CSMA wireless networks, in: 2011 IEEE International Conference on Communications, 2011, pp. 1–6.
- [8] A. Zocca, S.C. Borst, J.S.H. van Leeuwen, Mixing properties of CSMA networks on partite graphs, in: Proc. ValueTools Conf., 2012, pp. 117–126.
- [9] D. Shah, D.N.C. Tse, J.N. Tsitsiklis, Hardness of low delay network scheduling, *IEEE Transactions on Information Theory* 57 (12) (2011) 7810–7817.
- [10] L. Jiang, M. Leconte, J. Ni, R. Srikant, J. Walrand, Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling, in: Proc. IEEE INFOCOM Conf., 2011, pp. 371–375.
- [11] V. Subramanian, M. Alanyali, Delay performance of CSMA in networks with bounded degree conflict graphs, in: Proc. IEEE ISIT Conf., 2011, pp. 2373–2377.
- [12] J.L. van den Berg, J.E. Steif, Percolation and the hard-core lattice gas model, *Stochastic Processes and their Applications* 49 (2) (1994) 179–197.
- [13] C. Borgs, J.T. Chayes, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, V.H. Vu, Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics, in: Proc. 40th Annual Symposium on Foundations of Computer Science, IEEE, 1999, pp. 218–229.
- [14] D. Galvin, Sampling independent sets in the discrete torus, *Random Structures Algorithms* 33 (3) (2008) 356–376.
- [15] D. Randall, Slow mixing of Glauber dynamics via topological obstructions, in: Proc. 17th ACM–SIAM Symposium on Discrete Algorithm, 2006, pp. 870–879.
- [16] F. Manzo, F.R. Nardi, E. Olivieri, E. Scoppola, On the essential features of metastability: tunnelling time and critical configurations, *Journal of Statistical Physics* 115 (1) (2004) 591–642.
- [17] E. Olivieri, M.E. Vares, *Large Deviations and Metastability*, Cambridge University Press, 2005.
- [18] A. Blanca, D. Galvin, P. Tetali, D. Randall, Phase coexistence and slow mixing for the hard-core model on \mathbb{Z}^2 . Preprint at arxiv.org/abs/1211.6182.
- [19] L. Jiang, D. Shah, J. Shin, J. Walrand, Distributed random access algorithm: scheduling and congestion control, *IEEE Transactions on Information Theory* 56 (12) (2010) 6182–6207.
- [20] S. Rajagopalan, D. Shah, J. Shin, Network adiabatic theorem: an efficient randomized protocol for contention resolution, *Performance Evaluation Review* 37 (2009) 133–144.
- [21] D. Shah, J. Shin, Randomized scheduling algorithm for queueing networks, *The Annals of Applied Probability* 22 (2012) 128–171.
- [22] D. Shah, J. Shin, P. Tetali, Medium access using queues, in: Proc. FOCS Conf., 2011, pp. 698–707.
- [23] N. Bouman, S.C. Borst, J.S.H. van Leeuwen, Delays and mixing times in random-access networks, in: Proc. ACM SIGMETRICS, 2013, pp. 117–128.
- [24] J. Ghaderi, S.C. Borst, P.A. Whiting, Queue-based random-access algorithms: fluid limits and stability issues. Preprint at arxiv.org/abs/1302.5945.



Alessandro Zocca received his Bachelor's degree in mathematics from the University of Padua, Italy, in 2010, and the Master's degree in mathematics from the University of Cambridge, UK, in 2011. He is currently working as a Ph.D. student at Eindhoven University of Technology under the supervision of Sem Borst and Johan van Leeuwen. His interests are in theoretical and applied probability, in particular in stochastic networks and statistical mechanics. His Ph.D. research focuses on the performance analysis of spatial stochastic models, in which notions such as hitting times, conductance, and mixing times play a key role.



Sem Borst received his M.Sc. degree in applied mathematics from the University of Twente, The Netherlands, in 1990, and his Ph.D. degree from the University of Tilburg, The Netherlands, in 1994. During the fall of 1994, he was a visiting scholar at the Statistical Laboratory of the University of Cambridge, UK. In 1995, Sem joined the Mathematics of Networks and Systems research department of Bell Laboratories in Murray Hill, USA. From 1998 to 2006 he was a Senior Researcher at the Center for Mathematics and Computer Science (CWI) in Amsterdam, and since 1998 he also has a (part-time) appointment as professor in Stochastic Operations Research in the Department of Mathematics and Computer Science at Eindhoven University of Technology. Sem's main research areas are performance evaluation and resource allocation for stochastic systems, in particular computer-communication networks. He has published over 130 refereed journal articles and papers in conference proceedings, and holds 16 patents in a variety of areas. He was (co-)recipient of the best-paper awards at Sigmetrics/Performance 1992 and IEEE Infocom 2003, the 2001 Yosef Levy Prize, and the 2005 Van Dantzig Prize. Sem serves or has served on the editorial boards of several journals, such as IEEE/ACM Transactions on Networking, Mathematical Methods of Operations Research, Operations Research Letters, Performance Evaluation, Queueing Systems, and Wireless Networks, and has been a program committee member of numerous conferences.



Johan van Leeuwen (1978; Ph.D. in mathematics 2005, Eindhoven) is a full Professor of Stochastic Networks in the Stochastics section of the Department of Mathematics and Computer Science in Eindhoven University of Technology. His main research interests are in probability theory, stochastic networks, enumerative combinatorics, and complex and asymptotic analysis. Applications that motivate his research are congestion phenomena (queueing), resource allocation in communication networks, distributed control of complex networks, random graphs, and the spread of epidemics over networks. Johan was visiting professor at Columbia University and adjunct professor at NYU. He serves on the editorial boards of three international journals, and is a council member of the INFORMS Applied Probability Society. Johan is the recipient of an ERC Starting Grant (2010), and winner of the 2008 INFORMS Telecommunication Dissertation Award and the 2012 Erlang Prize for his contributions to applied probability.



Francesca Nardi has a tenure track position (since 2006) in the Stochastics section of the Department of Mathematics and Computer Science in Eindhoven University of Technology. She received her M.Sc. degree (1995) and Ph.D. degree (2001) in mathematics from the University of Rome "Tor Vergata", Italy. Francesca had post-doc and teaching positions at Eurandom Eindhoven University of Technology (2000–2003), University of Rome "La Sapienza" (2003–2005), University of Rome "Roma tre" (2005–2006). Her main research interests are in probability theory, statistical mechanics, and mathematical physics. Applications that motivate her research are statistical physics of equilibrium: phase diagrams and phase transitions; statistical physics of non-equilibrium: meta-stability; the spread of epidemics over network for moving populations; random walks in random environment; cut-off phenomena and random graphs. Francesca is the recipient of a STAR grant (2011) and an Aspasia grant (2011) from the Netherlands Organization for Scientific Research (NWO).