




Multicoloured Hardcore Model: Fast Mixing and Its Applications as a Scheduling Algorithm

Sam Olesker-Taylor   

Department of Statistics, University of Warwick, UK

Abstract

In the hardcore model, certain vertices in a graph are *active*: the active vertices must form an *independent set*. We extend this to a *multicoloured* version: instead of simply being active or not, the active vertices are assigned a colour; active vertices *of the same colour* must not be adjacent.

This models a scenario in which two neighbouring resources may *interfere* when active – eg, short-range radio communication. However, there are multiple *channels* (colours) available; they only interfere if both use the *same* channel. Other applications include routing in fibreoptic networks.

We analyse Glauber dynamics. Vertices update their status at random times, at which a uniform colour is proposed: the vertex is assigned that colour if it is available; otherwise, it is set inactive.

We find conditions for *fast mixing* of these dynamics. We also use them to model a queueing system: vertices only serve customers in their queue whilst active. The mixing estimates are applied to establish positive recurrence of the queue lengths, and bound their expectation in equilibrium.

2012 ACM Subject Classification Mathematics of computing → Probabilistic algorithms

Keywords and phrases mixing time, queueing theory, hardcore model, proper colourings, independent set, data transmission, randomised algorithms, routing, scheduling, multihop wireless networks

Digital Object Identifier 10.4230/LIPIcs.AofA.2024.20

Acknowledgements The author would like to thank Frank Kelly for multiple detailed discussions on this topic, as well as reading through the paper. Additionally, thanks go to Perla Sousi and Luca Zanetti, who also read the paper. Their feedback has been invaluable in preparing this manuscript.

1 Introduction and Main Results

We extend the hardcore model, used for sampling independent sets, to a multicoloured version. Given a graph $G = (V, E)$, our objective is to colour a subset $U \subseteq V$ of the vertices such that if $u, u' \in U$ satisfy $\{u, u'\} \in E$, then u and u' are painted with different colours. If there is only one colour, then this condition requires that there is no pair of mutually adjacent vertices. This is the definition of an *independent set*, so we recover the usual *hardcore model*.

We allow an arbitrary number $K \in \mathbb{N}$ of possible colours. If we required *all* vertices to be selected – ie, $U = V$ – then the condition is that no edge in the graph is *monochromatic*: the endpoints must receive different colours. We thus recover the *proper colouring* model. Our model models these two, sampling a properly coloured subset of vertices, or subgraph.

The motivation for this model comes from a desire for a decentralised (and randomised) algorithm for resource sharing. Two examples of this are short-range radio communication, where nearby agents on the same frequency interfere, and routing in fibreoptic networks. Both K and G are given parameters, depending on the particular engineering set-up.

A popular method for sampling proper colourings or independent sets is via *Glauber dynamics*. Our main result is on the *mixing time* of Glauber dynamics for the multicoloured hardcore model, defined precisely below. We then use the system to model a *queueing system*.

- Customers (eg, data packets) arrive to vertices at some (vertex-dependent) rate.
- Coloured vertices are *active*: they serve their customers at some (vertex-dependent) rate.
- Uncoloured vertices are *inactive*: they do not serve, but their queue can still grow.

We apply the mixing-time result to control the queue lengths, under certain conditions.



© Sam Olesker-Taylor;

licensed under Creative Commons License CC-BY 4.0

35th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AofA 2024).

Editors: Cécile Mailler and Sebastian Wild; Article No. 20; pp. 20:1–20:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A Glauber Dynamics

Let $G = (V, E)$ be a graph and $K \in \mathbb{N}$. Let $n := |V|$; write $[K]_0 := \{0, 1, \dots, K\}$. The state space Ω of the system is a subset of configurations $[K]_0^V = \{(\omega_v)_{v \in V} \mid \omega_v \in [K]_0 \forall v \in V\}$.

► **Definition A.1** (State Space). *Let $\Omega := \{\omega \in [K]_0^V \mid \omega \text{ is proper}\}$, where $\omega \in [K]_0^V$ is proper if*

$$\omega_u \neq \omega_v \quad \text{whenever} \quad \{u, v\} \in E \quad \text{and} \quad \omega_u + \omega_v > 0.$$

In a proper configuration, the colour of one vertex must be different to that of all its neighbours, except that colour 0 is exempt from this condition. We view colour 0 as inactive. A configuration is proper if the subgraph induced by its active vertices is properly coloured.

► **Definition A.2** (Glauber-Type Dynamics). *Let $\lambda = (\lambda_v) \in (0, \infty)^V$ and $\mathbf{p} = (p_v) \in [0, 1]^V$. We analyse the following continuous-time Markov chain, which we denote $\text{MCH}_\Omega(\lambda, \mathbf{p})$.*

- *Choose vertex $v \in V$ to update at rate λ_v , simultaneously over all vertices.*
- *Once vertex $v \in V$ is chosen, toss a p_v -biased coin: $C \sim \text{Bern}(p_v)$.*
 - *If $C = 1$, then choose a (non-zero) colour $k \in [K]$ uniformly at random. If colour k is available for v – ie, no neighbour of v has colour k – then paint v with colour k .*
 - *Otherwise, deactivate v – ie, colour 0 – whether or not it was active before.*

Denote the equilibrium distribution by π . The equilibrium active time, or service rate, is

$$s_v := \sum_{\omega \in \Omega: \omega_v \neq 0} \pi(\omega) \quad \text{for} \quad v \in V.$$

The usual Glauber dynamics for proper colourings proposes a colour chosen uniformly amongst available colours. However, this requires whoever is making the colour choice to know which colours are available for that vertex. This is unreasonable in the context of routing algorithms in fiberoptic networks, for example. It is often much faster to check if a single proposed colour is available than to find out which colours are available.

Our main theorem establishes *fast mixing*. First, we define mixing times precisely.

► **Definition A.3** (Mixing Times). *The total-variation distance between distributions μ and π is*

$$\|\mu - \pi\|_{\text{TV}} := \max_{A \subseteq \Omega} |\mu(A) - \pi(A)| = \frac{1}{2} \sum_{\omega \in \Omega} |\mu(\omega) - \pi(\omega)|.$$

The mixing time of a Markov chain $X = (X^t)_{t \geq 0}$ on Ω with invariant distribution π is

$$t_{\text{mix}}(\varepsilon) := \inf\{t \geq 0 \mid \max_{x \in \Omega} \|\mathbb{P}_x[X^t \in \cdot] - \pi\|_{\text{TV}} \leq \varepsilon\} \quad \text{for} \quad \varepsilon \in (0, 1).$$

► **Theorem A** (Fast Mixing). *Suppose that there exists $\beta > 0$ such that*

$$\frac{1}{K} \sum_{u \in V: \{u, v\} \in E} p_u \lambda_u / \lambda_v \leq 1 - \beta \quad \text{for all} \quad v \in V.$$

Let $\lambda_{\min} := \min_{v \in V} \lambda_v$. If $X, Y \sim \text{MCH}_\Omega(\lambda, \mathbf{p})$, then

$$\max_{x, y \in \Omega} \|\mathbb{P}_x[X^t \in \cdot] - \mathbb{P}_y[Y^t \in \cdot]\|_{\text{TV}} \leq \min\{2ne^{-\beta \lambda_{\min} t}, 1\}.$$

In particular,

$$t_{\text{mix}}(\varepsilon) \leq (\beta \lambda_{\min})^{-1} \log(2n/\varepsilon) \quad \text{for all} \quad \varepsilon \in (0, 1).$$

Remark A (Fast-Mixing Condition). The condition in Theorem A arises from requiring the Wasserstein distance between X and Y to contract in a single step, uniformly. Distance is measured vertex-wise: $d(x, y) := \sum_{v \in V} \mathbf{1}\{x_v \neq y_v\}$ for $x, y \in \Omega$. Namely, we prove that if configurations x and y differ only in that vertex v is active in one but not the other, then

$$\frac{d}{dt} \mathbb{E}_{(x,y)}[d(X^t, Y^t)]|_{t=0} \leq \lambda_v \left(\frac{1}{K} \sum_{u \in V: \{u,v\} \in E} p_u \lambda_u / \lambda_v - 1 \right)$$

under some coupling. The given condition ensures this is negative, uniformly in $x, y \in \Omega$. A standard application of *path coupling* [6] extends this uniform contraction to all $x, y \in \Omega$. \triangle

The graph G and number K of colours are given by the application. In contrast, the parameters $(\lambda_v, p_v)_{v \in V}$ may be chosen by the operator. There are good heuristics for taking

$$\lambda_v \propto d_v \quad \text{and} \quad p_v \propto (K/d_v) \wedge 1.$$

In short, high-degree nodes have more impact on their neighbours, and hence should be updated faster: so, take $\lambda_v \propto d_v$. Further, if v is active with probability p'_v , then it remove a total of $p'_v d_v$ colour choices in expectation (from its neighbours). There are K colours, so vertices shouldn't remove more than K in expectation: hence, $s_v d_v \propto K$; so, take $p_v \propto K/d_v$.

We work in continuous time, so scaling all the rates λ inversely scales the mixing time. We choose the normalisation $\sum_{v \in V} \lambda_v = n$; so, vertices each update at rate 1 on average.

► **Corollary A** (Heuristic-Driven Choice). *Suppose that $\lambda_v = d_v/\bar{d}$ and $p_v \leq \frac{2}{3}K/d_v$ for all $v \in V$, where $\bar{d} := \frac{1}{n} \sum_{v \in V} d_v$ is the average degree. Let $X, Y \sim \text{MCH}_\Omega(\lambda, p)$. Then,*

$$\max_{x,y \in \Omega} \left\| \mathbb{P}_x[X^t \in \cdot] - \mathbb{P}_y[Y^t \in \cdot] \right\|_{\text{TV}} \leq \min\{2ne^{-(\delta/\bar{d})t/3}, 1\},$$

where $\delta := \min_{v \in V} d_v$. In particular,

$$t_{\text{mix}}(\varepsilon) \leq 3(\bar{d}/\delta) \log(2n/\varepsilon) \quad \text{for all } \varepsilon \in (0, 1).$$

It is standard, or, at least, very common, in the hardcore-model ($K = 1$) literature to require $p_v = p < 1/\Delta$ for all v , where $\Delta := \max_{v \in V} d_v$ is the maximum degree; see, eg, [2, 15, 7] or [14, Theorem 5.9]. We take more care, requiring only $p_v < K/d_v$ for each v ; [10] have a similar improvement, but restricted to the usual hardcore model ($K = 1$).

A consequence of requiring $p_v = p < 1/\Delta$ is that the mixing time is often proportional to Δ . Ours is proportional to \bar{d}/d_{\min} , which is often significantly smaller.

The bound $p < 1/\Delta$ is natural, up to a factor e . Indeed, for the (usual) hardcore model, it has been known since Kelly [13] that the infinite Δ -regular tree has a critical threshold at $p_c(\Delta) \approx e/\Delta$, for large Δ : the corresponding Gibbs distribution is unique if and only if $p < p_c(\Delta)$. When $p < p_c(\Delta)$, known as the *uniqueness regime* to physicists, the “influence” of one vertex on another decays exponentially in their relative distance. On the other hand, long-range correlations persist when $p > p_c(\Delta)$. See [1, §1.2] for more discussion on this.

Based on this, it appears that we should be able to only require $p_v \leq (1 - \eta)eK/d_v$ and still obtain fast mixing. This would be a natural extension of the critical threshold: $p_c(\Delta, K) := eK/\Delta$. We demonstrate this via some simulations at the end of the paper.

We also investigate the proportion of time that vertices are active in equilibrium.

► **Proposition A** (Equilibrium Service Rates). *Suppose that $p_v \leq \frac{1}{3}K/\tilde{d}_v$ for all $v \in V$, where $\tilde{d}_v := \max\{d_u \mid u \sim v \text{ or } u = v\}$ is the maximal degree in the neighbourhood of $v \in V$. Then,*

$$\frac{1}{3}p_v \leq s_v \leq p_v \quad \text{for all } v \in V.$$

Our proof is quite flexible, allowing more general p_v . We discuss how to generalise it, and tighten the bounds, after its proof. Again, we expect that we only need $p_v \leq (1 - \eta)eK/d_v$.

B Queueing Network

Our next results concerns queue length in a network. The proof relies on fast mixing.

► **Definition B** (Queueing Network). Let $\lambda, \nu, \mu \in (0, \infty)^V$ and $\mathbf{p} \in [0, 1]^V$. Let $X \sim \text{MCH}_\Omega(\lambda, \mathbf{p})$. The state space of the queueing network is \mathbb{N}^V . For $q \in \mathbb{N}^V$ and $v \in V$, let

$$q_u^{v, \pm} := q_u \pm \mathbf{1}\{u = v\} \quad \text{for } u \in V;$$

that is, $q^{v, \pm}$ adds/removes one from the v -th queue. The transition rates given $X = x$ are

$$q \rightarrow \begin{cases} q^{v, +} & \text{at rate } \nu_v \\ q^{v, -} & \text{at rate } \mu_v \mathbf{1}\{x_v \neq 0\} \end{cases} \quad \text{for each } v \in V;$$

that is, the v -th queue has arrivals at rate ν_v always and services at rate μ_v provided v is active. We denote the law of this queueing network by $\text{QMCH}_\Omega(\lambda, \mathbf{p}; \nu)$.

We show that the queues are jointly *positive recurrent* – ie, the expected time until all queues are simultaneously empty is finite – under the fast-mixing conditions of Theorem A and the assumption that the arrival rate ν_v is smaller than the equilibrium service rate s_v .

► **Theorem B** (Stable Queues). Suppose that there exists $\beta > 0$ such that

$$\frac{1}{K} \sum_{u \in V: \{u, v\} \in E} p_u \lambda_u / \lambda_v \leq 1 - \beta \quad \text{for all } v \in V.$$

Suppose also that $\nu_v < s_v$ for all $v \in V$. If $Q \sim \text{QMCH}_\Omega(\lambda, \mathbf{p}; \nu)$, then Q is positive recurrent:

$$\tau := \inf\{t \geq 0 \mid Q^t = 0\} \quad \text{satisfies } \mathbb{E}_q[\tau] < \infty \quad \text{for all } q \in \mathbb{N}^V.$$

Moreover, if Q^0 is in equilibrium, then, writing $\lambda_{\min} := \min_{v \in V} \lambda_v$,

$$\mathbb{E}[Q_v^0] \leq \frac{6n \log(2n/e)}{\beta \lambda_{\min} (s_v - \nu_v)^2} \quad \text{for all } v \in V.$$

We now evaluate this under the heuristic-driven choice from Corollary A.

► **Corollary B** (Heuristic-Driven Choice). Suppose that $\lambda_v = d_v / \bar{d}$ and $p_v \leq \frac{2}{3} K / d_v$ for all $v \in V$, where $\bar{d} := \frac{1}{n} \sum_{v \in V} d_v$ is the average degree. Let $\delta := \min_{v \in V} d_v$. Suppose also that $\nu_v < s_v$ for all $v \in V$. Let $Q \sim \text{QMCH}_\Omega(\lambda, \mathbf{p}; \nu)$. Then, in equilibrium,

$$\mathbb{E}[Q_v^0] \leq \frac{18 \bar{d} n \log(2n/e)}{\delta (s_v - \nu_v)^2} \quad \text{for all } v \in V.$$

A related result was proved by Jiang et al [10] for the usual hardcore model (one colour). Also, they restrict to the special case $p_v = p < 1/\Delta$, where Δ is the maximum degree.

2 Motivation and Related Work

Fibreoptic Routing Application

Our original motivation was to create a *fully decentralised* random access scheme for resource sharing in fibreoptic routing networks. There, nodes are connected by *links*, and they communicate with each other along *routes*, which are sequences of links. Multiple routes may share a subset of links; such routes *interfere*. Each link has a collection of *frequencies* available.

A naive approach has the source node send the data to the first intermediary node on the route, along with instructions of where to send on. That intermediary node processes the data and sends it onto the next node. This continues until the data reaches its target destination.

It is possible for different frequencies to be used along the route, due to the intermediary processing. When checking whether it is possible for a certain collection of routes to be active simultaneously, it is enough to check that no individual link is overloaded. However, the intermediary processing adds overhead. If the time it takes to transmit the data along the link is larger than the processing time, then the overhead is unimportant. However, in fibreoptic networks, data is sent along links extremely quickly, and the processing overhead becomes the performance bottleneck.

Instead of processing and resending the data at an intermediary node, an *optical switch* is configured. This switch is like a prism: light coming from a single source is sent in different directions, depending on its colour. This allows a *light path* to be set up, removing the processing overhead; however, the *same* frequency must be used throughout the entire route.

The difficulty is in choosing the frequency (colour) of the light path. Now, it is not enough to simply check that each link is not overloaded marginally, as the colours are correlated. In the set-up of the multicoloured hardcore model, the vertices correspond to routes, and two routes (vertices) are adjacent, forming an edge, if they interfere – ie, share a link. Certainly, not all routes will be able to be on simultaneously; an access scheme must be devised.

I originally learnt of this model from a talk by Walker [22] at the *Algorithms and Software for Quantum Computers* event at the Isaac Newton Institute. There, the speaker was looking to quantum computation for solutions. I, as a probabilist, took a randomised approach.

The multicoloured hardcore model has the significant benefit of *decentralisation*. All decisions made can be made by the individual vertices, without any need for synchronisation or knowledge of the state of the other routes. A vertex can even request a light path blindly [11]: the path is set up if it does not conflict with any other already-active paths; otherwise, an error is returned to the initiator. Moreover, optical-switch reconfiguration is fast and easy.

The hardcore model is a popular and well-studied model for random access schemes where there is only a single frequency: ON or OFF. A toy model for this is local radio communication: vertices represent pairs of agents who wish to communicate; nearby pairs of agents cannot communicate simultaneously. Quite separately, Glauber dynamics are used to sample proper colourings on a graph. It seems natural to combine these two, yielding a multicoloured hardcore model which can model more complex interference situations, such as when multiple independent radio frequencies are available. However, to the best of my knowledge, this multicoloured hardcore set-up has not been studied before in the context of routing.

Multihop Wireless Networks

Another application of this type of random routing scheme is to *multihop wireless networks*. In cellular and wireless local area networks, wireless communication only occurs on the last link between a base station and the wireless end system. In *multihop* wireless networks, there are one or more intermediate nodes along the path; these receive and forward packets via the wireless links. There are several benefits to the multihop approach, including extended coverage and improved connectivity, higher transfer rates and the avoidance of wide deployment of cables. Unfortunately, protocols, particularly those for routing, developed for fixed or cellular networks, or the Internet, are not optimal for these, more complicated, multihop wireless networks; see, eg, [5].

A highly prominent example of multihop wireless networks is in the development and deployment of 5G cellular networks [21]. Conventional cellular networks employ well-planned deployment of tower-mounted base stations. They are undergoing a fundamental change to deployment of smaller base stations. Multihop relaying can be instrumental for tation. See [8, §4.1] for more details, from which part of this paragraph was paraphrased.

A multihop network with a single transmission frequency falls precisely into the framework of the (usual) hardcore model. Glauber dynamics is a powerful tool used to generated randomised, approximate solutions to combinatorially difficult problems. Moreover, it often has natural decentralised implementation. It has already been used in the past to design and analyse distributed scheduling algorithms for multihop wireless networks; see, particularly, [10, 4], from which this paragraph was paraphrased, as well as [16, 17, 3, 9, 19].

Multihop wireless networks with *multiple* transmission frequencies correspond precisely to our model. To the best of our knowledge, it has received little attention. However, with technological and engineering advances, it may become an important extension in the future.

(A)synchronicity

One aspect to point out is our lack of synchronicity: we use *continuous time*, so sites update one at a time. In practice, engineering implementations often prefer synchronised updates. This is the case in [10], where the (usual) hardcore model is analysed and an *independent set* of vertices – ie, a set of vertices with no edges between them – is updated simultaneously. It is crucial that it is an *independent* set: the changes to one vertex in the set do not affect the other vertices, and the updates can be done independently, in a parallel, distributed manner.

The (independent) set of vertices still needs to be chosen in each step. In [10], the authors simply prescribe a distribution q over the collection of all independent sets; no comment is made on *how* to sample one. In principle, this distribution is very complicated, and perhaps even needs approximating – eg, via Glauber dynamics for the (usual) hardcore model.

The path coupling technique that we use, and is used in [10], is robust to parallel updates, provided one update does not affect the others – as for updating an independent set of vertices. If N is the expected size of the independent set chosen – ie, $N := \mathbb{E}_{S \sim q}[|S|]$ – then the mixing bound behaves as if time is sped up by a factor N . We consider single-site, continuous-time updates for simplicity; but, our analysis extends to the parallel set-up, too.

Spin Systems in Statistical Mechanics

Spin systems are widely studied in statistical mechanics, crossing combinatorics, probability and physics: these involve a graph $G = (V, E)$ and a discrete set \mathcal{K} of *spins*; each vertex $v \in V$ is assigned a spin $k \in \mathcal{K}$. Adjacent vertices interact with each other. A zoo of examples of spin systems is discussed extensively in the very recent paper by Peled and Spinka [18].

- In proper colourings, $\mathcal{K} = \{1, \dots, K\}$ and the constraint is hard: adjacent vertices *must not* have the same colour. The hardcore model is similar with $\mathcal{K} = \{0, 1\}$.
- In the Ising model, $\mathcal{K} = \{\pm 1\}$ and the constraint is soft: vertices prefer to be aligned with their neighbours, with strength controlled by the *inverse temperature* $\beta \geq 0$.

The multicoloured hardcore model is discussed in [18, §3.2.2]. It was originally introduced by Runnels and Lebowitz [20] in the context of lattice gases.

The results of [20, 18] are specialised to \mathbb{Z}^d . The latter is most interested in the case where the dimension d is much larger than the number K of colours. The motivating example for this paper is the fiberoptic routing, for which the lattice \mathbb{Z}^d – particularly in high dimensions – is not an appropriate model. Our results appear to be the first on general graphs.

Notation

We briefly recall some notation which is used throughout the paper.

- The underlying graph is $G = (V, E)$. Let $n := |V|$ denote its number of vertices, and write $u \sim v$ if $\{u, v\} \in E$. The degree of $v \in V$ is $d_v := \sum_{u \in V} \mathbf{1}\{u \sim v\} = |\{u \in V \mid u \sim v\}|$.
- There are $K \in \mathbb{N}$ colours, and we abbreviate $[K]_0 := \{0, 1, \dots, K\}$.
- The update rates and probabilities are $\lambda \in (0, \infty)^V$ and $p \in [0, 1]^V$, respectively.
- The state space is $\Omega := \{\omega \in [K]_0^V \mid \omega \text{ is proper}\}$, where $\omega \in [K]_0^V$ is *proper* if

$$\omega_u \neq \omega_v \quad \text{whenever} \quad \{u, v\} \in E \quad \text{and} \quad \omega_u + \omega_v > 0.$$

- The multicoloured hardcore model is denoted $\text{MCH}_\Omega(\lambda, p)$; its equilibrium distribution π .
- For $\text{QMCH}_\Omega(\lambda, p; \nu)$, the arrival rates are $\nu \in (0, \infty)^V$ and equilibrium service rates

$$s_v := \sum_{\omega \in \Omega: \omega_v \neq 0} \pi(\omega) \quad \text{for} \quad v \in V.$$

3 Proofs of Main Theorems

A Mixing

In this section, we use the classical path coupling argument of Bubley and Dyer [6] to upper bound the mixing time. Throughout, $X, Y \sim \text{MCH}_\Omega(\lambda, p)$, under the “natural” coupling:

- the vertex-update clocks are coupled, so the same vertex is chosen at the same time;
- the subsequent coin toss and colour selection are also coupled.

This coupling is clearly *coalescent*:

$$X^t = Y^t \quad \text{implies} \quad X^s = Y^s \quad \text{for all} \quad s \geq t.$$

Proof of Theorem A. We use path coupling, so must define a path space. We say that $x, y \in [K]_0^V$ are *adjacent* if there is a unique $v \in V$ such that $x_v \neq y_v$ and $0 \in \{x_v, y_v\}$. In other words, our path space is generated by activating an inactive vertex or deactivating an active vertex; changing the colour of an already active vertex *is not* permitted. This space is connected: let $d(x, y)$ denote the distance between two configurations $x, y \in [K]_0^V$; then, $\mathbf{1}\{x \neq y'\} \leq d(x, y) \leq 2n$ for all $x, y \in [K]_0^V$, going via the empty configuration $(0, \dots, 0) \in \Omega$.

For $v \in V$ and $x \in [K]_0^V$, denote the *available colours at v in x* by

$$\mathcal{A}_v(x) := \{1, \dots, K\} \setminus \cup_{u \in V: \{u, v\} \in E} \{x_u\} = \{k \in \{1, \dots, K\} \mid x_u \neq k \forall u \sim v\}.$$

Suppose that $(X^0, Y^0) = (x, y) \in \Omega^2$ with $d(x, y) = 1$; say, $0 = x_v \neq y_v$. Consider the first step of the process from these states. Suppose that vertex $u \in V$ updates.

- Suppose that $u \not\sim v$. Then, $\mathcal{A}_x(u) = \mathcal{A}_y(u)$, since $x_w = y_w$ for all $w \sim u$. Hence, we can perform the same update in both X and Y . The relative distance is unchanged, unless $u = v$, in which case the two coalesce.
- Suppose that $u \sim v$; in particular, $u \neq v$. We may not have $\mathcal{A}_u(x) = \mathcal{A}_u(y)$, but always

$$\mathcal{A}_u(x) \cup \{x_u\} = \mathcal{A}_u(y) \cup \{y_u\}.$$

Hence, $|\mathcal{A}_u(x) \triangle \mathcal{A}_u(y)| \leq 1$. So, the probability that a proposed colour is valid for one and not the other is at most $1/K$. If this is the case, then the relative distance increases by 1; otherwise, it remains unchanged. The probability *some* colour is proposed is p_u .

20:8 Multicoloured Hardcore Model: Fast Mixing and Application

It is in this last step that the assumption $0 \in \{x_v, y_v\}$ is used: without it, the symmetric difference could be of size 2, giving a probability $2/K$. Summing over $u \in V$, the relative distance increases by 1 at rate at most $\frac{1}{K} \sum_{u: u \sim v} p_u \lambda_u$ and decreases by 1 at rate λ_v . Hence,

$$\frac{d}{dt} \mathbb{E}_{x,y}[d(X^t, Y^t)]|_{t=0} \leq \lambda_v \left(\frac{1}{K} \sum_{u: u \sim v} p_u \lambda_u / \lambda_v - 1 \right) \leq -\beta \lambda_v,$$

with the last inequality using the (main) assumption of the theorem. This can be extended to general $x, y \in [K]_0^V$ – ie, not requiring $d(x, y) = 1$ – by looking at contraction along geodesics, in the usual manner for path coupling. Hence, recalling that $\lambda_{\min} = \min_v \lambda_v$,

$$\max_{x,y \in [K]_0^V} \frac{d}{dt} \mathbb{E}_{x,y}[d(X^t, Y^t)]|_{t=0} \leq -\beta \lambda_{\min}.$$

By the Grönwall inequality, integrating this and using $\mathbf{1}\{x \neq y\} \leq d(x, y) \leq 2n$, we obtain

$$\max_{x,y \in [K]_0^V} \mathbb{P}_{x,y}[X^t \neq Y^t] \leq \max_{x,y \in [K]_0^V} \mathbb{E}_{x,y}[d(X^t, Y^t)] \leq 2ne^{-\beta t}.$$

Finally, the coupling representation of total-variation distance implies that

$$\max_{x,y \in \Omega} \|\mathbb{P}_x[X^t \in \cdot] - \mathbb{P}_y[Y^t \in \cdot]\|_{\text{TV}} \leq \min\{2ne^{-\beta \lambda_{\min} t}, 1\}. \quad \blacktriangleleft$$

Remark. If preferred, instead of using a continuous-time version of path coupling, discretise time: let $\tilde{X}^\ell := X^{\delta\ell}$ and $\tilde{Y}^\ell := Y^{\delta\ell}$, where δ is some very small real number. Then,

$$\mathbb{E}_{x,y}[d(\tilde{X}^1, \tilde{Y}^1)] \leq (1 - \beta \lambda_{\min} \delta + o(\delta))d(x, y) \quad \text{uniformly,}$$

using the fact that the diameter is finite to obtain a uniform $o(\delta)$ term. Path coupling gives

$$\mathbb{E}_{x,y}[d(\tilde{X}^\ell, \tilde{Y}^\ell)] \leq 2n(1 - \beta \lambda_{\min} \delta + o(\delta))^\ell \leq 2ne^{-\beta \lambda_{\min} \delta \ell + o(\delta \ell)} n.$$

Given $t \geq 0$, let $\ell := \lfloor t/\delta \rfloor \geq t/\delta - 1$. Then,

$$\mathbb{E}_{x,y}[d(X^t, Y^t)] \leq \mathbb{E}_{x,y}[d(\tilde{X}^\ell, \tilde{Y}^\ell)] \leq 2ne^{-\beta t + o(1)}.$$

Finally, taking $\delta \downarrow 0$, we deduce the same bound as before.

We close this section with a discussion of the equilibrium service rates. Here, we assume

$$p_v \leq \frac{1}{3}K/\tilde{d}_v \quad \text{where} \quad \tilde{d}_v := \max\{d_u \mid u \sim v \text{ or } u = v\} \quad \text{for } v \in V.$$

Proof of Proposition A. The quantity we estimate is the proportion of colours available at a vertex. This allows estimation of the probability an attempted colouring is successful.

Clearly, in equilibrium, each neighbour u of v is active with probability at most $p_u = \frac{1}{3}K/\tilde{d}_u$; in particular, $s_v \leq p_v$. Hence, if N_v is the number of colours available at v , then

$$N_v \lesssim \text{Bin}(d_v, \frac{1}{3}K/d_v) \quad \text{in equilibrium.}$$

It can be shown that $\mathbb{P}[\text{Bin}(d, \frac{1}{3}k/d) \geq \frac{1}{2}k] \leq \frac{1}{3}$ whenever $k \leq 3d$. This implies that

$$\mathbb{P}[N_v \geq \frac{1}{2}K] \leq \frac{1}{3}.$$

Hence, upon refreshing, at least $\frac{1}{2}$ of the colours are available with probability at least $\frac{2}{3}$. So, the probability that the proposed colour is accepted is at least $\frac{1}{3}$. Thus, $s_v \geq \frac{1}{3}p_v$. \blacktriangleleft

We discuss briefly extensions of this proof, including heuristics for an upper bound on s_v .

Remark. If we require $p_v \leq (1 - \delta)K/\tilde{d}_v$, then the above argument says that at least a proportion δ of the colours are free in expectation. If K (and \tilde{d}_v) are large, then the Binomial concentrates. There is then a probability δ that a uniformly proposed colour is available.

We can extend this, heuristically at least. If $u, u' \sim v$, then the colours at u and u' should be approximately independent if K is large and the graph has few triangles. If $k_1, \dots, k_K \sim^{\text{iid}} \text{Unif}([K])$, then $\frac{1}{K}|\{k_1, \dots, k_K\}| \approx 1/e$, suggesting that, in fact, a proportion $1/e$ are available after K choices. This would suggest $s_v \geq p_v/e$.

We can also try to iterate this argument. Instead of upper bounding the expected number of colours taken by $\sum_{u:u \sim v} p_u$, we can bound by $\sum_{u:u \sim v} s_u$. Suppose that s_v does not vary much over the vertices: $s_v \approx \bar{s} := \frac{1}{n} \sum_u s_u$, the average of s ; see, eg, Figure 2 later. Also, assume graph regularity: $d_v = d$, and $p_v = p$, for all v . Then, $\sum_{u:u \sim v} s_u \approx d\bar{s}$. This imposes

$$\bar{s} \leq p(1 - d\bar{s}/K); \quad \text{ie,} \quad \bar{s} \leq p/(1 + pd/K).$$

Including the factor $1/e$ from the previous heuristic improves this to $\bar{s} \approx p/(1 + e^{-1}pd/K)$. \triangle

B Queues

Next, we investigate the *stability* of the queueing network: ie, its positive recurrence (or lack thereof) and expected queue length in equilibrium. The end goal is Theorem B. Similar properties for a related model are established in [10, §V], using the usual *Lyapunov function*

$$L^t := \sum_{v \in V} (Q_v^t)^2 \quad \text{for } t \geq 0 \quad \text{where } Q = (Q^t)_{t \geq 0} \sim \text{QMCH}_\Omega(\lambda, \mathbf{p}; \nu).$$

There, the model is slightly simpler, with unit service times, rather than Exponentials. Moreover, they require $p_v = p \leq 1/\Delta$ for all $v \in V$, where $\Delta := \max_v d_v$ is the maximum degree of the graph $G = (V, E)$, and treat Δ as a constant, which is absorbed into a final, unquantified constant. For a sequence $(G_n)_{n \in \mathbb{N}}$ of graphs, this implicitly assumes bounded degrees: $\sup_{n \in \mathbb{N}} \Delta_n < \infty$. We allow much greater generality, both in G and in \mathbf{p} .

We denote by τ the first time the queue is empty:

$$\tau := \inf\{t \geq 0 \mid Q^t = 0, \cup_{s \in [0, t]} Q^s \neq \{0\}\}.$$

Positive recurrence is equivalent to having $\mathbb{E}_q[\tau] < \infty$ for some, and hence all, $q \neq 0$.

Proof of Theorem B. We establish negative drift for an appropriate Lyapunov function L :

$$L^t := \frac{1}{2} \sum_{v \in V} (Q_v^t)^2 \quad \text{for } t \geq 0. \tag{1}$$

We fix some notation and conventions. By the memoryless property of the service times, we may assume that the vertices are *always* providing service, but that a service attempt is rejected if the vertex is inactive at the time of the attempt. Then, the arrivals and attempted services form Poisson processes, independent of each other and the underlying MCH process.

Fix $v \in V$ and $t, T \geq 0$. Write $\hat{S}_v[T, T+t)$ for the number of attempted services by vertex v between times T and $T+t$, and write $\hat{s}_v := \hat{s}_v[T, T+t) := \frac{1}{t} \hat{S}_v[T, T+t)$ for the average (attempted) service rate in this interval. Similarly, write $\hat{A}_v[T, T+t)$ and $\hat{a}_v := \hat{a}_v[T, T+t) := \frac{1}{t} \hat{A}_v[T, T+t)$ for the number of arrivals and average service rate, respectively, between T and $T+t$.

Using these definitions, we have the following simple inequality:

$$Q_v^{T+t} \leq [Q_v^T - \hat{S}_v[T, T+t)]_+ + \hat{A}_v[T, T+t) = [Q_v^T - t\hat{s}_v]_+ + t\hat{a}_v,$$

20:10 Multicoloured Hardcore Model: Fast Mixing and Application

where $[\alpha]_+ := \max\{\alpha, 0\}$ for $\alpha \in \mathbb{R}$. Hence, using $[Q_v^t - t\hat{s}_v]_+ \leq Q_v^t$,

$$\begin{aligned} (Q_v^{T+t})^2 &\leq (Q_v^T - t\hat{s}_v)^2 + 2t[Q_v^T - t\hat{s}_v]_+ \hat{a}_v + t^2 \hat{a}_v^2 \\ &\leq (Q_v^T)^2 + 2tQ_v^T(\hat{a}_v - \hat{s}_v) + t^2(\hat{a}_v^2 + \hat{s}_v^2). \end{aligned} \quad (2)$$

Plugging this into the definition (1) of L bounds its random increment:

$$L^{T+t} - L^T \leq t \sum_{v \in V} Q_v^T(\hat{a}_v - \hat{s}_v) + \frac{1}{2}t^2 \sum_{v \in V}(\hat{a}_v^2 + \hat{s}_v^2). \quad (3)$$

Now, if $\hat{\tau}_v$ is the proportion of time during $[T, T+t]$ that vertex v is active, then

$$t\hat{a}_v = \hat{A}_v[T, T+t] \sim \text{Pois}(t\nu_v) \quad \text{and} \quad t\hat{s}_v = \hat{S}_v[T, T+t] \sim \text{Pois}(t\hat{\tau}_v).$$

To emphasise, the implicit Poisson variables are independent of the MCH process. Recall that

$$\text{if } P \sim \text{Pois}(\mu), \quad \text{then } \mathbb{E}[P] = \mu \quad \text{and} \quad \mathbb{E}[P^2] = \mu + \mu^2.$$

Now, $\nu_v < s_v$, by assumption, and $s_v \leq p_v \leq 1$; also, $\hat{\tau}_v \leq 1$. Hence,

$$\mathbb{E}[\hat{a}_v] = \nu_v, \quad \mathbb{E}[\hat{a}_v^2] \leq 2, \quad \mathbb{E}[\hat{s}_v] \leq 1 \quad \text{and} \quad \mathbb{E}[\hat{s}_v^2] \leq 2.$$

Plugging these into (3) bounds the (expected) drift:

$$\mathbb{E}[L^{T+t} - L^T \mid (X^T, Q^T)] \leq t \sum_{v \in V} Q_v^T(\nu_v - \mathbb{E}[\hat{s}_v \mid X^T]) + \frac{3}{2}nt^2; \quad (4)$$

the (attempted) service rate $\hat{s}_v[T, T+t]$ depends only on X^T , not Q^T .

It remains to handle $\mathbb{E}[\hat{s}_v \mid X^T]$. The attempted services are a thinned Poisson process. So,

$$\mathbb{E}[\hat{s}_v \mid X^T] = \mathbb{E}[\hat{\tau}_v \mid X^T] \quad \text{and} \quad \tau_v = \frac{1}{t} \sum_T^{T+t} \mathbf{1}\{X_v^s \neq 0\} ds.$$

So, if we write $\mu_{x,s}$ for the law of X^s given $X^0 = x$, then

$$\mathbb{E}[\hat{s}_v \mid X^T] = \frac{1}{t} \int_T^{T+t} \mathbb{P}[X^s \neq 0 \mid X^T] ds \frac{1}{t} \int_0^t \mu_{X^T, s}(\{\omega \in \Omega \mid \omega_v \neq 0\}) ds.$$

This is very similar to the equilibrium (attempted) service rate

$$s_v = \sum_{\omega \in \Omega: \omega_v \neq 0} \pi(\omega) = \pi(\{\omega \in \Omega \mid \omega_v \neq 0\});$$

in fact, by the ergodic theorem, $\hat{s}_v[T, T+t] \rightarrow s_v$ as $t \rightarrow \infty$. Quantitatively,

$$\begin{aligned} |\mathbb{E}[\hat{s}_v \mid X^T] - s_v| &= \left| \frac{1}{t} \int_0^t \mu_{X^T, s}(\{\omega \in \Omega \mid \omega_v \neq 0\}) ds - \pi(\{\omega \in \Omega \mid \omega_v \neq 0\}) \right| \\ &\leq \frac{1}{t} \int_0^t |\mu_{X^T, s}(\{\omega \in \Omega \mid \omega_v \neq 0\}) - \pi(\{\omega \in \Omega \mid \omega_v \neq 0\})| \\ &\leq \frac{1}{t} \int_0^t \|\mu_{X^T, s} - \pi\|_{\text{TV}} ds. \end{aligned}$$

It is here that we apply the mixing result, Theorem A: for any $x \in \Omega$ and $s \geq 0$,

$$\|\mu_{x,s} - \pi\|_{\text{TV}} \leq \min\{2ne^{-\beta\lambda_{\min}s}, 1\};$$

note that the first hypothesis of Theorem B is precisely that required for Theorem A. Then,

$$\begin{aligned} \int_0^t \|\mu_{X^T, s} - \pi\|_{\text{TV}} ds &\leq t_0 + n \int_{t_0}^{t \vee t_0} e^{-\beta\lambda_{\min}s} ds \\ &\leq t_0 + (\beta\lambda_{\min})^{-1} =: t_1 \quad \text{where} \quad t_0 := (\beta\lambda_{\min})^{-1} \log(2n). \end{aligned}$$

In particular, this is independent of t , so vanishes once divided by t and $t \rightarrow \infty$:

$$|\mathbb{E}[\hat{s}_v | X^T] - s_v| \leq t_1/t \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

We want to plug this bound into (4). Let $\varepsilon_v := \frac{1}{2}(s_v - \nu_v)$ and $t_v := t_1/\varepsilon_v$. Then,

$$|\mathbb{E}[\hat{s}_v | X^T] - s_v| \leq \varepsilon_v \quad \text{whenever } t \geq t_v \quad \text{for all } v \in V.$$

Set $t_\star := \max_v t_v$, so $t_\star \geq t_v$. Plugging this into (4),

$$\begin{aligned} \mathbb{E}[L^{T+t} - L^T | (X^T, Q^T)] &\leq -t \sum_{v \in V} Q_v^T (s_v - \nu_v - \varepsilon_v) + \frac{3}{2}nt^2 \\ &\leq -\frac{1}{2}t \sum_{v \in V} Q_v^T (s_v - \nu_v) + \frac{3}{2}nt^2 \quad \text{whenever } t \geq t_\star. \end{aligned} \quad (5)$$

This expression is negative for large enough $\|Q^T\|$. This establishes negative drift of L . Hence, by the Foster–Lyapunov criterion (eg, [12, Proposition D.1]), $(Q^t)_{t \geq 0}$ is positive recurrent.

It remains to control the expected queue length in equilibrium. We start in equilibrium and take the expectation of the increment $(Q_v^{t_v})^2 - (Q_v^0)^2$. By stationarity and (2),

$$0 = \mathbb{E}[(Q_v^{t_v})^2 - (Q_v^0)^2] \leq -t_v \mathbb{E}[Q_v^0] (s_v - \nu_v - \varepsilon_v) + \frac{3}{2}nt_v^2,$$

using the same manipulations as before. Rearranging,

$$\mathbb{E}[Q_v^0] \leq \frac{3}{2}nt_v / (s_v - \nu_v - \varepsilon_v) \leq 6nt_1 / (s_v - \nu_v)^2.$$

Finally, $t_1 = (\beta\lambda_{\min})^{-1}(\log(2n) + 1) = (\beta\lambda_{\min})^{-1} \log(2n/e)$. ◀

4 Simulations: Queue Lengths and Equilibrium Service Rate

We close the paper with a short discussion of some simulations. Specifically, we investigate the queue lengths and the proportion of time that a vertex is active as a rolling average – namely,

$$\hat{Q}_v^t := \frac{1}{t} \sum_{s=0}^{t-1} Q_v^s \quad \text{and} \quad \hat{s}_v^t := \frac{1}{t} \sum_{s=0}^{t-1} \mathbf{1}\{X_v^s \neq 0\} \quad \text{for } t \geq 0.$$

Then, $\hat{Q}_v^t \rightarrow \mathbb{E}_\pi[Q_v^0]$ and $\hat{s}_v^t \rightarrow s_v$, the expected equilibrium queue length and service rate.


Our choice of parameters is driven by the same heuristics as for Corollaries A and B:

$$\lambda_v := d_v/\bar{d}, \quad p_v := \min\{\frac{4}{5}eK/d_v, \frac{3}{4}\} \quad \text{and} \quad \nu_v := \frac{1}{3}p_v \quad \text{for } v \in V.$$


Notice the prefactor in p_v : it is $\frac{4}{5}e > 2$, rather than $\frac{1}{3}$ or $\frac{2}{3}$. This is to emphasise the fact that we really can take p_v close to eK/d_v , yet still get high, and stable, service rates s_v .

Figure 1 show the time-averaged queue lengths and service rates when the underlying graph is an Erdős–Rényi graph. Figure 2 show the same for a random regular graph. The average degree is 30 and $K = 10$ colours are used; so, almost all vertices satisfy $p_v = \frac{4}{5}eK/d_v \approx 0.5$. A collection of 10 vertices with typical degrees to be displayed are chosen randomly. Time is scaled so that the average vertex update-rate is 1 – ie, scaled by $\frac{1}{n} \sum_v (\lambda_v + \nu_v + 1)$.

We see that the empirical service rates settle down really quite quickly, and appear to be remain stable. Moreover, the values s_v to which they converge appear to be on the same order as the proposal probabilities p_v . This suggests many proposals are accepted, but not *too many*: if $s_v \approx p_v$, then perhaps a higher proposal probability p_v could have been used. In particular, we found that the normalised difference $|s_v - p_v|/p_v$ averaged around 60%.



■ **Figure 1** The underlying graph is Erdős–Rényi with $n = 500$ vertices and edge probability $40/n$.



■ **Figure 2** The underlying graph is drawn uniformly over 40-regular graphs on $n = 500$ vertices.

The queue lengths, on the other hand, fluctuate a more. They are a bit more stable in the random regular graph (Figure 2) compared with the Erdős–Rényi graph (Figure 1), perhaps due to inhomogeneities. It is not even completely clear what they are converging to.

We suggest that this is likely caused by the inhomogeneities in the graph along with the fact that we take $\nu_v = p_v/3 \approx 0.33p_v$, which is pretty close to $s_v \approx 0.4p_v$. Indeed, the same calculations (not shown) with $\nu_v = 0.2p_v$ result in much more stable queues.

The primary objective is to get as large an equilibrium service rate s_v as possible, or at least its average $\bar{s} = \frac{1}{n} \sum_v s_v$. Since the 60% above is still quite a large rejection rate, we also tested a slightly smaller value of p_v : namely, we used $p_v = \frac{2}{3}eK/d_v \approx 0.45$. However, we found that \bar{s} was about 10% smaller for these parameters, for both random graph models.

A random d -regular graph locally looks like a d -regular tree, so it is not reasonable to expect better than $eK/d = p_c(\Delta, K)$, the earlier critical threshold. Similarly, a sparse Erdős–Rényi graph locally looks like a Bienaymé–Galton–Watson tree with $\text{Pois}(\bar{d})$ degrees.

References

- 1 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science*, pages 1319–1330. IEEE Computer Soc., Los Alamitos, CA, 2020. doi:10.1109/FOCS46700.2020.00125.
- 2 Nayantara Bhatnagar, Allan Sly, and Prasad Tetali. Decay of correlations for the hardcore model on the d -regular random graph. *Electron. J. Probab.*, 21:Paper No. 9, 42 pp., 2016. doi:10.1214/16-EJP3552.

- 3 R. Boorstyn, A. Kershbaum, B. Maglaris, and V. Sahin. Throughput analysis in multihop csma packet radio networks. *IEEE Transactions on Communications*, 35(3):267–274, March 1987. doi:10.1109/TCOM.1987.1096769.
- 4 Torsten Braun, Andreas Kasser, Maria Kihl, Veselin Rakocevic, Vasilios Siris, and Geert Heijenk. Multihop wireless networks. In Vasilios Siris, Torsten Braun, Francisco Barcelo-Arroyo, Dirk Staehle, Giovanni Giambene, and Yevgeni Koucheryavy, editors, *Traffic and QoS Management in Wireless Multimedia Networks: COST 290 Final Report*, pages 201–265. Springer US, Boston, MA, 2009.
- 5 Raffaele Bruno and Maddalena Nurchis. Survey on diversity-based routing in wireless mesh networks: Challenges and solutions. *Computer Communications*, 33(3):269–282, February 2010. doi:10.1016/j.comcom.2009.09.003.
- 6 Ross Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, FOCS '97, pages 223–, Washington, DC, USA, 1997. IEEE Computer Society. doi:10.1109/SFCS.1997.646111.
- 7 Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures Algorithms*, 45(1):78–110, 2014. doi:10.1002/rsa.20479.
- 8 Song-Nam Hong and Ivana Marić. Multihop wireless backhaul for 5g. In Ivana Marić, Shlomo Shamai (Shitz), and Osvaldo Simeone, editors, *Information Theoretic Perspectives on 5G Systems and Beyond*, pages 131–165. Cambridge University Press, Cambridge, 2022. doi:10.1017/9781108241267.004.
- 9 L. Jiang and J. Walrand. A distributed csma algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking*, 18(3):960–972, June 2010. doi:10.1109/TNET.2009.2035046.
- 10 Libin Jiang, Mathieu Leconte, Jian Ni, Rayadurgam Srikant, and Jean Walrand. Fast mixing of parallel glauher dynamics and low-delay csma scheduling. *IEEE Transactions on Information Theory*, 58(10):6541–6555, October 2012. doi:10.1109/TIT.2012.2204032.
- 11 Frank Kelly. Private communication, 2018.
- 12 Frank Kelly and Elena Yudovina. *Stochastic Networks*, volume 2 of *Institute of Mathematical Statistics Textbooks*. Cambridge University Press, Cambridge, 2014. doi:10.1017/CB09781139565363.
- 13 Frank P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society*, 47(3):379–395, 415–428, 1985. arXiv:2345773.
- 14 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI, USA, second edition, 2017. doi:10.1090/mbk/107.
- 15 Pinyan Lu, Kuan Yang, and Chihao Zhang. Fptas for hardcore and ising models on hypergraphs. In *33rd Symposium on Theoretical Aspects of Computer Science*, volume 47 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 51, 14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.
- 16 J. Ni and R. Srikant. Distributed csma/ca algorithms for achieving maximum throughput in wireless networks. In *2009 Information Theory and Applications Workshop*, pages 250–250, 2009. doi:10.1109/ITA.2009.5044953.
- 17 Jian Ni, Bo Tan, and R. Srikant. Q-csma: Queue-length based csma/ca algorithms for achieving maximum throughput and low delay in wireless networks. In *Proceedings of the 29th Conference on Information Communications*, INFOCOM'10, pages 271–275, Piscataway, NJ, USA, 2010. IEEE Press.
- 18 Ron Peled and Yinon Spinka. Long-range order in discrete spin systems. *arXiv:2010.03177 [math-ph]*, January 2024. doi:10.48550/arXiv.2010.03177.

20:14 Multicoloured Hardcore Model: Fast Mixing and Application

- 19 Shreevatsa Rajagopalan, Devavrat Shah, and Jinwoo Shin. Network adiabatic theorem: An efficient randomized protocol for contention resolution. *SIGMETRICS Perform. Eval. Rev.*, 37(1):133–144, June 2009. doi:10.1145/2492101.1555365.
- 20 L. K. Runnels and J. L. Lebowitz. Analyticity of a hard-core multicomponent lattice gas. *Journal of Statistical Physics*, 14(6):525–533, 1976. doi:10.1007/BF01012851.
- 21 Pol Torres Compta, Frank H. P. Fitzek, and Daniel E. Lucani. Network coding is the 5g key enabling technology: Effects and strategies to manage heterogeneous packet lengths. *Transactions on Emerging Telecommunications Technologies*, 26(1):46–55, January 2015. doi:10.1002/ett.2899.
- 22 Nigel Walker. Design challenges in BT’s network. www.newton.ac.uk/seminar/23532, 2018.